

Inference of the low-energy constants in Δ -full chiral effective field theory including a correlated truncation error

Isak Svensson , Andreas Ekström, and Christian Forssén 

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



(Received 27 April 2023; revised 7 March 2024; accepted 11 April 2024; published 18 June 2024)

We sample the posterior probability distributions of the low-energy constants (LECs) in Δ -full chiral effective field theory (χ EFT) up to third order. We use eigenvector continuation for fast and accurate emulation of the likelihood and Hamiltonian Monte Carlo to draw effectively independent samples from the posteriors. Our Bayesian inference is conditioned on the Granada database of neutron-proton (np) cross sections and polarizations. We use priors grounded in χ EFT assumptions and a Roy-Steiner analysis of pion-nucleon scattering data. We model correlated EFT truncation errors using a two-feature Gaussian process, and find correlation lengths for np scattering energies and angles in the ranges 45–83 MeV and 24–39 degrees, respectively. These correlations yield a nondiagonal covariance matrix and reduce the number of independent scattering data with factors of 8 and 4 at the second and third chiral orders, respectively. The relatively small difference between the second- and third-order predictions in Δ -full χ EFT suppresses the marginal variance of the truncation error and the effects of its correlation structure. Our results are particularly important for analyzing the predictive capabilities in *ab initio* nuclear theory.

DOI: [10.1103/PhysRevC.109.064003](https://doi.org/10.1103/PhysRevC.109.064003)

I. INTRODUCTION

A chiral effective field theory (χ EFT) description of the nuclear interaction [1–5] is endowed with a power counting (PC) to organize the order-by-order contributions of the strong-interaction dynamics to nuclear observables. This kind of organization, regardless of the particulars of the adopted PC [6], is a hallmark of EFT [7] and *ab initio* [8] approaches to nuclear theory since it promises a handle on the theoretical uncertainty coming from truncating the EFT expansion [9]. Accounting for the truncation error is key to mitigating overfitting of the low-energy constants (LECs) as well as assessing the importance of discrepancies. Indeed, the modeling of EFT truncation errors can play a significant role in the robustness of LEC inferences and ensuing nuclear predictions [10–16]. Melendez *et al.* [17] have proposed a Bayesian model for the truncation error that accounts for finite correlations across independent variables, e.g., the scattering energy and angle for nucleon-nucleon (NN) scattering cross sections and polarizations. To date, LEC inference in χ EFT typically accounts for truncation errors in the fully correlated or uncorrelated limits [18–22], and the robustness with respect to the correlation structure is not well known.

In this paper, we quantify a correlated truncation error and analyze its effects on a Bayesian estimation of the LEC

posteriors for a Δ -full χ EFT description of the neutron-proton (np) interaction up to next-to-next-to-leading order (NNLO) [23,24]. This extends our previous work on Bayesian LEC estimation in Δ -less χ EFT where we employed an uncorrelated truncation error [14,21,22]. We also use eigenvector continuation (EC) [25], i.e., a reduced basis method [26], to efficiently and accurately emulate the scattering amplitudes entering the np scattering-data likelihood. Following our previous publications [21,22], we employ Hamiltonian Monte Carlo (HMC) [27] to draw effectively independent samples from the LEC posteriors.

The $\Delta(1232)$ resonance plays an important role in nuclear physics since it represents a rather low excitation energy and couples strongly to the pion-nucleon (πN) system. This was recognized already in early χ EFT descriptions of the NN interaction [28–30], and several modern χ EFT interactions incorporate the Δ as well [31–34]. In Δ -full χ EFT there are four subleading πN LECs up to NNLO, usually denoted c_1, c_2, c_3, c_4 . They govern the strength of subleading 2π -exchange diagrams of the NN interaction and the leading three-nucleon (NNN) 2π exchange with an intermediate Δ excitation, the so-called Fujita-Miyazawa force [35,36]. A Roy-Steiner analysis of πN scattering amplitudes by Hoferichter *et al.* [37] has enabled a determination of the πN LECs. Unfortunately, the relatively unknown value of the axial $\pi N \Delta$ coupling h_A propagates to approximately five times greater uncertainties for c_2, c_3, c_4 (compared to a determination in Δ -less χ EFT) when matching Δ -full χ EFT to the subthreshold amplitudes in the Roy-Steiner formalism [38]. A well-founded truncation error would therefore pave the way for learning more about the strength of subleading 2π exchange, and leading NNN forces, from NN scattering data.

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. Funded by Bibsam.

This paper is organized as follows: In Sec. II we present our statistical model for linking experiment and χ EFT. In Sec. III we discuss our priors and likelihood, and in Sec. IV we introduce the two-feature Gaussian-process model of the correlated EFT truncation error. In Sec. V we discuss the training of EC np scattering emulators. In Sec. VI we present the results from HMC sampling of the LEC posteriors. A summary and an outlook are given in Sec. VII.

II. LINKING EXPERIMENT AND THEORY

Following our previous papers [14,21,22], we relate an experimental measurement y_{exp} of some scattering observable with a theoretical prediction $y_{\text{th}}^{(k)}$, up to chiral order k , using an additive model to account for the respective uncertainties, δy_{exp} and $\delta y_{\text{th}}^{(k)}$:

$$y_{\text{exp}}(\vec{x}) = y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x}) + \delta y_{\text{exp}}(\vec{x}) + \delta y_{\text{th}}^{(k)}(\vec{x}). \quad (1)$$

The theory prediction y_{th} depends deterministically on the vector of LECs $\vec{\alpha}$ and the independent variable $\vec{x} = (T_{\text{lab}}, \theta)$, where T_{lab} denotes the kinetic energy of the incoming nucleon in the laboratory frame and θ denotes the scattering angle in the center-of-mass frame. For the total scattering cross section we have $\vec{x} = (T_{\text{lab}})$ as this observable is integrated over all θ . We have suppressed the explicit $\vec{\alpha}$ dependence of $\delta y_{\text{th}}^{(k)}$ as we will assume a fixed parameter value in the model of the theory error.

The composition of the LEC vector $\vec{\alpha}$ depends on the chiral order. For the np potentials in this paper, the LEC vector up to each order is given by

$$\vec{\alpha}_{\text{LO}} = (\tilde{C}_{1S0}, \tilde{C}_{3S1}), \quad (2)$$

$$\vec{\alpha}_{\text{NLO}} = (\tilde{C}_{1S0}^{np}, \tilde{C}_{3S1}, C_{1S0}, C_{3P0}, \times C_{1P1}, C_{3P1}, C_{3S1}, C_{3S1-3D1}, C_{3P2}), \quad (3)$$

$$\vec{\alpha}_{\text{NNLO}} = (\tilde{C}_{1S0}^{np}, \tilde{C}_{3S1}, C_{1S0}, C_{3P0}, C_{1P1}, C_{3P1}, \times C_{3S1}, C_{3S1-3D1}, C_{3P2}, c_1, c_2, c_3, c_4), \quad (4)$$

where LO is leading order ($k = 0$), NLO is next-to-leading order ($k = 2$), and NNLO is next-to-next-to-leading order ($k = 3$).¹ We employ units and a notation linked to the momentum partial-wave basis; see Refs. [4,23] for details. The potential is nonlocally regulated with a regulator cutoff $\Lambda = 450$ MeV as in Ref. [24]. In the following we will only refer to the generic vector $\vec{\alpha}$, while the specific chiral order, if important, should be obvious from the context.

The power counting of Δ -full χ EFT allows us to express a prediction $y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x})$ up to chiral order k as a sum of order-by-order contributions,

$$y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x}) = y_{\text{ref}}(\vec{x}) \sum_{i=0}^k c^{(i)}(\vec{\alpha}; \vec{x}) Q^i, \quad (5)$$

where y_{ref} is some characteristic scale to be decided, and $c^{(i)}$ are dimensionless expansion coefficients. The dimensionless

expansion parameter $Q = f(m_\pi, \delta, p)/\Lambda_\chi$ is a ratio composed of a soft scale given by some function of the pion mass (m_π), the ΔN mass splitting δ , the (external) relative NN momentum p , and a hard scale Λ_χ for which we assume a point-estimate value of 600 MeV [18,39]. We treat $\delta \sim m_\pi$ as a small scale [40], although it does not vanish in the chiral limit. We resum the potential to all orders which obfuscates the form of f , and in line with [17] we therefore assume the following functional form:

$$f(m_\pi, p) = \frac{p^8 + m_\pi^8}{p^7 + m_\pi^7}, \quad (6)$$

which facilitates a smooth transition across the soft scale m_π . We find that the exact form of this function does not impact our inference significantly, and reverting to a canonical estimate $f = \max(m_\pi, p)$ does not change any of our results. The upshot of having an order-by-order description of y_{th} , as in Eq. (5), is a handle on the theoretical uncertainty via

$$\delta y_{\text{th}}^{(k)}(\vec{x}) = y_{\text{ref}} \sum_{i=k+1}^{\infty} c^{(i)}(\vec{x}) Q^i. \quad (7)$$

Clearly, we have not explicitly extracted EFT coefficients $c^{(i>k)}$ for any \vec{x} and as a consequence we are uncertain about their values. However, if we assume naturalness to hold, it is reasonable to have $c^{(i)} \sim \mathcal{O}(1)$ for *all* expansion coefficients, including the unseen ones. Furthermore, contributions can be either positive or negative. We will therefore consider the coefficients as drawn from some underlying probability distribution bounded by a finite variance \bar{c}^2 expected to be of order 1. In addition, there is most likely information about $y_{\text{th}}(\vec{x})$ at nearby values of the independent variable. As such, there is likely nonzero covariance between expansion coefficients $c^{(i)}(\vec{x}_m)$ and $c^{(i)}(\vec{x}_n)$ at different values \vec{x}_m and \vec{x}_n . That correlation structure is of primary interest in this paper.

III. SETTING UP THE LEC PARAMETER INFERENCE

We seek a posterior probability density function (PDF) $\text{pr}(\vec{\alpha}|\mathcal{D}, I)$ for the LECs $\vec{\alpha}$ in Δ -full χ EFT up to NNLO conditioned on selected np scattering data \mathcal{D} from the Granada database [41,42] and other information I as specified below. In the following, we use bold symbols to denote quantities that depend on a range of values for the independent variable \vec{x} , e.g., $\mathbf{x} = (\vec{x}_1, \vec{x}_2, \dots)$ and $\mathcal{D} = (y_{\text{exp}}(\vec{x}_1), y_{\text{exp}}(\vec{x}_2), \dots)$, where we let y denote observable type (differential cross section, vector polarization, etc.).

Using Bayes's rule we can express the posterior in terms of a likelihood $\text{pr}(\mathcal{D}|\vec{\alpha}, I)$, prior $\text{pr}(\vec{\alpha}|I)$, and evidence $\text{pr}(\mathcal{D}|I)$ as

$$\text{pr}(\vec{\alpha}|\mathcal{D}, I) = \frac{\text{pr}(\mathcal{D}|\vec{\alpha}, I)\text{pr}(\vec{\alpha}|I)}{\text{pr}(\mathcal{D}|I)}. \quad (8)$$

We ignore the evidence term in this paper as it provides overall normalization and does not impact the shape of the posterior.

¹Due to symmetries [3–5] the order $k = 1$ vanishes.

A. Prior

We assume independent priors for the LECs of the NN contact and πN potentials, i.e.,

$$\text{pr}(\vec{\alpha}|I) = \text{pr}(\vec{\alpha}_{NN}|I) \times \text{pr}(\vec{\alpha}_{\pi N}|I), \quad (9)$$

where $\vec{\alpha}_{\pi N} = (c_1, c_2, c_3, c_4)$ denotes² the subleading πN LECs. For these, we place a normal and correlated prior based on the (nondiagonal) covariance matrix $\Sigma_{\pi N}$ obtained in the Roy-Steiner analysis of πN scattering data by Siemens *et al.* [38],

$$\text{pr}(\vec{\alpha}_{\pi N}|I) = \mathcal{N}(\vec{\mu}_{\pi N}, \Sigma_{\pi N}), \quad (10)$$

with

$$\Sigma_{\pi N} = \begin{bmatrix} +6.11 & -0.63 & +6.26 & +0.25 \\ -0.63 & +277.43 & -359.84 & +174.20 \\ +6.26 & -359.84 & +474.01 & -224.92 \\ +0.25 & +174.20 & -224.92 & +119.16 \end{bmatrix} \times \frac{10^{-4}}{\text{GeV}^2} \quad (11)$$

and

$$\vec{\mu}_{\pi N} = [-0.74, -0.49, -0.65, +0.96] \text{ GeV}^{-1}. \quad (12)$$

Along the lines of naturalness, we place a normal and uncorrelated prior on the LECs $\vec{\alpha}_{NN}$ of the contact potential,

$$\text{pr}(\vec{\alpha}_{NN}|I) = \mathcal{N}(\vec{0}, \Sigma_{NN}), \quad (13)$$

where Σ_{NN} is a diagonal covariance matrix with standard deviations $5 \times 10^4 \text{ GeV}^{-(k+2)}$ for the contact LECs first appearing at orders $k = 0, 2$. This is the same prior as in our previous papers [21,22].

