DIRECT AND INDIRECT INVERSE-LAPLACE STRATEGIES FOR SPECTRAL RECONSTRUCTION

Marco Aliberti [1,2] , Francesco Di Renzo[1,2], Petros Dimopoulos[1,2], **Demetrianos Gavriel [*,1,2]** ,
Jose Alejandro Vidal Maxia[1,2]

[1]University of Parma, [2]INFN Parma, [*]**Speaker**

Istituto Nazionale di Fisica Nucleare

UNIVERSITÀ DI PARMA

SPECTRAL DENSITIES

Euclidean lattice correlators are related to spectral densities via **a Laplace transform**:

$$C(t) = \int_0^\infty dE\, e^{-Et}\, \rho(E)$$

At finite volume, the energy spectrum is discrete:

$$\rho(E) = \sum_n Z_n\, \delta(E - E_n)\ , \qquad Z_n \equiv |\langle 0|\mathcal{O}|n\rangle|^2$$

SPECTRAL DENSITIES

Euclidean lattice correlators are related to spectral densities via **a Laplace transform**:

$$C(t) = \int_0^\infty dE\, e^{-Et}\, \rho(E)$$

At finite volume, the energy spectrum is discrete:

$$\rho(E) = \sum_n Z_n\, \delta(E - E_n)\ , \qquad Z_n \equiv |\langle 0|\mathcal{O}|n\rangle|^2$$

Challenges

- The Laplace kernel **suppresses high-energy contributions** and makes **nearby states difficult to distinguish**.
- Lattice correlators contain only a **finite number of noisy Euclidean-time samples**.

Reconstructing the underlying spectral information is therefore a severely **ill-posed inverse problem**.

TWO RECONSTRUCTION STRATEGIES

1. Direct reconstruction

Reconstruct a regularized, finite-resolution spectral observable.

$$C(t) \longrightarrow \widehat{\rho}_\sigma(E)$$

- The result depends on **the resolution kernel and regularization choice.**
- Improving resolution generally increases sensitivity to statistical noise.
- Localized structures or peaks may be identified afterward.

TWO RECONSTRUCTION STRATEGIES

1. Direct reconstruction

Reconstruct a regularized, finite-resolution spectral observable.

$$C(t) \longrightarrow \widehat{\rho}_\sigma(E)$$

- The result depends on **the resolution kernel and regularization choice.**
- Improving resolution generally increases sensitivity to statistical noise.
- Localized structures or peaks may be identified afterward.

2. Indirect spectral construction

Infer a finite-state representation over a selected Euclidean-time window.

$$C(t) \longrightarrow \left\{\widehat{E}_n, \widehat{Z}_n\right\}_{n=1}^{N} \longrightarrow \quad \widehat{\rho}(E) = \sum_{n=1}^{N} \widehat{Z}_n\, \delta\left(E - \widehat{E}_n\right)$$

- The **extracted energies and weights** define an effective discrete spectral model.
- Performance depends on state separation, model order, and noise.
- A smeared spectrum can subsequently be obtained by convolution with a resolution kernel

INDIRECT STRATEGY

EXPONENTIAL-FITTING METHODS

Over a selected Euclidean-time window, approximate the correlator by a finite sum of exponentials:

$$C(t) \simeq \sum_{n=1}^{N} Z_n e^{-E_n t} \,, \qquad Z_n \geq 0$$

Several approaches have been introduced, including Prony, matrix-pencil, Lanczos, and GEVP-based methods.

They exploit the **same finite-exponential spectral structure**, but differ in how they construct and truncate the reduced eigenvalue problem and control noise-induced modes. [1]

[1] M. Fischer *et al.*, Eur. Phys. J. A **56**, 206 (2020)

EXPONENTIAL-FITTING METHODS

Over a selected Euclidean-time window, approximate the correlator by a finite sum of exponentials:

$$C(t) \simeq \sum_{n=1}^{N} Z_n e^{-E_n t} \,, \qquad Z_n \geq 0$$

Several approaches have been introduced, including Prony, matrix-pencil, Lanczos, and GEVP-based methods.

They exploit the **same finite-exponential spectral structure**, but differ in how they construct and truncate the reduced eigenvalue problem and control noise-induced modes. [1]

Reference indirect method in our benchmarks: Truncated Hankel Correlator method [2]

The THC method first constructs a **low-rank approximation of** the maximal correlator **Hankel matrix** by retaining its most significant modes.

A **Prony generalized eigenvalue problem** is then solved to extract approximations to the energy levels.

[1] M. Fischer *et al.*, Eur. Phys. J. A **56**, 206 (2020)

[2] J. Ostmeyer and C. Urbach, arXiv:2510.15500 (2025)

Z-RCNW: RATIONAL INITIALIZATION IN Z-SPACE

For equally spaced Euclidean times, $$C_n \equiv C(n\Delta t) \simeq \sum_{k=1}^{K} Z_k q_k^n \ , \qquad q_k = e^{-E_k \Delta t}$$

The finite z-transform is $$G_N(z) \equiv \sum_{n=0}^{N-1} C_n z^{-n} \simeq \sum_{k=1}^{K} Z_k \frac{1-(q_k/z)^N}{1-q_k/z}$$

