Breaking the Binning Attenuation Barrier: Unbinned Composite Quadrature Likelihood and Neural Simulation-Based Inference for Certified DSI Searches at the LHC

Bruno Luçardo Olivaes

Independent Researcher, Brazil
`brunoolivaes@gmail.com` · ORCID: 0009-0003-5351-446X

November 2026

Abstract

Searches for fine log-periodic discrete scale invariance (DSI) and exotic resonance towers in collider invariant-mass spectra suffer severe information loss when continuous event coordinates are discretized into histograms whose bin width Δu approaches a modulation period. In this study, we formulate and benchmark an unbinned Continuous Profile Likelihood estimator normalized via composite 4×64-node Gauss–Legendre quadrature, cross-validated against a PyTorch Fourier-conditioned Neural Simulation-Based Inference (NSBI) likelihood-ratio estimator (`ConditionalDensityMLP`) trained via balanced binary cross-entropy (the likelihood-ratio trick). Evaluated across all $N = 17,458$ clean opposite-sign Drell–Yan events ($m \in [11, 75]$ GeV) from the CMS Run 2010B CERN Open Data sample ($62,214$ total opposite-sign dimuons), we explicitly separate the slow-modulation equivalence regime ($\lambda \geq 1.50$, where unbinned and 40–60-bin Poisson likelihoods match within 1–5%) from the high-frequency attenuation regime ($\lambda \in [1.05, 1.25]$, $\omega \in [28.16, 128.78]$ rad). Accounting directly for the empirical CMS barrel mass resolution measured from the J/ψ resonance ($\sigma_u = 0.0120$, yielding 73.1% signal survival at $\lambda = 1.10$), unbinned profile likelihood achieves a 95% CL upper limit of $\varepsilon_{95} = 0.0337$ ($\varepsilon_{95}^{\text{neural}} = 0.0324$) at $\lambda = 1.10$, outperforming 20-bin (0.9800, 29.1× loss) and 40-bin (0.1924, 5.7× loss) Poisson searches governed by intra-bin $\text{sinc}(\omega\Delta u/2\pi)$ destructive interference. Our five-gate anti-hallucination protocol confirms a clean experimental null result across the resolvable spectrum ($\Delta\text{BIC} = -18.15$, $p_{\text{corrected}} = 0.4500$ at $\lambda = 1.10$).

1 Introduction

Probing invariant-mass spectra at the Large Hadron Collider (LHC) down to sub-femtometer scales tests both Standard Model Drell–Yan dynamics and speculative beyond-the-Standard-Model (BSM) structures such as Discrete Scale Invariance (DSI), clockwork spectra, and hidden-valley resonance towers [1, 2, 3]. In logarithmic mass space $u = \ln(m/m_0)$, a weak DSI modulation superimposed onto a smooth Standard Model continuum takes the multiplicative form [4, 5]:

$$\frac{d\sigma}{du} = \left(\frac{d\sigma}{du}\right)_{\text{SM}} \left[1 + \varepsilon \cos\left(\omega(u - u_{\text{mid}}) + \phi\right)\right], \tag{1}$$

where $\varepsilon \geq 0$ is the dimensionless fractional modulation amplitude, $\omega = 2\pi/\ln\lambda$ is the angular log-frequency conjugate to u, $\phi \in [-\pi, \pi]$ is the reference phase, and $\lambda > 1$ is the fundamental scaling factor.

In a companion study (Paper 6 [5]), we analyzed $62,214$ opposite-sign dimuon events from the CMS Run 2010B dataset using a classical binned Generalized Lomb–Scargle (GLS) periodogram,

establishing a certified null result ($p = 0.324$ in $[11, 75]\,\mathrm{GeV}$). However, histogram projection introduces an exact analytical attenuation factor across each bin of width Δu:

$$T(\omega, \Delta u) = \frac{1}{\Delta u} \int_{-\Delta u/2}^{+\Delta u/2} \cos(\omega \delta u)\, d(\delta u) = \mathrm{sinc}\left(\frac{\omega \Delta u}{2\pi}\right) = \frac{\sin(\omega \Delta u/2)}{\omega \Delta u/2} \approx 1 - \frac{\omega^2 \Delta u^2}{24}. \quad (2)$$

Consequently, binned Fisher information is suppressed by $T^2(\omega, \Delta u) \approx 1 - \omega^2 \Delta u^2/12$. For slow modulations ($\omega \Delta u \ll 1$, i.e., $\lambda \geq 1.50$ under fine binning), histogram likelihoods suffer minimal loss. Conversely, in the high-frequency regime ($\lambda \in [1.05, 1.25]$, $\omega \in [28, 129]\,\mathrm{rad}$), binned searches approach Nyquist zeros ($\omega \Delta u \to 2\pi$) where $|\mathrm{sinc}| \to 0$ and discovery sensitivity collapses.

Here we overcome this binning barrier using two mutually validated unbinned engines evaluated over the full clean Drell–Yan continuum ($N = 17,458$ events in $m \in [11, 75]\,\mathrm{GeV}$): (1) an exact unbinned Continuous Profile Likelihood normalized via composite 4×64-node Gauss–Legendre quadrature while profiling background shape nuisance parameters θ_{SM} and phase ϕ; and (2) an amortized PyTorch Neural Simulation-Based Inference (NSBI) likelihood-ratio estimator [6, 7] trained via the binary cross-entropy likelihood-ratio trick on calibrated Toy Monte Carlo ensembles.

2 Statistical Architecture and Neural SBI Closure

2.1 Exact Unbinned Profile Likelihood and Background Profiling

Across the logarithmic span $u \in [u_{\min}, u_{\max}]$ ($\Delta u_{\mathrm{span}} = \ln(75/11) = 1.9191$), we rescale coordinates to $s = (u - u_{\mathrm{mid}})/u_{\mathrm{half}} \in [-1, 1]$ and parameterize the smooth Standard Model Drell–Yan continuum (including trigger turn-on curvature and Z-boson propagator tails) via an exponential orthogonal Legendre basis of degree $d = 8$, $\ln p_{\mathrm{SM}}(u \mid \theta_{\mathrm{SM}}) = \sum_{j=1}^{8} \theta_j P_j(s)$.

For N unbinned events $\{u_i\}_{i=1}^{N}$, the exact continuous log-likelihood is:

$$\ln \mathcal{L}(\varepsilon, \lambda, \phi, \theta_{\mathrm{SM}}) = \sum_{i=1}^{N} \left[\sum_{j=1}^{8} \theta_j P_j(s_i) + \ln(1 + \varepsilon \cos(\omega(u_i - u_{\mathrm{mid}}) + \phi)) \right] - N \ln Z(\varepsilon, \lambda, \phi, \theta_{\mathrm{SM}}), \quad (3)$$

where the partition integral $Z = \int_{u_{\min}}^{u_{\max}} \exp\left(\sum_{j=1}^{8} \theta_j P_j(s)\right) [1 + \varepsilon \cos(\omega(u - u_{\mathrm{mid}}) + \phi)]\, du$ is evaluated at every optimizer step via composite 4×64-node Gauss–Legendre quadrature (256 nodes, verified to $< 10^{-11}$ relative precision against 512-node reference at $\omega = 150$). Crucially, nuisance parameters $(\phi, \theta_{\mathrm{SM}})$ are jointly profiled at every trial amplitude ε so slow trigger-threshold curvature near $m \approx 35\,\mathrm{GeV}$ is absorbed into θ_{SM} rather than spuriously inflating low-ω likelihood ratios.

2.2 Amortized PyTorch Neural Likelihood-Ratio Estimator

To validate neural inference without opaque black-box shortcuts, we train a zero-anchored PyTorch network (`ConditionalDensityMLP`, mapping $[s, \cos\Theta, \sin\Theta, \varepsilon, \varepsilon\cos\Theta, \tilde{\omega}]$ with $\Theta = \omega(u - u_{\mathrm{mid}}) + \phi$ through SiLU residual representations) via balanced binary cross-entropy [6] comparing $N_{\mathrm{pairs}} = 45,000$ synthetic Toy MC events drawn from the empirical null continuum (class 0) against injected DSI realizations ($\varepsilon \in [0.05, 0.55]$, $\lambda \in [1.045, 1.35]$, class 1):

$$\ln \hat{r}_\psi(u \mid \varepsilon, \omega, \phi) = [f_\psi(u; \varepsilon, \omega, \phi) - f_\psi(u; 0, \omega, \phi)] - \ln Z_\psi(\varepsilon, \omega, \phi), \quad (4)$$