B. Likelihood

We consider only strong interaction effects in elastic reaction channels, and neglect all electromagnetic interactions in our calculations. Thus, we condition our LEC inferences on low-energy np scattering data below the pion-production threshold $T_{\text{lab}} = 290 \text{ MeV}$, and omit np scattering data with $T_{\text{lab}} < 30 \text{ MeV}$ as the nonzero magnetic moments of neutrons and protons can distort the low-energy np scattering amplitude [43,44]. We partition the experimental data set as $\mathcal{D} = \{\mathcal{D}_y\}_{y=1}^{N_y}$, where $N_y = 18$ is the number of unique types of scattering cross sections and polarizations in the considered energy range. In Fig. 1 we summarize in tabular form the data of measured np scattering observables that we include, as well as some of the results that will be discussed below.

We assume that experimental and theoretical errors are independent of each other such that the total covariance matrix can be written as

$$\Sigma = \Sigma_{\text{exp}} + \Sigma_{\text{th}}. \quad (14)$$

The (diagonal) covariance matrix Σ_{exp} is provided by the experimenters and we employ normalization factors from the

Granada database. We model the covariance of the EFT errors independently for each type of np scattering observable y , and make specific assumptions and choices I_y per observable type. The (block-diagonal) covariance matrix Σ_{th} is given by

$$\Sigma_{\text{th}} = \begin{bmatrix} \Sigma_{\text{th},1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \Sigma_{\text{th},2} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \Sigma_{\text{th},N_y} \end{bmatrix}, \quad (15)$$

where the (nondiagonal) covariance matrix $\Sigma_{\text{th},y}$ at chiral order k contains elements

$$(\Sigma_{\text{th},y})_{mm} = \text{cov}[\delta y_{\text{th}}^{(k)}(\vec{x}_m), \delta y_{\text{th}}^{(k)}(\vec{x}_n)] \quad (16)$$

that we model using a Gaussian process; see Sec. IV.

In accordance with the principle of maximum entropy we employ a normally distributed data likelihood, which factorizes as

$$\text{pr}(\mathcal{D}|\vec{\alpha}, I) = \prod_{y=1}^{N_y} \text{pr}(\mathcal{D}_y|\vec{\alpha}, I_y), \quad (17)$$

where

$$\text{pr}(\mathcal{D}_y|\vec{\alpha}, I_y) \propto \exp\left[-\frac{\mathbf{r}_y^T(\vec{\alpha})(\Sigma_{\text{exp},y} + \Sigma_{\text{th},y})^{-1}\mathbf{r}_y(\vec{\alpha})}{2}\right]. \quad (18)$$

Here, $\mathbf{r}_y(\vec{\alpha}) = [r_{y,1}(\vec{\alpha}), r_{y,2}(\vec{\alpha}), \dots, r_{y,N_{d,y}}(\vec{\alpha})]$ is a (column) vector of $N_{d,y}$ residuals, each one given by the difference between experimental and order- k theoretical prediction values for the independent variable \vec{x} , i.e.,

$$r_{y,j}(\vec{\alpha}) = [y_{\text{exp}}(\vec{x}_j) - y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x}_j)]. \quad (19)$$

All cross sections are computed from amplitudes obtained numerically via the nonrelativistic Lippmann-Schwinger equation [12,21]. This approach is virtually exact and we do not account for any method uncertainties of numerical origin in this paper. Note that in Eq. (19) we assume that all errors taken into account have mean zero, which will be the case in this paper.

IV. CORRELATING THE EFT TRUNCATION ERROR IN NN SCATTERING

Following the suggestion by Melendez *et al.* [17], we model the EFT truncation error using Gaussian processes (\mathcal{GP} s) [46,47]. This results in a covariance matrix Σ_{th} from Eq. (15) with off-diagonal elements. In this section, we limit our discussion to the \mathcal{GP} modeling of a single covariance submatrix $\Sigma_{\text{th},y}$ for a specific observable type y , e.g., the differential cross section $\sigma(\theta)$, and therefore omit the observable index y in the following.

We develop a \mathcal{GP} model to handle two features, i.e., $\mathcal{GP} : \mathbb{R}^2 \rightarrow \mathbb{R}^1$, for correlating the distribution of expansion coefficients at different values of the scattering energy *and* angle. For total cross sections we only operate with the scattering energy, as all angles are integrated over. We expect unseen coefficients $c^{(i)}(\mathbf{x}) \equiv \{c^{(i)}(\vec{x}_j)\}_j$ to be distributed as

$$c^{(i)}(\mathbf{x}) | m, \ell, \bar{c}^2 \sim \mathcal{GP}[m, \bar{c}^2 \kappa(\mathbf{x}', \mathbf{x}; \ell)], \quad (20)$$

²Not to be confused with the EFT expansion coefficients $c^{(0)}, c^{(2)}, \dots$ in Eq. (5).

Notation	Definition		Acronym	$N_{d,y}$	$N_{T_{\text{lab}},y}$	n_{eff}	$\widehat{\ell}_{T_{\text{lab}}}^{\text{MAP}}$ (MeV)	$\widehat{\ell}_{\theta}$ (deg)	\widehat{c}^2
σ_{tot}	total cross section		SGT	119	113	30.8	49		0.56^2
σ_T	$\sigma_{\text{tot}}(\uparrow\downarrow) - \sigma_{\text{tot}}(\uparrow\uparrow)$		SGTT	3	3	—	—		—
σ_L	$\sigma_{\text{tot}}(\rightrightarrows) - \sigma_{\text{tot}}(\leftrightharpoons)$		SGTL	4	4	3.8	47		1.95^2
Notation	Tensor	Illustration	Acronym	$N_{d,y}$	$N_{T_{\text{lab}},y}$	n_{eff}	$\widehat{\ell}_{T_{\text{lab}}}^{\text{MAP}}$ (MeV)	$\widehat{\ell}_{\theta}$ (deg)	\widehat{c}^2
$\sigma(\theta)$	I_{0000}		DSG	1207	68	352.9	73	39	0.61^2
$A(\theta)$	D_{s0k0}		A	5	1	5.0	68	37	0.65^2
$A_t(\theta)$	K_{0ks0}		AT	30	2	25.5	74	33	0.49^2
$A_{yy}(\theta)$	A_{00nn}		AYY	58	4	13.7	66	36	1.27^2
$A_{zz}(\theta)$	A_{00kk}		AZZ	45	2	16.7	70	28	0.98^2
$D(\theta)$	D_{n0n0}		D	13	1	4.5	70	27	0.65^2
$D_t(\theta)$	K_{0nn0}		DT	36	5	31.8	57	29	0.67^2
$D_x^z(\theta)$	D_{0s0k}		D0SK	8	1	3.1	60	29	0.53^2
$A_y(\theta)$	P_{n000}		P	503	28	298.7	59	28	0.31^2
$N_{zz}^y(\theta)$	N_{0nkk}		NNKK	8	1	8.0	45	24	0.22^2
$N_{zy}^x(\theta)$	N_{0skn}		NSKN	12	1	10.9	81	35	0.87^2
$N_{xy}^x(\theta)$	N_{0ssn}		NSSN	4	1	4.0	76	31	0.62^2
$R(\theta)$	D_{s0s0}		R	5	1	5.0	74	32	0.67^2
$R_t(\theta)$	K_{0ss0}		RT	29	3	25.2	83	26	0.55^2
$R_t'(\theta)$	K_{0sk0}		RPT	1	1	1.0	61	35	0.53^2

FIG. 1. The database \mathcal{D} of np scattering observables with $30 \leq T_{\text{lab}} \leq 290$ MeV with notation, definition, and acronym in the Granada database [41,42], number of data $N_{d,y}$ and measurement energies $N_{T_{\text{lab}},y}$ per observable type y , effective dimension n_{eff} given theory error correlations, maximum *a posteriori* (MAP) values for \mathcal{GP} correlation lengths across scattering energy $\widehat{\ell}_{T_{\text{lab}}}^{\text{MAP}}$ and angle $\widehat{\ell}_{\theta}$, and marginal variance \widehat{c}^2 . The spin-dependent integrated cross sections σ_T and σ_L have polarizations parallel and antiparallel to each other and transverse (T) or longitudinal (L) to the beam direction. The σ_T observable was omitted in our analysis; see the text for details. For spin polarizations we list the polarization tensors, X_{srbt} , using the notation of [45] where subscripts denote the polarization of the (s)catter, (r)eccoil, (b)eam, and (t)arget particles in the direction of the incoming projectile (\hat{k}), normal to the scattering plane (\hat{n}), $\hat{s} = \hat{n} \times \hat{k}$, and 0 is an unpolarized state. The illustrations show initial and final spin-polarizations in the laboratory frame using gray arrows for reactants, blue arrows for the polarization direction in the scattering plane, and circled dots for the outward-pointing polarization vector. The projectile impinges from the left, the target (not indicated) is located in the center, and the projectile (recoil) scatters upwards (downwards) at an angle indicated by the arc (gray). In this table, and in the text, we leave the energy dependence of all observables implicit.

where m is assumed to be a constant mean function and ℓ , \bar{c}^2 denote the hyperparameters for the \mathcal{GP} correlation length(s) and marginal variance, respectively. The calibration of these parameters will be discussed in Sec. IV A. The marginal standard deviation \bar{c} is a scale factor that quantifies the average distance of the \mathcal{GP} away from its mean, and the correlation length quantifies the typical extent of the \mathcal{GP} fluctuations. Elastic np scattering observables typically exhibit a smooth behavior as a function of energy and angle. Therefore, we expect smooth variations of the expansion coefficients. Assuming stationarity, we employ a squared-exponential kernel to parametrize the correlation structure,

$$\kappa(\vec{x}_i, \vec{x}_j; \ell) = \exp\left[-\frac{(\vec{x}_i - \vec{x}_j)^T \ell^{-1} (\vec{x}_i - \vec{x}_j)}{2}\right], \quad (21)$$

where the length scale(s) are given by

$$\ell = \begin{bmatrix} \ell_{T_{\text{lab}}}^2 & 0 \\ 0 & \ell_{\theta}^2 \end{bmatrix}, \quad (22)$$

in an obvious notation.

We are mainly interested in the quantified hyperparameters ℓ and \bar{c}^2 that characterize the finite correlation length and variance of the EFT error. A \mathcal{GP} is a linear-Gaussian model for which there are closed-form expressions [46,47] for the distribution of predicted data conditioned on calibration data. Once the posterior distribution, or point estimates, of the hyperparameters are determined it is straightforward to quantitatively use the \mathcal{GP} to predict EFT expansion coefficients exhibiting a variance and correlation consistent with the calibration data. Furthermore, one can show [17] that the \mathcal{GP} for the EFT truncation error pertaining to a scattering observable y_{th} is given by

$$\delta y_{\text{th}}^{(k)}(\mathbf{x}) | m, \ell, \bar{c}^2, y_{\text{ref}}, Q \sim \mathcal{GP}[M^{(k)}(\mathbf{x}), \bar{c}^2 K^{(k)}(\mathbf{x}', \mathbf{x}; \ell)], \quad (23)$$

where the respective mean (M) and kernel (K) functions are given by

$$M^{(k)}(\mathbf{x}) = y_{\text{ref}}(\mathbf{x}) \frac{[Q(\mathbf{x})]^{k+1}}{1 - Q(\mathbf{x})} m, \quad (24)$$

and

$$K^{(k)}(\mathbf{x}', \mathbf{x}; \ell) = y_{\text{ref}}(\mathbf{x}') y_{\text{ref}}(\mathbf{x}) \times \frac{[Q(\mathbf{x}')Q(\mathbf{x})]^{k+1}}{1 - Q(\mathbf{x}')Q(\mathbf{x})} \kappa(\mathbf{x}', \mathbf{x}; \ell). \quad (25)$$

This lays the foundation for how we model the covariance matrix in Eq. (15) and sample the LEC posterior $\text{pr}(\vec{\alpha} | \mathcal{D}, I)$ via the likelihood and prior defined in Secs. III A and III B. As will be discussed shortly, we set $m = 0$ in this paper.

