Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Reading Primes from Riemann Zeros: Finite-Height Noise and Detection Laws for the Zeta Echo Spectrum

Eng Ebrahem

habibamabrouk815@gmail.com

Mansoura University

Research Article

Keywords: Riemann zeta zeros, spectral form factor, Landau formula, von Mangoldt function, prime detection, Loschmidt echo, finite-height error

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11849960/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

The first N non-trivial zeros of the Riemann zeta function define a spectral coherence $C_N(u) = |N^{-1}\Sigma\, e^{i\gamma_n u}|^2$, the survival probability of the uniformly aligned state under evolution in $u = \log x$. By Landau's formula this coherence has echoes at $u = \log n$ exactly when n is a prime power. We turn this classical fact into a quantitative finite-height readout theory. Using the first 1000 zeros ($\gamma_{1000} = 1419.42$), the estimator $\Lambda_N(n) = -\sqrt{n}\ \mathrm{Re}\ S_N(\log n)/(T/2\pi)$ recovers the von Mangoldt function $\Lambda(n)$ with a noise level that follows $\sigma_\Lambda(n) \approx 0.675\, n/(T/2\pi)$ (fitted exponent 0.999, bootstrap interval 0.95–1.06). Prime powers are separable from other integers (AUC ≥ 0.90) up to a limit X satisfying $X/\log X \approx 0.97\, T/2\pi$, which corresponds to roughly $\log(T/2\pi e)$ zeros per readable prime and lies far beyond the naive resolution scale $T/2\pi$. The readout depends on the arithmetic positions of the zeros: the AUC is 0.997 for the zeta zeros and 0.50 ± 0.04 (Poisson) and 0.50 ± 0.05 (GUE) for surrogates of equal density. Errors in the ordinates attenuate the echo by an exact Debye–Waller factor $\exp(-\delta^2 \log^2 n/2)$. A smooth taper gives no consistent gain. All results are reproducible from one script.

1. Introduction

The explicit formula of Riemann and von Mangoldt expresses the prime-counting functions through a sum over the non-trivial zeros $\rho = 1/2 + i\gamma$ of the zeta function [1–4]. Each zero contributes a term proportional to $x^\rho = x^{1/2} e^{i\gamma \log x}$, an oscillation of frequency γ in the variable $u = \log x$. In this sense the zeros are the frequencies from which the primes are synthesised. The converse statement is Landau's formula [2, 3]: for fixed $x > 1$ the sum of x^ρ over zeros up to height T equals $-(T/2\pi)\Lambda(x)$ up to a lower-order error, where Λ is the von Mangoldt function. A sum over zeros therefore has a sharp response exactly at prime powers [4–6].

In the language of quantum chaos the ordinates γ_n form a spectrum, the exponential sum over them is the amplitude of the spectral form factor, and its squared modulus is the survival probability (Loschmidt echo) of a uniformly aligned state [6–10]. Montgomery's pair-correlation theorem [4] and Odlyzko's numerical work [5] place the zeros in the universality class of the Gaussian unitary ensemble (GUE) [11–15], and Berry's diagonal approximation explains the linear ramp of the form factor as a sum over prime-power echoes [6, 7]. These facts are classical.

What is less often stated is quantitative information for a *finite* set of zeros. If only the first N zeros are available, how large is the error made when primes are read from them, how does that error depend on the integer being tested and on the height $T = \gamma_N$, how many primes can be read as a function of N, and does the reading depend on the arithmetic positions of the zeros or only on their density? This paper measures these quantities on the first 1000 zeros and states them as compact empirical laws.

Contributions. (i) Exact properties of the finite coherence C_N (Section 2) and of its response to errors in the ordinates, a Debye–Waller factor (Section 3.3). (ii) A calibrated readout estimator Λ_N and measurements of its accuracy (Sections 3 and 5.1). (iii) An empirical noise law $\sigma_\Lambda(n) \approx \kappa n/(T/2\pi)$

(Section 5.2). (iv) An empirical readout limit and the resulting cost of about log(*T*/2πe) zeros per readable prime (Section 5.3). (v) A surrogate test showing that the readout is specific to the arithmetic structure of the zeros, and a negative result for tapered estimators (Sections 5.4 and 5.5). Statements are marked as Proposition (proved), Heuristic (argued, not proved) or Observation (numerical, with the stated range of validity).