For finite N, $\boldsymbol{G_N(z)}$ **is a polynomial in** $\boldsymbol{z^{-1}}$. RCNW (our version of rational approximation) fits a rational model to the sampled transform and retains poles satisfying

$$0 < q_k < 1 \ , \qquad E_k = -\frac{\log q_k}{\Delta t}$$

Z-RCNW: RATIONAL INITIALIZATION IN Z-SPACE

For equally spaced Euclidean times, $$C_n \equiv C(n\Delta t) \simeq \sum_{k=1}^{K} Z_k q_k^n \ , \qquad q_k = e^{-E_k \Delta t}$$

The finite z-transform is $$G_N(z) \equiv \sum_{n=0}^{N-1} C_n z^{-n} \simeq \sum_{k=1}^{K} Z_k \frac{1-(q_k/z)^N}{1-q_k/z}$$

For finite N, **$G_N(z)$ is a polynomial in z^{-1}**. RCNW (our version of rational approximation) fits a rational model to the sampled transform and retains poles satisfying

$$0 < q_k < 1 \ , \qquad E_k = -\frac{\log q_k}{\Delta t}$$

Reconstruction workflow of z-RCNW:

$$C_n \xrightarrow{\text{finite } z\text{-transform}} G_N(z) \xrightarrow{\text{RCNW}} \{q_k\} \xrightarrow{E_k = -\log q_k/\Delta t} \{E_k\} \xrightarrow{\text{time-domain refinement}} \{E_k, Z_k\}$$

DIRECT STRATEGY

SMEARED SPECTRAL-DENSITY RECONSTRUCTION

Direct methods reconstruct a finite-resolution spectral observable rather than attempting to identify individual delta-function states:

$$\rho_\sigma(\omega) \equiv \int_0^\infty dE\, \Delta_\sigma(\omega, E)\, \rho(E)\ , \qquad \widehat{\rho}_\sigma(\omega) = \sum_t g_t(\omega)\, C(t)$$

- **Smearing stabilizes the inverse problem** by sacrificing pointwise spectral resolution.
- The achievable **resolution is limited** by the Euclidean-time range, statistical noise, and regularization.
- The reconstructed observable should therefore be interpreted together with its achieved resolution kernel.

SMEARED SPECTRAL-DENSITY RECONSTRUCTION

Direct methods reconstruct a finite-resolution spectral observable rather than attempting to identify individual delta-function states:

$$\rho_\sigma(\omega) \equiv \int_0^\infty dE\, \Delta_\sigma(\omega, E)\, \rho(E)\ , \qquad \widehat{\rho}_\sigma(\omega) = \sum_t g_t(\omega)\, C(t)$$

- **Smearing stabilizes the inverse problem** by sacrificing pointwise spectral resolution.
- The achievable **resolution is limited** by the Euclidean-time range, statistical noise, and regularization.
- The reconstructed observable should therefore be interpreted together with its achieved resolution kernel.

Reference direct method in our benchmark: Hansen–Lupo–Tantalo method [3]

HLT chooses a target smearing kernel in advance and determines the coefficients $g_t(\omega)$ that best reproduce it, while balancing kernel resolution against statistical uncertainty.

[3] M. Hansen, A. Lupo and N. Tantalo, Phys. Rev. D 99, 094508 (2019).

ILDF: INVERSE-LAPLACE DOMAIN FITTING

First construct a regularized inverse-Laplace representation of the correlator and of a single trial exponential:

$$\mathbf{A}_{\text{data}} \equiv \mathcal{T}_\lambda[C]\ , \qquad \mathbf{A}(E) \equiv \mathcal{T}_\lambda\left[e^{-Et}\right]$$

Fit $\boldsymbol{A}_{data}$ with a finite positive spectral model

$$\mathbf{A}_{\text{model}} = \sum_{k=1}^{K} Z_k\,\mathbf{A}(E_k)\ , \qquad Z_k \geq 0$$

- Search and refine the energies by minimizing the transformed-domain residual.
- For fixed energies, determine the amplitudes by non-negative least squares.

ILDF **directly reconstructs a regularized inverse-Laplace representation**, from which candidate energy peaks may subsequently be identified.

ILDF: INVERSE-LAPLACE DOMAIN FITTING

First construct a regularized inverse-Laplace representation of the correlator and of a single trial exponential:

$$\mathbf{A}_{\text{data}} \equiv \mathcal{T}_\lambda[C] \,, \qquad \mathbf{A}(E) \equiv \mathcal{T}_\lambda\left[e^{-Et}\right]$$

Fit $\boldsymbol{A}_{data}$ with a finite positive spectral model

$$\mathbf{A}_{\text{model}} = \sum_{k=1}^{K} Z_k \, \mathbf{A}(E_k) \,, \qquad Z_k \geq 0$$

- Search and refine the energies by minimizing the transformed-domain residual.
- For fixed energies, determine the amplitudes by non-negative least squares.

ILDF **directly reconstructs a regularized inverse-Laplace representation**, from which candidate energy peaks may subsequently be identified.