A. Optimizing the \mathcal{GP} hyperparameters

We set up individual \mathcal{GP} s for the block-diagonal covariance matrices $\Sigma_{\text{th},y}$ in Eq. (15). The hyperparameters for each observable type y are optimized separately and we find maximum *a posteriori* (MAP) values for them using GSUM [17]. For each \mathcal{GP} we use a set of training data \mathbf{c}^* consisting of N_{train}

expansion coefficients obtained as

$$c^{(k)}(\vec{x}_j) = \frac{y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x}_j) - y_{\text{th}}^{(k-1)}(\vec{\alpha}; \vec{x}_j)}{y_{\text{ref}}[Q(\vec{x}_j)]^k}, \quad (26)$$

for a training range \mathbf{x}^* of values for the independent variable. We pick $N_{\text{train}} \approx 25$ –50 points uniformly distributed across approximately 4–6 energies $T_{\text{lab}} \in [30, 290]$ MeV and 5–10 angles $\theta \in [0, 180]$ degrees. The training grid varies somewhat depending on the observable type. The order-by-order differences in Eq. (26) are computed from theoretical predictions based on a MAP estimate for the LECs, $\vec{\alpha} = \vec{\alpha}^*$, obtained in line with Ref. [21] in the uncorrelated limit using $\bar{c} = 1$.

The LO predictions $y_{\text{th}}^{(0)}$ in Δ -full χ EFT are identical to the Δ -less predictions and as such they are rather poor, as expected in Weinberg PC [12,48]. Thus, the $c^{(2)}$ expansion coefficients incur rather unnatural values. One could argue that a Bayesian approach, with a prior for the expansion parameters, should be able to properly handle such outliers. However, a (likely) deficient LO will conflict with our initial assumptions in Eq. (7). Instead of using the truncation error to absorb deficiencies rooted in the PC, we decided to neglect LO predictions in this paper. Thus, all training data consist of $c^{(3)}$ expansion coefficients only. Despite this limited amount of data, we expect the available coefficients to carry valuable information about the correlated truncation error and we proceed with a Bayesian \mathcal{GP} model. For total and differential cross sections we set y_{ref} to the predictions of the Idaho-N3LO(500) potential [49], which is almost the same as using $y_{\text{ref}} = y_{\text{exp}}$. We encounter a zero crossing in y_{ref} for σ_T . This leads to a discontinuity in the $c^{(3)}$ coefficients as a function of T_{lab} ; see Fig. 2. Such small-scale structures are virtually impossible to faithfully represent using an infinitely smooth squared-exponential kernel. As there are only three experimental data points for this observable type, we decided to exclude them from \mathcal{D} rather than choosing an exceptional y_{ref} for this one case. For polarization data we encounter an excess of small-scale structures when setting y_{ref} to the predictions of Idaho-N3LO(500). Using $y_{\text{ref}} = 0.5$ removes them entirely and we are spared further kernel design. We choose 0.5 as a representative value for these observables. Other choices could be $y_{\text{ref}} = 1.0$ [18] or an average of experimental values [21].

All training data are pruned for outliers using a three-interquartiles threshold [21]. In the end, all $c^{(3)}(\mathbf{x})$ coefficients are of natural order, and a vast majority of them pass the outlier filter. Operating with a \mathcal{GP} allows us to incorporate known symmetry constraints on the polarization observables. Indeed, some polarizations are identically equal to zero at $\theta = 0$ and/or $\theta = 180$ degrees. We impose such boundary constraints by adding zero-valued expansion coefficients for the endpoint(s) of the angular direction to the training data. In a future application we will formally incorporate continuous boundary constraints in the \mathcal{GP} [50].

We employ a bounded and uniform prior for the length scales ℓ of the form

$$\text{pr}(\ell_{\theta} | I) = \mathcal{U}(\epsilon, 180) \text{ degrees} \quad (27)$$

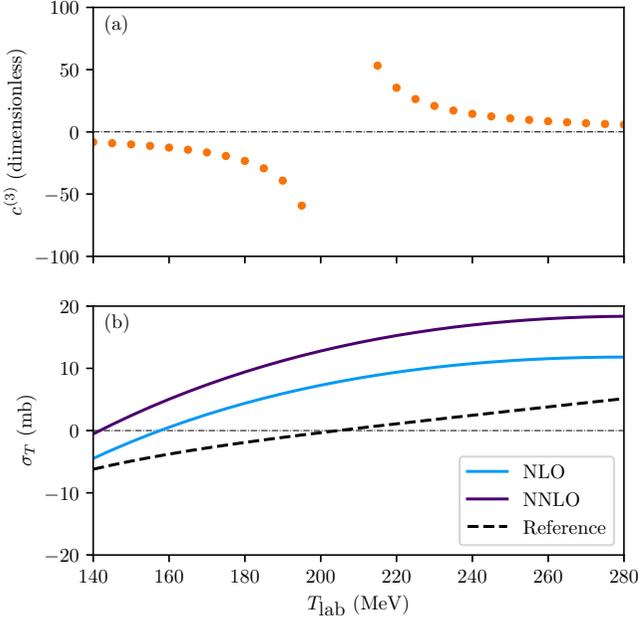


FIG. 2. (a) Calculated expansion coefficients for σ_T . (b) The NLO, NNLO, and reference values from which the coefficients are calculated. Note the zero crossing for the reference value at $T_{\text{lab}} \approx 205$ MeV.

and

$$\text{pr}(\ell_{T_{\text{lab}}}|I) = \mathcal{U}(\epsilon, 290) \text{ MeV}, \quad (28)$$

where $\epsilon = 10^{-5}$ is introduced to avoid numerical issues in the optimization of the kernel (the exact value of ϵ is unimportant in this case since the posterior values are never close to the edge of the prior). We place a conjugate inverse- χ^2 prior on the variance \bar{c}^2 according to

$$\text{pr}(\bar{c}^2|I) = \chi^{-2}(v_0 = 1, \tau_0^2 = 1) \quad (29)$$

and a Dirac delta prior $\text{pr}(m) = \delta(m)$ on the mean m . This encodes our expectation of natural \bar{c}^2 values, although the heavy tail in our prior allows for some unnaturalness, and expansion coefficients symmetrically distributed with mean $m \equiv 0$. Our chosen prior for \bar{c}^2 is shown in orange in Fig. 3. We do not allow the mean to vary owing to the limited amount of information—one chiral order—we have available to learn from. This strict prior on m could be substituted for, e.g., a Gaussian prior if more information is available. If so, Eq. (19) must be updated to include the systematic contribution of a truncation error with nonzero mean.

The conjugacy of the $\chi^{-2}(\bar{c}^2)$ prior implies that the posterior $\text{pr}(\bar{c}^2|m=0, \ell, \mathbf{c}^*, \mathbf{x}, I)$ is also an inverse- χ^2 PDF with updated parameters [17]. The prior for the length scale ℓ is not conjugate, and to find the posterior MAP values ($\hat{\ell}$, $\hat{\bar{c}}^2$) of

$$\begin{aligned} \text{pr}(\ell, \bar{c}^2|m=0, \mathbf{c}^*, \mathbf{x}^*, I) &\propto \text{pr}(\mathbf{c}^*|m=0, \ell, \bar{c}^2, \mathbf{x}^*, I) \\ &\times \text{pr}(\bar{c}^2|m=0, I)\text{pr}(\ell|I) \end{aligned} \quad (30)$$

we use the numerical optimization routine L-BFGS-B [51] with multiple restarts. We include a so-called nugget, set to 10^{-10} , to avoid numerical issues. The resulting MAP

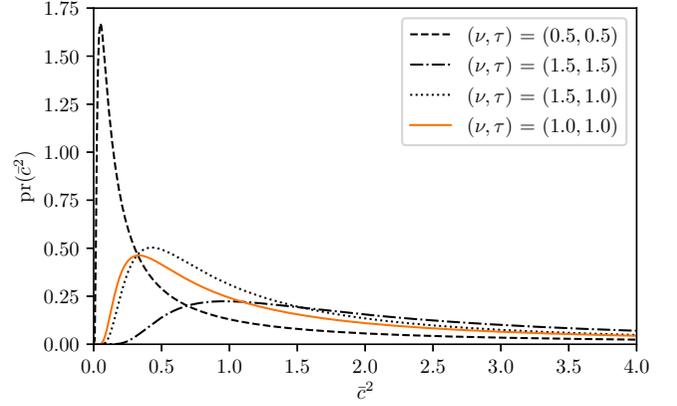


FIG. 3. The inverse- χ^2 prior $\text{pr}(\bar{c}^2|I)$ (orange) for the \mathcal{GP} hyperparameter \bar{c}^2 . Shown in black are alternative priors that we use to show the robustness of the inference.

estimates of the hyperparameters are listed in the table in Fig. 1. The delta prior on m obviously yields a corresponding MAP value $\hat{m} = 0$. The \mathcal{GP} predictions for expansion coefficients $c^{(3)}$ for three of the most abundant observable types are shown in Fig. 4.

The amount of training data is sufficient to make the posterior likelihood-dominated in most cases. We explore a number of priors, shown in Fig. 3, and find that the inference is rather robust as long as we allow for very small values $\bar{c}^2 \gtrsim 0.2^2$. Indeed, a prior that is too restrictive in this regard can bias the result for total cross sections, where the training data are one dimensional.

B. Validating the \mathcal{GP} hyperparameters

We validate the \mathcal{GP} hyperparameters for each observable type y using a set of N_{val} complementary validation data $\tilde{\mathbf{c}}$. The validation data are generated by the same method as the training data, but with energies and angles shifted so that the validation values $\tilde{\mathbf{x}}$ for the independent variable do not overlap with ones used during training. In addition to visual inspections, which are certainly useful, we also employ a set of diagnostics inspired by Refs. [17,52]. These diagnostics are the Mahalanobis distance, pivoted Cholesky errors, and the credible interval diagnostic. These, and a few more, are thoroughly discussed in Ref. [17].

The (squared) Mahalanobis distance, D_{MD}^2 , is a multivariate analog to the sum of squared residuals. Here, it is computed as

$$D_{\text{MD}}^2(\tilde{\mathbf{c}}; \hat{\mathbf{m}}, \tilde{\Sigma}_{\mathcal{GP}}) = [\tilde{\mathbf{c}} - \hat{\mathbf{m}}]^T \tilde{\Sigma}_{\mathcal{GP}}^{-1} [\tilde{\mathbf{c}} - \hat{\mathbf{m}}], \quad (31)$$

where $\hat{\mathbf{m}}$ is the \mathcal{GP} mean at $\tilde{\mathbf{x}}$ —equivalent to zero with our choice of prior—and

$$\tilde{\Sigma}_{\mathcal{GP}} = \hat{\bar{c}}^2 \kappa(\tilde{\mathbf{x}}', \tilde{\mathbf{x}}; \hat{\ell}). \quad (32)$$

This distance is commonly used to quantify the deviation of a prediction compared to data in a correlated setting. Here, we use it to diagnose whether a set of validation data could reasonably have been drawn from a \mathcal{GP} with covariance according to Eq. (32) and mean $m = \hat{m} = 0$. Either too large or

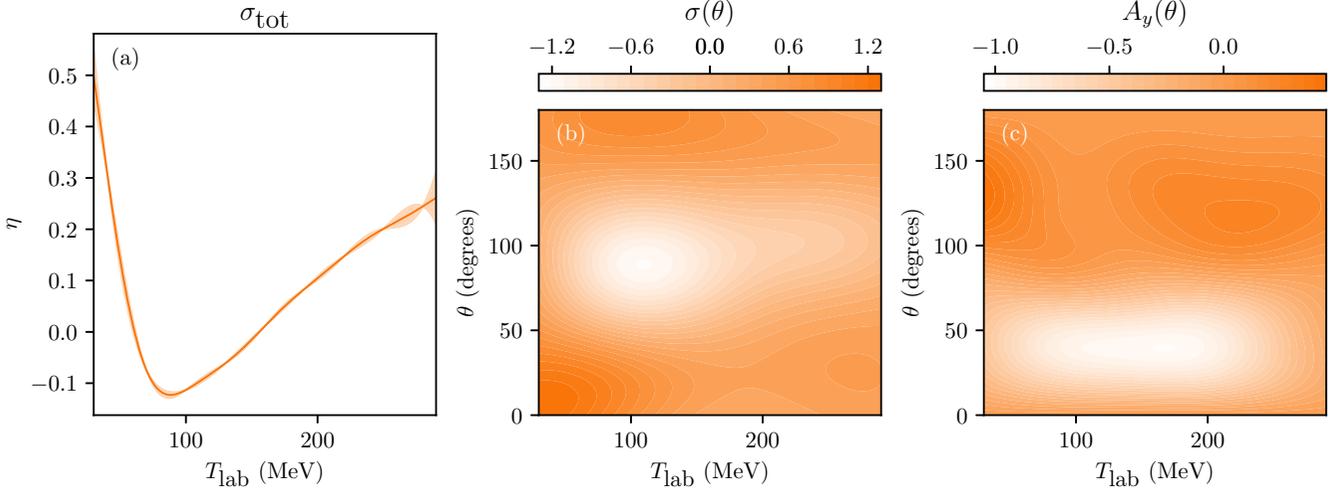


FIG. 4. Mean \mathcal{GP} predictions η for (dimensionless) expansion coefficients for (a) σ_{tot} , (b) $\sigma(\theta)$, and (c) $A_y(\theta)$. The light colored band in (a) indicates a 95% credible interval.

small values of D_{MD}^2 , compared to the reference χ^2 distribution with N_{val} degrees of freedom, point to a possible conflict between the \mathcal{GP} and the validation data. The Mahalanobis distance is a scalar measure and does not provide detailed insight about the possible tension with respect to the validation data.