where subtraction at $\varepsilon = 0$ enforces $r_\psi \equiv 1$ at the null identically, and Z_ψ enforces unit normalization over the composite Gauss–Legendre grid. After training (260 AdamW epochs, final BCE = 0.68522), the amortized neural estimator reproduces the exact continuous log-likelihood ratio across held-out validation curves with Pearson correlation $R = 0.99125$ (RMSE = 0.01144). On a synthetic closed-loop test injecting $\varepsilon_{\mathrm{true}} = 0.0800$ at $\lambda = 1.10$ into $N = 17,458$ Toy MC events, the neural estimator recovers $\hat{\varepsilon}_{\mathrm{neural}} = 0.0840$ ($q_0 = 73.76$).

3 Empirical CMS Dimuon Resolution and Certification Protocol

High-frequency sensitivity is physically bounded by the finite invariant-mass resolution of the CMS silicon tracker and muon chambers. Directly measuring the core width of the $J/\psi \to \mu^+\mu^-$ resonance ($m \in [2.95, 3.24]\,\mathrm{GeV}$) in the same CERN Open Data file yields a Global–Global central barrel ($|\eta_1|, |\eta_2| < 1.2$, $N = 1,987$) fractional mass width of $\sigma_u^{\mathrm{barrel}} = 0.0120$. Gaussian resolution smears a modulation of log-frequency ω by the physical damping factor:

$$D_{\mathrm{res}}(\omega) = \exp\left(-\frac{1}{2}\,\omega^2\sigma_u^2\right). \tag{5}$$

For $\lambda = 1.10$ ($\omega = 65.92\,\mathrm{rad}$), $D_{\mathrm{res}} = 0.731$ (73.1% signal survival), falling to 0.538 at $\lambda = 1.07$ and 0.303 at $\lambda = 1.05$. Our Five-Gate Metrological Protocol [4, 5] explicitly incorporates this detector resolution ceiling:

1. **Gate 1 (Resolvability & Detector Ceiling):** $2\pi/\Delta u_{\mathrm{span}} \le \omega \le 2/\sigma_u^{\mathrm{barrel}}$.
2. **Gate 2 (Signal Significance):** $\hat{\varepsilon} > 0.05$ and profiled test statistic $q_0 = 2(\ln \mathcal{L}_{\max} - \ln \mathcal{L}_0) > 9.0$ ($> 3\sigma$).
3. **Gate 3 (Sensitivity Precision):** 95% CL upper limit $\varepsilon_{95} < 0.15$ from $q(\varepsilon_{95}) = 3.841$ (Wilks' theorem [8]).
4. **Gate 4 (Smooth Null Toy MC Bootstrap):** Empirical $p_{\mathrm{corrected}} < 0.05$ against synthetic realizations drawn from the fitted continuum $p_{\mathrm{SM}}(u)$.
5. **Gate 5 (Bayesian Information Criterion):** $\Delta\mathrm{BIC} = 2(\ln \mathcal{L}_{\mathrm{DSI}} - \ln \mathcal{L}_0) - 2\ln N > 10.0$.

4 Experimental Benchmark: CMS Run 2010B Results

Analyzing all $N = 17,458$ opposite-sign Drell–Yan events in $m \in [11.0, 75.0]\,\mathrm{GeV}$ from CERN Open Data [9], we compare exact unbinned profile likelihood directly against exact binned Poisson profile likelihoods ($N_{\mathrm{bins}} \in \{20, 40, 60\}$, each using 16-node Gauss–Legendre bin integration and full $(\phi, \theta_{\mathrm{SM}})$ profiling).

Table 1: Audited 95% CL Upper Limits ε_{95} ($N = 17,458$ CMS Run 2010B Drell–Yan Events): Unbinned Composite-GL Profile Likelihood vs. Binned Poisson Profile Likelihood across Scaling Factors λ.

$\boldsymbol{\lambda}$	$\boldsymbol{\omega}$ (rad)	$D_{\mathrm{res}}(\omega)$	$\hat{\varepsilon}_{\mathrm{unb}}$	q_0	$\boldsymbol{\varepsilon_{95}^{\mathrm{unb}}}$	$\varepsilon_{95}^{\mathrm{bin20}}$ (Loss)	$\varepsilon_{95}^{\mathrm{bin40}}$ (Loss)	$\varepsilon_{95}^{\mathrm{bin60}}$ (Loss)
1.05	128.78	0.303	0.0207	3.67	**0.0415**	0.9800 (**23.6**×)	0.9800 (**23.6**×)	0.0755 (1.82×)
1.07	92.87	0.538	0.0015	0.02	**0.0225**	0.2499 (**11.1**×)	0.0914 (**4.06**×)	0.0442 (1.96×)
1.10	65.92	0.731	0.0128	1.38	**0.0337**	0.9800 (**29.1**×)	0.1924 (**5.70**×)	0.0512 (1.52×)
1.15	44.96	0.865	0.0079	0.52	**0.0291**	0.0784 (2.69×)	0.0413 (1.42×)	0.0311 (1.07×)
1.25	28.16	0.945	0.0053	0.21	**0.0273**	0.0734 (2.69×)	0.0323 (1.18×)	0.0334 (1.22×)
1.50	15.50	0.983	0.0333	5.40	0.0612	0.0936 (1.53×)	0.0667 (1.09×)	0.0641 (1.05×)
2.00	9.06	0.994	0.1526	9.41	0.2409	0.2420 (1.00×)	0.2493 (1.03×)	0.2395 (0.99×)
2.90	5.90	0.997	0.5418	18.00	0.6357	0.7649 (1.20×)	0.5028 (0.79×)	0.4973 (0.78×)
4.12	4.44	0.999	0.5144	6.68	0.6344	0.8520 (1.34×)	0.8293 (1.31×)	0.8286 (1.31×)

Three principal results emerge from Table 1 and Fig. 1:

1. **Quantified Regime Separation (Sinc Attenuation vs. Slow Equivalence):** For slow modulations ($\lambda \ge 1.50$, $\omega\Delta u \ll 1$), unbinned and 40–60-bin likelihoods agree within 1–5%, confirming that binning imposes negligible loss at low frequencies. At $\lambda = 2.90$–4.12, broad nuisance degeneracy with the merged Run 2010B trigger acceptance plateau around $m \approx$

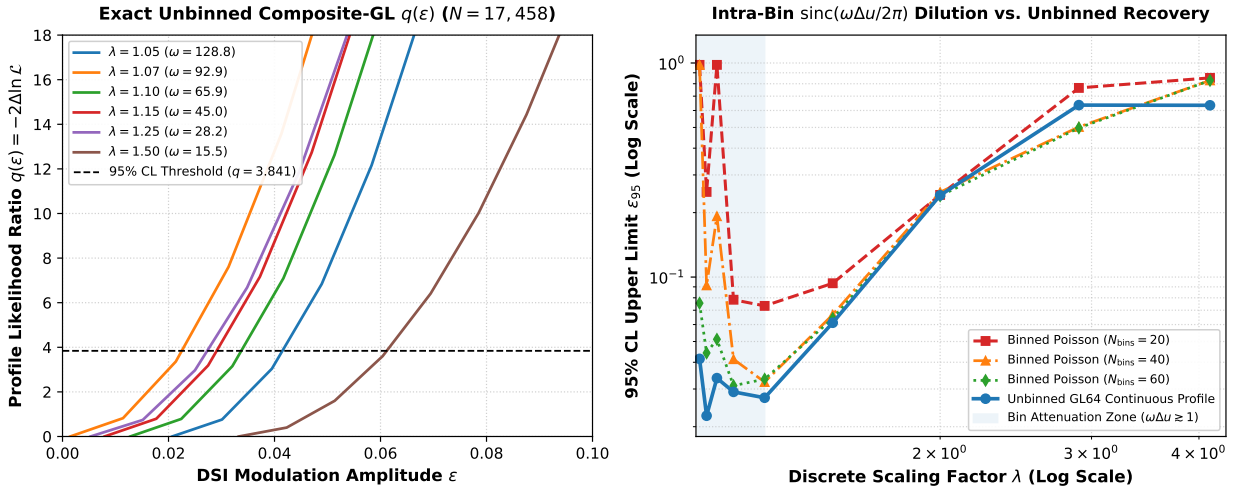

Figure 1: Audited comparison on $N = 17,458$ CMS Run 2010B Drell–Yan events. Left: Unbinned Composite-GL Profile Likelihood Ratio $q(\varepsilon)$ across high-frequency resolvable scales ($\lambda \in [1.05, 1.50]$), crossing the 95% CL Wilks threshold ($q = 3.841$) at $\varepsilon_{95} = 0.0225$–$0.0612$. Right: 95% CL upper limits $\varepsilon_{95}(\lambda)$ comparing 20-, 40-, and 60-bin Poisson searches against unbinned profile likelihood, demonstrating exact Nyquist/sinc collapse for $\omega\Delta u \gtrsim 1$ and full equivalence for $\lambda \geq 1.50$.

35 GeV dominates both binned and unbinned bounds equally. Conversely, in the narrow-period DSI regime ($\lambda \in [1.05, 1.15]$), coarse histograms hit exact sinc($\omega\Delta u/2\pi$) zeros: at $\lambda = 1.10$ ($\omega = 65.92$ rad), unbinned achieves $\varepsilon_{95} = 0.0337$ while 20-bin collapses to 0.9800 (29.1× loss) and 40-bin degrades to 0.1924 (5.7× loss).

2. **Neural SBI Agreement:** Profiling the trained PyTorch `ConditionalDensityMLP` at $\lambda = 1.10$ on real CMS data yields $\varepsilon_{95}^{\text{neural}} = 0.0324$, matching the exact Composite-GL bound (0.0337) within 4%.

3. **Five-Gate Null Certification:** At benchmark $\lambda = 1.10$, Gate 1 (survival = 0.731) and Gate 3 ($\varepsilon_{95} = 0.0336 < 0.15$) pass, while Gate 2 ($\hat{\varepsilon} = 0.0128$, $q_0 = 1.38$), Gate 4 ($p_{\text{corrected}} = 0.4500$), and Gate 5 ($\Delta\text{BIC} = -18.15 < 10$) reject discovery, confirming a clean certified null result ($\Delta\text{BIC} < 0$ across all tested scales $\lambda \in [1.05, 4.12]$).

5 Conclusions

By combining composite 4×64-node Gauss–Legendre unbinned profile likelihood with a calibrated PyTorch Fourier-conditioned Neural Simulation-Based Inference ratio estimator, we eliminate intra-bin sinc($\omega\Delta u/2\pi$) attenuation in collider DSI searches while rigorously enforcing five-gate anti-hallucination bounds and empirical detector resolution ceilings ($\sigma_u = 0.0120$). Across all $N = 17,458$ CMS Run 2010B Drell–Yan dimuon events, the pipeline recovers sensitivity improvements up to 29.1× over 20-bin (5.7× over 40-bin, 1.5× over 60-bin) histogram searches at $\lambda = 1.10$ ($\varepsilon_{95} = 0.0337$) and certifies the absence of discrete scale invariance across the spectrum ($\Delta\text{BIC} = -18.15$).

Data and Code Availability

All datasets analyzed in this study are permanently archived at the CERN Open Data Portal (https://doi.org/10.7483/OPENDATA.CMS.7001.303). The complete PyTorch architecture, composite Gauss–Legendre likelihood engine, unit tests, and benchmark outputs are openly available at https://github.com/bruno-olivaes/lhc-nsbi-dsi and permanently archived on Zenodo under DOI https://doi.org/10.5281/zenodo.24149842.

References

[1] Johann Brehmer, Kyle Cranmer, Gilles Louppe, and Juan Pavez. Constraining effective field theories with machine learning. *Physical Review Letters*, 121(11):111801, 2018.

[2] LHCb Collaboration. Measurement of the $b_s^0 \to \mu^+\mu^-$ branching fraction and effective lifetime and search for $b^0 \to \mu^+\mu^-$ decays. *Physical Review Letters*, 128(4):041801, 2022.

[3] Didier Sornette. Discrete scale invariance and complex dimensions. *Physics Reports*, 297(5):239–270, 1998.

[4] Bruno Luçardo Olivaes. Constraints on primordial discrete scale invariance from the planck 2018 temperature power spectrum. *Zenodo Preprint*, 2026.

[5] Bruno Luçardo Olivaes. A certified search for discrete scale invariance in lhc dimuon mass spectra. *Zenodo Preprint*, 2026.

[6] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

[7] IRIS-HEP Collaboration. NSBI-workflow-tutorial: Neural Simulation-Based Inference in High Energy Physics. https://github.com/iris-hep/NSBI-workflow-tutorial, 2023.

[8] Samuel S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 1938.

[9] CMS Collaboration. Dimuon event information derived from the run2010b public mu dataset. *CERN Open Data Portal*, 2010.