

LA-UR-24-27320

Accepted Manuscript

Emulators for Scarce and Noisy Data: Application to Auxiliary-Field Diffusion Monte Carlo for Neutron Matter

Armstrong, Cassandra Lee

Somasundaram, Rahul

Giuliani, Pablo

Godbey, Kyle

Tews, Ingo

Provided by the author(s) and the Los Alamos National Laboratory (2025-11-13).

To be published in: Physical Review Letters

DOI to publisher's version: 10.1103/9928-wyjm

Permalink to record:

<https://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-24-27320>



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Emulators for Scarce and Noisy Data: Application to Auxiliary-Field Diffusion Monte Carlo for Neutron Matter

Cassandra L. Armstrong,¹ Pablo Giuliani,² Kyle Godbey,² Rahul Somasundaram,^{3,4} and Ingo Tews^{3,*}

¹*Intelligence and Space Research Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

²*Facility for Rare Isotope Beams, Michigan State University, East Lansing, Michigan 48824, USA*

³*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

⁴*Department of Physics, Syracuse University, Syracuse, New York 13244, USA*

(Dated: October 28, 2025)

Understanding the equation of state (EOS) of pure neutron matter is necessary for interpreting multimessenger observations of neutron stars. Reliable data analyses of these observations require well-quantified uncertainties for the EOS input, ideally propagating uncertainties from nuclear interactions directly to the EOS. This, however, requires calculations of the EOS for a prohibitively large number of nuclear Hamiltonians, solving the nuclear many-body problem for each one. Quantum Monte Carlo methods, such as auxiliary-field diffusion Monte Carlo (AFDMC), provide precise and accurate results for the neutron matter EOS, but they are very computationally expensive, making them unsuitable for the fast evaluations necessary for uncertainty propagation. Here, we employ parametric matrix models to develop fast emulators for AFDMC calculations of neutron matter and use them to directly propagate uncertainties of coupling constants in the Hamiltonian to the EOS. As these uncertainties include estimates of the effective field theory truncation uncertainty, this approach provides robust uncertainty estimates for use in astrophysical data analyses. This Letter will enable novel applications such as using astrophysical observations to put constraints on coupling constants for nuclear interactions.

Introduction. The equation of state (EOS) of pure neutron matter (PNM) is a key input to constrain properties of neutron stars (NSs), such as their masses, radii, and tidal deformabilities [1–4]. This is because neutron-star properties are dominated by the neutron-star core where proton fractions are small. Using observatories such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) [5, 6] or NASA’s Neutron Star Interior Composition Explorer (NICER) [7–11], one can extract neutron-star properties, and hence, constrain the EOS [10, 12–20]. The coming decades will provide even more astrophysical data from current and next-generation observatories, such as Cosmic Explorer [21, 22] and the Einstein Telescope [23, 24], ushering in a data-rich era for neutron stars that will allow us to further constrain the EOS. To properly analyze and extract robust nuclear-physics information from this multimessenger data, it is important to provide statistically meaningful EOS priors with a strong foundation in uncertainty quantification.

At low densities, the neutron-matter EOS can be calculated using nuclear Hamiltonians from chiral effective field theory (EFT). Chiral EFT is a systematic expansion for nuclear forces and describes nuclear interactions in neutron matter up to densities of about two times the nuclear saturation density, $n_{\text{sat}} \sim 0.16 \text{ fm}^{-3}$ [25–28]. Chiral EFT Hamiltonians depend on a set of parameters, called low-energy couplings (LECs), that are usually calibrated to scattering data and properties of atomic nuclei [29–33]. Using chiral EFT Hamiltonians, one can then employ a quantum many-body method to

solve for the EOS of PNM. Here, we employ accurate and precise quantum Monte Carlo (QMC) methods [34], and in particular, auxiliary-field diffusion Monte Carlo (AFDMC) [35], to obtain the energy per particle in PNM. Importantly, the use of chiral EFT interactions enables uncertainty estimates due to the truncation of the EFT expansion, reflecting our incomplete knowledge of nuclear interactions. [27, 36]. Currently, these uncertainties are estimated *a posteriori*: calculations are performed order-by-order in the EFT expansion, and an uncertainty band is estimated from these results [27, 36]. Recent AFDMC results for the EOS of PNM have employed such *a posteriori* estimates for theoretical uncertainties [37]. Nevertheless, the statistical interpretation of these error bands is either unclear [36] or distributions are assumed to be Gaussian [27].

The truncation uncertainty of the EFT expansion can also be accounted for during the fit of the LECs [38]. Then, the distributions of the LECs include not only statistical uncertainties stemming from the fit data, but also an estimate for the systematic EFT uncertainty. A robust statistical interpretation of such uncertainty bands can also be obtained by propagating these uncertainties from the LECs [38] directly to the EOS [39, 40]. In the case of the EOS of PNM, this would allow us to extract a consistent shape of the EOS uncertainty distribution at each density given an ansatz for the distribution of the LECs. This approach, however, is currently not feasible for the EOS of PNM because AFDMC calculations are very computationally expensive. For example, the calculation of the energy per particle at one fixed density and for one LEC set costs on the order of several 100,000 CPU-hours. In this Letter, we address this problem and present emulators for AFDMC calculations

* itews@lanl.gov




FIG. 1. Full propagated error bands for the energy per particle of PNM as function of number density at 68% and 90% confidence intervals. The red band shows the mean and 68% uncertainty of Ref. [37] using a Gaussian Process error prescription [27].

of pure neutron matter based on the parametric matrix model (PMM) [33, 41]. These emulators allow us to speed up AFDMC calculations by a factor of 10^8 , enabling us to rapidly carry out the many thousands of calculations necessary for error propagation from LECs to the EOS. Using these emulators, we provide reliable uncertainty bands for the EOS of PNM, and show our main result in Fig. 1.

Methods—The PMM [41] is a general purpose machine-learning algorithm that has shown excellent performance on various tasks, including image classification, curve extrapolation, and surrogate modeling [33, 42]. For the latter, it has been especially powerful for emulating eigenvalue equations such as the Schrödinger equation. In this context, it can be interpreted as a data-driven variant of the reduced-basis method that has been recently adopted by the low-energy nuclear-theory community [43–50]. The reduced-basis method identifies the physically relevant subspace spanned by solutions for different LECs through a basis expansion informed by previous high-fidelity solutions—solutions found in the full model space that are here given by highly precise and accurate results from AFDMC—and then projects the necessary equations describing the system into this subspace. In contrast, the PMM approach implicitly identifies a relevant effective low-dimensional space by exploiting the eigen-system of matrices whose elements are fit to data, which here are the high-fidelity AFDMC energies. In this Letter, we construct a PMM inspired by the structure of the chiral EFT Hamiltonian up to a chosen order in the EFT expansion, that depends on the LECs in an affine way. Generally, we define

$$\hat{H}(\{c_i\}) = H_0 + c_1 \cdot H_1 + c_2 \cdot H_2 + \dots, \quad (1)$$

where $\{c_i\}$ are our control parameters, i.e., the LECs of the Hamiltonian, and all the involved matrices have dimensions $N_{\text{dim}} \times N_{\text{dim}}$. For example, in the $N_{\text{dim}} = 2$ case and for two control parameters c_1 and c_2 we have:

$$\hat{H} = \begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix}, H_0 = \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}, H_1 = \begin{bmatrix} c & d \\ d & e \end{bmatrix}, H_2 = \begin{bmatrix} f & g \\ g & h \end{bmatrix}, \quad (2)$$

with real matrix elements $a-h$, with H_0 being a diagonal matrix containing all terms that are not dependent on the LECs, and H_1 and H_2 being symmetric matrices. The matrix elements are found through fitting a selected eigenvalue of \hat{H} to data for the observable we want to emulate, in this case, AFDMC calculations of the PNM energy. The matrices in Eq. (2) can be extended to higher matrix dimensions N_{dim} , which can improve the PMM emulator at the cost of more PMM parameters to fit. This general structure for the PMM was chosen based on previous experience with the deuteron [33] and with a toy problem discussed in the Supplemental Material [51].

To generate the training data we employ the AFDMC method. AFDMC is a QMC algorithm that solves the Schrödinger equation by performing an evolution of a trial wave function $|\psi_T\rangle$ in imaginary time τ to project out the ground state of the system $|\psi_0\rangle$ [34, 52],

$$|\psi_0\rangle = \lim_{\tau \rightarrow \infty} e^{-H\tau} |\psi_T\rangle. \quad (3)$$

The direct computation of the propagator $\exp\{-H\tau\}$ for arbitrary τ is typically impossible, but the propagator can be calculated for small imaginary times $\delta\tau = \tau/N$ with large N using Monte Carlo methods [34]. For this, the nuclear input Hamiltonian has to be local, as is the case for the various chiral EFT interactions that have been developed in the past decades for use in QMC methods [38, 53–60]. Using perturbation theory, however, it is also possible to treat small nonlocalities [61]. Here, we use the local high-cutoff chiral EFT interactions of Ref. [38] with $R_0 = 0.6$ fm at leading order (LO) with 2 LECs, next-to-leading order (NLO), and next-to-next-to-leading order (N²LO), with 9 LECs each. These interactions were fit to nucleon-nucleon scattering data over a range of cutoffs using a Bayesian fitting protocol that explicitly accounts for truncation uncertainties of the EFT expansion and provides us with full posterior distributions for all LECs. The three-nucleon interactions at N²LO are determined by pion-nucleon LECs that we take from the Roy-Steiner analysis [62] and do not vary here. In pure neutron matter, the shorter-range three-nucleon interactions do not contribute for nonlocal interactions [63], and are negligible for local interactions at high cutoffs [37]. For these interactions, we generated full AFDMC solutions for pure neutron matter for several LEC sets. At LO (NLO), we performed AFDMC simulations with 14 neutrons at nuclear saturation density for 30 (60) LEC sets. At N²LO, we have calculated 30 high-fidelity data points using simulations with 66 neutrons at 0.12 fm^{-3} , 0.16 fm^{-3} , 0.24 fm^{-3} , 0.32 fm^{-3} , and

0.48 fm⁻³, following the final setup of Ref. [37]. Importantly, the chosen cutoff $R_0 = 0.6$ fm removes regulator artifacts [57, 64]. The generation of all training data cost ~ 22 million CPU-hours.

PMMs for AFDMC—For propagation of LEC distributions to PNM, full AFDMC simulations are prohibitively expensive. Here, we apply the PMM, focusing on the lowest eigenvalue, to construct a surrogate model of AFDMC instead. We constructed PMMs at LO and NLO to understand the best setup and study the necessary amount of training data to achieve a good reproduction of AFDMC energies. We present more details for the PMM at these orders in the Supplemental material [51] but summarize the most important results here. At LO, where the PMM is described by the three terms in Eq. (1), we split the 30 high-fidelity data points into a training set with 20 points and validation set with the remaining 10 points. The matrix elements of the PMM are then fit to the training energies using SciPy’s Dual Annealing optimization routine [65], and the resulting PMM is used to predict the PNM energies for the validation LEC sets. We note that, especially for larger N_{dim} , we typically encounter a situation with a larger number of unknowns than training data. The Dual Annealing approach allows us to find solutions to this ill-posed problem that nevertheless have good extrapolation power. However, when initially setting up the PMM, we found that, instead of the error decreasing or leveling off with an increasing number of training points N_{train} , there appeared spikes with unexpectedly high percent errors for some values of N_{train} . This increase in percent error is caused by a level crossing in the resulting PMM fit that occurs at energies above those in the training set. In these cases, the training data is described by the lowest eigenvalue of \hat{H} , but the validation data is described in part by the second eigenvalue. To circumvent this problem and avoid the level crossing in the PMM fit, we tested various fitting strategies. The fitting approach that led to the best and most consistent results is defined by ensuring that the maximum energy in the training set is always included when training, regardless of the number of training points, and that the fit for $N_{\text{train}} + 1$ training points starts from the result of the fit for N_{train} . We provide more details in the Supplemental Material [51]. With the specified fitting routine, we tested the validation error for the PMM as a number of different N_{train} and N_{dim} . At LO, for $N_{\text{dim}} = 2$ the PMM improves as a function of N_{train} until $N_{\text{train}} = 4$ where the error settles around 0.03%. The error is similarly small for different $N_{\text{dim}} > 2$.

At NLO and N²LO, the number of LECs describing the nuclear Hamiltonian increases from 2 to 9. Using the fitting protocol specified above, we first built PMMs at NLO as calculations are cheaper. At NLO, we compared the performance of a PMM with 9 operator LECs and a PMM with the 6 relevant spectral LECs, see Ref. [54] for details. With 9 operator LECs, the PMM at NLO is set




FIG. 2. Convergence of the relative error of the PMM reproduction of the AFDMC validation energies with the number of training points N_{train} for $N_{\text{dim}} = 1 - 4$ at N²LO at nuclear saturation density, $n_{\text{sat}} = 0.16$ fm⁻³. The training and validation sets are the same for all N_{dim} .

up as

$$\begin{aligned} \hat{H} = & H_0 + C_S \cdot H_1 + C_T \cdot H_2 + C_1 \cdot H_3 \\ & + C_2 \cdot H_4 + C_3 \cdot H_5 + C_4 \cdot H_6 \\ & + C_5 \cdot H_7 + C_6 \cdot H_8 + C_7 \cdot H_9, \end{aligned} \quad (4)$$

and with 6 spectral LECs, the PMM is

$$\begin{aligned} \hat{H} = & H_0 + d_{11} \cdot H_1 + d_{22} \cdot H_2 + d_3 \cdot H_3 \\ & + d_4 \cdot H_4 + d_6 \cdot H_5 + d_7 \cdot H_6, \end{aligned} \quad (5)$$

where d_i are functions of C_i of Eq. (4). We found that the final performance of both PMMs at NLO is comparable but that the PMM that uses 6 spectral LECs converges faster and more consistently. This can be understood as fewer matrix elements need to be fit to data. We explored bounds on the matrix elements when fitting and found that setting the bounds on each matrix element to $[-500, 500]$ fm⁻³⁽⁻⁵⁾ for LO (NLO) contacts ensures that the eigenvalues are well reproduced and reasonable. We tested the effect of varying N_{dim} and N_{train} on the PMM performance at NLO and found that the percent error levels off at about 1% error for $N_{\text{train}} \gtrsim 20$ when $N_{\text{dim}} \geq 2$. This lower bound on the error is probably due to the statistical uncertainty of the AFDMC calculations of the same order and is small compared to uncertainties due to the truncation of the chiral expansion, which is of the order of 10% at saturation density [37]. Therefore, we do not expect this emulator uncertainty to affect the conclusions in this Letter. Results at NLO are given in the Supplemental Material [51].

With the fitting protocol described above and using a PMM for spectral LECs [see Eq. (5)], we now turn to N²LO. The calculations at this order were done with a larger number of neutrons and AFDMC walkers for 30 LEC sets. We split these 30 calculations into N_{train}

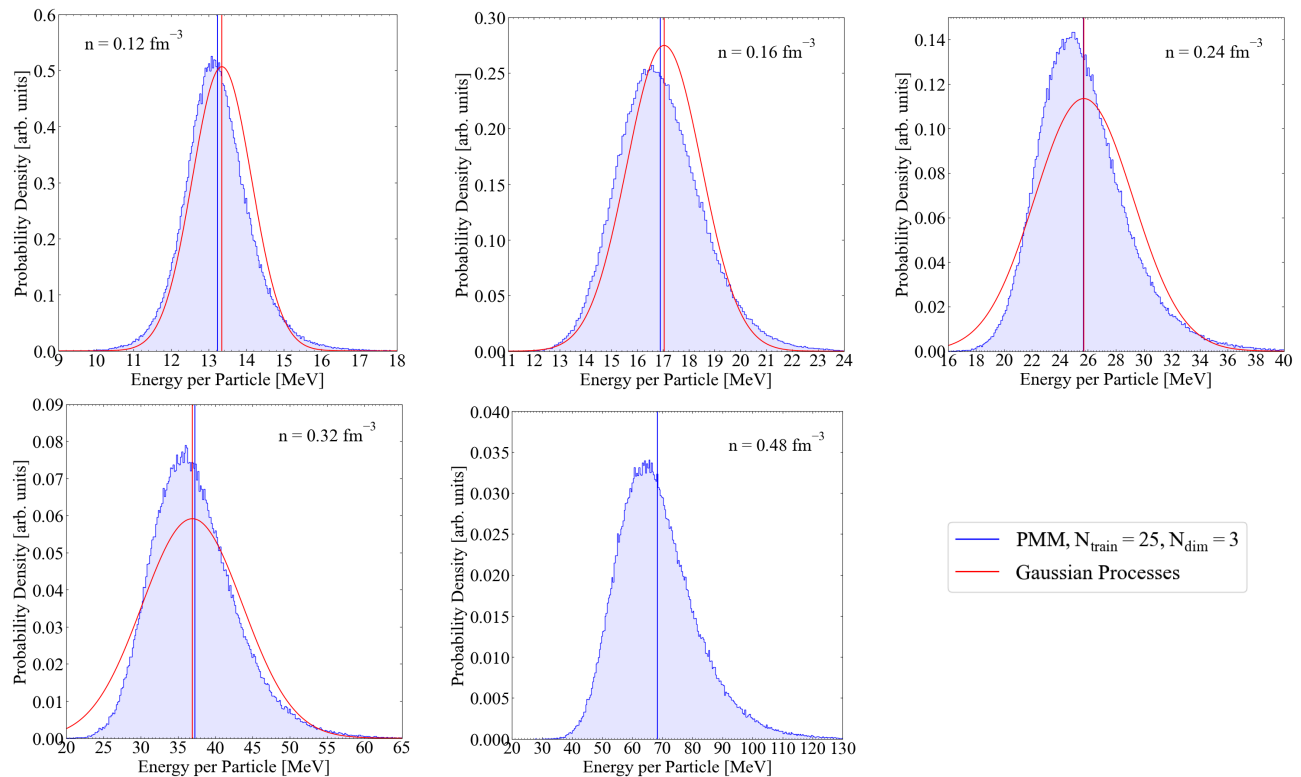


FIG. 3. Probability density functions (PDFs) for the neutron-matter energy per particle at 5 different densities when propagating full posterior distributions for all LECs. Where available, we compare the PDFs to the previous uncertainty estimation method using Gaussian Processes [27, 37]. For both distributions, we specify the mean by a vertical line.

training data and use the remaining points for validation. At $N^2\text{LO}$, we find that the percent errors again level off at around 1% for $N_{\text{train}} \gtrsim 15$ for $N_{\text{dim}} = 1 - 4$, as shown in Figure 2 for nuclear saturation density (results are similar at other densities). We found that the PMM for $N_{\text{dim}} = 5$ does not produce a good fit for a small number of training points. Because the quality of the PMM does not considerably improve for $N_{\text{dim}} > 2$, in the following we use PMMs for 25 training points and $N_{\text{dim}} = 2$ at each density. We tested these choices when propagating 250,000 LEC samples and found only negligible differences in the predicted distributions of the energy per particle for larger number of training data or $N_{\text{dim}} > 2$ except at $3n_{\text{sat}}$, where more training data led to the appearance of a level crossing. However, this density is likely above the range of validity of chiral EFT even for high-cutoff interactions, but we include it here for future checks of the range of validity of chiral EFT. We note that the PMMs are trained at each density independently. We also stress that while $N_{\text{dim}} = 1$ leads to a surprisingly good validation error in Fig. 2, this is generally not the case and depends on the order of the calculations as well as on the randomly chosen training and validation data sets. In the Supplemental Material [51], we show the case at NLO where $N_{\text{dim}} = 1$ leads to a much worse validation error. However, this is indicating that at $N^2\text{LO}$ the dependence of the energy on the LECs

is approximately linear.

EOS of PNM—Using the PMMs as specified above, we propagate 250,000 samples from the full LEC posteriors at $N^2\text{LO}$ to distributions of energies per particle in PNM at $0.75n_{\text{sat}}$, n_{sat} , $1.5n_{\text{sat}}$, $2n_{\text{sat}}$, and $3n_{\text{sat}}$. For each LEC set, we interpolate between these 5 density points using the simple function proposed in Ref. [66] and use this to compute an error band in the full density range of $1 - 3n_{\text{sat}}$. Our main result is shown in Fig. 1 up to $2n_{\text{sat}}$, where we compare to previous results using a Gaussian process (GP) to estimate the truncation error. We find that the uncertainty bands are comparable but that the GP approach leads to a slightly larger 68% confidence interval. Furthermore, we find that the error is not symmetric.

To illustrate this point, we show histograms of the energy per particle at the 5 training densities in Figure 3. We find that the full distribution of energies is mostly Gaussian at lower densities but with a slight skew toward higher energies. At higher densities, the skewness in the distribution increases. We note that the LEC distributions resemble skewed Gaussians [38]. Hence, our results show that the shape of the distribution is not conserved when propagating errors from LECs to PNM.

While the predicted EOS does not include an explicit model for truncation uncertainties, we stress that these uncertainties are absorbed in part by the LEC distribu-




FIG. 4. PDFs for the pressure in neutron-star matter at twice saturation density when propagating full posterior distributions for all LECs. We also show the result of the LIGO collaboration of Ref. [67].

tions. We acknowledge that this approach is not fully comprehensive, especially at lower orders in the EFT. In particular, we estimate these uncertainties in two-nucleon scattering and absorb them only in short-range LECs while keeping the pion-exchange terms fixed. Furthermore, our approach does not know about the structure of higher-order contributions and, therefore, fails if this structure can only be poorly absorbed by the truncated interaction. This is especially true at LO (see Ref. [38]), but improves at higher orders, where corrections can be expected to be small; see Fig. 1. Other approaches [27, 36] rely on different assumptions about the structure of these omitted terms and estimate errors based on expansion coefficients that are only informed by the truncated interactions. Similarly, these estimates are also problematic at lower orders in the EFT [27, 36]. A strength of our approach is that it naturally enables us to account for correlated uncertainties across different nuclear systems because uncertainties are estimated from distributions of the Hamiltonian directly and not system by system. While no approach is without limitations, we find that different estimates agree at N^2 LO, supporting the validity of our framework.

Finally, we study the impact of our framework on neutron stars, and compute the pressure in β -equilibrium at twice nuclear saturation density, $p_\beta(2n_{\text{sat}})$. Our error estimates for this quantity are obtained from our PNM results by using the metamodel of Refs. [68, 69] to account for the small proton and lepton fractions present in neutron stars. We extract a pressure of $p_\beta(2n_{\text{sat}}) = 16.8_{-5.7}^{+6.0}$ MeV/fm³ at 90% confidence level and show the resulting histogram in Fig. 4. This is in excellent agreement with the determination from gravitational-wave data of neutron-star merger GW170817 by the LIGO collaboration [67] of $p_\beta(2n_{\text{sat}}) = 21.8_{-10.6}^{+16.9}$ MeV/fm³ at 90%

confidence level. Our calculations demonstrate consistency between chiral interactions and gravitational-wave neutron star observations where uncertainties in the former are fully propagated to the EOS using emulators.

Conclusion—We have built and applied fast and accurate surrogate models to AFDMC for pure neutron matter. These emulators are based on the PMM approach and allow us to propagate uncertainties from LECs directly to the EOS of PNM. We have found that the PMM is an excellent approach to emulating AFDMC for dense matter, faithfully reproducing our high-fidelity calculations with only a small number of high-fidelity training points. However, care is necessary when building these emulators as level crossings can appear in the final fit results.

The LEC uncertainties generated in the fit of Ref. [38] include an estimate of the truncation uncertainty of the chiral EFT expansion. Using our PMM emulators, we have propagated LEC posteriors accounting for this source of uncertainty to the EOS of neutron matter at multiple densities and found distributions of the energy per particle that are skewed at each density, which the skewness increasing with density. Our final EOS uncertainty band differs from the previous uncertainties found using GPs. While producing the training data for the emulators cost 22.4 million CPU-hours (30-60 LEC sets depending on the order), the final results are obtained for 250,000 LEC sets in ~ 6.5 seconds on commodity hardware.

Our results are important for understanding astrophysical observations of neutron stars and their mergers and provide robust nuclear-theory input for analyses of multi-messenger data. As we will explore in future work, the use of fast emulators will allow us to sample over LECs in inferences of astrophysical data, enabling fits of LECs directly from this data [70].

ACKNOWLEDGMENTS

We thank E. Bonilla, J. Carlson, S. Gandolfi, and B. Reed for useful discussions. This document has been approved for unlimited release, and was assigned LA-UR-24-27320.

C.L.A. was supported by the Laboratory Directed Research and Development (LDRD) program of Los Alamos National Laboratory (LANL) under project number 20230315ER and by LANL through its Center for Space and Earth Science, which is funded by LANL’s LDRD program under project number 20240477CR. P.G. and K.G. were supported by the National Science Foundation CSSI program under award No. OAC-2004601 (BAND Collaboration). R.S. acknowledges support from the Nuclear Physics from Multi-Messenger Mergers (NP3M) Focused Research Hub which is funded by the National Science Foundation under Grant Number 21-16686, and from the LDRD program of LANL under project number 20220541ECR. I.T. was supported by the U.S. De-

partment of Energy, Office of Science, Office of Nuclear Physics, under contract No. DE-AC52-06NA25396, by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) NUCLEI program, and by the LDRD program of LANL under project numbers 20220541ECR and 20230315ER.

Computational resources have been provided by the

Los Alamos National Laboratory Institutional Computing Program, which is supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001, and by the National Energy Research Scientific Computing Center (NERSC), which is supported by the U.S. Department of Energy, Office of Science, under contract No. DE-AC02-05CH11231.

-
- [1] J. M. Lattimer and M. Prakash, *Astrophys. J.* **550**, 426 (2001), arXiv:astro-ph/0002232.
- [2] F. Özel and P. Freire, *Ann. Rev. Astron. Astrophys.* **54**, 401 (2016), arXiv:1603.02698 [astro-ph.HE].
- [3] J. M. Lattimer, *Ann. Rev. Nucl. Part. Sci.* **71**, 433 (2021).
- [4] K. Chatziioannou, H. T. Cromartie, S. Gandolfi, I. Tews, D. Radice, A. W. Steiner, and A. L. Watts, (2024), arXiv:2407.11153 [nucl-th].
- [5] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].
- [6] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, *Phys. Rev. Lett.* **121**, 091102 (2018), [Erratum: *Phys. Rev. Lett.* **121**, 259902 (2018)], arXiv:1804.08583 [astro-ph.HE].
- [7] T. E. Riley *et al.*, *Astrophys. J. Lett.* **887**, L21 (2019), arXiv:1912.05702 [astro-ph.HE].
- [8] T. E. Riley *et al.*, *Astrophys. J. Lett.* **918**, L27 (2021), arXiv:2105.06980 [astro-ph.HE].
- [9] M. C. Miller *et al.*, *Astrophys. J. Lett.* **887**, L24 (2019), arXiv:1912.05705 [astro-ph.HE].
- [10] M. C. Miller *et al.*, *Astrophys. J. Lett.* **918**, L28 (2021), arXiv:2105.06979 [astro-ph.HE].
- [11] D. Choudhury *et al.*, *Astrophys. J. Lett.* **971**, L20 (2024), arXiv:2407.06789 [astro-ph.HE].
- [12] C. D. Capano, I. Tews, S. M. Brown, B. Margalit, S. De, S. Kumar, D. A. Brown, B. Krishnan, and S. Reddy, *Nature Astron.* **4**, 625 (2020), arXiv:1908.10352 [astro-ph.HE].
- [13] T. Dietrich, M. W. Coughlin, P. T. H. Pang, M. Bulla, J. Heinzel, L. Issa, I. Tews, and S. Antier, *Science* **370**, 1450 (2020), arXiv:2002.11355 [astro-ph.HE].
- [14] P. Landry, R. Essick, and K. Chatziioannou, *Phys. Rev. D* **101**, 123007 (2020), arXiv:2003.04880 [astro-ph.HE].
- [15] H. Tan, J. Noronha-Hostler, and N. Yunes, *Phys. Rev. Lett.* **125**, 261104 (2020), arXiv:2006.16296 [astro-ph.HE].
- [16] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nissanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021), arXiv:2105.06981 [astro-ph.HE].
- [17] I. Legred, K. Chatziioannou, R. Essick, S. Han, and P. Landry, *Phys. Rev. D* **104**, 063003 (2021), arXiv:2106.05313 [astro-ph.HE].
- [18] T. Gorda, O. Komoltsev, and A. Kurkela, *Astrophys. J.* **950**, 107 (2023), arXiv:2204.11877 [nucl-th].
- [19] H. Koehn *et al.*, *Phys. Rev. X* **15**, 021014 (2025), arXiv:2402.04172 [astro-ph.HE].
- [20] N. Rutherford *et al.*, *Astrophys. J. Lett.* **971**, L19 (2024), arXiv:2407.06790 [astro-ph.HE].
- [21] D. Reitze *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019), arXiv:1907.04833 [astro-ph.IM].
- [22] M. Evans *et al.*, “A Horizon Study for Cosmic Explorer: Science, Observatories, and Community,” (2021), arXiv:2109.09882 [astro-ph.IM].
- [23] M. Punturo *et al.*, *Class. Quant. Grav.* **27**, 194002 (2010).
- [24] M. Maggiore *et al.*, *JCAP* **03**, 050 (2020), arXiv:1912.02622 [astro-ph.CO].
- [25] I. Tews, J. Carlson, S. Gandolfi, and S. Reddy, *Astrophys. J.* **860**, 149 (2018), arXiv:1801.01923 [nucl-th].
- [26] R. Essick, I. Tews, P. Landry, S. Reddy, and D. E. Holz, *Phys. Rev. C* **102**, 055803 (2020), arXiv:2004.07744 [astro-ph.HE].
- [27] C. Drischler, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *Phys. Rev. Lett.* **125**, 202702 (2020), arXiv:2004.07232 [nucl-th].
- [28] C. Drischler, P. G. Giuliani, S. Bezoui, J. Piekarewicz, and F. Viens, *Phys. Rev. C* **110**, 044320 (2024).
- [29] P. Navrátil, V. G. Gueorguiev, J. P. Vary, W. E. Ormand, and A. Nogga, *Phys. Rev. Lett.* **99**, 042501 (2007), arXiv:nucl-th/0701038.
- [30] 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, *Phys. Rev. C* **91**, 051301 (2015), [Erratum: *Phys. Rev. C* **109**, 059901 (2024)], arXiv:1502.04682 [nucl-th].
- [31] P. Reinert, H. Krebs, and E. Epelbaum, *Eur. Phys. J. A* **54**, 86 (2018), arXiv:1711.08821 [nucl-th].
- [32] D. R. Entem, R. Machleidt, and Y. Nosyk, *Phys. Rev. C* **96**, 024004 (2017), arXiv:1703.05454 [nucl-th].
- [33] R. Somasundaram, C. L. Armstrong, P. Giuliani, K. Godbey, S. Gandolfi, and I. Tews, *Physics Letters B*, 139558 (2025).
- [34] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa, *Rev. Mod. Phys.* **87**, 1067 (2015), arXiv:1412.3081 [nucl-th].
- [35] K. E. Schmidt and S. Fantoni, *Phys. Lett. B* **446**, 99 (1999).
- [36] E. Epelbaum, H. Krebs, and U. G. Meißner, *Eur. Phys. J. A* **51**, 53 (2015), arXiv:1412.0142 [nucl-th].
- [37] I. Tews, R. Somasundaram, D. Lonardonì, H. Götting, R. Seutin, J. Carlson, S. Gandolfi, K. Hebeler, and A. Schwenk, *Phys. Rev. Res.* **7**, 033024 (2025), arXiv:2407.08979 [nucl-th].
- [38] R. Somasundaram, J. E. Lynn, L. Huth, A. Schwenk, and I. Tews, *Phys. Rev. C* **109**, 034005 (2024), arXiv:2306.13579 [nucl-th].
- [39] W. G. Jiang, C. Forssén, T. Djärv, and G. Hagen, *Phys. Rev. C* **109**, L061302 (2024), arXiv:2212.13203 [nucl-th].
- [40] W. G. Jiang, C. Forssén, T. Djärv, and G. Hagen, *Phys. Rev. C* **109**, 064314 (2024), arXiv:2212.13216 [nucl-th].
- [41] P. Cook, D. Jammooa, M. Hjorth-Jensen, D. D.

- Lee, and D. Lee, *Nature Commun.* **16**, 5929 (2025), arXiv:2401.11694 [cs.LG].
- [42] B. T. Reed, R. Somasundaram, S. De, C. L. Armstrong, P. Giuliani, C. Capano, D. A. Brown, and I. Tews, *The Astrophysical Journal* **974**, 285 (2024).
- [43] T. Duguet, A. Ekström, R. J. Furnstahl, S. König, and D. Lee, *Reviews of Modern Physics* **96** (2024), 10.1103/revmodphys.96.031002.
- [44] P. Giuliani, K. Godbey, E. Bonilla, F. Viens, and J. Piekarewicz, *Frontiers in Physics* **10** (2023), 10.3389/fphy.2022.1054524.
- [45] C. Drischler, J. A. Melendez, R. J. Furnstahl, A. J. Garcia, and X. Zhang, *Frontiers in Physics* **10** (2023), 10.3389/fphy.2022.1092931.
- [46] E. Bonilla, P. Giuliani, K. Godbey, and D. Lee, *Phys. Rev. C* **106**, 054322 (2022).
- [47] R. Cheng, K. Godbey, Y. Niu, Y. Ma, W. He, and S. Wang, arXiv preprint arXiv:2411.15492 (2024).
- [48] D. Odell, P. Giuliani, K. Beyer, M. Catacora-Rios, M.-H. Chan, E. Bonilla, R. J. Furnstahl, K. Godbey, and F. M. Nunes, *Physical Review C* **109**, 044612 (2024).
- [49] D. Frame, R. He, I. Ipsen, D. Lee, D. Lee, and E. Rrapaj, *Physical review letters* **121**, 032501 (2018).
- [50] B. Hu, W. Jiang, T. Miyagi, Z. Sun, A. Ekström, C. Forssén, G. Hagen, J. D. Holt, T. Papenbrock, S. R. Stroberg, *et al.*, *Nature Physics* **18**, 1196 (2022).
- [51] See Supplemental Material for additional details on the PMM setup for different models and interactions.
- [52] J. E. Lynn, I. Tews, S. Gandolfi, and A. Lovato, *Ann. Rev. Nucl. Part. Sci.* **69**, 279 (2019), arXiv:1901.04868 [nucl-th].
- [53] A. Gezerlis, I. Tews, E. Epelbaum, S. Gandolfi, K. Hebeler, A. Nogga, and A. Schwenk, *Phys. Rev. Lett.* **111**, 032501 (2013), arXiv:1303.6243 [nucl-th].
- [54] A. Gezerlis, I. Tews, E. Epelbaum, M. Freunek, S. Gandolfi, K. Hebeler, A. Nogga, and A. Schwenk, *Phys. Rev. C* **90**, 054323 (2014), arXiv:1406.0454 [nucl-th].
- [55] M. Piarulli, L. Girlanda, R. Schiavilla, R. Navarro Pérez, J. E. Amaro, and E. Ruiz Arriola, *Phys. Rev. C* **91**, 024003 (2015), arXiv:1412.6446 [nucl-th].
- [56] I. Tews, S. Gandolfi, A. Gezerlis, and A. Schwenk, *Phys. Rev. C* **93**, 024305 (2016), arXiv:1507.05561 [nucl-th].
- [57] J. E. Lynn, I. Tews, J. Carlson, S. Gandolfi, A. Gezerlis, K. E. Schmidt, and A. Schwenk, *Phys. Rev. Lett.* **116**, 062501 (2016), arXiv:1509.03470 [nucl-th].
- [58] M. Piarulli, L. Girlanda, R. Schiavilla, A. Kievsky, A. Lovato, L. E. Marcucci, S. C. Pieper, M. Viviani, and R. B. Wiringa, *Phys. Rev. C* **94**, 054007 (2016), arXiv:1606.06335 [nucl-th].
- [59] M. Piarulli *et al.*, *Phys. Rev. Lett.* **120**, 052503 (2018), arXiv:1707.02883 [nucl-th].
- [60] M. Piarulli and I. Tews, *Front. in Phys.* **7**, 245 (2020), arXiv:2002.00032 [nucl-th].
- [61] R. Curry, R. Somasundaram, S. Gandolfi, A. Gezerlis, and I. Tews, *Phys. Rev. C* **111**, 015801 (2025), arXiv:2409.16365 [nucl-th].
- [62] M. Hoferichter, J. Ruiz de Elvira, B. Kubis, and U.-G. Meißner, *Phys. Rept.* **625**, 1 (2016), arXiv:1510.06039 [hep-ph].
- [63] K. Hebeler and A. Schwenk, *Phys. Rev. C* **82**, 014314 (2010), arXiv:0911.0483 [nucl-th].
- [64] L. Huth, I. Tews, J. E. Lynn, and A. Schwenk, *Phys. Rev. C* **96**, 054003 (2017), arXiv:1708.03194 [nucl-th].
- [65] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, *Nature Methods* **17**, 261 (2020).
- [66] S. Gandolfi, J. Carlson, and S. Reddy, *Phys. Rev. C* **85**, 032801 (2012), arXiv:1101.1921 [nucl-th].
- [67] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. Lett.* **121**, 161101 (2018), arXiv:1805.11581 [gr-qc].
- [68] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, *Phys. Rev. C* **97**, 025805 (2018), arXiv:1708.06894 [nucl-th].
- [69] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, *Phys. Rev. C* **97**, 025806 (2018), arXiv:1708.06895 [nucl-th].
- [70] R. Somasundaram, I. Svensson, S. De, A. E. Deneris, Y. Dietz, P. Landry, A. Schwenk, and I. Tews, (2024), arXiv:2410.00247 [nucl-th].

SUPPLEMENTAL MATERIAL

In the following, we will give additional details on the PMMs constructed in this work.

PMM for a toy model




FIG. 1. Accuracy for reproducing eigenvalues of a randomly generated 10-dimensional eigenvalue equation. We show the logarithm of the relative error between the full 10-dimensional model and PMM emulation for different numbers of training points. From left to right, the PMM is trained on 5, 10, and 15 points which are shown in red. Each subsequent row increases the number of dimensions in the matrices fit starting from 2-dimensional matrices for the top row and ending with 5-dimensional matrices for the bottom row.




FIG. 2. The evolution of the average validation error as function of $N_{\text{train}} = 1 - 100$ for $N_{\text{dim}} = 2$ for a toy problem with 9 control parameters.

To test the PMM approach, better understand the required number of training data for PMMs of different dimension, and assess how well PMMs extrapolate with scarce training data, we have set up a toy problem. We have randomly generated a 10-dimensional toy Hamiltonian with two parameters, similar to Eq. (1), and exactly solved for its eigenvalues as function of the two parameters C_S and C_T . We have then constructed PMMs with two to five dimensions and studied the relative errors of the PMM reproduction of the full 10-dimensional problem. We show the obtained relative errors in Fig. 1 for 5, 10, and 15 training points.

We find that the PMM with two parameters generally works very well already with a small number of training data in parts of the parameter space that are much larger than the phase space covered by the training data. In the 2-dimensional case, for which there are 8 unknown matrix elements, the PMM saturates between 5-10 training points as expected. For higher-dimensional PMMs, though, the relative uncertainties are very small in large parts of the parameter space even if there are fewer training data than matrix elements to determine. In case of a 5-dimensional PMM (35 unknowns), the parameter space is well recovered even with 15 training data. This highlights the power of the PMM as an efficient emulator for interpolation and extrapolation.

We have also set up a similar toy problem for a 9-parameter PMM to test the optimal number of training data to construct PMMs resembling the case at NLO and $N^2\text{LO}$ with operator LECs. We show the relative error for a two-dimensional PMM (29 unknown matrix elements) as function of number of training points N_{train} in Fig. 2. We find that the error decreases and levels off at 35-55 training data, but slowly increases again for larger number of training data, likely due to overfitting. We conclude that N_{train} should be similar to the number of free matrix elements, which suggests that we need ~ 20 training data for our PMM fits at NLO and $N^2\text{LO}$ with six spectral LECs (20 matrix elements for two-dimensional PMM).

A. Testing the PMM setup at LO

As discussed in the main text, at LO we generated 30 high-fidelity data points. Initially, we randomly split them into a training set with 20 data points and a validation set with 10 validation points. We then trained a two-dimensional PMM for various N_{train} keeping the random order of training points fixed and obtained the results in Fig. 3. We find that the relative error decreases to sub-percent uncertainties for $N_{\text{train}} = 6$ but at larger values of training points the error can suddenly increase to much larger values, sometimes even to 20 – 30% at higher orders. We found that this is due to level crossings of the two lowest eigenvalues of the PMM matrix that appear at energies above the




FIG. 3. The relative error of a 2-dimensional PMM when reproducing the AFDMC validation energies as function of number of training points N_{train} . For all lines, the AFDMC energies are first split into a training set and a validation set using different prescriptions. For the blue line, the training points are randomly chosen and ordered. For the orange line, the training set includes the global highest-energy AFDMC result in the overall data set at first position, so all N_{train} include this maximum. For the green line, the training set is randomly chosen but with the highest AFDMC energy in the training set at first position, so all N_{train} include this point. For the red line, the training set is randomly ordered, but the fit results for each N_{train} is used as a starting point for the subsequent fit with $N_{\text{train}} + 1$. For the purple line, all N_{train} include the highest energy in the training set and the fit for each N_{train} is used as a starting point for the subsequent fit at $N_{\text{train}} + 1$.

ones included in the training set but intersects with the validation set, see Fig. 4. In these cases, the correct PMM prediction would sometimes be given by the second eigenvalue, and the relative uncertainties for these validation points blow up the average uncertainty. This can happen even if one additional training point is included and the PMM fit shifts accordingly, leading to the spikes observed in Fig. 3. This problem also appears at larger N_{dim} , see Figure 5 for $N_{\text{dim}} = 3$ and $N_{\text{dim}} = 4$.

We then explored several options to avoid these level crossings. A reasonable solution is to choose training points so that they span the full energy range. The simplest option is to include the highest-energy AFDMC result in the training set at first position, which effectively pushes higher eigenvalues to high energies for all N_{train} . From Fig. 3, we see this eliminates the issue and relative errors level off at very low values already for $N_{\text{train}} \geq 4$. When generating training data for the PMM, we cannot *a priori* ensure that we sample parameter sets that lead to sufficiently high energies. Hence, in a second approach we first randomly split the data set into a training and a validation set, as before, but move the highest-energy datum in the training set on position one of the training set, i.e., it is always included in training. This leads to the green curve, partially reducing the spikes in relative error. Finally, as these spikes are related to the fit result for the matrix elements and suddenly appear at certain number of training data, it is reasonable to try to avoid them by starting the fit for $N_{\text{train}} + 1$ from the previous fit result for N_{train} for which the level crossing did not appear. Testing this (red line), we find a similarly stable behavior as for the fit that always includes the global energy maximum, albeit with slightly larger uncertainties. Finally, we combined this approach with the one always including the training set maximum and obtain the purple curve. This approach leads to stable results with small errors, and we have decided to employ this to generate the PMMs in this paper.

We have then employed this approach at LO at nuclear saturation density. Using a PMM with 5 training points and $N_{\text{dim}} = 2$, we have propagated $\sim 200,000$ LEC samples to the energy per particle, see Fig. 6. We found that the final distribution changes only negligibly for different N_{train} and larger N_{dim} . We find that the uncertainty distribution at LO is well described by a Gaussian.




FIG. 4. Correlation of AFDMC energies with $C_S - 3C_T$, C_S , and C_T for $N_{\text{train}} = 5$ and $N_{\text{dim}} = 2$. All rows are for different PMM fits and show all PMM eigenvalues if 100k LEC sets are propagated. The upper row is for a PMM with a level crossing and a gap in the energy spectrum. The middle row is for a PMM with level crossing but without a gap in the energy spectrum. The lower row is for a PMM without a level crossing.

B. PMM at NLO

Finally, we discuss our PMM at NLO. We have extensively tested the PMM at NLO as the number of LECs at this order is the same as at $N^2\text{LO}$. At NLO, we have followed the same fit prescription as at LO. As discussed in the main text, we have set up the NLO PMM in terms of 6 spectral LECs [38] to reduce the number of fit parameters. We have explicitly checked that the final errors are similar to PMMs defined in terms of 9 operator LECs, but the PMM convergence is faster and more stable. In Fig. 6, we show the distribution of energies per particle when propagating $\sim 200,000$ LEC sets at NLO. We find that the uncertainty distribution at NLO is also well described by a Gaussian. We conjecture that the different shape at $N^2\text{LO}$ is due to the impact of three-nucleon forces.

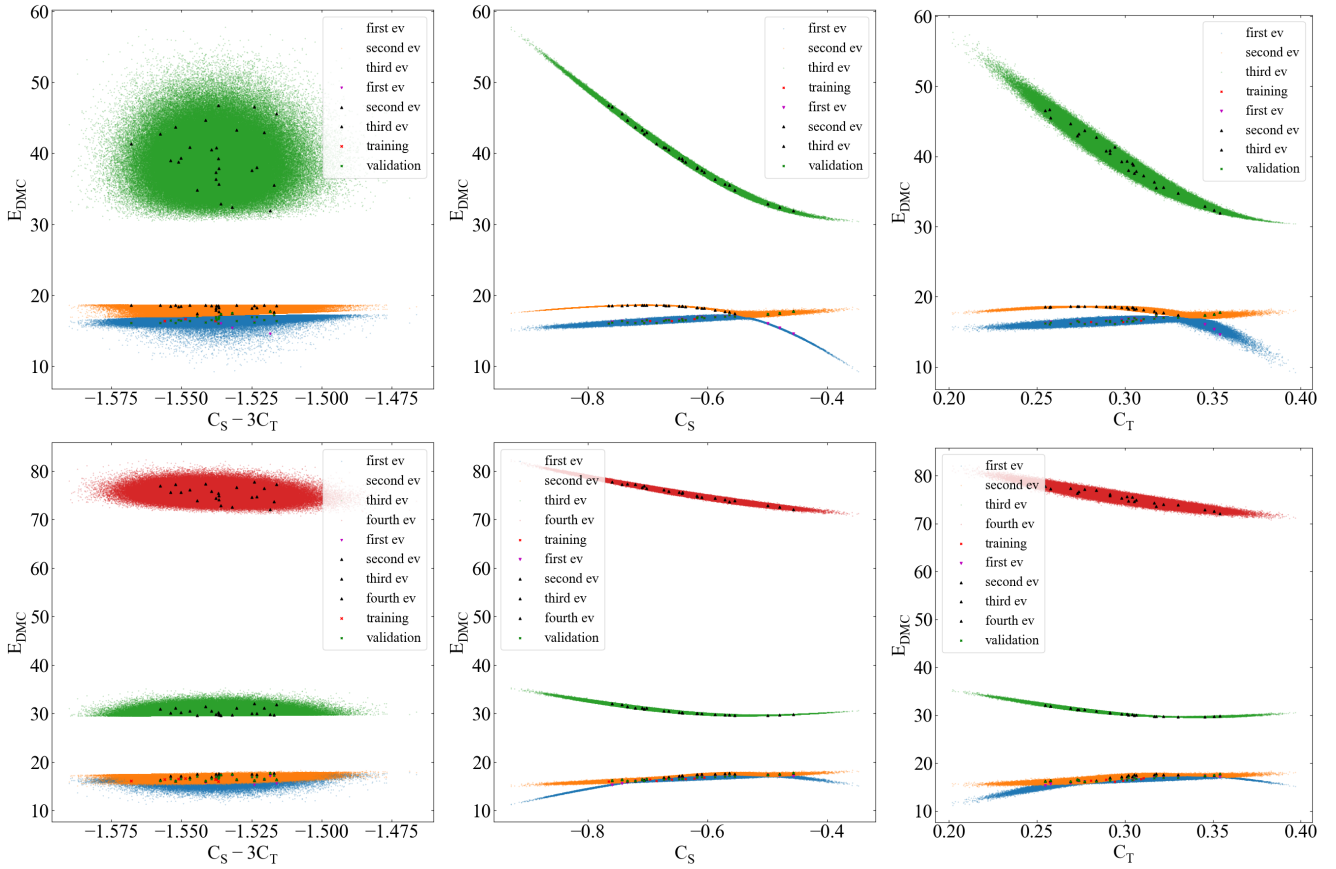


FIG. 5. Same as Fig. 4 but for $N_{\text{dim}} = 3$ (upper row) and $N_{\text{dim}} = 4$ (lower row).

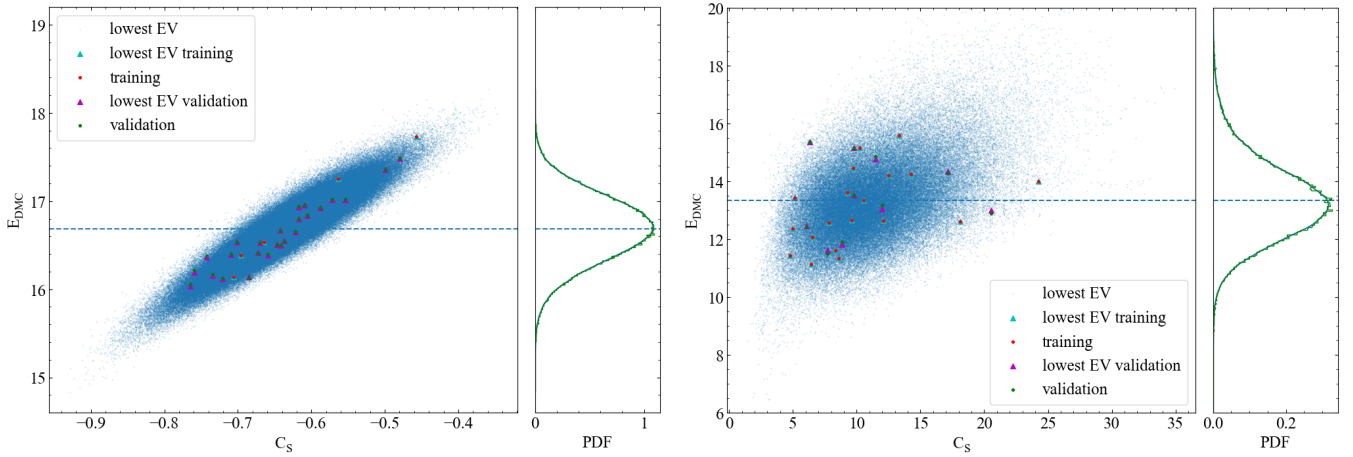


FIG. 6. Left: Distribution of energies per particle at LO using a PMM fit with $N_{\text{train}} = 5$ and $N_{\text{dim}} = 2$. The distribution is obtained by propagating uncertainties on C_S and C_T to the observable. We also show the 1-D histogram and its Gaussian KDE fit. The average error for this fit is 0.3%. Right: Similar, but for NLO with $N_{\text{train}} = 20$ and $N_{\text{dim}} = 3$. The average error for this fit is 0.5%.

We show the convergence of the relative validation error as function of N_{train} in the left panel of Fig. 7. In contrast to the LO case, we find that the convergence is slower, and stabilizes for $\sim 20 - 25$ training points. We also find that the error improves up to $N_{\text{dim}} = 4$ but generally the error is comparable for $N_{\text{dim}} = 2 - 5$ for $N_{\text{train}} = 25$ and is of the order of 1%.

In the right panel of Fig. 7, we show the detailed evolution of the validation errors for the two-dimensional PMM

for 20-50 training points. We find that the relative validation errors can vary widely between validation data points, but that the uncertainty stabilizes at rather similar values if more and more training points are included.

Finally, during the construction of our PMMs we have realized that for different random starting guesses of the matrix elements and for different sets of training and validation data, that the quality of the fit, i.e., the final relative validation error, can vary slightly. To quantify the variation, we have performed 20 different PMM fits for the NLO data. For each of these fits, we selected 10 random validation points and present a histogram of the validation error for all validation points in Fig. 8. We find that the average error over all 20 PMMs is about 1.9%, which is close to the error in Fig. 7.

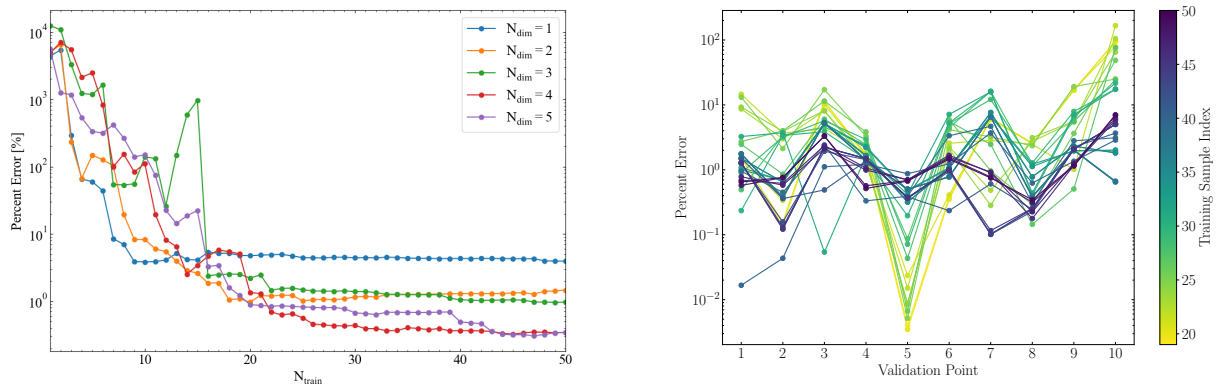


FIG. 7. Left: The relative error of the PMM reproduction of the AFDMC validation energies at NLO as function of number of training points N_{train} for $N_{\text{dim}} = 1 - 5$ at nuclear saturation density. The training and validation sets are the same for all N_{dim} . Right: Percent errors for the 10 validation points for the PMM fits at NLO with $N_{\text{dim}} = 2$ and $N_{\text{train}} = 20 - 50$.




FIG. 8. Distribution of PMM validation errors for all validation points at NLO using 6 spectral LECs and Dual Annealing and $N_{\text{dim}} = 3$. We have performed 20 different PMM fits for $N_{\text{train}} = 25$ randomly selected training points and 10 randomly selected validation points.