More about the method and its recent developments in the upcoming talks **[Di Renzo 9:40, Vidal Maxia 10:00].**

CONTROLLED BENCHMARKS

THREE CONTROLLED TOY MODELS

The exact correlators are generated from $$C(t) = \sum_{k=1}^{K} Z_k e^{-E_k t} \ , \qquad Z_k > 0 \ , \qquad t = 0, 1, \ldots, 32$$

THREE CONTROLLED TOY MODELS

The exact correlators are generated from
$$C(t) = \sum_{k=1}^{K} Z_k e^{-E_k t} \; , \qquad Z_k > 0 \; , \qquad t = 0, 1, \ldots, 32$$

Easy model: Three well-separated states, $\Delta E_{min}^{easy} = 0.35$

Medium model: Nearby low-energy pair and weaker excited structures, $\Delta E_{min}^{medium} = 0.08$

Hard model: Six crowded states , $\Delta E_{min}^{hard} = 0.03$

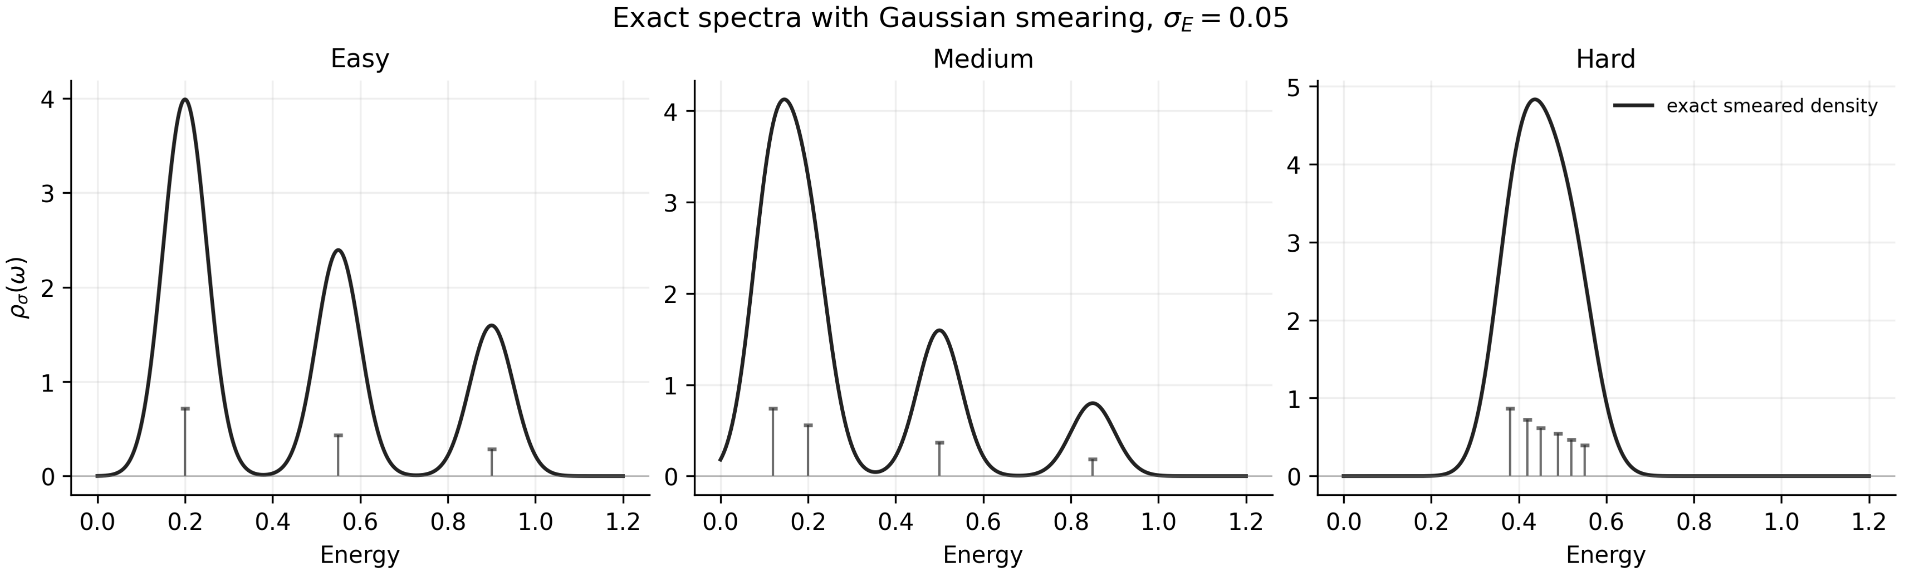

COMMON BENCHMARK PROTOCOL

Apply 2 indirect methods: **THC** (GitHub implementation)**, z-RCNW**

and 2 direct methods: **HLT** (GitHub implementation)**, ILDF**

COMMON BENCHMARK PROTOCOL

Apply 2 indirect methods: **THC** (GitHub implementation)**, z-RCNW**

and 2 direct methods: **HLT** (GitHub implementation)**, ILDF**

Noise input

- The full-sample mean C(t) is the central noisy correlator.
- **100 bootstrap means** are passed identically to THC, z-RCNW, and HLT.
- HLT regularization is selected once and held fixed across bootstrap replicas.

Finite-resolution diagnostic

$$\Delta_\rho(\sigma_E) = \int dE \, \left|\rho^{\mathrm{true}}_{\sigma_E}(E) - \rho^{\mathrm{rec}}_{\sigma_E}(E)\right|$$

- All methods are compared after applying the same Gaussian resolution.
- HLT is evaluated only as a direct smeared-density reconstruction.

RECONSTRUCTION FROM NOISELESS DATA

Easy and medium: the expected states and smeared spectra are recovered.

Hard: some nearby levels merge or shift, yet the smeared spectrum remains well reconstructed.

HLT: evaluated only as a smeared-density method, with no discrete state count.

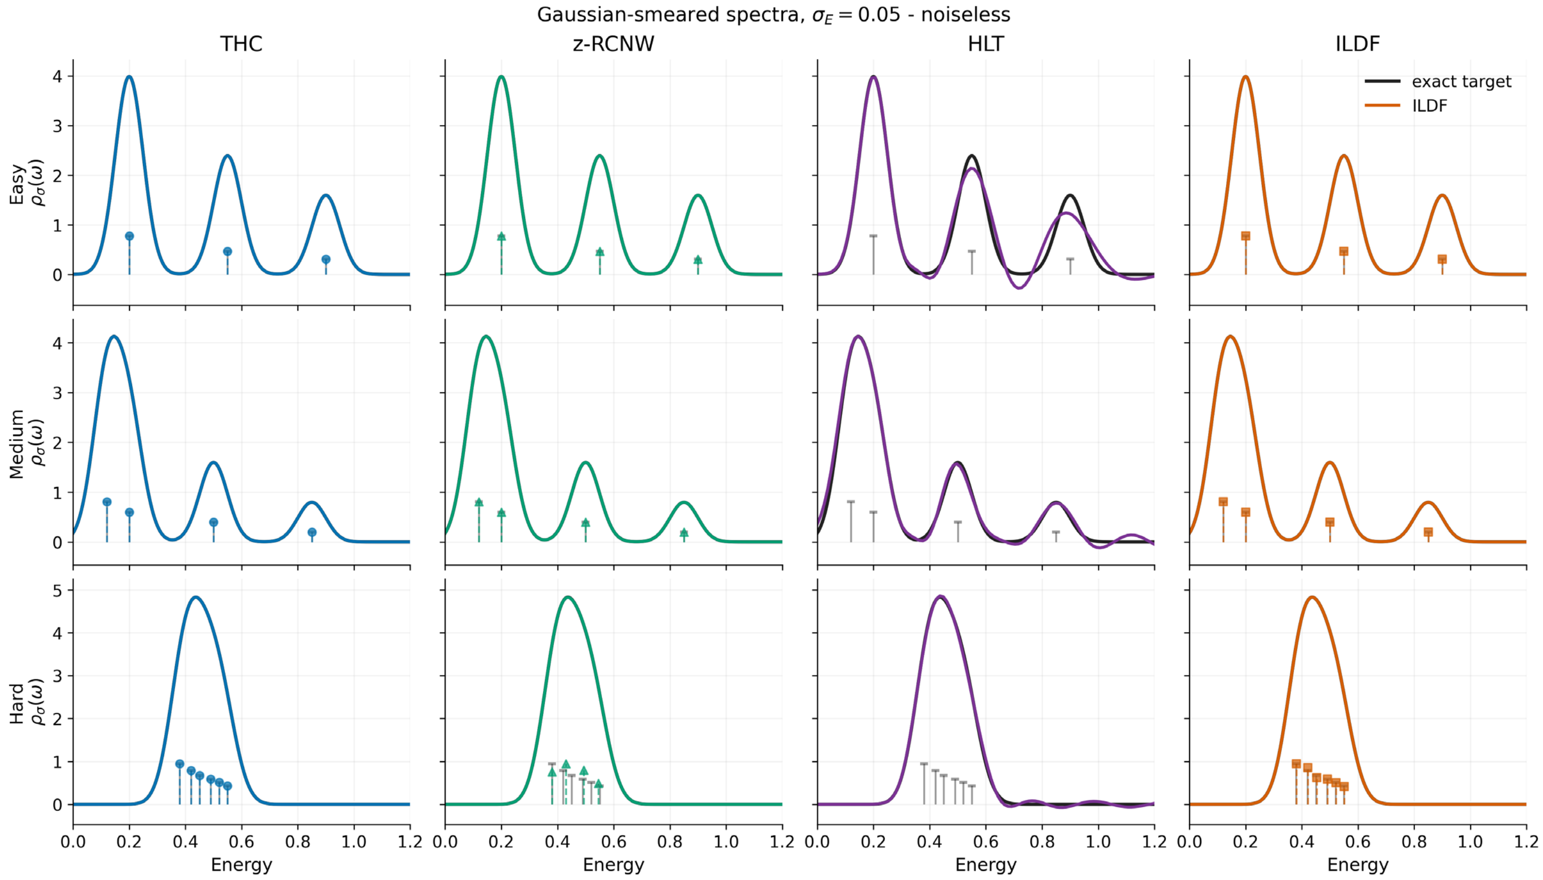

RECONSTRUCTION WITH 1% NOISE

At 1% uncertainty on the correlator mean, the focus shifts to bootstrap stability and finite-resolution consistency.

THC: median of two effective states for all models.

z-RCNW: median of three states for easy and medium, and two for hard.

HLT: direct smeared-density estimate with broad bands

ILDF: no noisy-data analysis. Work in progress.

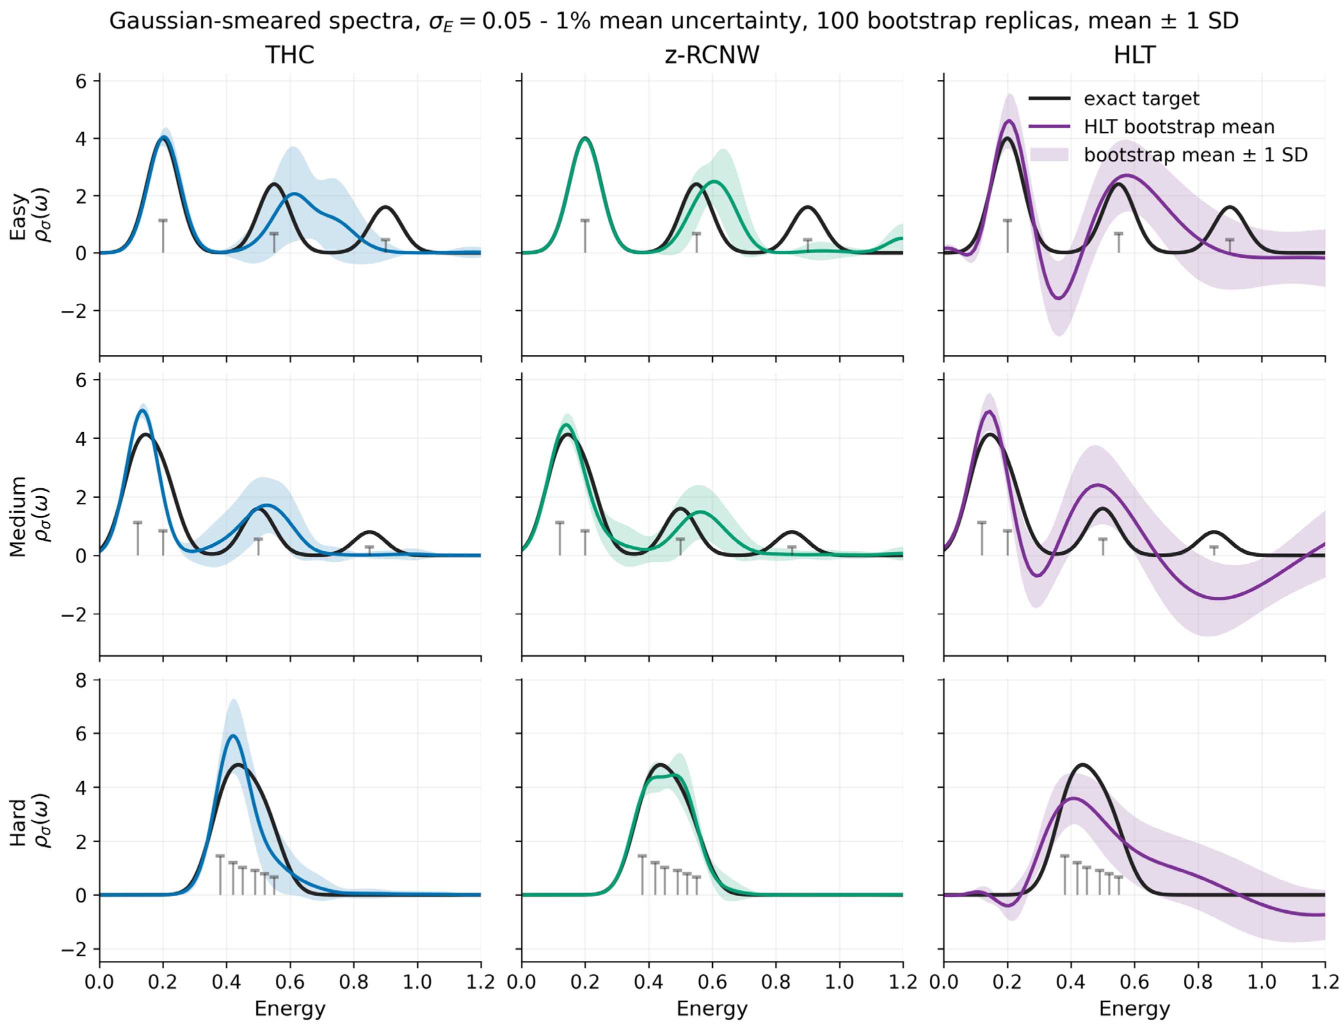

RECONSTRUCTION WITH 1% NOISE

At 1% uncertainty on the correlator mean, the focus shifts to bootstrap stability and finite-resolution consistency.

THC: median of two effective states for all models.

z-RCNW: median of three states for easy and medium, and two for hard.

HLT: direct smeared-density estimate with broad bands

ILDF: no noisy-data analysis. Work in progress.

Individual states may be unresolved, while **the smeared spectral structure remains reasonably reproduced**.

This provides a preliminary comparison of the direct and indirect strategies.

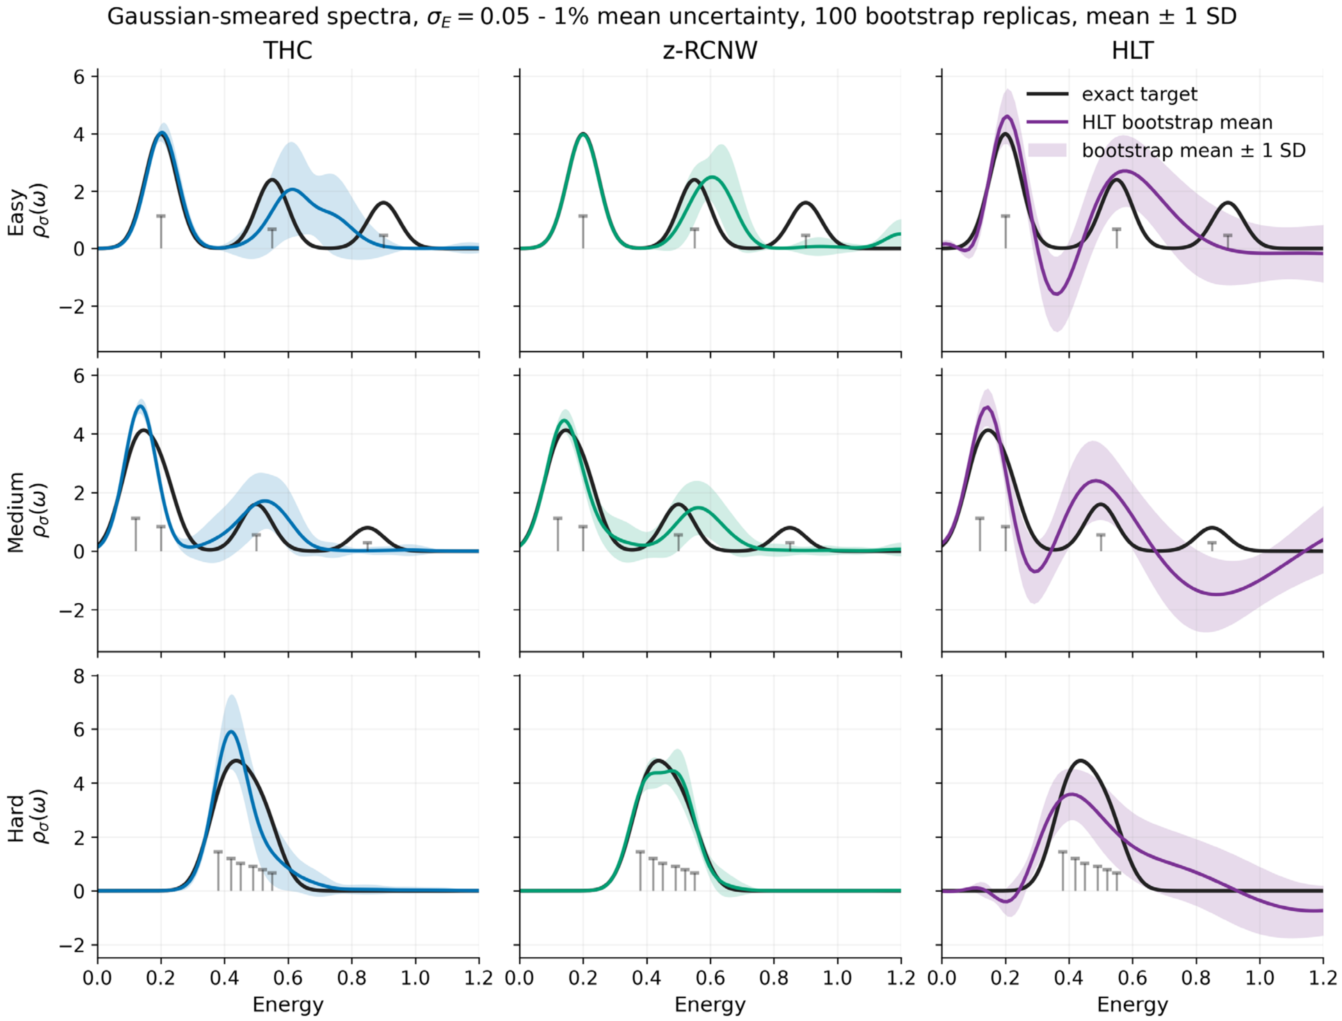

APPLICATION TO LATTICE-QCD DATA

NPLQCD DATA AND ANALYSIS SETUP

Dataset

Five NPLQCD correlator constructions corresponding to the five rows of Table I in Detmold et al.[4] :

$$N, \qquad NN\ (I = 1{,}1D), \qquad NN\ (I = 0{,}1D), \qquad NN\ (I = 1{,}3D), \qquad NN\ (I = 0{,}10D).$$

Only z-RCNW analysis

- Data-derived windows: $t_{min} \in \{1,2,3,4\},\ 8 \le N_t \le 16$.
- Full delete-one-block jackknife with 327 replicas.
- Window selection and reconstruction are repeated in every replica.

[4] W. Detmold et al., arXiv:2601.22273 (2026).

NPLQCD DATA AND ANALYSIS SETUP

Dataset

Five NPLQCD correlator constructions corresponding to the five rows of Table I in Detmold et al.[4] :

$$N, \qquad NN\ (I = 1,1D), \qquad NN\ (I = 0,1D), \qquad NN\ (I = 1,3D), \qquad NN\ (I = 0,10D).$$

Only z-RCNW analysis

- Data-derived windows: $t_{min} \in \{1,2,3,4\},\ 8 \le N_t \le 16.$
- Full delete-one-block jackknife with 327 replicas.
- Window selection and reconstruction are repeated in every replica.

Assessment

Lowest-energy consistency, correlator reconstruction, and stability under window and jackknife variations.

Reference: published block-Lanczos $E_0 a$ with statistical and bound widths.

[4] W. Detmold et al., arXiv:2601.22273 (2026).

NPLQCD GROUND-STATE COMPARISON

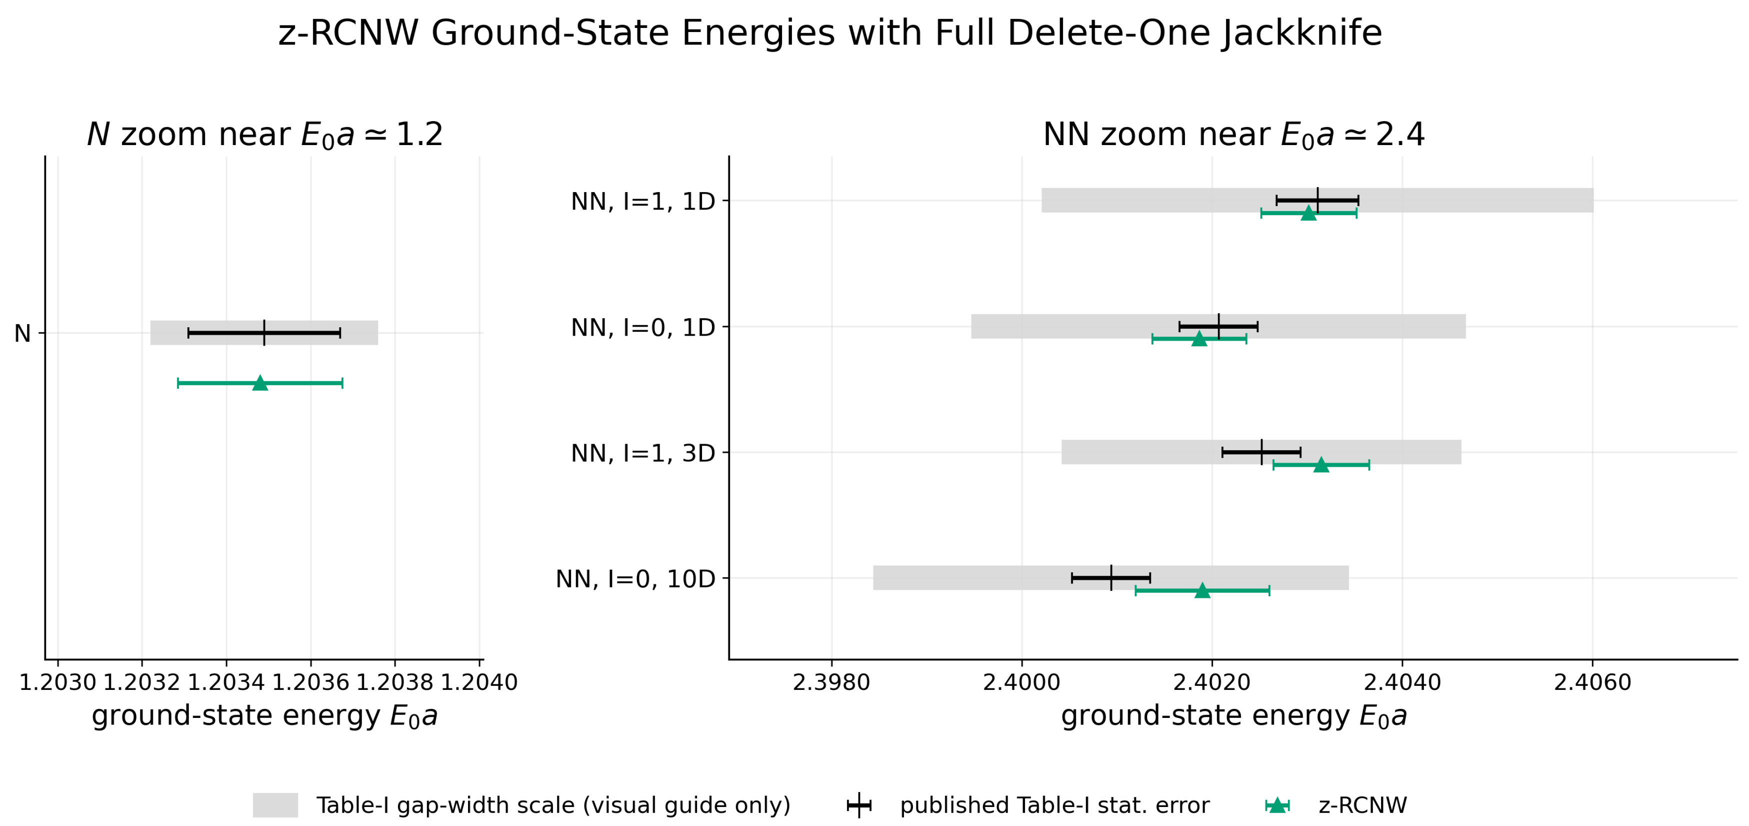

- Stable ground-state extraction is obtained for all five NPLQCD correlator constructions.
- The lowest-energy level is recovered in all 327 delete-one-block jackknife replicas.
- The reconstructed ground-state energies agree with the published Lanczos estimates within $|\Delta(E_0 a)| < 10^{-3}$

SUMMARY AND OUTLOOK

SUMMARY AND NEXT STEPS

Summary

- Direct methods recover smeared spectra; indirect methods recover effective states whose convolution reproduces the smeared structure.
- **Noiseless benchmark:** the expected structures are recovered for the easy and medium models. In the hard model, nearby levels may merge or shift into effective peaks, while the $\sigma_E = 0.05$ smeared spectrum remains well reproduced.
- **At 1% uncertainty on the correlator mean:** the recovered state count becomes unstable, but the finite-resolution spectral structure remains reasonably reconstructed.
- **NPLQCD data:** z-RCNW stably recovers the lowest-energy state in all five correlator constructions, consistently with the published Lanczos estimates.

Next steps

- Extend **ILDF to noisy correlators**, including bootstrap uncertainties and covariance weighting.
- Test the methods on additional published lattice-QCD datasets with different **channels, operators, volumes, and noise levels**.

THANK YOU

BACKUP: RCNW

Approximate sampled transform data by a compact rational model:

$$F(s_i) \simeq \widehat{F}(s_i) = \sum_{j=1}^{P} \frac{r_j}{s_i - p_j}$$

For an inverse-Laplace representation with physical decaying modes,

$$p_j = -E_j < 0 \,, \qquad f(t) = \sum_{j=1}^{P} r_j e^{-E_j t}$$

Workflow

Padé initialization provide candidate pole configurations $\longrightarrow$ noise-weighted residue fitting through regularized linear least-squares problem $\longrightarrow$ pruning and clustering of poles $\longrightarrow$ pole position refinement via variable projection $\longrightarrow$ Bayesian information criterion selection

BACKUP: BENCHMARK TABLES

Noiseless:

model	THC	z-RCNW	HLT	ILDF
easy	7.633997e-13	3.473289e-14	1.337098e-01	3.514857e-06
medium	9.587402e-13	2.427791e-13	6.774389e-02	7.057580e-06
hard	9.147813e-08	7.182133e-04	3.821673e-02	7.076600e-05

Noisy:

model	THC	z-RCNW	HLT
easy	7.038183e-01	4.928502e-01	1.165325e+00
medium	5.451476e-01	4.976796e-01	1.245215e+00
hard	4.534850e-01	1.210823e-01	1.020680e+00

BACKUP: JACKKNIFE-STABLE HIGHER-POLE CANDIDATES

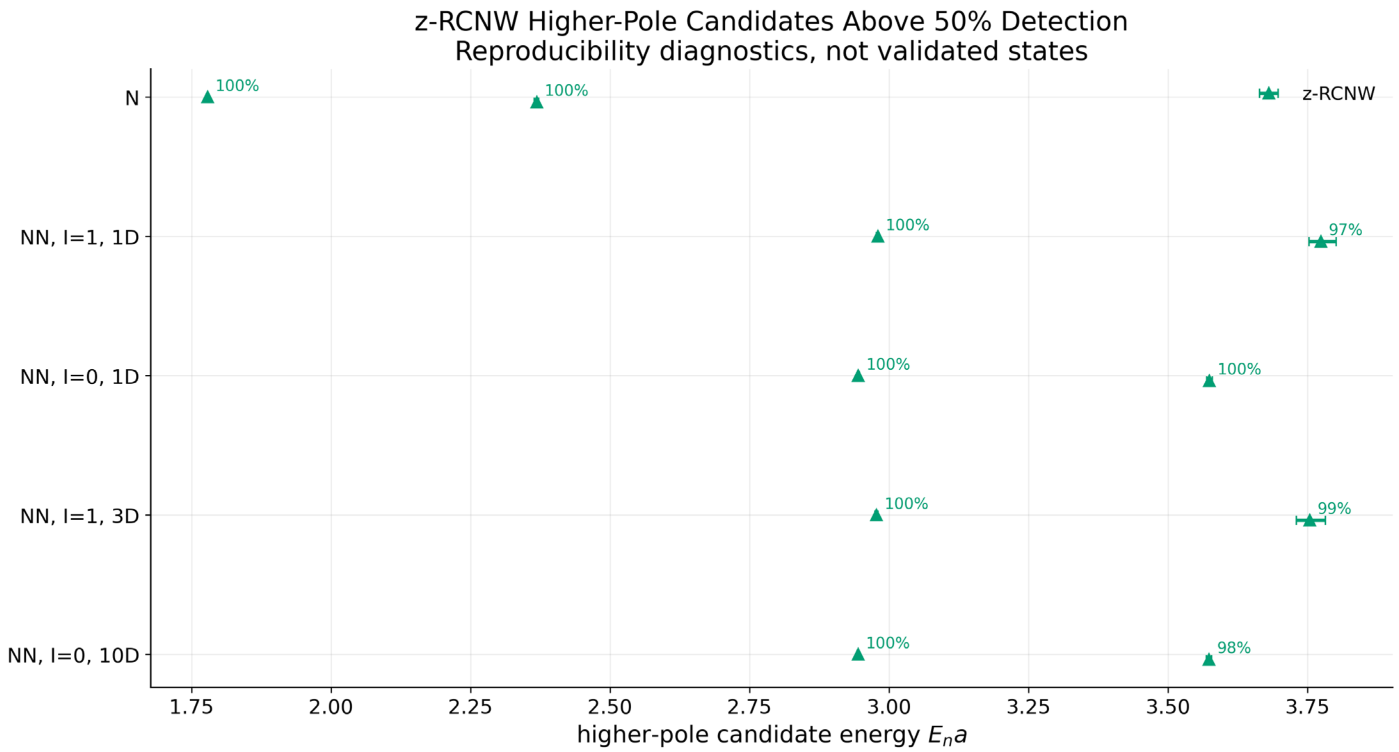

- Table I provides lowest-energy references only; no one-to-one higher-state reference exists for all five constructions.
- Candidates are retained when a compatible pole is detected in at least 50% of the 327 delete-one-block jackknife replicates; every candidate shown here is detected in 96.6%–100%.
- Horizontal intervals show the empirical 5th–95th percentile range; percentage labels give the detection fraction.
- Detection measures algorithmic reproducibility, not the probability that a candidate is a physical state.