Furthermore, it is instructive to study the \mathcal{GP} and the validation data point by point as a function of the independent variable. Such comparisons are not straightforward to interpret given the mutual correlation of the validation data. Therefore, we decorrelate and rescale the covariance matrix to independent unit variances using a Cholesky decomposition: $\tilde{\Sigma}_{\mathcal{GP}} = \mathbf{G}\mathbf{G}^T$ where \mathbf{G} is a triangular standard deviation matrix. From this we define the Cholesky errors

$$\mathbf{D}_G = \mathbf{G}^{-1}[\tilde{\mathbf{c}} - \hat{\mathbf{m}}]. \quad (33)$$

To order the vector \mathbf{D}_G in a meaningful way we pivot the decomposition in decreasing conditional variances. One should not detect any pattern when plotting the pivoted \mathbf{D}_G versus the index of the validation data. To reveal further information one can introduce a ratio scale also on the abscissa by plotting the pivoted Cholesky errors versus the conditional variances used for pivoting [53].

All but one of the \mathcal{GP} s readily pass our diagnostics, with D_{MD}^2 landing within 95% credible intervals of the respective reference distributions. We see little to no structure in the pivoted Cholesky decompositions, and empirical coverages that roughly match the corresponding credible intervals. In addition, visual inspections indicate that the inferred hyperparameters are plausible. Naturally, incorporating expansion coefficients from other orders would provide a stronger test of our \mathcal{GP} model. We only reject σ_T , and therefore also the three experimental data points. Neither D_{MD}^2 nor the \mathbf{D}_G exhibit significantly poor performance for this observable given our validation data. However, upon visual inspection of the $c^{(3)}$ curve for this observable, shown in Fig. 2, we find that it is discontinuous for $T_{\text{lab}} \approx 205$ MeV because the reference value crosses zero at this point, leading to extremely large

coefficients near the crossing point and an abrupt change in sign. Clearly, our \mathcal{GP} model, based on a squared-exponential kernel, is ill suited to handling this discontinuity, and the inferred length scale is strongly dependent on the chosen training data and will approach zero as we increase the number of training data. If it were not for the existence of an experimental data point in the vicinity of the problematic scattering energy, we would have considered excluding only the problematic region. This example clearly shows that one should not blindly trust the diagnostics as they are conditioned on the chosen training and validation data.

C. The structure of the correlated truncation error

Introducing a correlated truncation error reduces the number of independent directions in the data space. Indeed, two data residing within one correlation length of the scattering energy or angle carry joint information. To quantify the impact of this we compute the effective dimension n_{eff} , per observable type y , using a measure [54,55] defined as $n_{\text{eff}} = \prod_{i=1}^{N_{d,y}} \gamma_i^{-\gamma_i}$, where γ_i denotes a normalized eigenvalue $\gamma_i = \lambda_i / \text{tr}(\mathbf{C}_y)$ of the respective correlation matrix. Here, $\mathbf{C}_y = \mathbf{S}_y^{-1} \Sigma_y \mathbf{S}_y^{-1}$ and $\mathbf{S}_y = \sqrt{\text{diag}(\Sigma_y)}$ with $\Sigma_y = \Sigma_{\text{exp},y} + \Sigma_{\text{th},y}$. We could equally well compute n_{eff} for Σ_y directly. However, operating with the correlation matrix leads to the values $n_{\text{eff}} = 1$ and $n_{\text{eff}} = N_{d,y}$ in the limits of having a full off-diagonal correlation and zero off-diagonal correlation, respectively. Consequently, we interpret n_{eff} as the effective number of data, and note that the logarithm of n_{eff} corresponds to the discrete Shannon entropy of the spectrum of the normalized eigenvalues γ_i .

The resulting n_{eff} values at NNLO are summarized in the table in Fig. 1. They show that the correlations reduce the number of effective data from 2087 to 841, which is still plentiful. The main reason for the relatively weak impact of the correlated truncation error is the dominance of the experimental variances along the diagonal of the covariance matrix Σ . Indeed, for the correlation matrices of the \mathcal{GP} kernel for

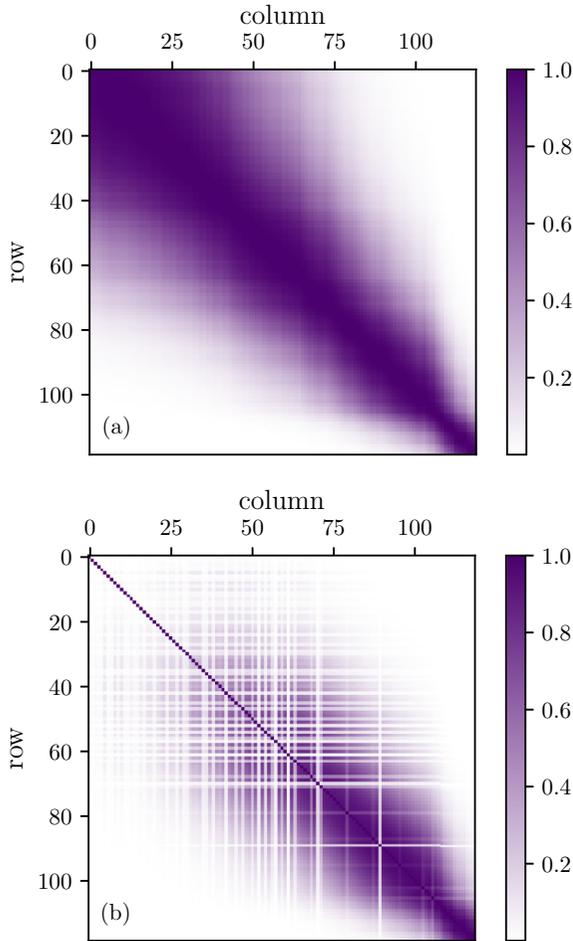


FIG. 5. Correlation matrices based on the NNLO covariance matrix $\Sigma_{\sigma_{\text{tot}}}$ of the 119 σ_{tot} data. In panels (a) and (b) we exclude and include, respectively, the experimental variances $\Sigma_{\text{exp},\sigma_{\text{tot}}}$. The data are sorted by energy in increasing order.

EFT truncation errors alone we find $n_{\text{eff}} \approx 5\text{--}10$. In Fig. 5 we show correlation matrices for the total cross section $C_{\sigma_{\text{tot}}}$ with and without the experimental covariance matrix accounted for. The diagonal dominance of the experimental errors is clearly visible. This dominance weakens for higher energies (higher indices in the figure) as the truncation error increases with the scattering energy T_{lab} . At NLO, where the truncation error is greater by one chiral order Q , the corresponding n_{eff} values are typically a factor 2 smaller as the correlations of the truncation error become relatively more important. The total number of effective data is 498. In Fig. 6(a) we compare the 68% credible interval $\Delta y_{\text{th}}^{(k)}$ of the EFT truncation error at NLO and NNLO with the experimental errors Δy_{exp} for all data.

The EFT truncation error, and its correlations, is expected to be more important for inferences conditioned on very precise pp data. To estimate this effect, we inspected the pp $A_y(\theta)$ data set more closely. With $T_{\text{lab}} \in [30, 290]$ MeV there are 496 data points, which is almost the same as in the np sector. Reducing the experimental variances of the np $A_y(\theta)$ data, to mimic the average level of pp variances, we observe $n_{\text{eff}} = 154$ and 91 at NNLO and NLO, respectively. For other

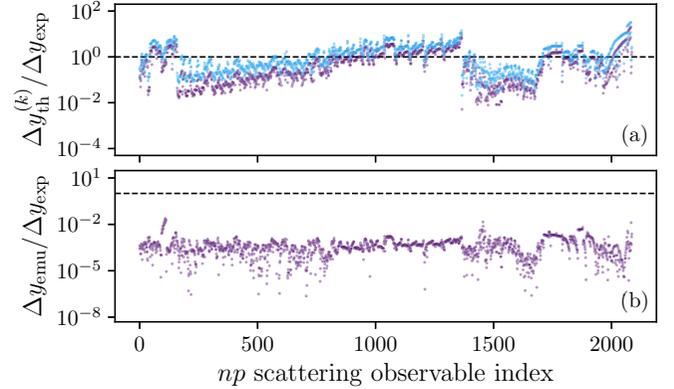


FIG. 6. The EFT truncation and EC emulator errors compared to experimental errors for all considered np scattering data. (a) The standard deviations of the EFT truncation errors at NLO and NNLO. (b) The NNLO EC emulator errors (Δy_{emu}). Colors as in Fig. 2. In both panels we divide by the the experimental 1σ errors (Δy_{exp}) and the horizontal dashed lines indicate a ratio of 1.

observable types, like $\sigma(\theta)$, the average pp variance is greater than the average np variance, and the opposite effect is likely observed. Note that the distribution of np and pp measurement energies and angles differ and we do not account for this in our estimate.

V. EMULATING np SCATTERING CROSS SECTIONS

We use an EC-based method [25] to construct accurate and efficient emulators for np scattering observables. Operating with emulators, instead of the exact simulators, helps reduce the computational cost of evaluating the likelihood. As an added bonus, once the emulators are trained they are straightforward to distribute in the scientific community.

We use Newton's functional formulation of on-shell np T -matrix elements for setting up the emulators [56]. Technically, this leads to one emulator per np partial-wave and unique scattering energy in the database \mathcal{D} (table in Fig. 1). Truncating the partial-wave expansion at a maximum np angular-momentum quantum number $J = 30$ leads to 182 partial waves. With $N_{T_{\text{lab}}} = 177$ energies in \mathcal{D} we end up with 32 214 T -matrix emulators per chiral order. The training values for the LECs are drawn according to a latin hypercube design in a sufficiently wide interval $[-4, +4]$ (see Sec. III A for units). To simplify the setup, we employ the same training protocol for all emulators and find it sufficient to use 7 (8) training points at NLO (NNLO). As there are at most 3 (7) relevant LECs acting per partial wave, this approach leads to very accurate emulation of all relevant scattering observables. We estimate that emulator errors Δy_{emu} are at least 10, and typically 1000, times smaller than experimental errors. This is quantified using the difference between emulated and exact results for all np observables in \mathcal{D} at ten different sets of LEC values following randomized latin hypercube designs within the training intervals specified above. In Fig. 6(b) we show the ratio $\Delta y_{\text{emu}}/\Delta y_{\text{exp}}$ at NNLO for a random set of LEC values. This is virtually identical to the NLO result. We conclude that the emulator errors are sufficiently small to neglect them in

our inferences. In the following we refer to the collection of all amplitude-emulators, per chiral order, as the *emulator* for scattering observables y .

VI. SAMPLING POSTERIOR DISTRIBUTIONS

We sample the LEC posteriors $\text{pr}(\vec{\alpha}|\mathcal{D}, I)$ at NLO and NNLO using HMC, an advanced Markov chain Monte Carlo (MCMC) algorithm suitable for high-dimensional sampling problems. We have detailed the use and performance of HMC in a previous paper [21] and further improved its performance in Ref. [22]. Here, we employ the same sampling strategy as in the latter paper. HMC needs tuning in order to perform well, and in particular a good so-called mass matrix is needed. The mass matrix should be a decent approximation to the inverse of the parameter covariance matrix. We extract such an approximation by maximizing the posterior through the standard BFGS algorithm [57–60].

HMC requires gradients of the (logarithm of the) posterior with respect to the sampled parameters. The underlying linearity of the emulators furnishes easy access to derivatives of T -matrix amplitudes with respect to the LECs. However, we employ the automatic differentiation (AD) and just-in-time (JIT) compilation tools for Python in the Google JAX [61] library to compute the required gradients and boost the execution speed of our Python code; it takes on the order of one second to evaluate the entire data likelihood, including derivatives. We opted for this approach due to its simplicity and speed. It also enables straightforward computation of derivatives of the posterior with respect to parameters other than the LECs, such as \mathcal{GP} hyperparameters. A threaded C-code implementation would likely be more efficient, but not pivotal for the present applications.

We diagnose the convergence of the MCMC chains using a criterion based on the integrated autocorrelation time τ_{int} [62]. This is a measure of the autocorrelation between MCMC samples, and it tends to be underestimated for short MCMC chains. In line with Ref. [21] we therefore declare convergence when (a) the estimation of τ_{int} has stabilized, and (b) when $N \geq 50\tau_{\text{int}}$, where N is the length of the MCMC chain. All our chains readily pass this test. Like other MCMC convergence tests, the τ_{int} criterion may falsely declare convergence, e.g., if the posterior is multimodal; we therefore search for local optima using BFGS optimization initialized at random positions. We have not detected signs of multimodality and are confident that our chains are converged.

A. LEC posteriors

In a first step we compare the LEC posteriors obtained with correlated and uncorrelated EFT truncation errors. In the uncorrelated limit we determined \bar{c}^2 from the root-mean-square value of the NLO-NNLO expansion coefficients for the σ_{tot} , $\sigma(\theta)$, $A_y(\theta)$ observables, i.e., omitting the LO contribution and discarding outliers. We employed the same grids of scattering energies and angles as in the training of the correlated \mathcal{GP} model for these observables; see Sec. IV A. This leads to a fixed $\bar{c}^2 = 0.42^2$.

We find that the introduction of a finite correlation length in the EFT truncation error does not affect the marginal LEC posteriors much; see Figs. 9–12 in Appendix A. However, there are some differences and similarities worth commenting on. Most noticeably, the marginal posteriors for all LECs are approximately twice as wide when including a correlated truncation error. This is consistent with the observed reduction in the effective number of data. Other than that, the posterior correlation structure remains the same and the respective locations of the modes are largely the same, except for a significant shift in the value of the $C_{3S1-3D1}$ contact and the C_{1P1} contact LEC becoming (almost) consistent with zero in the correlated (uncorrelated) limit at NLO (NNLO). At NNLO the MAP values for the subleading πN LECs are substantially shifted with respect to the mean of their prior values in Eq. (10). This result does not change when neglecting a correlated truncation error. In the case of a correlated truncation error we obtain a MAP value for the πN LECs,

$$\widehat{\vec{\alpha}}_{\pi N} = [-0.72(2), -0.96(5), -0.01(6), 0.69(5)] \text{ GeV}^{-1}, \quad (34)$$

with 68% credible intervals indicated by error bars in parentheses. The squared Mahalanobis distance of this point with respect to the mean and covariance of the prior is $D_{\text{MD}}^2(\widehat{\vec{\alpha}}_{\pi N}, \vec{\mu}_{\pi N}; \Sigma_{\pi N}) = 9.65$, which is just far enough for the posterior MAP value to be outside 95% of the prior probability mass. This can be cast in terms of a p value (0.047) and traditional significance testing, which would lead us to reject the correctness of the Roy-Steiner prior (null hypothesis) on the 5% significance level. However, we are inclined to place doubt on our model for the truncation error and in particular its variance. As a side remark, the D_{MD}^2 values for the MAP values of the πN LEC posteriors for Δ -less NNLO [21] and for the next order N3LO [22], inferred using an uncorrelated truncation error, are even greater and again point to significant tension between the respective LECs inferred using NN data and πN data.

We suspect that the truncation error is underestimated since the NLO-NNLO shift in the Δ -full theory is on the smaller side, which is not too surprising as the inclusion of the Δ -isobar pulls down higher-order contributions in the Δ -less theory. Also, we cannot rule out that the contributions at Δ -full N3LO are substantial owing to the introduction of a rich fourth-order contact potential. To shed some light on the possible underestimation of \bar{c}^2 we sampled the joint posterior $\text{pr}(\vec{\alpha}, \bar{c}^2|\mathcal{D}, I)$ in the uncorrelated limit of the truncation error. For this, we employed the same LEC prior as defined in Sec. III A, and assumed the same \bar{c}^2 for all observable types. For the latter parameter we employed an inverse- χ^2 prior with $\nu = 23.75$ and $\tau^2 = 0.19$, as obtained via conjugacy of an inverse- χ^2 (hyper)prior with $\nu_0 = \tau_0^2 = 1$ updated using training data from expansion coefficients $c^{(3)}$ for $\sigma(\theta)$, $A_y(\theta)$, σ_{tot} on the same grid of θ and T_{lab} values as used when inferring the \mathcal{GP} model parameters for the correlated truncation error. This prior is sharply peaked at $\bar{c}^2 \approx 0.45^2$. The marginal LEC values from the $\text{pr}(\vec{\alpha}, \bar{c}^2|\mathcal{D}, I)$ posterior are not noticeably different from before. However, the posterior modes for \bar{c}^2 are 0.96(5) and 3.5(2) at NLO, and NNLO, respectively, with

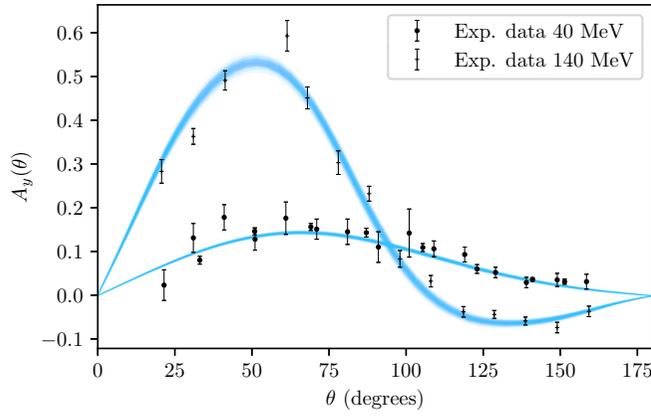


FIG. 7. 500 draws from the NLO PPDs for $A_y(\theta)$ at $T_{\text{lab}} = 40$ MeV and $T_{\text{lab}} = 140$ MeV using a correlated truncation error. The experimental data at 40 and 140 MeV are from Refs. [64,65] and [66], respectively.

68% credible intervals indicated in parentheses. Clearly, conditioning on \mathcal{D} has a significant impact and increases the EFT truncation error.

B. Posterior predictive distributions

In this section we quantify the posterior predictive distributions (PPDs) for selected np scattering observables at NLO ($k = 2$) and NNLO ($k = 3$). A PPD is a distribution of values for a, possibly unseen, observable y^* conditioned on experimental data [63]. Specifically, we will sample from $\text{pr}(y^*|\mathcal{D}, \vec{x}^*, I)$ with $y^* = y_{\text{th}}(\vec{x}^*) + \delta y_{\text{th}}(\vec{x}^*)$ at some value(s) \vec{x}^* for the independent variable. Note that we include our estimate of the truncation error in all PPDs. To achieve this, we sample $y_{\text{th}}^{(k)}(\vec{\alpha}; \vec{x}^*)$ for $\vec{\alpha} \sim \text{pr}(\vec{\alpha}|\mathcal{D}, I)$, i.e., we draw LEC values from the MCMC chains distributed according to the posteriors quantified in Sec. VI. To each sample we add a

normally distributed truncation error $\delta y_{\text{th}}^{(k)}(\vec{x}^*) \sim \mathcal{N}(0, \Sigma_{\text{th},y}^{(k)})$ where the covariance matrix is informed by a \mathcal{GP} with MAP values $\widehat{\ell}_\theta, \widehat{\ell}_{T_{\text{lab}}}, \widehat{c}^2$ according to the values listed in the table in Fig. 1.

As a model check we first inspect the PPD for some of the seen data, i.e., the data already in \mathcal{D} . In Fig. 7 we show 500 draws from the PPDs for the vector polarization $A_y(\theta)$ at NLO with $T_{\text{lab}} = 40$ MeV and $T_{\text{lab}} = 140$ MeV.

The model calculations were performed using EC emulators trained as above. The observed correlation length, $\widehat{\ell}_\theta = 28$ degrees, is reasonable, and we hit most of the data within the theory uncertainty. Due to the symmetries of the strong interaction we must have $A_y(\theta = 0) = A_y(\theta = 180) = 0$. As we incorporated this type of constraint during the training the \mathcal{GP} , the predictive variance indeed goes to zero for $\theta = 0, 180$ degrees. Given the rather long correlation lengths, the variance constraints at the angular endpoints appears to propagate to the interior to further suppress the truncation error. We have not studied the empirical coverage of the truncation error as we expect it to extrapolate rather poorly to unseen data at higher energies.

In Fig. 8 we quantify PPDs for unseen data and compare predictions based on correlated and uncorrelated truncation errors. We have drawn 500 samples from the PPDs for the differential cross section $\sigma(\theta)$ at $T_{\text{lab}} = 319$ MeV, the vector polarization $A_y(\theta)$ at $T_{\text{lab}} = 350$ MeV, and the total cross section σ_{tot} for $T_{\text{lab}} = 30\text{--}350$ MeV, where the latter is at NLO and the former two at NNLO using correlated as well as uncorrelated truncation errors.

The smoothing effect of the correlated truncation error is clearly visible in panels (a) and (c). For $\sigma(\theta)$, in panels (a) and (b), the predictions in the low- θ region fare significantly worse when including correlations. We speculate that this is connected to the reduced weight of high-energy np data residing within one correlation length of the truncation error. In panels (c) and (d) we compare $A_y(\theta)$ predictions at NNLO. Besides overall similarities, the predictive variance of $A_y(\theta)$ is finite

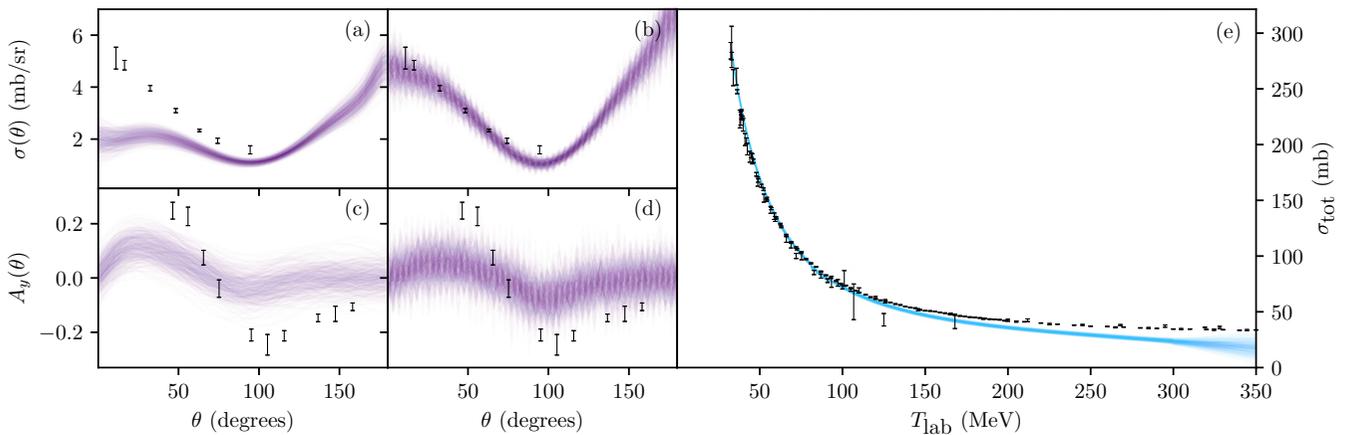


FIG. 8. 500 draws from the PPDs for unseen data for (a) $\sigma(\theta)$ at NNLO and $T_{\text{lab}} = 319$ MeV with a correlated error, and experimental data from Ref. [67]. (b) Same as (a) but with an uncorrelated error. (c) $A_y(\theta)$ at NNLO and $T_{\text{lab}} = 350$ MeV with a correlated error, and experimental data from Ref. [68]. (d) Same as (c) but with a correlated error. (e) σ_{tot} at NLO and $T_{\text{lab}} = 30\text{--}350$ MeV (see Ref. [42] for references to the experimental data).

at the edges even when we include a correlated truncation error for which we have imposed the relevant symmetry constraints. However, at $T_{\text{lab}} = 350$ MeV we are approximately one correlation length away from the training point where the constraint was imposed, and its effect has deteriorated. Finally, the PPD for the total cross section in panel (e) underestimates the experimental data for $T_{\text{lab}} > 100$ MeV, and the correlated truncation error appears to be somewhat too small as well.