2. Coherence of a finite set of zeros

Let $\gamma_1 < \gamma_2 < \dots$ be the positive ordinates of the non-trivial zeros. For $N \geq 1$ define the field and the spectral coherence [1–4], [14–16]

$$S_N(u) = \sum_{n=1}^{N} e^{i\gamma_n u}, \quad C_N(u) = |S_N(u)|^2 / N^2, \quad (1)$$

for real u (the variable u is log x in the explicit formula) [16]. All 1000 zeros used here lie on the critical line, in agreement with the verification of the Riemann Hypothesis to far greater height [17].

Proposition 1

(exact properties). (a) $0 \leq C_N(u) \leq 1$, $C_N(0) = 1$ and $C_N(-u) = C_N(u)$. (b) $N^2 C_N(u) = N + 2\sum_{n<m} \cos[u(\gamma_n - \gamma_m)]$, a sum over all pairs of zeros. (c) $C_N(u) = |\langle\varphi|e^{iD^*u}|\varphi\rangle|^2$, where $D = \mathrm{diag}(\gamma_1,\dots,\gamma_N)$ and $|\varphi\rangle$ is the uniform vector, so C_N is a survival probability. (d) If the γ_n are distinct, the long-time average of C_N is $1/N$. (e) $C_N(u) = 1 - \sigma^2 u^2 + O(u^4)$, where σ^2 is the variance of $\gamma_1,\dots,\gamma_N$; consequently the time needed to reach an orthogonal state is at least $\pi/(2\sigma)$.

Proof

(a) The triangle inequality gives $|S_N| \leq N$, and $S_N(-u)$ is the complex conjugate of $S_N(u)$. (b) Expand $|S_N|^2 = \sum_{n,m} e^{iu(\gamma_n - \gamma_m)}$. (c) $\langle\varphi|e^{iD^*u}|\varphi\rangle = S_N(u)/N$. (d) For $n \neq m$ the average of $e^{iu(\gamma_n - \gamma_m)}$ vanishes; the N diagonal terms give N/N^2. (e) Write $S_N/N = e^{i\mu u}(1 - \sigma^2 u^2/2 - i\kappa_3 u^3/6 + \dots)$ with μ the mean and κ_3 the third central moment; the modulus squared removes the phase and the cubic term enters at order u^6. The last statement is the Mandelstam–Tamm bound [10] for the spectral spread σ. ■

Item (b) is the sense in which pairs of zeros interact: the coherence is the average alignment of the pair phases $u(\gamma_n - \gamma_m)$. The scale $1/\sigma$ is the coherence length in u. Because σ grows roughly in proportion to γ_N, adding more zeros shortens the coherence length.

$$N^2 C_N(u) = N + 2\sum_{n<m} \cos[u(\gamma_n - \gamma_m)] \quad (2)$$

$$C_N(u) = 1 - \sigma^2 u^2 + O(u^4) \quad (3)$$

3. Prime readout from a finite set of zeros

3.1 Landau's formula and the estimator

Landau's formula states that for fixed $x > 1$ [2]

$\Sigma_{0<\gamma\le T} x^{\rho} = -(T/2\pi)\,\Lambda(x) + O(\log T)$, (4)

where $\Lambda(x) = \log p$ if x is a power of a prime p and $\Lambda(x) = 0$ otherwise [2]. On the critical line $x^{\rho} = \sqrt{x}\, e^{i\gamma \log x}$, so with $x = n$ an integer the sum equals $\sqrt{n}\, S_N(\log n)$ when $T = \gamma_N$. We therefore define the readout estimator [2]

$\Lambda_N(n) = -\sqrt{n} \cdot \mathrm{Re}\, S_N(\log n) / (T_N/2\pi),\ T_N = \gamma_N$. (5)

Taking the squared modulus and using the Riemann–von Mangoldt count $N \approx (T/2\pi)\log(T/2\pi e)$ gives the predicted height of the echo of the coherence at $u = \log n$:

$C_N(\log n) \approx [\,\Lambda(n) / (\sqrt{n} \cdot L_1)\,]^2,\ L_1 = \log(T/2\pi e)$. (6)

Equation (6) predicts that echoes occur only at prime powers, that their height is proportional to $\Lambda(n)^2/n$ (so the strongest echoes belong to primes near e^2 rather than to 2), and that all of them are of relative order $1/L_1^2$ or smaller compared with the main peak $C_N(0) = 1$.

Uniform versions of Landau's formula are due to Gonek [3]. Their error term contains a contribution of size x up to logarithmic factors. After multiplication by $\sqrt{n}/(T/2\pi)$ in (5) this becomes an error in Λ of order $n/(T/2\pi)$ up to logarithms. This is the scaling that is observed in Section 5.2; we regard it as consistent with, not derived from, Gonek's result.

3.2 Mean coherence: ramp and plateau

Heuristic 1 (diagonal approximation, after Berry [6]**).** Each echo contributes the area $(T/2\pi)\Lambda(n)^2/n$ to $|S_N|^2$ in the u-variable. Using $\Sigma_{n\le x}\Lambda(n)^2/n \sim \tfrac{1}{2}\log^2 x$, the mean of $|S_N(u)|^2$ over a window near u is about $(T/2\pi)u$. Dividing by N, the mean of $NC_N(u)$ rises linearly with slope $1/L_1$ and saturates at the plateau value 1 of Proposition 1(d) at the Heisenberg time $u \approx L_1$.

$\langle\, N\, C_N(u)\,\rangle \approx \min(u/L_1, 1)$. (7)

3.3 Robustness to errors in the ordinates

Proposition 2

(Debye–Waller factor). Let $\tilde{\gamma}_n = \gamma_n + \delta\xi_n$, where ξ_n are independent standard normal variables, and let $\tilde{S}_N$ be the corresponding field. Then for every real u [6–9]

$E[\,\tilde{S}_N(u)\,] = e^{-\delta^2 u^2/2}\, S_N(u)$, (8)

$$\mathrm{E}\,|\tilde{S}_N(u)|^2 = e^{-\delta^2 u^2}\,(\,|S_N(u)|^2 - N\,) + N. \quad (9)$$

Proof

$\mathrm{E}\,e^{i\delta\xi u} = e^{-\delta^2u^2/2}$, which gives (8). In $|\tilde{S}_N|^2$ the N diagonal terms are unaffected, and each of the off-diagonal terms is multiplied by $\mathrm{E}\,e^{i\delta(\xi_j-\xi_k)u} = e^{-\delta^2u^2}$, which gives (9). ■

At $u = \log n$ the expected echo is attenuated by $e^{-\delta^2\log^2 n/2}$. To read primes up to X the ordinates must therefore be known to within $\delta \ll 1/\log X$, a requirement that tabulated zeros, with errors far below 10^{-10}, satisfy by many orders of magnitude. For $\delta \to \infty$ Eq. (9) tends to N, i.e. C_N tends to the incoherent baseline $1/N$.

4. Data and methods

Data. The ordinates of the first 1000 zeros ($\gamma_1 = 14.134725$, $\gamma_{1000} = 1419.42248095$) were taken from a table and cross-checked against the arbitrary-precision routine zetazero of the library mpmath at $n = 1, 2, 10, 100, 250, 500, 750$ and 1000; the largest difference is 4.5×10^{-13}, the rounding of the table. The smooth Riemann–von Mangoldt count $(T/2\pi)\log(T/2\pi e) + 7/8$ evaluated at γ_{1000} equals 999.42, against the true count 1000, a difference of the size of the fluctuating term. Computations use double precision; a phase error $u\cdot 10^{-13}$ is negligible for the $u \le 10^4$ used here.

Subsets. To study dependence on N, the first N zeros are used with $N \in \{100, 150, 200, 300, 400, 500, 700, 1000\}$, and $T = \gamma_N$ in each case. Integers $n = 2,\ldots,2000$ are labelled as prime powers or not.

Noise. For each N, the root-mean-square value of $\Lambda_N(n)$ over non-prime-powers (where the true value is 0) is computed in nine logarithmically spaced bins of n, giving 64 points in total. The law $\log\sigma = \log\kappa + p\,\log[n/(T/2\pi)]$ is fitted by least squares, and a bootstrap over the 64 points gives an interval for p. The points of one N are not independent, so the interval is indicative.

Detection. Separation of prime powers from other integers is measured by the area under the receiver operating characteristic curve (AUC, computed through the Mann–Whitney statistic [16]) of $\Lambda_N(n)$ in sliding windows $(c/1.5, c]$. The AUC curve is smoothed with a three-point average, and the readout limits X_{90} and X_{95} are the window centres at which it crosses 0.90 and 0.95. An AUC of 0.5 means no information.

Surrogates. Surrogates of the same size $N = 1000$ and the same mean density are built in the unfolded variable and mapped to heights by inverting the smooth counting function. (a) Poisson: 1000 independent uniform points. (b) GUE: the bulk of the spectrum of a 2000×2000 GUE matrix, unfolded with the semicircle law. Twenty realisations of each are used. For Proposition 2, the ordinates are jittered with $\delta \in \{0.05, 0.1, 0.2, 0.3, 0.5\}$ and the readout at $n = 7$ is averaged over 60 realisations.

Reproducibility. Every number, table and figure below is generated by the accompanying script zeta_echo_reproduce.py from the table of ordinates (see Data and code availability).

5. Results

5.1 Accuracy of the readout

Table 1
Readout of the von Mangoldt function and echo heights from the first 1000 zeros.

n	$\Lambda(n)$	$\tilde{\Lambda}_{1000}(n)$	$C_{1000}(\log n)$ computed	$C_{1000}(\log n)$ Eq. (6)
2	0.693	0.696	0.01235	0.01230
3	1.099	1.086	0.02006	0.02059
4	0.693	0.680	0.00591	0.00615
5	1.609	1.609	0.02641	0.02652
6	0.000	-0.013	0.00000	0.00000
7	1.946	1.941	0.02746	0.02769
8	0.693	0.685	0.00299	0.00307
9	1.099	1.076	0.00657	0.00686
10	0.000	-0.039	0.00001	0.00000
11	2.398	2.380	0.02628	0.02675
13	2.565	2.522	0.02498	0.02590
16	0.693	0.610	0.00119	0.00154
25	1.609	1.507	0.00464	0.00530
49	1.946	1.841	0.00355	0.00396
97	4.575	4.490	0.01061	0.01104
101	4.615	4.537	0.01043	0.01079
127	4.844	4.563	0.00838	0.00946
128	0.693	0.246	0.00004	0.00019
211	5.352	4.969	0.00598	0.00695
229	5.434	4.690	0.00491	0.00660
289	2.833	2.185	0.00090	0.00142

Table 1 and Fig. 1 show that the estimator recovers $\Lambda(n)$ at small prime powers to within 0.05 ($n \le 13$), returns values within 0.04 of zero at the non-prime-powers 6 and 10, and gives 4.49 and 4.54 for the primes 97 and 101 (true values 4.575 and 4.615). The predicted echo height (6) agrees with the computed one to within 5% for primes up to 13 and within 4% at 97 and 101. For larger primes and for higher prime powers the computed echoes fall below the prediction (by 11% at 127 and 26% at 229); the power $128 = 2^7$, whose true value $\Lambda = 0.693$ is small, is at the noise level ($\hat{\Lambda} = 0.25$). The echo structure is therefore real and quantitatively predicted for n up to about 100, and degrades gradually beyond.

5.2 The noise law

Observation 1. For $100 \le N \le 1000$ and $8 \le n \lesssim 6T/2\pi$, the root-mean-square error of $\hat{\Lambda}_N(n)$ at non-prime-powers obeys

$$\sigma_\Lambda(n; T) \approx \kappa \cdot n \,/\, (T/2\pi), \quad \kappa = 0.675. \tag{10}$$

A free fit of the exponent gives $p = 0.999$ (bootstrap 95% interval 0.95 to 1.06); with p fixed at 1 the constant is $\kappa = 0.675$. The fit explains 97% of the variance of $\log \sigma$ over 64 points, with a residual scatter of a factor 1.27 (Fig. 2a). All eight values of N collapse onto one line when plotted against $n/(T/2\pi)$, so the error is controlled by the single ratio of the integer to the height. Its linear growth in n is not the $\sqrt{n}$ growth of a random-phase noise floor; it is the growth expected from the x-linear error term of the uniform Landau formula (Section 3.1).

5.3 The readout limit

Table 2
Readout limits: the largest integer X for which the smoothed AUC stays above 0.90 (X_{90}) or 0.95 (X_{95}).

N	γ_N	$T/2\pi$	X_{90}	X_{95}	$X_{90}/\log X_{90} \div (T/2\pi)$
100	236.5	37.6	149	65	0.79
150	318.9	50.7	309	148	1.06
200	396.4	63.1	367	231	0.99
300	541.8	86.2	544	139	1.00
400	679.7	108.2	701	450	0.99
500	811.2	129.1	870	572	1.00
700	1062.9	169.2	1154	694	0.97
1000	1419.4	225.9	1612	1215	0.97

Observation 2. For $100 \le N \le 1000$ the 0.90 readout limit satisfies

$X_{90} / \log X_{90} \approx 0.97 \cdot T/2\pi$ (scatter ± 0.07). (11)

Table 2 lists the limits. A free fit of the slope of $\log(X/\log X)$ against $\log(T/2\pi)$ gives 1.05 (bootstrap interval 0.93 to 1.18), compatible with 1 (Fig. 3). The 0.95 limit is less stable (ratio 0.59 ± 0.14, slope 1.25, and non-monotone at $N = 300$) because fewer prime powers fall in each window; we use the 0.90 limit. For $N = 1000$ the limits are $X_{90} \approx 1612$ and $X_{95} \approx 1215$.

Two consequences follow. First, since $\log X \approx \log p$ is the signal and (10) is the noise, equating signal to z_0 times noise gives $p/\log p = (T/2\pi)/(\kappa z_0)$, which reproduces (11) for $z_0 \approx 1.5$; the limit is therefore noise-limited. A naive resolution argument (neighbouring integers are separated by $1/n$ in u while the echo width is about $2\pi/T$) would give $n \lesssim T/2\pi \approx 226$ for $N = 1000$, yet the AUC remains 0.994 for $256 \le n < 512$. Second, as $X/\log X$ approximates the number of primes up to X, and $N \approx (T/2\pi)\log(T/2\pi e)$, Eq. (11) says that about $\log(T/2\pi e)/0.97 \approx 4.6$ zeros are needed per readable prime at $N = 1000$, and that the cost per prime grows only logarithmically. Extrapolating (outside the tested range) to $X = 10^4$ gives $T/2\pi \approx 1120$ and about 6.7×10^3 zeros.

5.4 Tapering does not help

Table 3
AUC of the readout in dyadic windows of n for uniform weights and for a Hann taper $\cos^2(\pi\gamma/2T)$ (normalised so that the leading term of Landau's formula is preserved).

***N* – weights**	**[2,16)**	**[16,32)**	**[32,64)**	**[64,128)**	**[128,256)**	**[256,512)**
100 – uniform	1.000	1.000	1.000	0.906	0.873	0.691
100 – Hann	1.000	1.000	0.986	0.851	0.808	0.651
300 – uniform	1.000	1.000	1.000	0.999	0.951	0.924
300 – Hann	1.000	1.000	1.000	0.997	0.922	0.837
1000 – uniform	1.000	1.000	1.000	1.000	0.992	0.994
1000 – Hann	1.000	1.000	1.000	1.000	1.000	0.985

A smooth taper lowers the sidelobes of the kernel of (4) but widens its main lobe. Table 3 shows no consistent gain: at $N = 100$ and $N = 300$ the uniform weights are better in every window beyond $n = 32$, and at $N = 1000$ the taper helps in [128,256) (1.000 against 0.992) but hurts in [256,512) (0.985 against 0.994). The sharp truncation at γ_N is retained throughout.

5.5 Arithmetic specificity and robustness

For $N = 1000$ and $16 \le n < 200$ the AUC is 0.997 for the zeta zeros, 0.503 ± 0.040 for Poisson surrogates and 0.504 ± 0.046 for GUE surrogates (20 realisations each; