

Hypernuclei with neural network quantum states

Andrea Di Donna^{1,2}, Lorenzo Contessi³, Alessandro Lovato^{2,4,5,*} and Francesco Pederiva^{1,2}

¹Dipartimento di Fisica, *University of Trento*, Via Sommarive 14, I-38123 Povo, Trento, Italy

²INFN-TIFPA Trento Institute of Fundamental Physics and Applications, Via Sommarive 14, I-38123 Povo, Trento, Italy

³Universite Paris-Saclay, CNRS-IN2P3, IJCLab, 91405 Orsay, France

⁴Physics Division, *Argonne National Laboratory*, Argonne, Illinois 60439, USA

⁵Computational Science Division, *Argonne National Laboratory*, Argonne, Illinois 60439, USA



(Received 7 August 2025; accepted 14 January 2026; published 11 February 2026)

Leveraging complementary machine-learning-based approaches, we compute properties of s - and p -shell Λ hypernuclei—including binding energies, single-particle densities, and radii—starting from the individual interactions among their constituents. These interactions are modeled based on a leading-order finite-cutoff pionless effective field theory expansion. The couplings and the range of the three-nucleon and Λ -nucleon-nucleon potentials are determined via a Gaussian process framework anchored on virtually exact few-body techniques. We solve the many-body Schrödinger equation using a variational Monte Carlo method based on neural network quantum states, extending it for the first time to include Λ particles alongside protons and neutrons. The predicted binding energies show remarkably good agreement with experimental results, given the simplicity of the input Hamiltonian. We also confirm the experimentally observed shrinkage of the proton radius in ${}^7_\Lambda\text{Li}$ compared to its parent nucleus, ${}^6\text{Li}$. This work paves the way for an *ab initio* description of medium-mass and heavy hypernuclei, which is critical for understanding the onset of strange degrees of freedom in the core of neutron stars.

DOI: [10.1103/wmxg-cnrg](https://doi.org/10.1103/wmxg-cnrg)

I. INTRODUCTION

Microscopic approaches to nuclear structure have witnessed tremendous progress over the last two decades. Two- and three-nucleon interactions are derived systematically within nuclear effective field theories [1–3], and quantitative estimates of their uncertainties can be obtained by leveraging Bayesian methods [4,5]. Using these interactions as input, quantum many-body approaches based on single-particle basis expansions—including coupled-cluster theory [6], the in-medium similarity renormalization group [7], and Gorkov-Green’s function methods [8,9]—are now routinely used to predict and describe properties of medium-heavy nuclei with controlled approximations. Quantum Monte Carlo methods have also advanced significantly, providing high-resolution descriptions of static and dynamic properties of nuclei with up to $A \simeq 20$ nucleons, including electromagnetic transitions, beta-decay rates, and electroweak response functions, with high accuracy [10,11].

The frontier of nuclear physics now lies in exotic systems. A prime example is neutron-rich nuclei, which are being explored at facilities such as FAIR, FRIB, JLab, J-PARC, and RIBF. These experiments provide invaluable information on the limits of nuclear stability and on elemental abundances. The same theoretical tools that push us toward the

neutron drip line also open the way to another class of exotic, strangeness-bearing systems. In this context, quantum many-body calculations based on systematically improvable interactions are indispensable for delivering predictions—with quantified theoretical uncertainties—for hypernuclei, self-bound systems that contain at least one hyperon in addition to protons and neutrons. The study of hypernuclei has a long experimental history: The first emulsion experiment was reported at a meeting of the Polish Academy of Sciences in 1952 [12]. Since then, accelerator-based experiments have focused extensively on measuring ground-state properties, excitation energies, transition strengths, and charge symmetry-breaking effects [13–19].

This work focuses on finite hypernuclei, which provide an excellent testbed for constraining hypernuclear interactions. These interactions play a fundamental role in understanding the phenomenology of dense nuclear matter in the interior of neutron stars [20,21]. At densities greater than twice nuclear saturation, the conversion of nucleons into hyperons may become energetically favorable, and Pauli blocking would prevent hyperons from decaying by limiting the phase space available to nucleons. Hyperon formation would lead to a reduction in the Fermi pressure exerted by the baryons and a softening of the equation of state, which would not be able to support the most massive neutron stars observed [22–25]. This apparent inconsistency is a longstanding problem known as the *hyperon puzzle* [26,27]. Its solution requires elucidating hypernuclear interactions through a combination of microscopic analysis and thorough experimental investigation.

For a long time, microscopic calculations of hypernuclei were limited to systems containing up to four baryons,

*Contact author: lovato@anl.gov

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

relying on phenomenological potentials [28–32]. More recently, few-body techniques combined with pionless effective field theory (EFT) interactions have been successfully employed to address the long-standing overbinding problem in light Λ hypernuclei [33]. Quantum Monte Carlo methods have enabled a consistent description of a wide range of hypernuclei and neutron-star matter, though they still rely on somewhat simplified phenomenological interactions [26,34,35]. Additionally, the no-core shell model has been applied to perform microscopic calculations of s - and p -shell hypernuclei, starting from two- and three-baryon interactions consistently derived within chiral EFT [36–39]. These calculations contribute to constraining the hyperon-nucleon interaction, using experimental ground-state and spectroscopic data for selected p -shell hypernuclei [40]. Even more recently, Ref. [41] reported the first calculation of light Λ hypernuclei with up to $A = 7$ based on two- and three-body hypernuclear potentials, both derived within chiral EFT.

In this work, we combine complementary machine-learning-based approaches to compute the properties of s - and p -shell single- Λ hypernuclei, paving the way for calculations of medium-mass systems with quantified uncertainties. Nuclear and hypernuclear forces are modeled based on a leading-order finite-cutoff pionless EFT [42] expansion, with low-energy constants (LECs) determined through a Gaussian-process (GP) framework informed by highly accurate few-body calculations. Using this Hamiltonian as input, we compute the ground-state properties of selected s - and p -shell Λ hypernuclei by extending the state-of-the-art variational Monte Carlo method based on neural network quantum states (VMC-NQS) [43–46] to include Λ particles alongside protons and neutrons. We present results for binding energies, Λ -separation energies, and nuclear and hypernuclear radii and densities, including an analysis of the experimentally observed shrinkage of the proton radius in ${}^7_\Lambda\text{Li}$. Our work lays the foundation for accurate calculations of medium-mass and heavy hypernuclei, as well as hypernuclear matter, with implications for resolving the long-standing hyperon puzzle.

II. HAMILTONIAN

The starting point of our analysis is a nonrelativistic Hamiltonian describing single- Λ hypernuclei ${}^A_\Lambda\text{Z}$, where A is the total number of baryons, which reads

$$H = - \sum_i \frac{\nabla_i^2}{2m_N} + V_{\text{NN}} + V_{\text{NNN}} - \frac{\nabla_\Lambda^2}{2m_\Lambda} + V_{\text{N}\Lambda} + V_{\text{NN}\Lambda}, \quad (1)$$

where Latin indices denote nucleons only, and we set $\hbar = 1$. The first three terms represent the Hamiltonian of the core nucleus ${}^A\text{Z}$, while the last three terms model the kinetic energy of the Λ particle, the nucleon- Λ (N Λ) potential, and the nucleon-nucleon- Λ (NN Λ) potential.

The nucleon-nucleon (NN) potential is expressed as the sum of a charge-independent term and the electromagnetic interaction: $V_{\text{NN}} = V_{\text{NN}}^{\text{CI}} + V_{\text{NN}}^{\text{EM}}$. Following Ref. [47], we assume $V_{\text{NN}}^{\text{EM}}$ to include only the Coulomb repulsion between finite-size protons. We use the parametrization described in Ref. [48], which accounts for finite-size effects through dipole form factors for the interacting protons. Other terms

proportional to the first and second powers of the fine-structure constant—such as two-photon interactions, Darwin-Foldy corrections, vacuum polarization, and magnetic interactions—are neglected.

We model the strong components of both nuclear and hypernuclear forces in a consistent manner inspired by a leading-order pionless EFT expansion [33,47]. The charge-independent components of the nucleon-nucleon and hyperon-nucleon potentials are given by

$$V_{\text{NN}}^{\text{CI}} = \sum_{\text{IS}} C_\lambda^{\text{IS}} \sum_{i<j} \mathcal{P}_{\text{IS}}(ij) \delta_\lambda(\mathbf{r}_{ij}), \quad (2)$$

$$V_{\text{N}\Lambda} = \sum_{\text{IS}} \tilde{C}_\lambda^{\text{IS}} \sum_i \mathcal{P}_{\text{IS}}(i\Lambda) \delta_\lambda(\mathbf{r}_{i\Lambda}), \quad (3)$$

where $\mathcal{P}_{\text{IS}}$ are projection operators for NN and N Λ pairs with total isospin I and spin S , respectively, while C_λ^{IS} and $\tilde{C}_\lambda^{\text{IS}}$ are the corresponding two-body low-energy constants, whose values depend upon the cutoff. We assume the presence of a single shallow virtual state in the $N - \Lambda$ system, restricted to each s -wave channel only, an approach that has proven successful in the past. The relevant isospin-spin combinations in the NN sector are $\text{IS} = \{(10), (01)\}$, corresponding to the 1S_0 and 3S_1 channels. Since strong and electromagnetic interactions do not mix nucleons with Λ particles, the latter are distinguishable and can be taken to have isospin zero. Hence, the s -wave channels for the N Λ potential are $\text{IS} = \{(\frac{1}{2}0), (\frac{1}{2}1)\}$.

In coordinate space, contact interactions are typically regularized by introducing a local Gaussian regulator, specified by its cutoff λ :

$$\delta_\lambda(\mathbf{r}) = \exp\left(-\frac{\lambda^2}{4}\mathbf{r}^2\right). \quad (4)$$

We note that the Gaussian regulator defined above is not normalized; its overall normalization is absorbed into the LECs, a convention consistent with established practice [49,50]. At leading order, the three-body interaction consists of a single spin-isospin-independent NNN term and three NNA contributions associated with the $\text{IS} = \{(0\frac{1}{2}), (1\frac{1}{2}), (0\frac{3}{2})\}$ channels, with explicit forms given by

$$V_{\text{NNN}} = D_\lambda \sum_{i<j<k} \left(\sum_{\text{cyc.}} \delta_\lambda(\vec{r}_{ik}) \delta_\lambda(\vec{r}_{jk}) \right), \quad (5)$$

$$V_{\text{NNA}} = \sum_{\text{IS}} \tilde{D}_\lambda^{\text{IS}} \sum_{i<j} \mathcal{Q}_{\text{IS}}(ij\Lambda) \delta_\lambda(\vec{r}_{i\Lambda}) \delta_\lambda(\vec{r}_{j\Lambda}). \quad (6)$$

The sum in Eq. (5) runs over all NNN triplets, where $\sum_{\text{cyc.}}$ denotes cyclic permutations among the particles within each triplet. Conversely, the sum in Eq. (6) runs over all NN pairs, with the operators $\mathcal{Q}_{\text{IS}}$ projecting the interaction onto baryon triplets with specific isospin I and spin S . Finally, D_λ and $\tilde{D}_\lambda^{\text{IS}}$ represent the LECs for NNN and NNA interactions, respectively, for a given value of the cutoff λ and spin-isospin channel. Note that the same cutoff value is used across all three spin-isospin channels defining the NNA potential.

A. Fitting procedure

In pionless EFT, the independence of observables with respect to the cutoff λ is ensured by allowing the LECs

TABLE I. Numerical values for the two-body LECs and cutoffs adjusted to reproduce the scattering lengths and effective ranges extracted from scattering experiments in the spin-singlet and spin-triplet channels. The np values are taken from Ref. [51], and the p Λ ones from Ref. [52].

Channel	IS	C_{IS} (MeV)	λ_{IS} (fm $^{-1}$)	a_{exp} (fm)	r_{exp} (fm)
np	10	-31.0633	1.101 17	-23.7148(43)	2.750(18)
	01	-68.3747	1.305 12	5.4112(15)	1.7436(19)
	$\frac{1}{2}0$	-33.5417	1.547 20	-1.80	2.80
p Λ	$\frac{1}{2}1$	-25.3115	1.413 79	-1.60	3.30

to run with λ , and by taking the cutoff to be much larger than the breaking scale of the theory [3,53]. However, standard applications of pionless EFT to many-body systems are severely limited by the instability of multifermion states at leading order [49,54,55]. In particular, taking a too large cutoff renders ^{16}O unstable with respect to breakup into four ^4He nuclei [49,54]. To overcome this issue, still motivated by pionless EFT but not claiming formal LO renormalization, we adopt a pragmatic *finite-cutoff* scheme introduced in Refs. [42,47,56,57], which captures the essential elements of nuclear binding while preserving formal consistency. Instead of taking λ to infinity, we fix its value separately in the NN and N Λ channels to reproduce the experimentally extracted effective ranges. The cutoffs for the NNN and NNA three-body forces are likewise tuned to experimental input, as discussed below.

As in Ref. [47], we determine the LECs and cutoff values of the NN potential by fitting to the neutron-proton (np) scattering lengths and effective ranges in the 1S_0 and 3S_1 channels, corresponding to isospin states $\text{IS} = \{(10), (01)\}$, respectively. For the experimental input, we adopt the values reported in Ref. [51]. Similarly, the LECs and cutoffs for the N Λ interaction are fixed to reproduce the scattering lengths and effective ranges inferred from low-energy Λp scattering cross sections measured in Ref. [52], in both the 1S_0 and 3S_1 channels, which correspond to isospin states $\text{IS} = \{(\frac{1}{2}0), (\frac{1}{2}1)\}$. The numerical fits are carried out using the *variable phase approach* [58]. The results are summarized in Table I, which lists the fitted LECs, cutoffs, and corresponding scattering parameters. The LECs and cutoffs obtained for the np channel are in close agreement with those reported in Ref. [47], with only minor deviations arising from slight differences in the scattering lengths and effective ranges used in the fits.

In principle, the coupling constant D_λ and the cutoff λ defining the NNN potential could be determined by fitting the binding energies, denoted as B_E , of ^3H , ^3He , and ^4He . However, owing to the strong correlation between these observables, many different combinations of LECs and cutoffs can reproduce them. In pionless EFT at LO, such correlations reflect renormalization-group invariance and the absence of a four-body force. In our finite-cutoff scheme, we observe the same pattern empirically; however, excessively large cutoffs can induce instabilities in heavier systems, and we therefore include the binding energy of ^{16}O as an additional fitting constraint. The experimental values and their uncertainties used in

TABLE II. Experimental ground-state binding energies used to constrain the three-body NNN interaction taken from Ref. [59]. The uncertainties were inflated to 2.5% of their respective values.

System	B_E (MeV)	σ (MeV)
^3H	8.48	0.21
^3He	7.80	0.20
^4He	28.30	0.71
^{16}O	127.6	3.2

the fit are reported in Table II. To improve the stability of the fitting algorithm, we have artificially inflated the uncertainties of the binding energies to 2.5% of their respective values, well below the theoretical uncertainty associated with a leading-order pionless EFT interaction [5,60].

Similarly, we determine the low-energy constants and cutoffs of the NNA potential by reproducing Λ -separation energies, defined as the energy difference between the hypernucleus and its parent nucleus:

$$B_\Lambda = E_{A-1Z} - E_{\Lambda Z}^\Lambda,$$

for $^3_\Lambda\text{H}$, $^4_\Lambda\text{H}_{S=0}$, $^4_\Lambda\text{H}_{S=1}$, and $^5_\Lambda\text{He}$; as in the nuclear sector, we also include $^{16}_\Lambda\text{O}$ to lift the residual few-body degeneracy and constrain the interaction range. Experimental values and their uncertainties are listed in Table III. Our calibration proceeds in two stages: (1) fit the NNN LEC and cutoff to $\{^3\text{H}, ^3\text{He}, ^4\text{He}, ^{16}\text{O}\}$ and (2) with the nuclear interaction fixed, fit the NNA couplings and cutoff to $\{^3_\Lambda\text{H}, ^4_\Lambda\text{H}_{S=0,1}, ^5_\Lambda\text{He}, ^{16}_\Lambda\text{O}\}$. This separation is justified by the mild sensitivity of B_Λ to the purely nucleonic Hamiltonian [34] and is discussed further in Sec. IV; a fully rigorous uncertainty quantification would jointly sample all LECs, e.g., via parametric matrix models [61], eigenvector continuation [62], or foundation NQS [63].

To map energies to parameters and quantify uncertainties, we use Gaussian-process emulators for $B_E(D_\lambda, \lambda)$ and $B_\Lambda(\tilde{D}_\lambda^{\text{IS}}, \tilde{\lambda})$ trained on few-body data generated with stochastic variational method (SVM) [64], while for ^{16}O and $^{16}_\Lambda\text{O}$ we use VMC-NQS (see Sec. III). The specific choice of the GP kernel, the assumed posterior statistical model, and the Markov chain Monte Carlo (MCMC) sampler are summarized in Appendix A. Here, we analyze the sampled posterior distributions for energies, LECs, and cutoffs, em-

TABLE III. Experimental Λ -separation energies used to constrain the three-body NNA interaction taken from the hypernuclear database [65]. The listed values correspond to the database version accessed in April 2025.

System	B_Λ (MeV)	σ (MeV)	IS dependence
$^3_\Lambda\text{H}$	0.164	0.043	$(0\frac{1}{2})$
$^4_\Lambda\text{H}_{S=0}$	2.169	0.042	$(0\frac{1}{2}), (1\frac{1}{2})$
$^4_\Lambda\text{H}_{S=1}$	1.081	0.046	$(0\frac{1}{2}), (1\frac{1}{2}), (0\frac{3}{2})$
$^5_\Lambda\text{He}$	3.102	0.030	$(0\frac{1}{2}), (1\frac{1}{2}), (0\frac{3}{2})$
$^{16}_\Lambda\text{O}$	13.00	0.089	$(0\frac{1}{2}), (1\frac{1}{2}), (0\frac{3}{2})$




FIG. 1. Joint distributions of the NNN potential LEC (in MeV) and cutoff (in fm⁻¹). The vertical dashed lines denote the 16th, 50th, and 84th percentiles of marginal distribution. The orange cross indicates the projection of the median point of the distribution.

phasizing correlations and credible intervals. In Fig. 1, we show corner plots displaying their marginal and joint posterior distributions for the NNN potential, while the joint posterior distributions of nuclear binding energies are shown in Fig. 2. The plots include histograms and credible intervals at the 16th, 50th, and 84th percentiles (corresponding to $\pm 1\sigma$ for a normal distribution). The overlaying best-fit values and corresponding energies are indicated by orange solid lines. The corresponding posteriors for the NNA sector and the Λ -hypernuclear separation energies are displayed in Figs. 3 and 4. In the diagonal marginals of Figs. 2 and 4, the light blue shaded curves show the Gaussian experimental constraints, i.e., $\mathcal{N}(B_{\text{exp}}, \sigma_{\text{exp}})$ (with means and standard deviations taken from Tables II and III) and are overlaid to visually compare the posterior predictions with the experimental likelihood.

The joint posterior distributions highlight how correlations among observables emerge from the symmetries of the underlying Hamiltonian and propagate into relationships among the low-energy constants. In particular, the binding energies of ^3H , ^3He , and ^4He , shown in Fig. 2, lie on an almost one-dimensional manifold in the (λ, D_λ) plane, resulting in a strong positive correlation between these parameters. This behavior implies an approximate degeneracy when fitting only light nuclei, and the narrow ridge observed in Fig. 1 extends to very large values of λ . The inclusion of the heavier nucleus ^{16}O in the fit breaks this degeneracy and constrains λ to a finite interval. This effect is clearly visible in the same figure, where the strong correlations among the light nuclei are significantly mitigated by the addition of ^{16}O . In practice, we identified the physically stable cutoff window by first calibrating on few-body systems within narrow cutoff slices and then evaluating ^{16}O (and $^{16}_\Lambda\text{O}$) along the resulting few-body-compatible

manifold. The heavy systems exhibit a clear monotonic trend with the three-body range, which selects a restricted region of cutoff values used for the final calibration stage. We note that the inclusion of ^{16}O in the calibration of three-body forces was first introduced in the construction of the chiral EFT interaction NNLO_{sat}, which yields accurate binding energies and charge radii for nuclei up to ^{40}Ca [66].

The picture appears to be even richer in the NNA sector. Each B_Λ couples to a specific linear combination of the three spin-isospin channels: $(0\frac{1}{2})$, $(1\frac{1}{2})$, and $(0\frac{3}{2})$. Consequently, observables that probe different mixtures of these channels—such as $B_\Lambda(^4_\Lambda\text{H})$ versus $B_\Lambda(^5_\Lambda\text{He})$ —populate well-defined bands in the multidimensional space of LECs and range. Experimental uncertainties thus map directly onto the width of the allowed parameters' region: the more precise the B_Λ data, the narrower the credible intervals for the corresponding couplings and ranges. The joint posterior for $B_\Lambda(^4_\Lambda\text{H})$ and $B_\Lambda(^4_\Lambda\text{H}_{S=1})$ is markedly narrower than the individual experimental error bars reported in Table III. It is important to note that this tightening does not necessarily reflect an improved theoretical constraint on the experimental measurements, but rather the limited flexibility of our leading-order interaction model, which has to be verified upon inclusion of subleading orders. Although not as pronounced as in the NNN case, correlations among few-body observables in the hypernuclear sector broaden the convergence region of the fitted parameters into a continuous path in parameter space, which allows for very large values of λ . As in the nuclear case, including $^{16}_\Lambda\text{O}$ in the fit is essential to constrain the range of the NNA interaction. This observable breaks the remaining degeneracy and selects a relatively narrow region in parameter space consistent with the experimental B_Λ data.

III. HYPERNUCLEAR NEURAL NETWORK QUANTUM STATE ANSATZ

We solve the nuclear quantum many-body problem by extending the highly flexible Pfaffian-Jastrow ansatz, originally introduced to model ultracold Fermi gases in Ref. [45] and later generalized to nuclear systems in Ref. [46], to include hyperons in addition to protons and neutrons. Since the strong and electromagnetic interactions do not mix nucleons with the Λ hyperon, we treat the latter as a distinguishable particle, so that

$$\langle X | \Psi_{PJ}^\Lambda \rangle \equiv \Psi_{PJ}^\Lambda(X) = e^{J(X)} \times \psi_\Lambda(\mathbf{x}_\Lambda) \times \text{Pf}[\Phi(X_N)], \quad (7)$$

where $J(X)$ is the complex-valued permutation-invariant Jastrow factor, $\text{Pf}[\Phi(X_N)]$ the nucleonic Pfaffian, and $\psi_\Lambda(\mathbf{x}_\Lambda)$ a single-particle wave function for the Λ particle. For the s - and p -shell systems considered here, pairing is not expected to be a dominant feature of the wave function. Nevertheless, a Pfaffian ansatz offers several advantages: (1) with an appropriate choice of pairing orbital, it reduces to a single Slater determinant; (2) since the orbital takes as input both spatial and spin-isospin coordinates, it naturally captures independent-pair correlations [67] beyond pairing; (3) it enables an efficient description of open-shell systems using a single orbital rather than multiple Slater determinants; and (4) it scales favorably with particle number, as a single orbital suffices irrespective of system size. To enhance the




FIG. 2. Joint posterior distributions of nuclear binding energies (MeV) obtained by varying the NNN potential LEC and cutoff. In the diagonal panels, vertical dashed lines mark the 16th, 50th, and 84th percentiles of each marginal distribution, and the light blue shaded areas represent Gaussian distributions centered at the experimental binding energies, with widths given by the corresponding experimental uncertainties in Table II. Orange markers and connecting lines indicate the binding energy computed with the median potential from Fig. 1.

expressiveness of the ansatz, we apply a permutation-equivariant backflow transformation to the single-particle coordinates. This idea dates back to Ref. [68] and has recently emerged as a powerful and widely adopted strategy to augment the expressivity of antisymmetric NQS [69–71]. Following Ref. [46], we employ the message-passing neural network developed in Refs. [45,72], which has proven to efficiently capture correlations in both the homogeneous electron gas and ultracold Fermi gases.

Concretely, we evaluate the hypernuclear ansatz on learned one-body and pair features, rather than on raw coordinates, namely, $\mathbf{g}_i(\mathbf{x}_i; \mathbf{x}_{k \neq i})$ and $\mathbf{g}_{ij}(\mathbf{x}_i, \mathbf{x}_j; \mathbf{x}_{k \neq i, j})$, respectively. In order for the NQS ansatz to be antisymmetric, these transformations are permutation-equivariant, so that a relabeling of particle indices induces the corresponding relabeling of the features. For the single-particle feature $\mathbf{g}_i$, the exchange $i \leftrightarrow j$ yields $\mathbf{g}_i \rightarrow \mathbf{g}_j$, while permutations among particles other than

i leave $\mathbf{g}_i$ unchanged. For the pair feature, an exchange $i \leftrightarrow j$ yields $\mathbf{g}_{ij} \rightarrow \mathbf{g}_{ji}$; swapping i with a third index k gives $\mathbf{g}_{ij} \rightarrow \mathbf{g}_{kj}$, and analogously exchanging j with k gives $\mathbf{g}_{ij} \rightarrow \mathbf{g}_{ik}$. On the other hand, $\mathbf{g}_{ij}$ is invariant under permutations that do not involve i or j (e.g., $k \leftrightarrow \ell$ with $k, \ell \notin i, j$). To let the backflow distinguish between nucleons and Λ hyperons—and to make it fully expressive—we assign the latter the auxiliary isospin value $t_{z, \Lambda} = 0$. This backflow construction is completely general and can, in principle, accommodate interactions that mix nucleons and Λ hyperons. Implementation details of the permutation-equivariant backflow, the Jastrow and Pfaffian parametrizations, and all architectural hyperparameters are provided in Appendix B.

A. Benchmark with SVM

The purely nucleonic Pfaffian-Jastrow ansatz of Eq. (B2) was recently validated by benchmarking the ground-state




FIG. 3. Joint distributions of the LECs (in MeV) and cutoff (in fm^{-1}) entering NNA potential. Plotting conventions are the same as in Fig. 1.

energies of selected nuclei with up to $A = 6$ nucleons, computed using the SVM method [73]. Agreement at better than the percent level was found using as input an improved-action pionless EFT Hamiltonians across a wide range of cutoff values.

Here, we focus on the hypernuclear sector. To gauge the accuracy of the NQS ansatz in Eq. (7), we performed benchmark calculations using the SVM method for light hypernuclei, finding excellent agreement, as shown in Table IV. In both methods, we use as input the median NNN and NNA potentials from Figs. 1 and 3, respectively.

Note that to accelerate convergence to the ground-state energy, we enforce that the NQS be an eigenstate of the parity operator by symmetrizing the wave function as $\Psi_{PJ}^\Lambda(X) \pm \Psi_{PJ}^\Lambda(PX)$, where the P flips the spatial coordinates of all the particles, including the Λ :

$$\mathcal{P}(\mathbf{r}_1, \dots, \mathbf{r}_{A-1}, \mathbf{r}_\Lambda) = (-\mathbf{r}_1, \dots, -\mathbf{r}_{A-1}, -\mathbf{r}_\Lambda).$$

IV. RESULTS

To set the stage, Fig. 5 displays the binding energy per nucleon computed for selected nuclei that serve as a baseline for hypernuclear calculations. The shaded gray regions highlight the nuclei used to determine the λ and D_λ parameters of the

TABLE IV. Hypernuclear binding energies (in MeV) computed using the VMC algorithm with the Pfaffian-Jastrow NQS, compared with results from the SVM method. Both methods use as input the median NNN and NNA potentials from Figs. 1 and 3.

System	NQS	SVM
${}^3_\Lambda\text{H}$	2.399(2)	2.409
${}^4_\Lambda\text{H}_{S=0}$	10.606(3)	10.601
${}^4_\Lambda\text{H}_{S=1}$	9.529(4)	9.490
${}^5_\Lambda\text{He}$	31.086(3)	31.082




FIG. 4. Joint posterior distributions of the Λ -separation energies B_Λ (MeV) for ${}^3_\Lambda\text{H}$, ${}^4_{\Lambda S=0,1}$, ${}^5_\Lambda\text{He}$, and ${}^{16}_\Lambda\text{O}$, obtained by propagating the sampled NNA interaction parameters (Fig. 3) through the GP emulators. In the diagonal panels, the light blue shaded curves show the Gaussian experimental constraints $\mathcal{N}(B_\Lambda^{\text{exp}}, \sigma_{\text{exp}})$ with the values listed in Table III. The remaining plot elements (marginals, credible intervals, and best-fit markers) are defined as in Fig. 2.

NNN potential, as discussed in Sec. II. The uncertainties in the NQS calculations reflect both the statistical error from the variational VMC, σ_{VMC} , and the uncertainty associated with the NNN force, σ_{NNN} . The latter is estimated by sampling ten different realizations of the NNN potential, with λ and D_λ drawn from the joint posterior distribution shown in Fig. 1. As a conservative estimate, the total energy uncertainty is computed by summing the two contributions in modulus rather than in quadrature: $\sigma_{B_E} = \sigma_{\text{NQS}} + \sigma_{\text{NNN}}$. In principle, the truncation error of the theory should also be included. However, due to the improved nature of the interaction, a study of cutoff variations is not feasible, and the lack of experimental data further complicates an uncertainty estimate based on variations of the theory fit input.

Consistent with the findings of Refs. [47,57], fitting the np scattering length and effective range in the 3S_1 channel constrains the NN potential to accurately reproduce the deuteron binding energy. As for ${}^3\text{H}$, ${}^4\text{He}$, and ${}^{16}\text{O}$, the excellent

agreement with experiment is expected, as their binding energies were used to determine the coupling and range of the NNN force. It is important to note, however, that the NQS binding energies, computed using a simple potential inspired by a leading-order pionless EFT expansion, successfully reproduce experimental values across several light nuclei—consistent with the findings of Ref. [57]. For the nuclei not included in the fit, we observe an average deviation of approximately 5% between the calculated B_E and experimental data. The largest discrepancies, around 9%, occur for ${}^{11}\text{B}$ and ${}^{11}\text{C}$, whose structure is strongly influenced by p -shell orbitals. This underbinding is likely due to the absence of explicit p -wave terms in the interaction [74]. Notably, the differences between theory and experiment are much smaller than the $\sim 30\%$ uncertainty typically expected at leading order in pionless EFT [75]. This improved agreement likely benefits from finite-range NN and NNN terms that can mimic some subleading effects.




FIG. 5. Binding energy per nucleon for selected nuclei up to ^{16}O . NQS results (orange) are compared with experimental data (blue). The shaded bands reflect both Monte Carlo statistical errors and model uncertainties in the coupling and range of the NNN force. Gray regions highlight the nuclei used to constrain the three-body force parameters.

Figure 6 illustrates the Λ -separation energies of selected hypernuclei up to $^{16}_{\Lambda}\text{O}$. The NQS calculations are benchmarked against experimental values from the hypernuclear database [65] (blue), as well as recent theoretical predictions [40,41,76] (green, purple, and red, respectively), which differ both in the input Hamiltonians and in the numerical methods used to solve the nuclear Schrödinger equation. Specifically, Ref. [40] employs the importance-truncated no-core shell model (IT-NCSM) with reoptimized chiral YN interactions at LO (as opposed to the original Jülich YN constrained only by YN scattering), together with an N^3LO chiral EFT NN interaction; no explicit NNA three-body terms are included. On the other hand, Ref. [76] uses nuclear lattice effective field theory (NLEFT) with N^3LO nucleon-nucleon interactions and leading-order s -wave $\text{N}\Lambda$ plus NNA forces constrained by $A = 4, 5$ hypernuclei. Finally, Ref. [41] reports NCSM calculations based on a chiral expansion of both $\text{N}\Lambda$ and NNA interactions up to N^2LO .

The NQS energies are supplemented by estimates of their theoretical uncertainties, which include both the statistical error from the VMC calculation, and uncertainties associated with the NNA potential: $\sigma_{B_{\Lambda}} = \sigma_{\text{NQS}} + \sigma_{\text{NNA}}$. The latter are




FIG. 6. Λ -separation energies for selected hypernuclei up to $^{16}_{\Lambda}\text{O}$. NQS results (orange) are compared with experimental data from the hypernuclear database (blue) and recent theoretical predictions from NLEFT (purple) and IT-NCSM (green). The horizontal error bars on the NQS points represent combined Monte Carlo statistical uncertainties and uncertainties in the NNA force.

obtained in the same manner as in Fig. 5, by considering ten different samples of the LECs and cutoffs entering the NNA force, drawn from the joint posterior distribution shown in Fig. 3. For simplicity, the coupling and range of the NNN force are kept fixed to their median values, indicated by the orange dots in Fig. 1. This approximation relies on the partial cancellation of purely nuclear effects in the energy difference defining B_{Λ} ; however, residual sensitivity to the nuclear Hamiltonian persists—particularly for simplified YN interactions—and has been estimated in s - and p -shell systems [77,78]. Within our finite-cutoff pionless EFT, this contribution is subdominant to our present theoretical uncertainties.

As expected, the NQS calculations for the systems used to calibrate the $N\Lambda$ interaction closely reproduce the experimental data. The agreement remains excellent—within theoretical and experimental uncertainties—even for the ${}^4_\Lambda\text{He}_{S=1}$ and ${}^4_\Lambda\text{He}_{S=0}$ states, which differ from their ${}^4_\Lambda\text{H}$ analogues only due to small charge-dependent and charge-symmetry-breaking effects. Interestingly, the B_Λ values of hypernuclei with $A \geq 6$ exhibit a nearly linear trend with the mass number A , up to ${}^{16}_\Lambda\text{O}$. Overall, the deviations from experimental data for hypernuclei not included in the fit remain small, with an average discrepancy below 9%. The largest deviations—significantly exceeding the combined theoretical and experimental uncertainties—are observed for ${}^{12}_\Lambda\text{B}$, ${}^{12}_\Lambda\text{C}$, and ${}^{13}_\Lambda\text{C}$. This pattern is consistent with the underbinding observed in the corresponding parent nuclei and likely stems from the absence of p -wave contributions in both the NN and $N\Lambda$ potentials. We expect these relatively small discrepancies to be resolved upon inclusion of such contributions in these potentials.

A. Hypernuclear densities and radii

Spatial distribution functions and the corresponding radii provide additional insight into the structure of hypernuclei. In particular, they highlight the smaller $N\Lambda$ scattering length compared to NN interactions and illustrate the role of the Pauli exclusion principle in shaping nuclear structure. Notably, unlike quantum many-body approaches based on harmonic-oscillator basis expansions—such as the IT-NCSM and NCSM—the shell structure of nuclei and hypernuclei in NQS calculations emerges naturally from the variational minimization of the Hamiltonian expectation value, rather than being encoded in the variational ansatz.

The single-particle radial density distributions denote the probability of finding a nucleon or a Λ particle at a distance r from the center of mass of the hypernucleus. They are defined as the ground-state expectation value

$$\rho_\tau(r) = \frac{1}{4\pi r^2} \frac{\langle \Psi_{PJ}^\Lambda | \sum_i \delta(r - |\vec{r}_i|) P_i^\tau | \Psi_{PJ}^\Lambda \rangle}{\langle \Psi_{PJ}^\Lambda | \Psi_{PJ}^\Lambda \rangle}. \quad (8)$$

In the above equation, P_i^τ denotes the proton, neutron, or Λ projector, and the index i runs over all A particles in the system, including both nucleons and the Λ hyperon. For illustrative purposes, we normalize the density distributions to unity, so that $\int 4\pi r^2 \rho_\tau(r) dr = 1$. In addition, to mitigate artifacts due to different spherical shell volumes, we multiply the density by a factor r^2 . The corresponding point radii are defined as

$$\langle r_\tau^2 \rangle = \frac{1}{N_\tau} \frac{\langle \Psi_{PJ}^\Lambda | \sum_i \vec{r}_i^2 P_i^\tau | \Psi_{PJ}^\Lambda \rangle}{\langle \Psi_{PJ}^\Lambda | \Psi_{PJ}^\Lambda \rangle}, \quad (9)$$

where N_τ denotes the number of protons or neutrons, and $N_\tau = 1$ for the Λ particle, since in this work we consider single- Λ hypernuclei.

To reduce statistical uncertainties, both the single-particle density distributions and the point radii are evaluated using approximately two million Monte Carlo configurations, binned into spherical shells of width $\Delta r = 0.04$ fm. We note that, in contrast to diffusion Monte Carlo calculations, no extrapolations are required to compute expectation values of operators

TABLE V. Point radii for protons, neutrons, and Λ hyperon in ${}^5_\Lambda\text{He}$, ${}^7_\Lambda\text{Li}$, and ${}^{16}_\Lambda\text{O}$, along with those of their corresponding parent nuclei. All values are reported in fm, with statistical uncertainties shown in parentheses, affecting the last digit.

System	$\sqrt{\langle r_p^2 \rangle}$	$\sqrt{\langle r_n^2 \rangle}$	$\sqrt{\langle r_\Lambda^2 \rangle}$
${}^5_\Lambda\text{He}$	1.57(1)	1.57(1)	2.32(1)
${}^4\text{He}$	1.42(1)	1.42(1)	—
${}^7_\Lambda\text{Li}$	2.30(1)	2.30(1)	2.30(2)
${}^6\text{Li}$	2.64(2)	2.63(2)	—
${}^{16}_\Lambda\text{O}$	2.48(1)	2.41(1)	2.15(1)
${}^{15}\text{O}$	2.50(1)	2.42(1)	—

that do not commute with the Hamiltonian, as NQS states provide pure estimators for these quantities.

The neutron, proton, and Λ -particle spatial density distributions of ${}^5_\Lambda\text{He}$, ${}^7_\Lambda\text{Li}_{\frac{1}{2}}$, and ${}^{16}_\Lambda\text{O}$ are displayed in the upper, middle, and lower panels of Fig. 7, respectively. In ${}^5_\Lambda\text{He}$, the Λ hyperon is weakly bound and tends to form an extended halo around the ${}^4\text{He}$ core. We observe similar behavior in other nuclei with $A < 7$, such as ${}^3_\Lambda\text{H}$, ${}^4_\Lambda\text{H}$, and ${}^4_\Lambda\text{He}$, indicating a significantly weaker attraction in the $N\Lambda$ potential compared to the NN interaction, consistent with the findings of Ref. [29]. Starting from ${}^7_\Lambda\text{Li}_{\frac{1}{2}}$, the spatial distribution of the Λ becomes more localized near the center of mass of the hypernucleus, compared to those of protons and neutrons. This behavior arises because the Λ , being distinguishable from nucleons, can continue to occupy inner s -shell orbitals, while protons and neutrons are subject to the Pauli exclusion principle and must populate higher p -shell states. This trend is further accentuated in ${}^{16}_\Lambda\text{O}$, where most protons and neutrons occupy p -shell states, while the Λ remains free to reside in the central region of the system.

Further insights can be gained by comparing the corresponding radii of the three hypernuclei discussed above and their parent nuclei, as reported in Table V. Consistent with the density distributions, the radius of the Λ hyperon is larger than that of the nucleons in ${}^5_\Lambda\text{He}$, becomes comparable in ${}^7_\Lambda\text{Li}_{\frac{1}{2}}$, and is significantly smaller in ${}^{16}_\Lambda\text{O}$. Even more revealing is the so-called “core shrinkage,” i.e., the relative change in the proton radius induced by the presence of the Λ hyperon [79,80]. To quantify this effect, we define the relative radius difference as

$$\delta r = \frac{\sqrt{\langle r_p^2 \rangle_\Lambda^A} - \sqrt{\langle r_p^2 \rangle^{A-1}}}{\sqrt{\langle r_p^2 \rangle^{A-1}}}, \quad (10)$$

where $\sqrt{\langle r_p^2 \rangle_\Lambda^A}$ is the proton point radius of the hypernucleus, and $\sqrt{\langle r_p^2 \rangle^{A-1}}$ is the corresponding radius for its parent nucleus.

We observe a two-regime behavior of δr as a function of A , which is consistent with trends observed in the spatial density distributions displayed in Fig. 7, where the proton density distribution of parent nuclei is given by a thick solid line.

For light systems such as ${}^3_\Lambda\text{H}$, ${}^4_\Lambda\text{H}$, ${}^4_\Lambda\text{He}$, and ${}^5_\Lambda\text{He}$, our NQS calculations yield $\delta r > 0$, indicating that the proton radius




FIG. 7. Neutron (orange squares), proton (green triangles), and Λ hyperon (blue circles) single-particle radial density distributions for ${}^5_\Lambda\text{He}$ (upper panel), ${}^7_\Lambda\text{Li}$ (middle panel), and ${}^{16}_\Lambda\text{O}$ (lower panel). The proton densities of the corresponding parent nuclei are also shown by the black solid line. For illustrative purposes, the distributions are normalized to unity and multiplied by a factor of r^2 , as discussed in the text.

in the hypernucleus is larger than in the parent nucleus. In these systems, the Λ hyperon is weakly bound and tends to form an extended halo around the core, effectively “pulling” the nucleons outward from the center of mass. For ${}^5_\Lambda\text{He}$, we

obtain $\delta r \simeq 0.11$, in qualitative agreement with the pioneering theoretical calculations of Ref. [29] for ${}^4_\Lambda\text{H}$ and ${}^4_\Lambda\text{He}$, despite different interactions and few-body methods are employed.

Starting from ${}^7_\Lambda\text{Li}$ —for which we find $\delta r \simeq -0.13$ —the proton radius of the hypernucleus becomes smaller than that of its parent nucleus, in qualitative agreement with the experimental findings of Ref. [14]. Within an extreme cluster model, ${}^6\text{Li}$ can be interpreted as having an α -deuteron clustered wave function, characterized by a relatively loose configuration and sizable separation between the clusters. When a Λ hyperon is added to form ${}^7_\Lambda\text{Li}$, it occupies the central $1s$ orbital and interacts attractively with nucleons in both the α and d clusters. Since the Λ hyperon is distinguishable from nucleons and does not experience Pauli repulsion, it acts as a “glue,” pulling the clusters closer together [79]. This behavior is reflected in the central panel of Fig. 7, where the proton density of ${}^6\text{Li}$ is broader than that of ${}^7_\Lambda\text{Li}$.

We observe that this gluelike effect gradually weakens with increasing A , decreasing to $\delta r \lesssim 0.01$ in ${}^{16}_\Lambda\text{O}$. This percent-level compaction agrees with relativistic mean-field calculations of Ref. [81], which predict only slight hyperon-induced modifications to the core radius in this and heavier mass regions. Because the Λ hyperon resides near the center of the nucleus, it probes the highest densities accessible in finite nuclei, close to saturation density. As a result, Λ separation energies are expected to provide meaningful constraints on the ΛNN interaction, which plays a critical role in the onset of hyperons in neutron-star matter.

V. CONCLUSIONS

In this work, we present calculations of single- Λ hypernuclei up to ${}^{16}_\Lambda\text{O}$ based on a finite-cutoff pionless EFT Hamiltonian. To solve the quantum many-body problem, we extend the Pfaffian-Jastrow NQS ansatz to include a single- Λ particle. For light systems, VMC calculations based on this ansatz are benchmarked against highly accurate SVM results, showing better than percent-level agreement.

The Hamiltonian includes NN and $\text{N}\Lambda$ two-body potentials, as well as NNN and $\text{NN}\Lambda$ three-body interactions. To avoid the instability of nuclei with more than $A = 4$ nucleons, the two-body LECs are fixed to the np and $p\Lambda$ s -wave scattering lengths and effective ranges in the 1S_0 and 3S_1 channels [47,73]. In the three-body sector, calibration leverages Gaussian-process emulators to interpolate nuclear binding energies and Λ -separation energies as functions of the potential ranges and couplings. Because few-body observables are strongly correlated, we find that fits based only on $A \leq 4$ data are nearly degenerate; adding ${}^{16}_\Lambda\text{O}$ and ${}^{16}_\Lambda\text{O}$ proves effective in constraining the NNN and $\text{NN}\Lambda$ potentials, thereby breaking this degeneracy.

As a first step, we validate the nuclear baseline: The binding energies of the parent nuclei in the s and p shells agree with experimental values at the $\sim 5\%$ level, despite not being used to determine the NNN force. Building on this agreement, we then compute the Λ -separation energies of s - and p -shell hypernuclei using samples from the calibrated ΛNN interactions, obtaining Λ -separation energies that fall within $\lesssim 9\%$ of experiment. The remaining deviations—larger for

$A = 12$ – 13 —point to missing p -wave components [74]; we expect leading p -wave terms in the NN and N Λ sectors to reduce them.

In addition to ground-state energies, NQS give us access to the wave functions of the systems, which allows us to compute the spatial distributions of nucleons and hyperons. Analyzing these distributions, we find that for $A \leq 5$ the Λ is weakly bound and forms an extended halo, leading to an expansion of the nuclear core. In contrast, in ${}^7_\Lambda\text{Li}$ the Λ acts as “glue,” shrinking the proton radius, consistent with experiment. These effects diminish as the number of nucleons increases and are barely noticeable in ${}^{16}_\Lambda\text{O}$. In this system, however, the hyperon, being distinguishable from the nucleons, occupies the central $1s$ orbital and thus experiences a density close to nuclear saturation, providing a valuable regime for calibrating NNA inputs of astrophysical relevance.

Looking ahead, the first extension of this work will entail including p -wave contributions in the nuclear interactions, which have proven essential for an accurate description of p -shell nuclei [74]. Second, we will perform a *joint* Bayesian inference over the combined set of NNN and NNA LECs. Moving beyond GPs, we plan on using fast multi-Hamiltonian emulators—such as parametric matrix models [61], eigenvector continuation [62], and foundation NQS [63]—to enable efficient posterior sampling across correlated nuclear and hypernuclear sectors. Finally, exploiting the favorable scalability of the VMC-NQS method with particle number, we will carry out calculations of medium-mass hypernuclei such as ${}^{40}_\Lambda\text{Ca}$ and ${}^{48}_\Lambda\text{Ca}$, where near-saturation densities probe a possible isospin dependence of the N Λ interaction [82]. At densities relevant to the onset of hyperons in neutron stars, we will transition to higher-resolution local chiral EFT NN and NNN interactions [83], potentially combining them with N Λ and NNA interactions as in Ref. [41]. The machine-learning methodology developed here will be particularly valuable, enabling efficient sensitivity analyses of chiral EFT nuclear and hypernuclear forces and streamlined calculations from light- to medium-mass hypernuclei.

ACKNOWLEDGMENTS

The authors acknowledge stimulating discussions with Alessandro Roggero. The present research was supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Contract No. DE-AC02-06CH11357 (A.L.), by the DOE Early Career Research Program (A.L.), by the STREAMLINE Collaboration Award No. DE-SC0024586 (A.L., F.P., and A.D.D.), and by the SciDAC-5 NUCLEI program (A.L.). Numerical calculations were performed using an ISCRE award for accessing the LEONARDO supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CINECA (Italy). We gratefully acknowledge the computing resources provided on Swing, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory. An award of computer time was provided by the U.S. Department of Energy’s (DOE) Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program. This research used resources from the Argonne Leadership Computing Facility, a U.S. DOE Office of Science user facility at Argonne National

Laboratory, which is supported by the Office of Science of the U.S. DOE under Contract No. DE-AC02-06CH11357.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: GAUSSIAN-PROCESS CALIBRATION AND VALIDATION

The fit for both the NNN and NNA potentials is efficiently performed by leveraging GPs [84] to interpolate binding energies and Λ -separation energies as continuous functions of the LECs and cutoffs. This approach is particularly advantageous when fitting the NNA force, as the relatively high dimensionality of the parameter space makes a direct grid search computationally prohibitive. First, we generate the datasets of binding energies

$$\{(D_\lambda, \lambda)_i, B_E(D_\lambda, \lambda)_i\}, \quad (\text{A1})$$

and Λ -separation energies

$$\{(\tilde{D}_\lambda^{\text{IS}}, \tilde{\lambda})_i, B_\Lambda(\tilde{D}_\lambda^{\text{IS}}, \tilde{\lambda})_i\}, \quad (\text{A2})$$

calculated at specific parameter values. We remark that the two sets correspond to the NNN and NNA potentials, respectively. For the GP calibration in the few-body sector, we solve light nuclei and hypernuclei using the highly accurate SVM [64], as it enables fast and precise calculations for few-body systems. However, with the exception of the figures displaying GP posteriors, all few-body results shown in the figures are obtained with VMC-NQS. The binding energy of ${}^{16}_\Lambda\text{O}$ and the Λ -separation energy of ${}^{16}_\Lambda\text{O}$, which are also used in the fit, are computed using the VMC-NQS method [43,44,46], extended in this work to the hypernuclear sector, as described in detail in Sec. III.

1. Training datasets and kernels

Due to its low computational cost, we employ the SVM to generate a set of 1000 training points for light nuclei and hypernuclei. For ${}^{16}_\Lambda\text{O}$ and ${}^{16}_\Lambda\text{O}$, which are more computationally intensive, we generate 50 training points each with the VMC-NQS method. These datasets are then used to train GP models, enabling smooth and analytic continuations of the nuclear binding energies and hypernuclear separation energies $B_E(D_\lambda, \lambda)$ and $B_\Lambda(\tilde{D}_\lambda^{\text{IS}}, \tilde{\lambda})$.

The GP interpolation not only predicts energies at untested parameter values but also quantifies their uncertainties, enhancing the robustness of the fitting procedure. The GP kernel imposes a prior on the mapping $(D, \lambda) \mapsto E$ by encoding our prior assumptions about the functional form of $E(D, \lambda)$. A GP model’s prior is fully determined by its second-order statistics via the kernel, which defines the pairwise covariance function between points in the input space. In our calculations, we use a combination of *rational quadratic* and *matern*32 kernels. As noted in Ref. [85], combining two kernels is analogous to a logical *or* operation: The resulting kernel will exhibit large values if either kernel independently has large values.

The rational quadratic kernel models infinitely differentiable functions with varying degrees of smoothness across the input space, while the matern32 kernel with $\nu = 3/2$ models smooth functions that may include inflection points where the derivative changes abruptly.

To accommodate differing length scales along the various axes of the parameter space, we apply the automatic relevance determination technique [86]. This requirement follows from two points: (1) the cutoff varies over a much narrower range than the LECs and (2) our analysis suggests that the hypernuclear $\tilde{D}_{0\frac{1}{2}}$ LEC, unlike the other two, varies only within a narrow region of the parameter space close to 0. The last observation stems from the fact that the binding energy of ${}^3_\Lambda\text{H}$ depends solely on $\tilde{D}_{0\frac{1}{2}}$ in the three-body sector, and the related separation energy is small—see Table III. Notably, the average distance between the Λ particle and the proton-neutron center of mass is about 10 fm, as reported in Ref. [87], a typical characteristic of halo hypernuclei. The hyperparameters of the kernel are optimized by maximizing the log-marginal likelihood. The posterior GP is derived using Algorithm 2.1 from Ref. [84], as implemented in the Python package GPy [88].

2. Sensitivity and correlation analysis

Assuming uniform priors on the LECs and the cutoff, the nuclear and hypernuclear posteriors over the joint parameter vector are proportional to a multivariate Gaussian likelihood:

$$\mathcal{N}(\vec{B}|\vec{D}, \lambda, \vec{\Sigma}) = \frac{e^{-\frac{1}{2}(\vec{B}-\vec{B}_{\text{exp}})\vec{\Sigma}^{-1}(\vec{B}-\vec{B}_{\text{exp}})}}{\sqrt{(2\pi)^k \det \vec{\Sigma}}}, \quad (\text{A3})$$

where $\vec{B} = \vec{B}(\vec{D}, \lambda)$ is the vector of the binding energies or separation energies for the NNN and NNA sector force, respectively, obtained through the GP. Analogously, $\vec{B}_{\text{exp}}$ are the k experimental values—four in the NNN sector and five in the NNA sector—and $\vec{\Sigma}$ is the $k \times k$ diagonal covariance matrix corresponding to the experimental uncertainties. The nuclear and hypernuclear parameter distributions corresponds to $\mathcal{N}(\vec{B}_E | D_\lambda, \lambda, \vec{\sigma})$ and $\mathcal{N}(\vec{B}_\Lambda | \tilde{D}_\lambda, \tilde{\lambda}, \vec{\sigma})$, respectively. Because the priors are uniform and sufficiently wide compared to the posterior support, they are effectively constant over the ranges shown in Figs. 1 and 3 and are therefore not overlaid in the corner plots.

To sample parameter distributions that accurately reflect experimental uncertainties, we employ MCMC techniques—specifically, a robust sampler optimized for highly anisotropic distributions—as implemented in the emcee package [89]. This approach enables a thorough exploration of the parameter space, yielding statistically robust estimates of uncertainties and correlations among the LECs, the cutoffs, and predicted energies.

APPENDIX B: NQS ARCHITECTURE, BACKFLOW TRANSFORMATIONS, AND TRAINING DETAILS

To clarify the notation, we denote the full set of coordinates as

$$X = (\mathbf{x}_1, \dots, \mathbf{x}_{A-1}, \mathbf{x}_A), \quad (\text{B1})$$

where each $\mathbf{x}_i = (\mathbf{r}_i, \mathbf{s}_i)$ comprises the spatial Cartesian coordinates $\mathbf{r}_i$ and the z components of the spin-isospin degrees of freedom, $\mathbf{s}_i = (s_{zi}, t_{zi})$, for the i th particle. As a convention, we take $t_z = 1$ for protons, $t_z = -1$ for neutrons, and $t_z = 0$ for the Λ particle. We first focus on the pure nucleonic sector and collectively denote the coordinates of the nucleons by

$$X_N = (\mathbf{x}_1, \dots, \mathbf{x}_{A-1}).$$

1. Jastrow and Pfaffian parametrizations

The amplitude of the Pfaffian-Jastrow ansatz for the nucleonic sector can be schematically written as [46]

$$\langle X_N | \Psi_{PJ} \rangle \equiv \Psi_{PJ}(X_N) = e^{J(X_N)} \times \text{Pf}[\Phi(X_N)], \quad (\text{B2})$$

where the complex-valued, permutation-invariant Jastrow factor is defined as

$$J(X_N) = a \tanh\left(\frac{U_J(X_N)}{a}\right) + i V_J(X_N). \quad (\text{B3})$$

Here, $U_J(X_N)$ and $V_J(X_N)$ are real-valued functions representing the logarithmic amplitude and phase, respectively. The parameter a acts as a cutoff that regularizes the growth of U_J and helps mitigate potential runaway instabilities. Following Ref. [90], we choose $a = 8$, which allows a maximum relative magnitude variation of approximately 10^{-7} . We enforce both $U_J(X_N)$ and $V_J(X_N)$ to be permutation invariant by employing the Deep Sets architecture [91,92] with *logsumexp* pooling:

$$U_J(X_N) = \rho_U \left[\log \left(\sum_{i \neq j=1}^{A-1} \exp(\phi_U(\mathbf{x}_i, \mathbf{x}_j)) \right) \right],$$

$$V_J(X_N) = \rho_V \left[\log \left(\sum_{i \neq j=1}^{A-1} \exp(\phi_V(\mathbf{x}_i, \mathbf{x}_j)) \right) \right], \quad (\text{B4})$$

where i runs over all the NN pairs. Both $\phi_{U,V}$ and $\rho_{U,V}$ are dense feedforward neural networks, each comprising two hidden layers with 32 nodes. The latent space—the output dimension of $\phi_{U,V}$ and the input dimension of $\rho_{U,V}$ —is set to 16; a detailed ablation analysis is presented below in Appendix B 4.

The antisymmetric part of the nucleonic wave function is given by the Pfaffian of a skew-symmetric matrix, which has proven effective in capturing strong pairing correlations across multiple spin-isospin channels in both ultracold Fermi gases [45] and low-density nuclear matter [46]. For nuclei with an even number of nucleons, we define the skew-symmetric matrix as

$$\Phi(X_N) = \begin{bmatrix} 0 & \phi(\mathbf{x}_1, \mathbf{x}_2) & \dots & \phi(\mathbf{x}_1, \mathbf{x}_{A-1}) \\ \phi(\mathbf{x}_2, \mathbf{x}_1) & 0 & \dots & \phi(\mathbf{x}_2, \mathbf{x}_{A-1}) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(\mathbf{x}_{A-1}, \mathbf{x}_1) & \phi(\mathbf{x}_{A-1}, \mathbf{x}_2) & \dots & 0 \end{bmatrix}. \quad (\text{B5})$$

To ensure skew symmetry, the pairing orbital is written as

$$\phi(\mathbf{x}_i, \mathbf{x}_j) = \eta(\mathbf{x}_i, \mathbf{x}_j) - \eta(\mathbf{x}_j, \mathbf{x}_i), \quad (\text{B6})$$

so that $\phi(\mathbf{x}_i, \mathbf{x}_j) = -\phi(\mathbf{x}_j, \mathbf{x}_i)$. In the above equation, $\eta(\mathbf{x}_i, \mathbf{x}_j)$ is a complex-valued function, whose logarithm is regularized as in Eq. (B3):

$$\eta(\mathbf{x}_i, \mathbf{x}_j) = \exp \left[a \tanh \left(\frac{u_\eta(\mathbf{x}_i, \mathbf{x}_j)}{a} \right) + i\pi v_\eta(\mathbf{x}_i, \mathbf{x}_j) \right]. \quad (\text{B7})$$

The real-valued feedforward neural networks u_η and v_η , which encode the logarithmic amplitude and phase of η , each include two hidden layers with 16 nodes, and their outputs are one dimensional.

Nuclei with an odd number of nucleons are treated by adding to the Pfaffian an unpaired single-particle orbital [45]. Thus, we extend $\Phi(X_N)$ to a $(2N+1) \times (2N+1)$ skew-symmetric matrix $\tilde{\Phi}(X_N)$ by introducing an additional row and column:

$$\tilde{\Phi} = \begin{bmatrix} \Phi(X_N) & \mathbf{u} \\ -\mathbf{u}^T & 0 \end{bmatrix}, \quad (\text{B8})$$

where the vector $\mathbf{u} \in \mathbb{C}^N$ is defined as

$$\mathbf{u} = \begin{bmatrix} \psi(\mathbf{x}_1) \\ \psi(\mathbf{x}_2) \\ \vdots \\ \psi(\mathbf{x}_N) \end{bmatrix}. \quad (\text{B9})$$

Consistent with the complex-valued functions above, we use two separate real-valued neural networks to parametrize the logarithm of $\psi(\mathbf{x}_i)$:

$$\psi(\mathbf{x}_i) = \exp \left[a \tanh \left(\frac{u_\psi(\mathbf{x}_i, \mathbf{x}_j)}{a} \right) + i\pi v_\psi(\mathbf{x}_i, \mathbf{x}_j) \right]. \quad (\text{B10})$$

We again set $a = 8$ to mitigate potential runaway instabilities. Both u_ψ and v_ψ are feedforward neural networks with two hidden layers containing 16 nodes each, and their outputs are one dimensional.

Since strong and electromagnetic interactions do not mix nucleons with Λ particles, the latter can be assumed to be distinguishable, removing the requirement for the total wave function to be antisymmetric under exchange $\mathbf{x}_i \leftrightarrow \mathbf{x}_\Lambda$. Hence, in analogy with the separation of spin-up and spin-down electrons in the presence of purely Coulomb interactions, the total wave function of single- Λ hypernuclei can be written as

$$\langle X | \Psi_{PJ}^\Lambda \rangle \equiv \Psi_{PJ}^\Lambda(X) = e^{J(X)} \times \psi_\Lambda(\mathbf{x}_\Lambda) \times \text{Pf}[\Phi(X_N)], \quad (\text{B11})$$

where $\psi_\Lambda(\mathbf{x}_\Lambda)$ is a generalized single-particle orbital for the Λ particle, whose logarithm is parametrized as in Eq. (B10). To incorporate correlations among the nucleons and the Λ , the permutation-invariant Jastrow factor now depends on the coordinates of all particles, and the sum in Eq. (B3) runs over both NN and N Λ pairs. The neural networks $U_J(X)$ and $V_J(X)$ can distinguish nucleons from the Λ because we assign the latter the auxiliary isospin value $t_z = 0$, which differs from the nucleon values $t_z = \pm \frac{1}{2}$.

2. Backflow transformations

We employ a message-passing neural network (MPNN) backflow transformation that takes as input single-particle “visible” features, defined as $\mathbf{v}_i = (\bar{\mathbf{r}}_i, \mathbf{s}_i)$. Unlike in periodic systems, where only the spin and isospin of particle i are used as inputs, here the Cartesian coordinates of the nucleons are also included. To automatically remove spurious center-of-mass contributions from all observables [93], we define the intrinsic spatial coordinates as $\bar{\mathbf{r}}_i = \mathbf{r}_i - \mathbf{R}_{\text{CM}}$, where $\mathbf{R}_{\text{CM}}$ denotes the center of mass of the nucleus (or hypernucleus). The “visible” pairwise features encode spatial as well as spin-isospin coordinates and are defined by $\mathbf{v}_{ij} = [(\mathbf{r}_i - \mathbf{r}_j), \sqrt{(\mathbf{r}_i - \mathbf{r}_j)^2}, \mathbf{s}_i, \mathbf{s}_j]$.

The initial hidden features for the nodes and the edges are obtained by concatenating the original and transformed single-particle and two-particle features, respectively, as

$$\mathbf{h}_i^0 = [\mathbf{v}_i, f_A(\mathbf{v}_i)], \quad \mathbf{h}_{ij}^0 = [\mathbf{v}_{ij}, f_B(\mathbf{v}_{ij})].$$

The main role of the two fully connected neural networks, f_A and f_B , is to preprocess the input coordinates. They also facilitate the implementation of the MPNN by ensuring that the dimensionality of the hidden features $\mathbf{h}_i^t$ and $\mathbf{h}_{ij}^t$ remains independent of t . These two networks each consist of two fully connected hidden layers with 22 nodes each.

The MPNN update is performed iteratively for $t = 0, \dots, T-1$ layers. Every time, information between the node and edge features is exchanged through a so-called message:

$$\mathbf{m}_{ij}^t = f_M^t(\mathbf{h}_i^{t-1}, \mathbf{h}_{ij}^{t-1}, \mathbf{h}_j^{t-1}).$$

For each particle i , the relevant messages are collected and pooled to eliminate any ordering with respect to the other particles $j \neq i$. As in Ref. [45], we use logsumexp pooling, a smooth alternative to max pooling:

$$\mathbf{m}_i^t = \log \left(\sum_{j \neq i} \exp(\mathbf{m}_{ij}^t) \right).$$

The hidden node and edge features are then updated as

$$\begin{aligned} \mathbf{h}_i^t &= (\mathbf{v}_i, f_F^t(\mathbf{h}_i^{t-1}, \mathbf{m}_i^t)), \\ \mathbf{h}_{ij}^t &= (\mathbf{v}_{ij}, f_G^t(\mathbf{h}_{ij}^{t-1}, \mathbf{m}_{ij}^t)). \end{aligned}$$

The functions f_M^t , f_F^t , and f_G^t are distinct fully connected neural networks with output dimensions matching those of f_A and f_B , and, like the latter, consist of two fully connected hidden layers with 22 nodes each. Including concatenated skip connections to the visible features ensures that the signal from the raw input remains accessible even as the MPNN depth T increases. Finally, after the T th iteration, we aggregate the hidden node and edge features into single-particle and pairwise feature vectors:

$$\begin{aligned} \mathbf{g}_{ij} &= (\mathbf{h}_i^T, \mathbf{h}_j^T, \mathbf{h}_{ij}^T), \\ \mathbf{g}_i &= \log \left(\sum_{j \neq i} \exp(\mathbf{g}_{ij}) \right). \end{aligned}$$

By construction, $\mathbf{g}_i$ is permutation equivariant with respect to the original input $\mathbf{x}_i$, meaning that it depends on $\mathbf{x}_i$ and is invariant under exchanges $\mathbf{x}_j \leftrightarrow \mathbf{x}_k$ for all $j, k \neq i$. Similarly,

$\mathbf{g}_{ij}$ is invariant under exchanges $\mathbf{x}_l \leftrightarrow \mathbf{x}_m$ for all $l, m \neq i, j$. We use $T \in \{1, 2\}$ message-passing layers. Consistent with related continuous-space studies [45], $T=1$ already captures most of the energy gain, while $T=2$ yields small, system-dependent improvements. Ablation studies on both the block depth and block width are discussed in Appendix B 4.

The pairwise feature vector $\mathbf{g}_{ij}$ serves as input to both the Pfaffian and the Jastrow as it replaces the original pair coordinates $(\mathbf{x}_i, \mathbf{x}_j)$ with $\mathbf{g}_{ij}$. Note that the Pfaffian includes only nucleon-nucleon pairs, while the Jastrow also includes nucleon- Λ pairs. Additionally, the permutation-equivariant output feature $\mathbf{g}_\Lambda$ is used as input to the Λ single-particle orbital Ψ_Λ , remaining invariant under nucleon permutations while not imposing $N\Lambda$ exchange symmetry, consistent with the Λ 's distinguishability.

3. Training algorithm

The optimal parameters entering the NQS are found exploiting the variational principle

$$\frac{\langle \Psi_V | H | \Psi_V \rangle}{\langle \Psi_V | \Psi_V \rangle} = E_V \geq E_0, \quad (\text{B12})$$

where E_0 is exact the ground-state energy of the system. The Metropolis-Hastings Monte Carlo algorithm is used to evaluate the variational energy E_V by sampling the spatial and spin-isospin coordinates [43]. As standard practice, to prevent the nucleons and the Λ particle from drifting away from the center of mass of the system, we add a mild Gaussian confinement:

$$\Psi_{p,j}^\Lambda(X) \rightarrow \Psi_{p,j}^\Lambda(X) \exp \left[-\alpha \sum_i |\mathbf{r}_i - \mathbf{R}_{\text{CM}}|^2 \right]. \quad (\text{B13})$$

We choose α small enough to stabilize the early stages of training without biasing observables, typically in the range 0.02–0.04.

To minimize the VMC energy, we use the stochastic-reconfiguration (SR) algorithm [94] with the RMSProp-inspired regularization introduced in Ref. [44]. At each iteration, the parameters are updated as

$$\mathbf{p} \leftarrow \mathbf{p} - \eta G^{-1} \nabla_{\mathbf{p}} E, \quad (\text{B14})$$

where η is a constant learning rate and G is the quantum geometric tensor [95].

To determine the convergence of the optimization, we monitor several complementary signals:

(1) We study block-averaged plateaus of the energy and of the local-energy variance, and require flat relative improvement $|\Delta E|/|E|$ over a sliding window. In practice, we require $|\Delta E|/|E| < 5 \times 10^{-4}$ over the last 200–500 SR steps and a variance plateau within 1σ .

(2) Using independent windows of the stored walk, we monitor improvements associated with the tentative update $\mathbf{p} + \Delta \mathbf{p}$. In particular, we compute the energy difference ΔE_{corr} , requiring it at convergence to be consistent zero within statistical errors. To reduce the latter, we estimate it by reweighting the same stored configurations with weights $w = |\Psi(\mathbf{p} + \Delta \mathbf{p})/\Psi(\mathbf{p})|^2$. Clearly, $\Delta E_{\text{corr}} \approx 0$ indicates no

TABLE VI. Ablation studies on ${}^7_\Lambda \text{Li}_{\frac{3}{2}}$ carry out with the same number of Monte Carlo configurations. The columns list the hidden neurons per layer for the Deep Sets/MPNN widths and the message-passing depth T .

Config	Deep Sets	MPNN	T	$E({}^7_\Lambda \text{Li})$ (MeV)
A	22	22	1	-35.9060 ± 0.0091
B	22	32	1	-35.8892 ± 0.0090
C	32	22	1	-35.8941 ± 0.0096
D	32	32	2	-35.9070 ± 0.0099

statistically significant change, while $\Delta E_{\text{corr}} < 0$ signals an energy decrease.

We additionally compute the quantum fidelity between the original and updated state:

$$F_{\text{dp}} = \frac{|\langle \Psi(\mathbf{p} + \Delta \mathbf{p}) | \Psi(\mathbf{p}) \rangle|^2}{\|\Psi(\mathbf{p} + \Delta \mathbf{p})\|^2 \|\Psi(\mathbf{p})\|^2}. \quad (\text{B15})$$

A necessary condition to achieve convergence is achieving $F_{\text{dp}} \gtrsim 0.9995$ in the last 200–500 SR steps.

(3) We monitor block-averaged expectation values of the targeted quantum numbers and require stationarity within statistical errors. For $Q \in \{S^2, L^2\}$, we demand (a) $|\langle Q \rangle - Q_{\text{target}}| < \varepsilon_Q$ and (b) a 1σ plateau over the last 200–500 SR steps, with typical tolerances $\varepsilon_{Q^2}/[Q(Q+1)] \lesssim 1\%$. We likewise track L_x, L_y, L_z , and radii $\langle r^2 \rangle$ (total, core, p , n , Λ) with criteria analogous to energy convergence.

(4) We require five independent initializations of the NQS to provide the same converged ground-state energy within statistical errors.

4. Ablation studies

We carried out ablation tests on the hypernuclei by varying: (1) the number of neurons per hidden layer in the Deep Sets heads ($22 \leftrightarrow 32$ at fixed latent dimension 16); (2) the number of neurons per hidden layer in the MPNN blocks ($22 \leftrightarrow 32$); and (3) the message-passing depth ($T=1 \leftrightarrow 2$). For the ${}^7_\Lambda \text{Li}_{\frac{3}{2}}$ study reported in Table VI, items (1), (2), and (3) were tested both individually and jointly. At fixed sampling (same walkers and steps), total energies remained within $1\text{--}2\sigma$ of each other; unless noted otherwise, we therefore retain the compact settings and use $T=2$ only when it yields a measurable decrease in the variational energy. These trends are consistent with related continuous-space NQS studies [45], which report weak sensitivity to moderate changes in depth/width.

[1] E. Epelbaum, H.-W. Hammer, and U.-G. Meissner, Modern theory of nuclear forces, *Rev. Mod. Phys.* **81**, 1773 (2009).

[2] R. Machleidt and D. R. Entem, Chiral effective field theory and nuclear forces, *Phys. Rep.* **503**, 1 (2011).

- [3] H. W. Hammer, S. König, and U. van Kolck, Nuclear effective field theory: Status and perspectives, *Rev. Mod. Phys.* **92**, 025004 (2020).
- [4] I. Svensson, A. Ekström, and C. Forssén, Bayesian parameter estimation in chiral effective field theory using the Hamiltonian Monte Carlo method, *Phys. Rev. C* **105**, 014004 (2022).
- [5] J. M. Bub, M. Piarulli, R. J. Furnstahl, S. Pastore, and D. R. Phillips, Bayesian analysis of nucleon-nucleon scattering data in pionless effective field theory, *Phys. Rev. C* **111**, 034005 (2025).
- [6] G. Hagen, T. Papenbrock, M. Hjorth-Jensen, and D. J. Dean, Coupled-cluster computations of atomic nuclei, *Rep. Prog. Phys.* **77**, 096302 (2014).
- [7] H. Hergert, S. K. Bogner, T. D. Morris, A. Schwenk, and K. Tsukiyama, The in-medium similarity renormalization group: A novel *ab initio* method for nuclei, *Phys. Rep.* **621**, 165 (2016).
- [8] W. H. Dickhoff and C. Barbieri, Selfconsistent Green's function method for nuclei and nuclear matter, *Prog. Part. Nucl. Phys.* **52**, 377 (2004).
- [9] V. Soma, C. Barbieri, and T. Duguet, *Ab initio* Gorkov-Green's function calculations of open-shell nuclei, *Phys. Rev. C* **87**, 011303(R) (2013).
- [10] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa, Quantum Monte Carlo methods for nuclear physics, *Rev. Mod. Phys.* **87**, 1067 (2015).
- [11] J. E. Lynn, I. Tews, S. Gandolfi, and A. Lovato, Quantum Monte Carlo methods in nuclear physics: Recent advances, *Annu. Rev. Nucl. Part. Sci.* **69**, 279 (2019).
- [12] M. Danysz and J. Pniewski, Delayed disintegration of a heavy nuclear fragment: I, *London, Edinburgh Dublin Philos. Mag. J. Sci.* **44**, 348 (1953).
- [13] M. May, *et al.*, Observation of hypernuclear gamma-ray transitions in ${}^7_\Lambda\text{Li}$ and ${}^9_\Lambda\text{Be}$, *Phys. Rev. Lett.* **51**, 2085 (1983).
- [14] K. Tanida, *et al.*, Measurement of the $B(E2)$ of hyper ${}^7_\Lambda\text{Li}$ and shrinkage of the hypernuclear size, *Phys. Rev. Lett.* **86**, 1982 (2001).
- [15] H. Akikawa, *et al.*, Hypernuclear fine structure in ${}^9_\Lambda\text{Be}$, *Phys. Rev. Lett.* **88**, 082501 (2002).
- [16] O. Hashimoto and H. Tamura, Spectroscopy of Λ hypernuclei, *Prog. Part. Nucl. Phys.* **57**, 564 (2006).
- [17] T. O. Yamamoto, *et al.* (J-PARC E13 Collaboration), Observation of spin-dependent charge symmetry breaking in ΛN interaction: Gamma-ray spectroscopy of ${}^4_\Lambda\text{He}$, *Phys. Rev. Lett.* **115**, 222501 (2015).
- [18] T. Gogami, *et al.*, High resolution spectroscopic study of ${}^{10}_\Lambda\text{Be}$, *Phys. Rev. C* **93**, 034314 (2016).
- [19] A. Gal, E. V. Hungerford, and D. J. Millener, Strangeness in nuclear physics, *Rev. Mod. Phys.* **88**, 035004 (2016).
- [20] I. Bombaci, The hyperon puzzle in neutron stars, *JPS Conf. Proc.* **17**, 101002 (2017).
- [21] K. A. Maslov, E. E. Kolomeitsev, and D. N. Voskresensky, Solution of the hyperon puzzle within a relativistic mean-field model, *Phys. Lett. B* **748**, 369 (2015).
- [22] P. Demorest, T. Pennucci, S. Ransom, M. Roberts, and J. Hessels, Shapiro delay measurement of a two solar mass neutron star, *Nature (London)* **467**, 1081 (2010).
- [23] J. Antoniadis, *et al.*, A massive pulsar in a compact relativistic binary, *Science* **340**, 1232322 (2013).
- [24] H. T. Cromartie, *et al.* (NANOGrav), Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar, *Nat. Astron.* **4**, 72 (2019).
- [25] E. Fonseca, *et al.*, Refined mass and geometric measurements of the high-mass PSR J0740+6620, *Astrophys. J. Lett.* **915**, L12 (2021).
- [26] D. Lonardoni, A. Lovato, S. Gandolfi, and F. Pederiva, Hyperon puzzle: Hints from quantum Monte Carlo calculations, *Phys. Rev. Lett.* **114**, 092301 (2015).
- [27] J. Haidenbauer, U. G. Meißner, N. Kaiser, and W. Weise, Lambda-nuclear interactions and hyperon puzzle in neutron stars, *Eur. Phys. J. A* **53**, 121 (2017).
- [28] A. Nogga, H. Kamada, and W. Gloeckle, The hypernuclei ${}^4_\Lambda\text{He}$ and ${}^4_\Lambda\text{H}$: Challenges for modern hyperon-nucleon forces, *Phys. Rev. Lett.* **88**, 172501 (2002).
- [29] E. Hiyama, M. Kamimura, T. Motoba, T. Yamada, and Y. Yamamoto, $\Lambda - \Sigma$ conversion in ${}^4_\Lambda\text{He}$ and ${}^4_\Lambda\text{H}$ based on a four-body calculation, *Phys. Rev. C* **65**, 011301(R) (2001).
- [30] H. Nemura, Y. Akaishi, and Y. Suzuki, *Ab initio* approach to s -shell hypernuclei ${}^3_\Lambda\text{H}$, ${}^4_\Lambda\text{H}$, ${}^4_\Lambda\text{He}$, and ${}^5_\Lambda\text{H}$ with a ΛN - ΣN interaction, *Phys. Rev. Lett.* **89**, 142504 (2002).
- [31] I. N. Filikhin and A. Gal, Faddeev-Yakubovsky search for ${}^4_{\Lambda\Lambda}\text{H}$, *Phys. Rev. Lett.* **89**, 172502 (2002).
- [32] H. Garcilazo, A. Valcarce, and T. Fernandez-Carames, ΛNN and ΣNN systems at threshold: II. The effect of D waves, *Phys. Rev. C* **76**, 034001 (2007).
- [33] L. Contessi, N. Barnea, and A. Gal, Resolving the Λ hypernuclear overbinding problem in pionless effective field theory, *Phys. Rev. Lett.* **121**, 102502 (2018).
- [34] D. Lonardoni, S. Gandolfi, and F. Pederiva, Effects of the two-body and three-body hyperon-nucleon interactions in Λ -hypernuclei, *Phys. Rev. C* **87**, 041303(R) (2013).
- [35] D. Lonardoni, F. Pederiva, and S. Gandolfi, Accurate determination of the interaction between Λ hyperons and nucleons from auxiliary field diffusion Monte Carlo calculations, *Phys. Rev. C* **89**, 014314 (2014).
- [36] R. Wirth, D. Gazda, P. Navrátil, A. Calci, J. Langhammer, and R. Roth, *Ab initio* description of p -shell hypernuclei, *Phys. Rev. Lett.* **113**, 192502 (2014).
- [37] R. Wirth and R. Roth, Induced hyperon-nucleon-nucleon interactions and the hyperon puzzle, *Phys. Rev. Lett.* **117**, 182501 (2016).
- [38] R. Wirth and R. Roth, Light neutron-rich hypernuclei from the importance-truncated no-core shell model, *Phys. Lett. B* **779**, 336 (2018).
- [39] R. Wirth and R. Roth, Similarity renormalization group evolution of hypernuclear Hamiltonians, *Phys. Rev. C* **100**, 044313 (2019).
- [40] M. Knöll and R. Roth, Hyperon-nucleon interaction constrained by light hypernuclei, *Phys. Lett. B* **846**, 138258 (2023).
- [41] H. Le, J. Haidenbauer, U.-G. Meißner, and A. Nogga, Light Λ hypernuclei studied with chiral hyperon-nucleon and hyperon-nucleon-nucleon forces, *Phys. Rev. Lett.* **134**, 072502 (2025).
- [42] L. H. Avraham and B. Bazak, Few-nucleon systems within finite-cutoff pionless EFT, *Phys. Lett. B* **872**, 140100 (2026).
- [43] C. Adams, G. Carleo, A. Lovato, and N. Rocco, Variational Monte Carlo calculations of $A \leq 4$ nuclei with an artificial neural-network correlator ansatz, *Phys. Rev. Lett.* **127**, 022502 (2021).

- [44] A. Lovato, C. Adams, G. Carleo, and N. Rocco, Hidden-nucleons neural-network quantum states for the nuclear many-body problem, *Phys. Rev. Res.* **4**, 043178 (2022).
- [45] J. Kim, G. Pescia, B. Fore, J. Nys, G. Carleo, S. Gandolfi, M. Hjorth-Jensen, and A. Lovato, Neural-network quantum states for ultra-cold Fermi gases, *Commun. Phys.* **7**, 148 (2024).
- [46] B. Fore, J. Kim, M. Hjorth-Jensen, and A. Lovato, Investigating the crust of neutron stars with neural-network quantum states, *Commun. Phys.* **8**, 108 (2025).
- [47] R. Schiavilla, L. Girlanda, A. Gnech, A. Kievsky, A. Lovato, L. E. Marcucci, M. Piarulli, and M. Viviani, Two- and three-nucleon contact interactions and ground-state energies of light- and medium-mass nuclei, *Phys. Rev. C* **103**, 054003 (2021).
- [48] N. Auerbach, J. Hüfner, A. K. Kerman, and C. M. Shakin, A theory of isobaric analog resonances, *Rev. Mod. Phys.* **44**, 48 (1972).
- [49] L. Contessi, A. Lovato, F. Pederiva, A. Roggero, J. Kirscher, and U. van Kolck, Ground-state properties of ^4He and ^{16}O extrapolated from lattice QCD with pionless EFT, *Phys. Lett. B* **772**, 839 (2017).
- [50] N. Barnea, L. Contessi, D. Gazit, F. Pederiva, and U. Van Kolck, Effective field theory for lattice nuclei, *Phys. Rev. Lett.* **114**, 052501 (2015).
- [51] R. W. Hackenburg, Neutron-proton effective range parameters and zero-energy shape dependence, *Phys. Rev. C* **73**, 044002 (2006).
- [52] G. Alexander, U. Karshon, A. Shapira, G. Yekutieli, R. Engelmann, H. Filthuth, and W. Lughofer, Study of the $\Lambda - N$ system in low-energy $\Lambda - p$ elastic scattering, *Phys. Rev.* **173**, 1452 (1968).
- [53] P. F. Bedaque and U. van Kolck, Effective field theory for few nucleon systems, *Annu. Rev. Nucl. Part. Sci.* **52**, 339 (2002).
- [54] A. Bansal, S. Binder, A. Ekström, G. Hagen, G. R. Jansen, and T. Papenbrock, Pion-less effective field theory for atomic nuclei and lattice nuclei, *Phys. Rev. C* **98**, 054301 (2018).
- [55] I. Stetcu, B. R. Barrett, and U. van Kolck, No-core shell model in an effective-field-theory framework, *Phys. Lett. B* **653**, 358 (2007).
- [56] B.-N. Lu, N. Li, S. Elhatisari, D. Lee, E. Epelbaum, and U.-G. Meißner, Essential elements for nuclear binding, *Phys. Lett. B* **797**, 134863 (2019).
- [57] A. Gnech, B. Fore, A. J. Tropiano, and A. Lovato, Distilling the essential elements of nuclear binding via neural-network quantum states, *Phys. Rev. Lett.* **133**, 142501 (2024).
- [58] F. Calogero, A novel approach to elementary scattering theory, *II Nuovo Cimento* **27**, 261 (1963).
- [59] International Atomic Energy Agency (IAEA), Live chart of nuclides, <https://www-nds.iaea.org/relnsd/vcharthtml/VChartHTML.html>.
- [60] A. Ekström and L. Platter, Quantifying the breakdown scale of pionless effective field theory, *Phys. Lett. B* **860**, 139207 (2025).
- [61] P. Cook, D. Jammooa, M. Hjorth-Jensen, D. D. Lee, and D. Lee, Parametric matrix models, *Nature Commun.* **16**, 5929 (2025).
- [62] D. Frame, R. He, I. Ipsen, D. Lee, D. Lee, and E. Rrapaj, Eigenvector continuation with subspace learning, *Phys. Rev. Lett.* **121**, 032501 (2018).
- [63] R. Rende, L. L. Viteritti, F. Becca, A. Scardicchio, A. Laio, and G. Carleo, Foundation neural-networks quantum states as a unified ansatz for multiple Hamiltonians, *Nat. Commun.* **16**, 7213 (2025).
- [64] Y. Suzuki and K. Varga, *Stochastic Variational Approach to Quantum Mechanical Few-Body Problems*, Lecture Notes in Physics New Series M, Monographs (Springer, Berlin, Heidelberg, New York, Barcelona, Hong Kong, London, Milan, Paris, Singapore, Tokyo, 1998), Vol. 54.
- [65] P. Eckert and P. Achenbach, Chart of hypernucleides—hypernuclear structure and decay data, Mainz Hypernuclei Database, <https://hypernuclei.kph.uni-mainz.de/>, accessed 1 April 2025.
- [66] A. Ekström, G. R. Jansen, K. A. Wendt, G. Hagen, T. Papenbrock, B. D. Carlsson, C. Forssén, M. Hjorth-Jensen, P. Navrátil, and W. Nazarewicz, Accurate nuclear radii and binding energies from a chiral interaction, *Phys. Rev. C* **91**, 051301(R) (2015); **109**, 059901(E) (2024).
- [67] J. C. Owen, Independent pair correlations in a cluster expansion formalism, *Phys. Lett. B* **77**, 9 (1978).
- [68] R. P. Feynman and M. Cohen, Energy spectrum of the excitations in liquid helium, *Phys. Rev.* **102**, 1189 (1956).
- [69] D. Luo and B. K. Clark, Backflow transformations via neural networks for quantum many-body wave functions, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [70] J. Hermann, Z. Schätzle, and F. Noé, Deep-neural-network solution of the electronic Schrödinger equation, *Nat. Chem.* **12**, 891 (2020).
- [71] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, *Ab initio* solution of the many-electron Schrödinger equation with deep neural networks, *Phys. Rev. Res.* **2**, 033429 (2020).
- [72] G. Pescia, J. Nys, J. Kim, A. Lovato, and G. Carleo, Message-passing neural quantum states for the homogeneous electron gas, *Phys. Rev. B* **110**, 035108 (2024).
- [73] L. Contessi, M. Schäfer, A. Gnech, A. Lovato, and U. van Kolck, A renormalizable theory for not-so-light nuclei, [arXiv:2505.09299](https://arxiv.org/abs/2505.09299).
- [74] M. Gattobigio, A. Kievsky, and M. Viviani, Embedding nuclear physics inside the unitary-limit window, *Phys. Rev. C* **100**, 034004 (2019).
- [75] S. König, Energies and radii of light nuclei around unitarity, *Eur. Phys. J. A* **56**, 113 (2020).
- [76] F. Hildenbrand, S. Elhatisari, Z. Ren, and U.-G. Meißner, Towards hypernuclei from nuclear lattice effective field theory, *Eur. Phys. J. A* **60**, 215 (2024).
- [77] H. Le, J. Haidenbauer, U.-G. Meißner, and A. Nogga, Separation energies of light Λ hypernuclei and their theoretical uncertainties, *Eur. Phys. J. A* **60**, 3 (2024).
- [78] D. Gazda, T. Y. Htun, and C. Forssén, Nuclear physics uncertainties in light hypernuclei, *Phys. Rev. C* **106**, 054001 (2022).
- [79] E. Hiyama, M. Kamimura, T. Motoba, T. Yamada, and Y. Yamamoto, Three body model study of $A = 6$, $A = 7$ hypernuclei: Halo and skin structures, *Phys. Rev. C* **53**, 2075 (1996).
- [80] E. Hiyama, M. Kamimura, K. Miyazaki, and T. Motoba, Gamma transitions in $A = 7$ hypernuclei and a possible derivation of hypernuclear size, *Phys. Rev. C* **59**, 2351 (1999).
- [81] J. M. Yao, Z. P. Li, K. Hagino, M. T. Win, Y. Zhang, and J. Meng, Impurity effect of lambda hyperon on collective excitations of atomic nuclei, *Nucl. Phys. A* **868–869**, 12 (2011).

- [82] D. Lonardoni and F. Pederiva, Medium-mass hypernuclei and the nucleon-isospin dependence of the three-body hyperon-nucleon-nucleon force, [arXiv:1711.07521](#).
- [83] M. Piarulli, *et al.*, Light-nuclei spectra from chiral dynamics, *Phys. Rev. Lett.* **120**, 052503 (2018).
- [84] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* Adaptive Computation and Machine Learning (MIT Press, Cambridge, Massachusetts, 2008), 3rd ed.
- [85] D. Duvenaud, Automatic model construction with Gaussian processes, Ph.D. thesis, University of Cambridge, 2014.
- [86] D. J. C. MacKay, Bayesian non-linear modeling for the prediction competition, in *Maximum Entropy and Bayesian Methods*, edited by G. R. Heidbreder (Springer Netherlands, Dordrecht, 1996), pp. 221–234.
- [87] S. Acharya, *et al.* (ALICE Collaboration), Measurement of the lifetime and Λ separation energy of $^3_\Lambda\text{H}$, *Phys. Rev. Lett.* **131**, 102302 (2023).
- [88] SheffieldML, GPy: A Gaussian processes framework, 2024, <https://github.com/SheffieldML/GPy>, accessed 10 October 2024.
- [89] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [90] M. Bukov, M. Schmitt, and M. Dupont, Learning the ground state of a non-stoquastic quantum Hamiltonian in a rugged neural network landscape, *SciPost Phys.* **10**, 147 (2021).
- [91] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov, and A. Smola, Deep sets, [arXiv:1703.06114](#).
- [92] E. Wagstaff, F. B. Fuchs, M. Engelcke, I. Posner, and M. Osborne, On the limitations of representing functions on sets, [arXiv:1901.09006](#).
- [93] P. Massella, F. Barranco, D. Lonardoni, A. Lovato, F. Pederiva, and E. Vigezzi, Exact restoration of Galilei invariance in density functional calculations with quantum Monte Carlo, *J. Phys. G* **47**, 035105 (2020).
- [94] S. Sorella, Wave function optimization in the variational Monte Carlo method, *Phys. Rev. B* **71**, 241103(R) (2005).
- [95] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum natural gradient, *Quantum* **4**, 269 (2020).