VII. SUMMARY AND OUTLOOK

We quantified a correlated truncation error for Δ -full χ EFT predictions of np scattering observables at NLO and NNLO. The correlation structure was modeled using a \mathcal{GP} with a squared-exponential kernel, and the resulting MAP values for the correlation lengths in the scattering energy and angle directions are in the ranges of 45–83 MeV and 24–39 degrees, respectively, for the 17 different observable types in the set of np scattering data that we consider. These are significant correlation lengths that, in principle, could reduce the effective number of independent np data in the likelihood by two orders of magnitude and therefore strongly impact the LEC inference. However, other than doubled widths of the univariate marginal posteriors for all LECs we find that the introduction of a correlated EFT truncation error does not change the structure of the LEC posteriors by much. This is explained well by the relatively small marginal variances of the truncation errors that we quantify in this paper. Indeed, we found effective dimensions of the data likelihood that were approximately 1/8 and 1/4 of the length of the experimental np database, at NLO and NNLO respectively.

The marginal variance of the uncorrelated truncation error increases up to four times when jointly sampling it with the LEC values. Thus, future inferences should attempt to marginalize predictions over the hyperparameters of the \mathcal{GP} model and if possible also the breakdown scale, or the expansion parameter Q . For this, the HMC algorithm with AD should provide the necessary leverage to enable sampling of a posterior with, at least, doubled dimensionality. There is most likely useful information to learn about the truncation error from theoretical predictions at N3LO in Δ -full χ EFT. Adding more chiral orders in the analysis of the truncation error might reveal challenging structures in the order-by-order data that call for more sophisticated \mathcal{GP} -kernel design than we employ in this paper.

It will be interesting in future studies to explore the performance of the inferred interactions in predictions of many-nucleon observables. The general *ab initio* modeling

capabilities of Δ -full chiral interactions at NNLO were studied recently in Refs. [15,16,69]. These papers targeted, in particular, predictions for heavy nuclei up to ^{208}Pb and infinite nuclear matter. A common strategy was to first use history matching to identify interaction parametrizations that give acceptable model predictions for a suite of low-energy nuclear observables. Energies and radii of light nuclei, as well as selected np phase shifts, were included sequentially in an implausibility measure that led to the removal of large regions of the LEC parameter space. Although a lower regulator cutoff $\Lambda = 394$ MeV was used in the history-matching analyses, it is still interesting to compare the identified regions with the parameter posterior PDF inferred in this paper. First, we find that the posterior mode from this paper projected on $c_{2,3,4}$ is situated just outside the non-implausible range. That is not surprising since the Roy-Steiner prior was employed without updates throughout history matching. Furthermore, the mode is located at the upper end of the non-implausible range in the C_{150} direction and at the lower edges of the C_{1P1} and C_{3P2} ranges. The significance and origin of this tension should be explored in a joint analysis of NN , few-, and many-nucleon observables.

The scattering emulators and HMC sampling presented here can be combined with history matching and reduced-basis methods for many-body observables to enable a joint analysis of chiral interactions. Such studies, however, will rely on an improved understanding of correlated EFT truncation errors and soft momentum scales in finite nuclei.

ACKNOWLEDGMENTS

We thank Jordan Melendez for providing detailed and insightful comments on the manuscript. This work was supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (Grant Agreement No. 758027), the Swedish Research Council (Grant No. 2021-04507). The computations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) partially funded by the Swedish Research Council through Grant Agreement No. 2018-05973.

APPENDIX: LEC POSTERIOR PDFS

In this Appendix we show (Figs. 9–12) the LEC posterior PDFs for the values of the LECs in Δ -full χ EFT up to NNLO, conditioned on np scattering data with $30 \leq T_{\text{lab}} \leq 290$ MeV and a truncation error in the correlated and uncorrelated limits. The corresponding MCMC chains are available upon request.

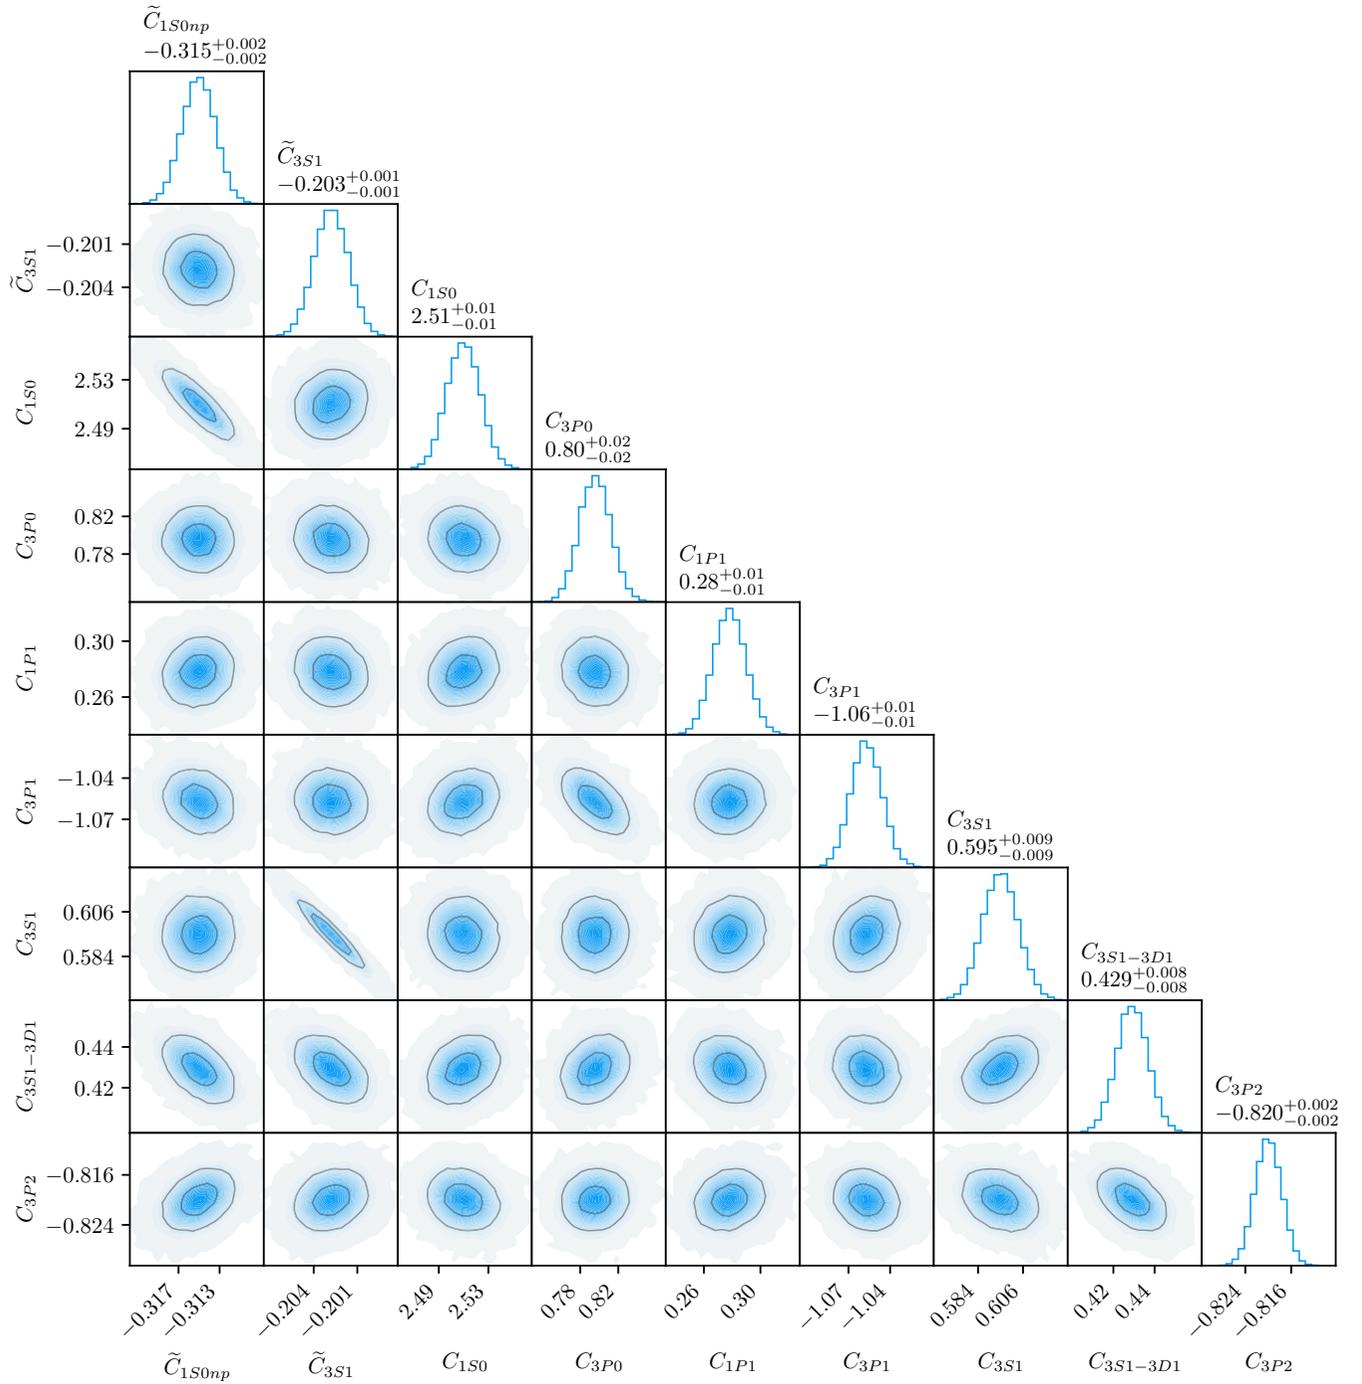


FIG. 9. Posterior PDF of the values for the LECs at NLO inferred using an EFT truncation error in the uncorrelated limit. The inner (outer) gray contour line encloses 39% (86%) of the probability mass. The dot-dashed vertical lines indicate a 68% credibility interval in the univariate marginals. See main text for units and conventions.

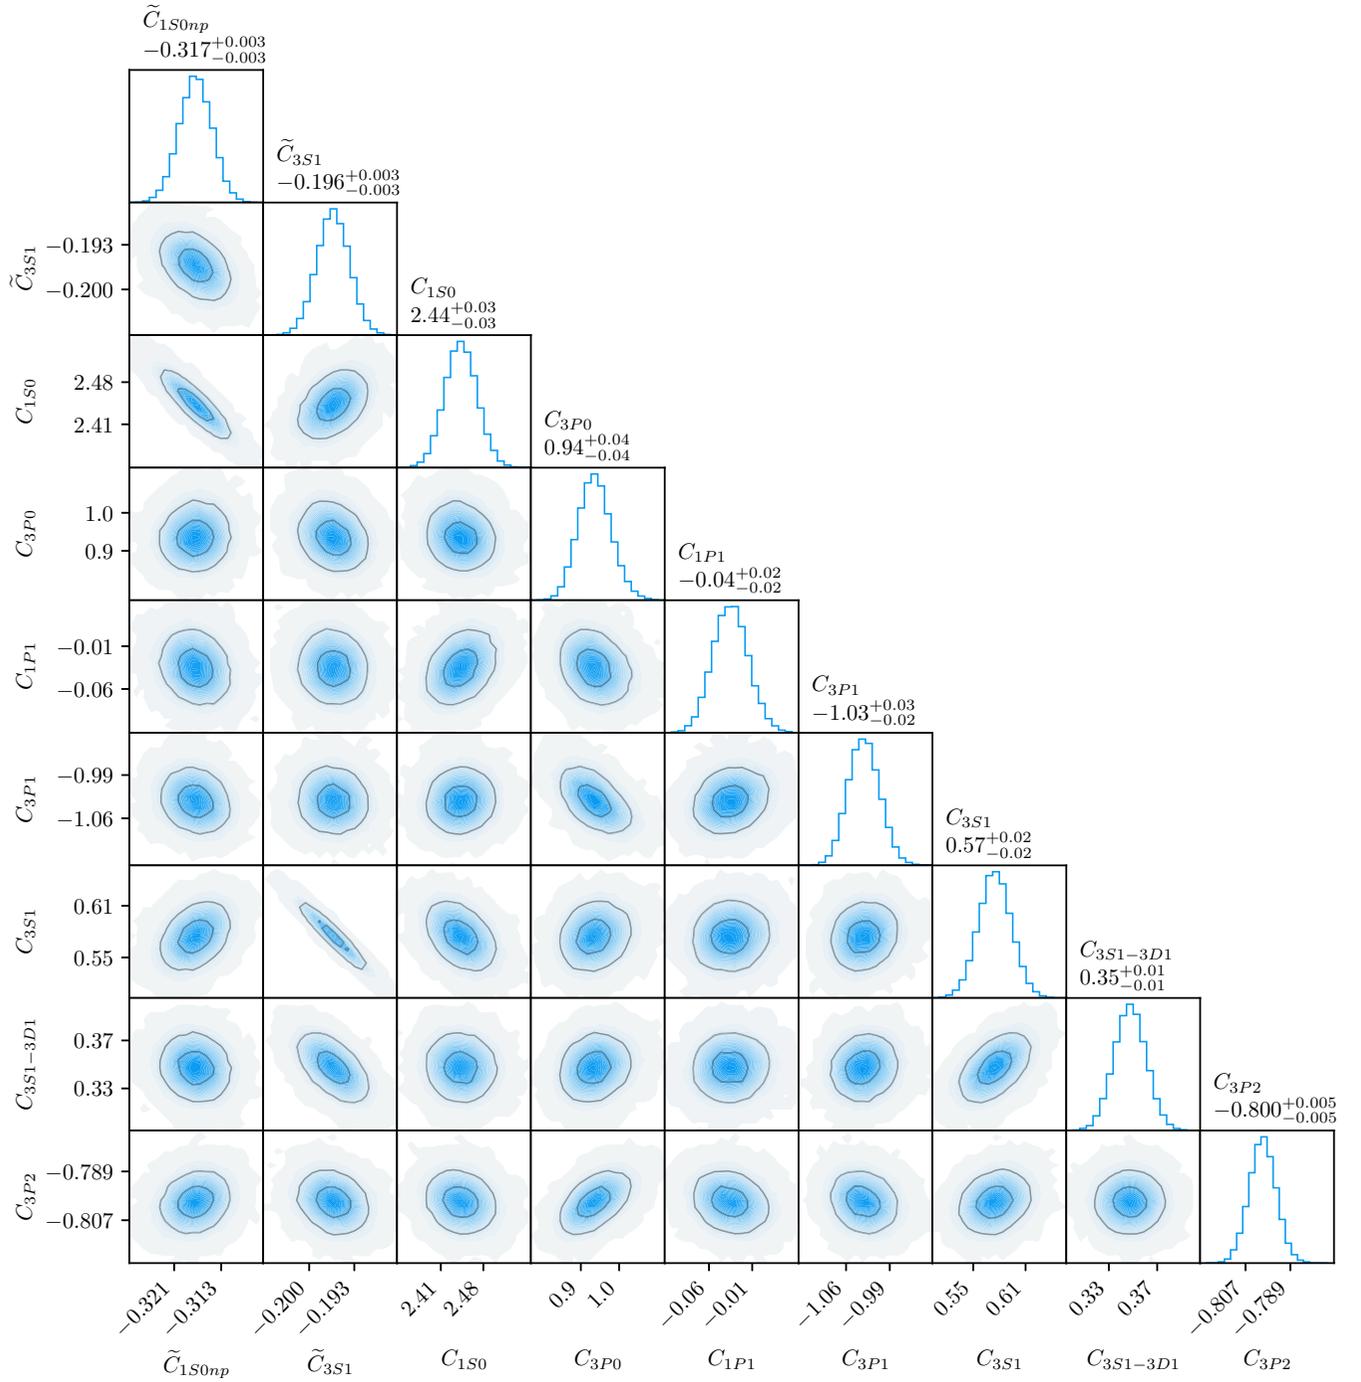


FIG. 10. Posterior PDF of the values for the LECs at NLO inferred using a correlated model for the EFT truncation error. See Fig. 9 for figure notation.

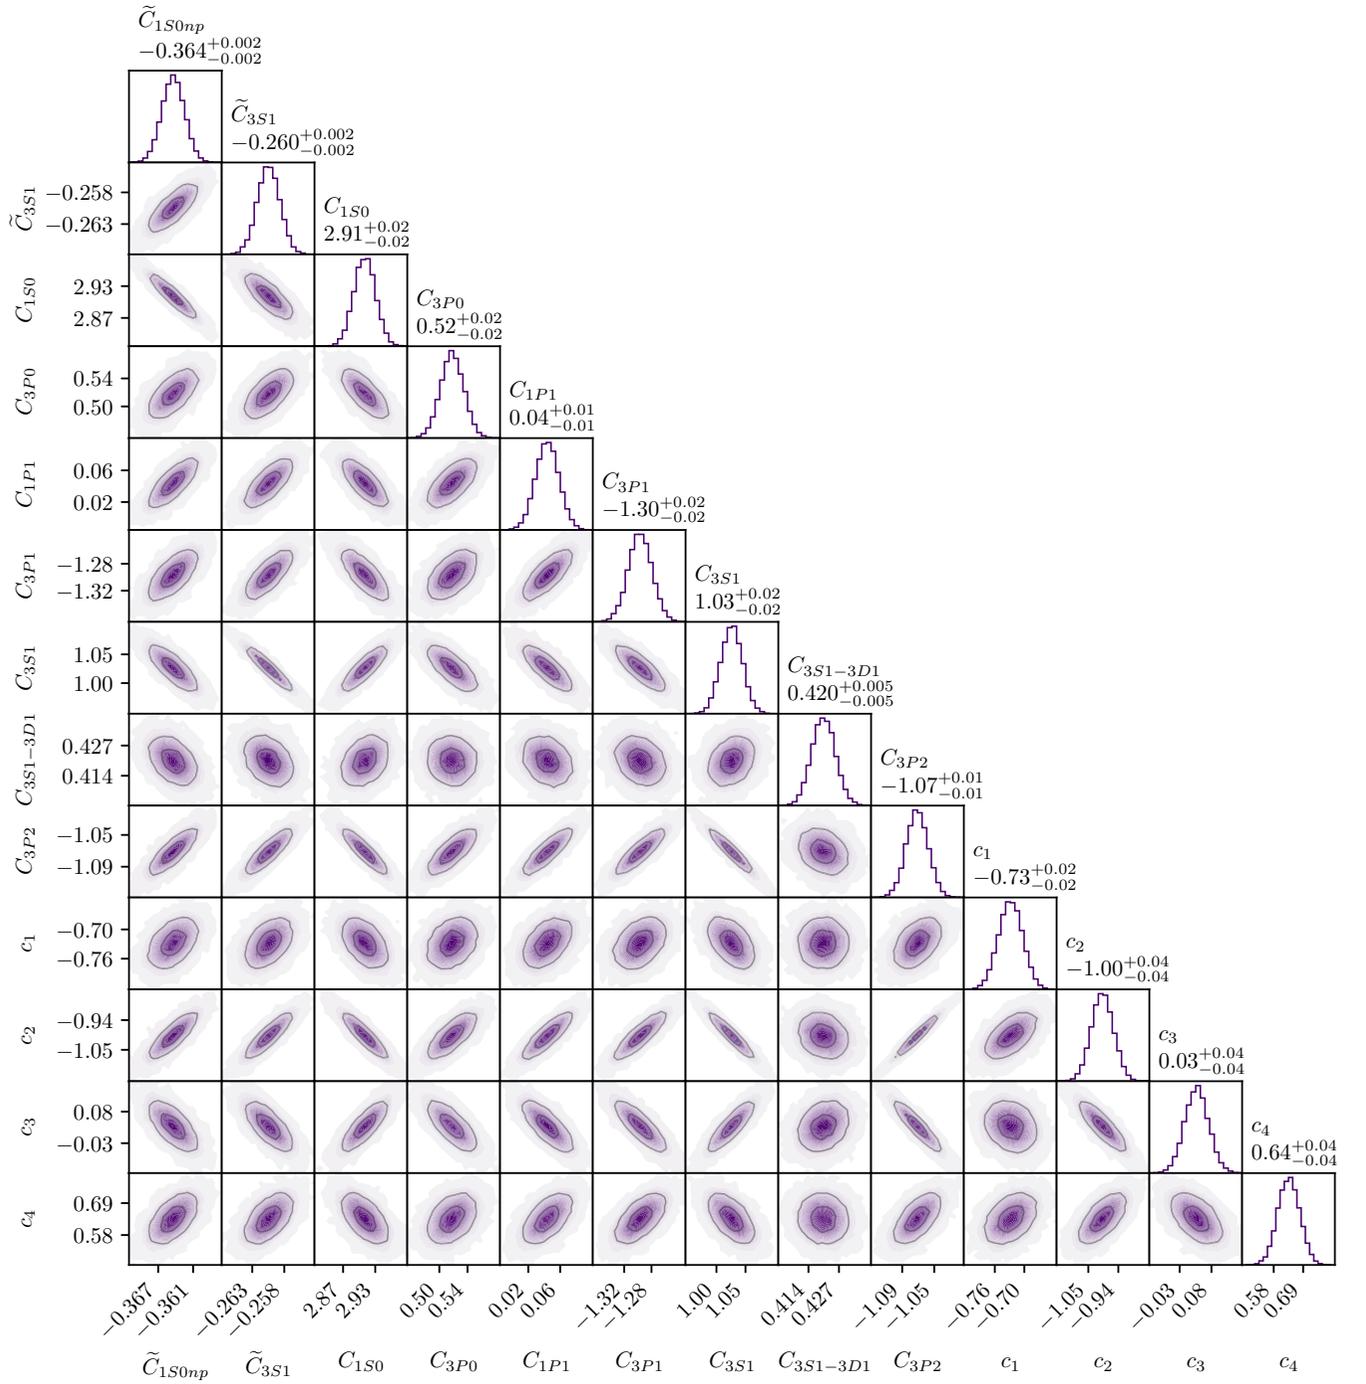


FIG. 11. Posterior PDF of the values for the LECs at NNLO inferred using an EFT truncation error in the uncorrelated limit. See Fig. 9 for figure notation.

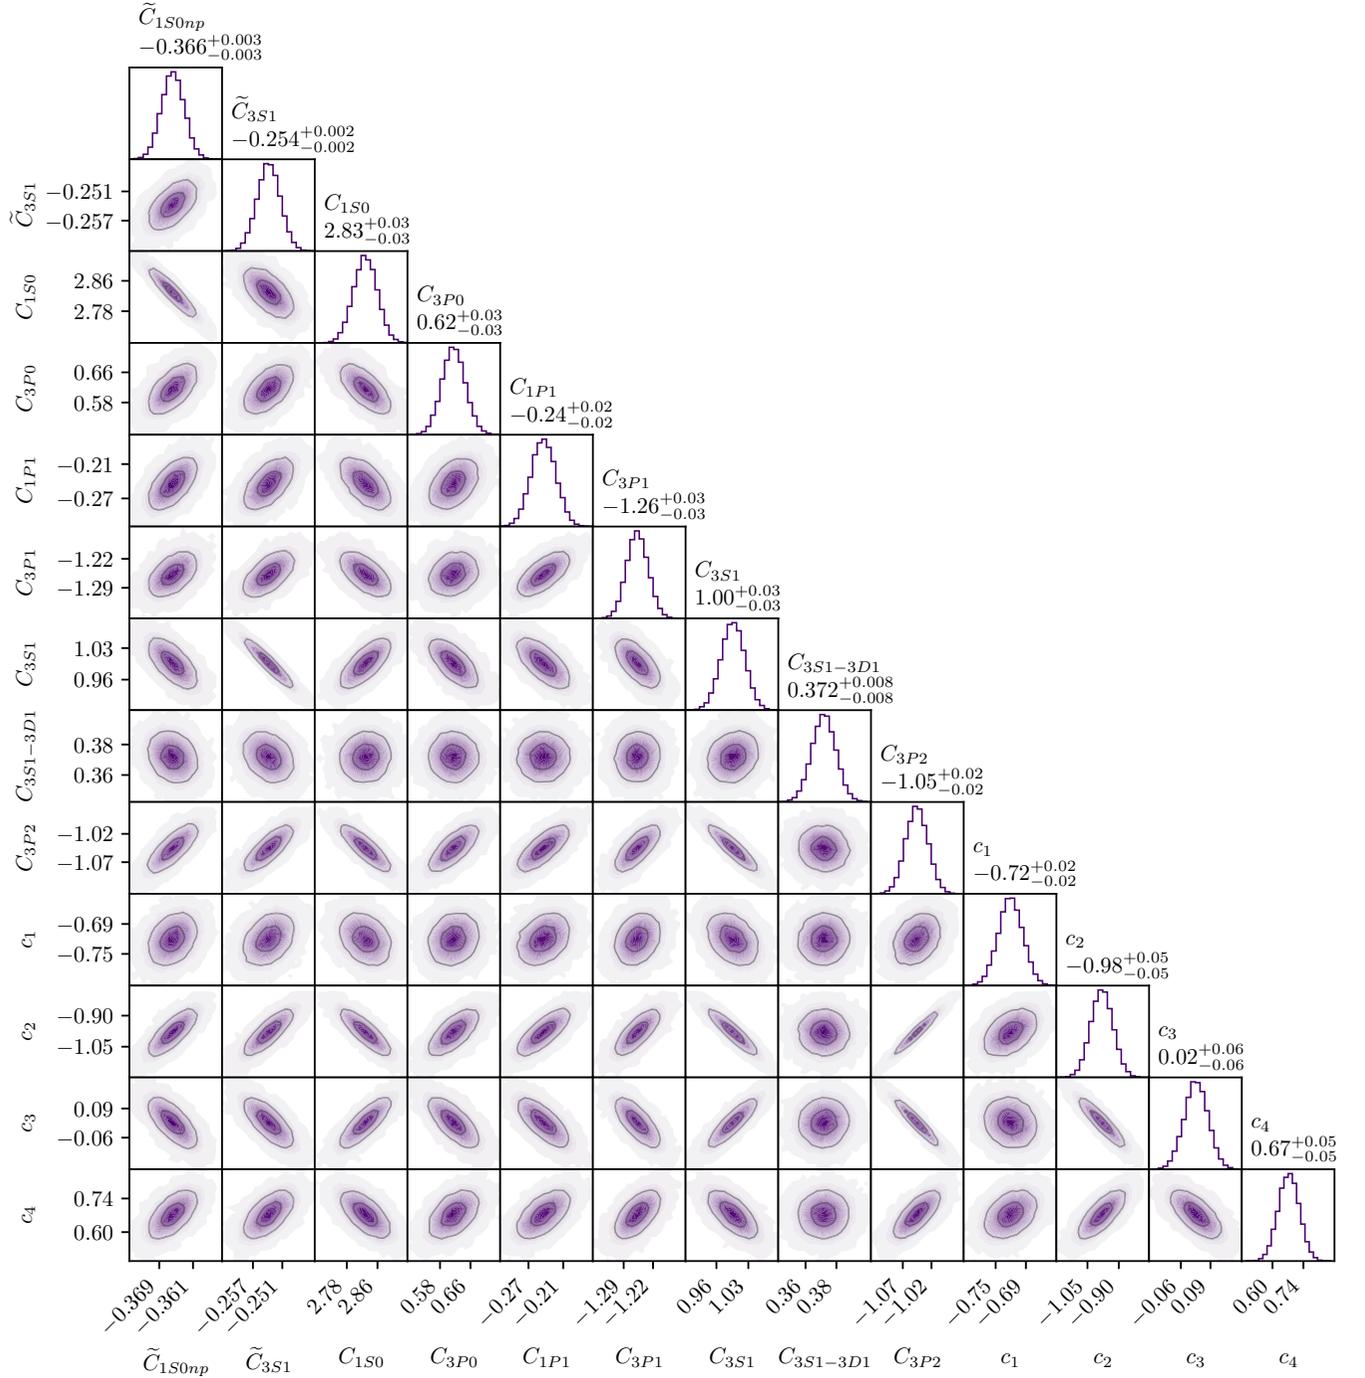


FIG. 12. Posterior PDF of the values for the LECs at NNLO inferred using a correlated model for the EFT truncation error. See Fig. 9 for figure notation.

- [1] S. Weinberg, *Phys. Lett. B* **251**, 288 (1990).
 [2] S. Weinberg, *Nucl. Phys. B* **363**, 3 (1991).
 [3] E. Epelbaum, H.-W. Hammer, and Ulf-G. Meißner, *Rev. Mod. Phys.* **81**, 1773 (2009).
 [4] R. Machleidt and D. R. Entem, *Phys. Rep.* **503**, 1 (2011).
 [5] H. W. Hammer, S. König, and U. van Kolck, *Rev. Mod. Phys.* **92**, 025004 (2020).
 [6] C. J. Yang, *Eur. Phys. J. A* **56**, 96 (2020).
 [7] S. Weinberg, *Physica A* **96**, 327 (1979).
 [8] A. Ekström, C. Forssén, G. Hagen, G. R. Jansen, W. Jiang, and T. Papenbrock, *Front. Phys.* **11**, 1129094 (2023).
 [9] R. J. Furnstahl, N. Klco, D. R. Phillips, and S. Wesolowski, *Phys. Rev. C* **92**, 024005 (2015).

- [10] 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(R) (2015).
- [11] E. Epelbaum, H. Krebs, and U. G. Meißner, *Eur. Phys. J. A* **51**, 53 (2015).
- [12] B. D. Carlsson, A. Ekström, C. Forssén, D. F. Strömberg, G. R. Jansen, O. Lilja, M. Lindby, B. A. Mattsson, and K. A. Wendt, *Phys. Rev. X* **6**, 011019 (2016).
- [13] S. R. Stroberg, J. D. Holt, A. Schwenk, and J. Simonis, *Phys. Rev. Lett.* **126**, 022501 (2021).
- [14] S. Wesolowski, I. Svensson, A. Ekström, C. Forssén, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *Phys. Rev. C* **104**, 064001 (2021).
- [15] W. G. Jiang, C. Forssén, T. Djärv, and G. Hagen, *Phys. Rev. C* **109**, L061302 (2024).
- [16] W. G. Jiang, C. Forssén, T. Djärv, and G. Hagen, *Phys. Rev. C* **109**, 064314 (2024).
- [17] J. A. Melendez, R. J. Furnstahl, D. R. Phillips, M. T. Pratola, and S. Wesolowski, *Phys. Rev. C* **100**, 044001 (2019).
- [18] J. A. Melendez, S. Wesolowski, and R. J. Furnstahl, *Phys. Rev. C* **96**, 024003 (2017).
- [19] S. Wesolowski, N. Klco, R. J. Furnstahl, D. R. Phillips, and A. Thapaliya, *J. Phys. G: Nucl. Part. Phys.* **43**, 074001 (2016).
- [20] S. Wesolowski, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *J. Phys. G: Nucl. Part. Phys.* **46**, 045102 (2019).
- [21] I. Svensson, A. Ekström, and C. Forssén, *Phys. Rev. C* **105**, 014004 (2022).
- [22] I. Svensson, A. Ekström, and C. Forssén, *Phys. Rev. C* **107**, 014001 (2023).
- [23] H. Krebs, E. Epelbaum, and U.-G. Meissner, *Eur. Phys. J. A* **32**, 127 (2007).
- [24] A. Ekström, G. Hagen, T. D. Morris, T. Papenbrock, and P. D. Schwartz, *Phys. Rev. C* **97**, 024332 (2018).
- [25] S. König, A. Ekström, K. Hebeler, D. Lee, and A. Schwenk, *Phys. Lett. B* **810**, 135814 (2020).
- [26] C. Drischler, J. A. Melendez, R. J. Furnstahl, A. J. Garcia, and X. Zhang, *Front. Phys.* **10**, 1092931 (2022).
- [27] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
- [28] C. Ordóñez, L. Ray, and U. van Kolck, *Phys. Rev. Lett.* **72**, 1982 (1994).
- [29] U. van Kolck, *Phys. Rev. C* **49**, 2932 (1994).
- [30] C. Ordóñez, L. Ray, and U. van Kolck, *Phys. Rev. C* **53**, 2086 (1996).
- [31] M. Piarulli, L. Girlanda, R. Schiavilla, R. Navarro Pérez, J. E. Amaro, and E. Ruiz Arriola, *Phys. Rev. C* **91**, 024003 (2015).
- [32] 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).
- [33] D. Logoteta, I. Bombaci, and A. Kievsky, *Phys. Rev. C* **94**, 064001 (2016).
- [34] W. G. Jiang, A. Ekström, C. Forssén, G. Hagen, G. R. Jansen, and T. Papenbrock, *Phys. Rev. C* **102**, 054301 (2020).
- [35] J. Fujita and H. Miyazawa, *Prog. Theor. Phys.* **17**, 360 (1957).
- [36] E. Epelbaum, H. Krebs, and U.-G. Meissner, *Nucl. Phys. A* **806**, 65 (2008).
- [37] M. Hoferichter, J. Ruiz de Elvira, B. Kubis, and U.-G. Meißner, *Phys. Rep.* **625**, 1 (2016).
- [38] D. Siemens, J. Ruiz de Elvira, E. Epelbaum, M. Hoferichter, H. Krebs, B. Kubis, and U. G. Meißner, *Phys. Lett. B* **770**, 27 (2017).
- [39] P. Reinert, H. Krebs, and E. Epelbaum, *Eur. Phys. J. A* **54**, 86 (2018).
- [40] T. R. Hemmert, B. R. Holstein, and J. Kambor, *J. Phys. G: Nucl. Part. Phys.* **24**, 1831 (1998).
- [41] R. Navarro Pérez, J. E. Amaro, and E. Ruiz Arriola, *Phys. Rev. C* **88**, 024002 (2013).
- [42] R. N. Pérez, J. E. Amaro, and E. R. Arriola, *Phys. Rev. C* **88**, 064002 (2013).
- [43] V. G. J. Stoks and J. J. de Swart, *Phys. Rev. C* **42**, 1235 (1990).
- [44] R. Machleidt, *Phys. Rev. C* **63**, 024001 (2001).
- [45] J. Bystricky, F. Lehar, and P. Winternitz, *J. Phys. (France)* **39**, 1 (1978).
- [46] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning Series (MIT Press, Cambridge, 2005).
- [47] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st ed. (Springer, Berlin, 2007).
- [48] C.-J. Yang, A. Ekström, C. Forssén, and G. Hagen, *Phys. Rev. C* **103**, 054304 (2021).
- [49] D. R. Entem and R. Machleidt, *Phys. Rev. C* **68**, 041001(R) (2003).
- [50] I. Vernon, S. E. Jackson, and J. A. Cumming, *SIAM/ASA J. Uncertain.* **7**, 838 (2019).
- [51] R. Byrd, P. Lu, J. Nocedal, and C. Zhu, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [52] L. S. Bastos and A. O'Hagan, *Technometrics* **51**, 425 (2009).
- [53] Y. Al-Taweel, Diagnostics and simulation-based methods for validating Gaussian process emulators, Ph.D. thesis, University of Sheffield, 2018.
- [54] M. O. Hill, *Ecology* **54**, 427 (1973).
- [55] M. D. Giudice, *Multivar. Behav. Res.* **56**, 527 (2021).
- [56] J. A. Melendez, C. Drischler, A. J. Garcia, R. J. Furnstahl, and X. Zhang, *Phys. Lett. B* **821**, 136608 (2021).
- [57] C. G. Broyden, *IMA J. Appl. Math.* **6**, 76 (1970).
- [58] R. Fletcher, *Comput. J.* **13**, 317 (1970).
- [59] D. Goldfarb, *Math. Comput.* **24**, 23 (1970).
- [60] D. F. Shanno, *Math. Comput.* **24**, 647 (1970).
- [61] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: composable transformations of Python+NumPy programs, <http://github.com/google/jax> (2018).
- [62] A. Sokal, Monte Carlo methods in statistical mechanics: Foundations and new algorithms, in *Functional Integration: Basics and Applications*, edited by C. DeWitt-Morette, P. Cartier, and A. Folacci (Springer, Boston, 1997), pp. 131–192.
- [63] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Rubin, A. Vehtari, and D. B. Dunson, *Bayesian Data Analysis*, Texts in Statistical Science Series (Chapman & Hall/CRC, London, 2014).
- [64] J. Wilczynski, J. Hansmeyer, F. Brady, P. Doll, W. Heeringa, J. Hiebert, H. Klages, and P. Plischke, *Nucl. Phys. A* **425**, 458 (1984).
- [65] A. Langsford, P. Bowen, G. Cox, G. Huxtable, and R. Riddle, *Nucl. Phys.* **74**, 241 (1965).
- [66] G. H. Stafford, S. Tornabene, and C. Whitehead, *Phys. Rev.* **106**, 831 (1957).

- [67] R. Keeler, R. Dubois, E. Auld, D. Axen, M. Comyn, G. Ludgate, L. Robertson, J. Richardson, D. Bugg, J. Edgington, W. Gibson, A. Clough, N. Stewart, and B. Dieterle, *Nucl. Phys. A* **377**, 529 (1982).
- [68] R. T. Siegel, A. J. Hartzler, and W. A. Love, *Phys. Rev.* **101**, 838 (1956).
- [69] B. Hu, W. Jiang, T. Miyagi, Z. Sun, A. Ekström, C. Forssén, G. Hagen, J. D. Holt, T. Papenbrock, S. R. Stroberg and I. Vernon, *Nat. Phys.* **18**, 1196 (2022).