Bayesian uncertainty quantification in ab initio nuclear theory

PAINT workshop, TRIUMF Vancouver, 27 Feb–1 Mar 2024

Isak Svensson

1 March 2024

SFB 1245

European Research Council
Established by the European Commission

TECHNISCHE UNIVERSITÄT DARMSTADT

CHALMERS

A little negativity

All is not yet well.

Based on:

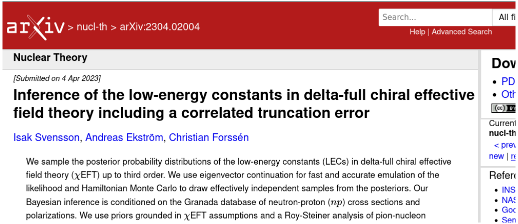

arXiv > nucl-th > arXiv:2304.02004

Search... All fi

Help | Advanced Search

Nuclear Theory

[Submitted on 4 Apr 2023]

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

Isak Svensson, Andreas Ekström, Christian Forssén

We sample the posterior probability distributions of the low-energy constants (LECs) in delta-full chiral effective field theory (χEFT) 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 χEFT assumptions and a Roy-Steiner analysis of pion-nucleon

Dov
- PD
- Oth

Current
nucl-th
< prev
new | re

Refere
- INS
- NAS
- Goo
- Sem

Still not published, despite very positive report . . .
Coming to Phys. Rev. C soon-ish.

Very recent similar paper by BUQEYE:

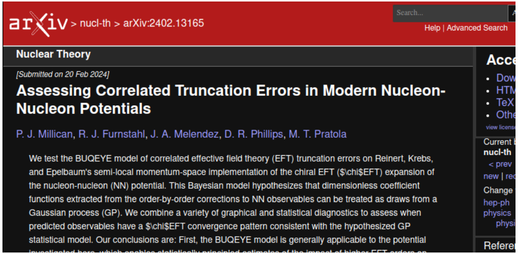

arXiv > nucl-th > arXiv:2402.13165

Search...

Help | Advanced Search

Nuclear Theory

[Submitted on 20 Feb 2024]

Assessing Correlated Truncation Errors in Modern Nucleon-Nucleon Potentials

P. J. Millican, R. J. Furnstahl, J. A. Melendez, D. R. Phillips, M. T. Pratola

We test the BUQEYE model of correlated effective field theory (EFT) truncation errors on Reinert, Krebs, and Epelbaum's semi-local momentum-space implementation of the chiral EFT (χEFT) expansion of the nucleon-nucleon (NN) potential. This Bayesian model hypothesizes that dimensionless coefficient functions extracted from the order-by-order corrections to NN observables can be treated as draws from a Gaussian process (GP). We combine a variety of graphical and statistical diagnostics to assess when predicted observables have a χEFT convergence pattern consistent with the hypothesized GP statistical model. Our conclusions are: First, the BUQEYE model is generally applicable to the potential

Acce
- Dow
- HTM
- TeX
- Othe

Current
nucl-th
< prev
new | rec

Change
hep-ph
physics
physi

Refere

Goal: to make Bayesian predictions of scattering observables in Δ-full χEFT with **correlated** EFT truncation errors

Ingredients:

❶ A method to compute observables (scattering emulators)

❷ Input data (NN scattering cross sections)

❸ A way to model correlated EFT truncation errors (Gaussian processes)

❹ Efficient MCMC sampling (Hamiltonian Monte Carlo)

Resulting in:

❺ Posteriors for LECs, predictions of observables

1. Emulators for NN scattering cross sections

We have developed a new code to compute scattering observables based on **eigenvector continuation (EC)** of T-matrix elements[1].

$$T \approx V(\vec{\alpha}) + \frac{1}{2}\vec{b}^T \mathbf{B}^{-1}\vec{b} \quad (\vec{\alpha} : \text{LECs})$$

where

$$\begin{aligned} b_i &= T_i G_0 V + V G_0 T_i \\ B_{ij} &= T_i G_0 T_j + T_j G_0 T_i - T_i G_0 V G_0 T_j - T_j G_0 V G_0 T_i. \end{aligned}$$

i, j indexes the *snapshot* (training point).

Implemented in Python. Currently no EM interactions, so limited to np scattering above 30 MeV[2]

[1] Melendez *et al.*, Phys. Lett. B **821** (2021)

[2] Stoks and De Swart, Phys. Rev. C **42** (1990) & Machleidt, Phys. Rev. C **63** (2001)

1. Emulation of a differential cross section

Emulator performance when varying just one LEC ($\widetilde{C}^{np}_{1S0}$):

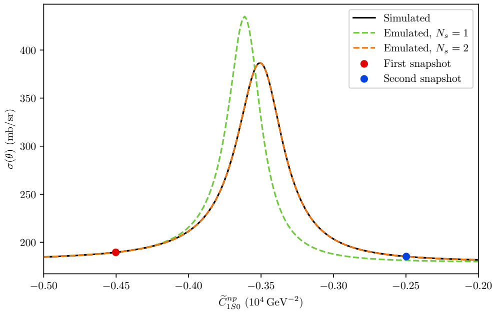

1.0 MeV, 25 degrees.

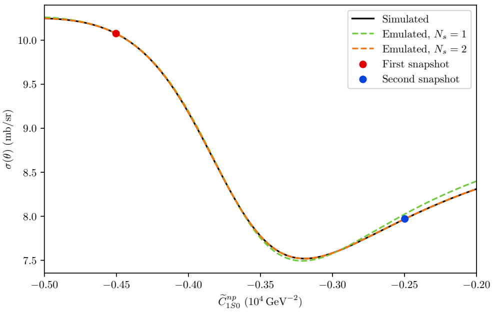

100 MeV, 25 degrees.

1. Emulator error compared to experimental error

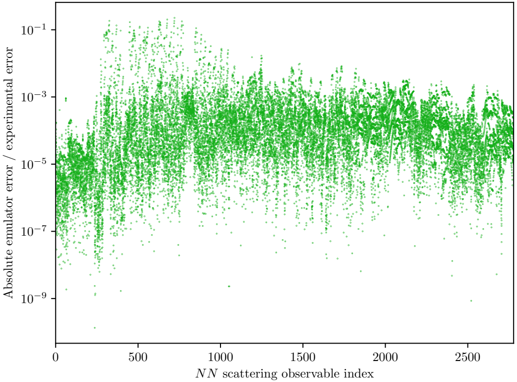

Emulator error compared to experimental errors for NNLO emulators with 8 snapshots. The x axis represents the world database of np scattering data with $0 < T_{\text{lab}} < 290$ MeV. LECs randomized in $[-4, 4]$ in appropriate units.

1. Speed and gradients through JAX

We need **gradients** of observables for HMC sampling.

- Automatic differentiation using JAX[3]
- JAX also provides just-in-time compilation, accelerating execution

With $J \leq 30$, we can compute 2,779 cross sections[4] in

- $\sim$ **0.4 s** without gradients
- $\sim$ **1.3 s** with gradients

at Δ-NNLO.

Includes evaluation of T-emulators in 87,906 PW channels.

[3] Bradbury *et al.*, `http://github.com/google/jax`

[4] Sub-30 MeV included for benchmarking. CPU: Intel Core i9-12900K.

2. Data

The data D we employ:

- neutron-proton scattering cross sections, $30 \leq T_{\mathrm{lab}} \leq 290$ MeV
- $T_{\mathrm{lab}} > 290$ MeV reserved for validation.

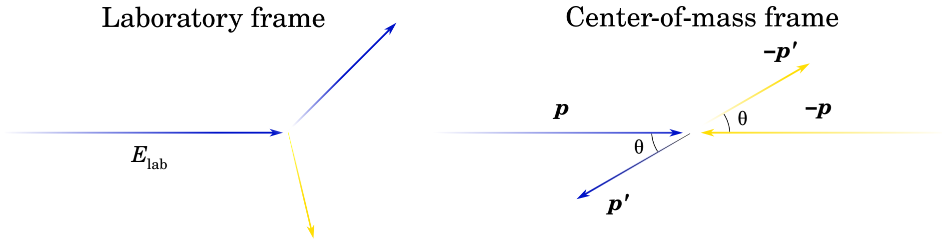

3. Bayesian inference and predictions

We make predictions of y using a **posterior predictive distribution (PPD)**:

$$\mathrm{pr}(y|D,I) = \int \mathrm{pr}(y|\vec{\alpha},I)\,\mathrm{pr}(\vec{\alpha}|D,I)\,d\vec{\alpha}$$

For this we need the joint **posterior** for the LECs $\mathrm{pr}(\vec{\alpha}|D,I)$. **Bayes' theorem**:

$$\underbrace{\mathrm{pr}(\vec{\alpha}|D,I)}_{\text{Posterior}} \propto \underbrace{\mathrm{pr}(D|\vec{\alpha},I)}_{\text{Likelihood}} \times \underbrace{\mathrm{pr}(\vec{\alpha}|I)}_{\text{Prior}}$$

Our priors are grounded in EFT. E.g., Roy-Steiner[5] for πN LECs.

[5] Hoferichter *et al.*, Phys. Rept. **625** (2016), Siemens *et al.*, Phys. Lett. B **770** (2017)

3. EFT expansion of an observable y

$$y = y_{\text{ref}} \sum_{i=0}^{k} c_i \Bigg(\underbrace{\frac{f(p, m_\pi)}{\Lambda_b}}_{\equiv Q} \Bigg)^i = y_{\text{ref}} \big(\underbrace{c_0}_{\text{LO}} + \underbrace{c_2 Q^2}_{\text{NLO}} + \underbrace{c_3 Q^3}_{\text{NNLO}} + \ldots + c_k Q^k \big)$$

Assume c_i are normally distributed:

$$\text{pr}(c_i | I) \sim \mathcal{N}(0, \bar{c}^2)$$

This leads to a normal distribution for δy[a]:

$$\delta y = \mathcal{N}(0, \sigma^2), \qquad \sigma^2 = \bar{c}^2 y_{\text{ref}}^2 \frac{Q^{2(k+1)}}{1 - Q^2}$$

[a] Wesolowski *et al.*, J. Phys. G **46** (2019), I. S., licentiate thesis (2021)

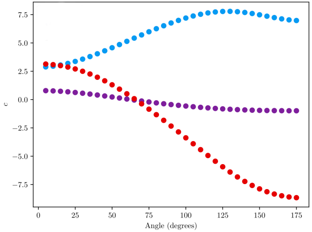

Expansion coefficients at different chiral orders. (From earlier work.)

3. Correlated EFT truncation uncertainties

The expansion coefficients c_i are **correlated** across scattering energy and angle with some finite correlation length.

We use Gaussian processes trained on known expansion coefficients to learn about **correlation lengths** and the **variance** $\bar{c}^2$.[6]

We find that accounting for correlated EFT uncertainties

- Decreases the effective number of data by a factor of 4–8 (NNLO/NLO)
- Increases the width of marginal LEC posteriors 2–3 times

We find length scales in the ranges **40–120 MeV & 25–45 degrees** and $\mathbf{\bar{c}^2 \sim 0.5^2}$ (probably underestimated)

[6] Melendez *et al.*, Phys. Rev. C **100** (2019)

Mean GP predictions of expansion coefficients for some observables

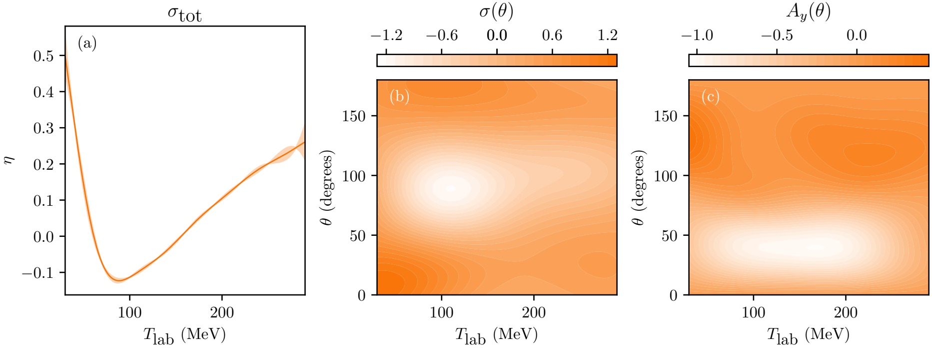

4. Sampling high-dimensional spaces with Hamiltonian Monte Carlo

Realistic χEFT potentials feature $\sim$ 15–30 LECs (or more) $\longrightarrow$ must use MCMC.

Random walk MCMC will not work very well in such high-dimensional spaces.
We obtain a guided walk using Hamiltonian Monte Carlo (HMC)[7].

HMC exploits gradients of the posterior to efficiently explore the parameter space.

JAX provides the gradients.

[7] Duane *et al.*, Phys. Lett. B **195**(2) (1987)

4. A computational balancing act

We must weigh computational cost versus quality of the samples:

- HMC increases the **computational cost** of each MCMC sample
- But also increases the **information density** (by decreasing autocorrelations)

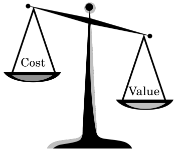

Weighing these factors we have found that[8]

HMC is ~ 6 times more efficient than `emcee`[9] in our applications.

[8] I.S., A. Ekström, C. Forssén, Phys. Rev. C **105** (2022)

[9] Foreman-Mackey *et al.*, PASP **125** (2013)

4. Autocorrelation function for the correlated Δ-NNLO sampling

The autocorrelation function $\rho(h)$ meaures the correlation between samples h MCMC steps apart

- *No* correlation: $\rho(h) = 0$
- *Full* correlation: $\rho(h) = 1$
- *Anti*-correlation: $\rho(h) < 0$

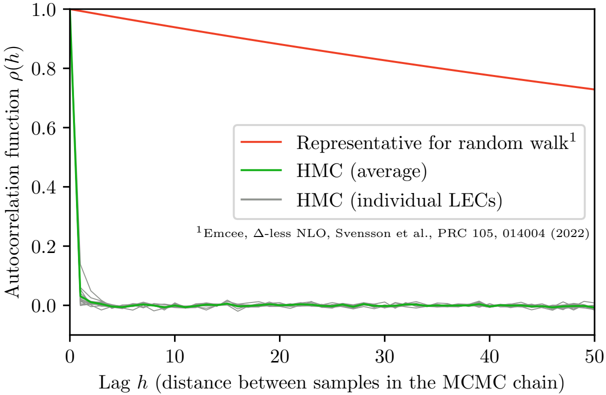

HMC provides uncorrelated samples here.

5. Posteriors for the LECs at NLO and NNLO

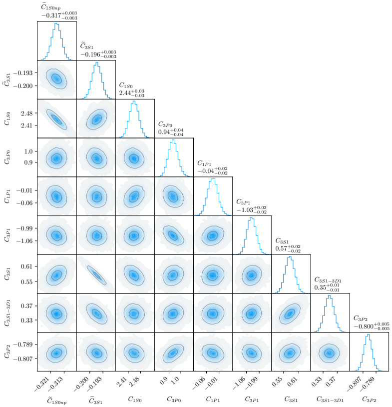

NLO posterior.

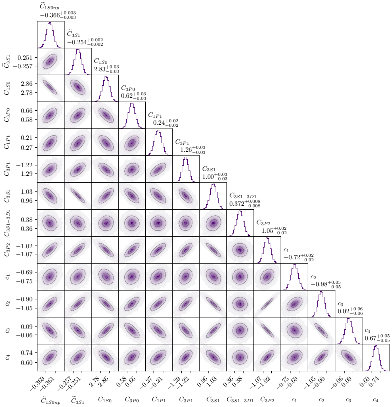

NNLO posterior.

Accounting for correlated errors **softens LEC correlations** somewhat and **increases the width** of the posteriors.

5. Tensions with Roy-Steiner prior

Comparison between R-S prior (purple contours) and MAP (black dot):

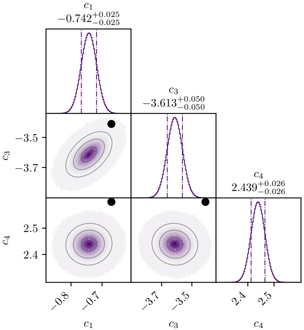

With uncorrelated error model (Δ-less χEFT)

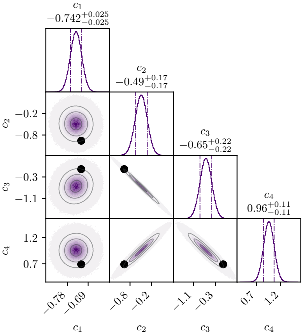

With correlated error model (Δ-full χEFT)

5. Predicting unseen data

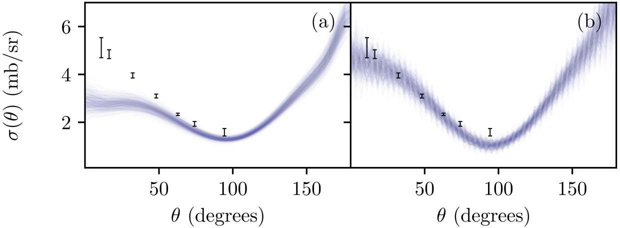

PPDs for np differential cross sections at laboratory energy 319 MeV. (a) Correlated error model. (b) Uncorrelated error model.

This prediction is better with an uncorrelated error model *but* several highly similar data sets are included in the calibration data. **Overfitting!**

Some earlier work

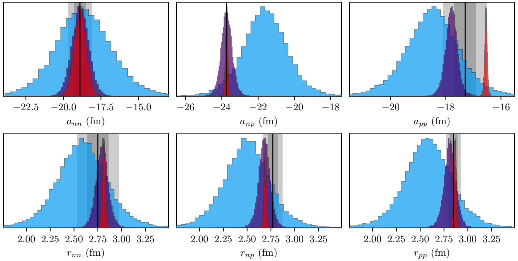

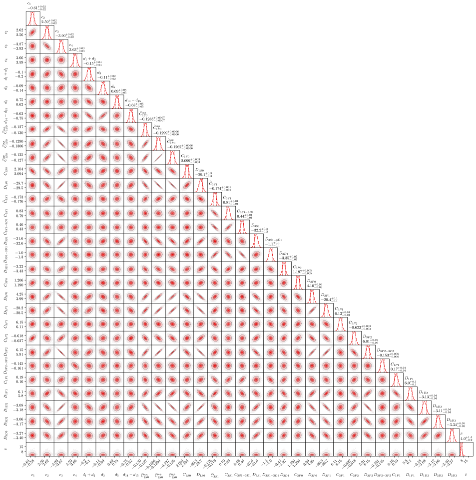

Outlook

There are lots of future prospects. To name a few:

- Improved EFT error modeling, e.g., accounting for uncertainties in more hyperparameters simultaneously.
- Simultaneous inference of two- and three-nucleon forces.
- Advances in HMC sampling technology, primarily for ease-of-use but also to tackle more challenging problems.
- Bayesian predictions of a wide variety of nuclear observables!

Thank you!

Merci!

Danke!

Tack!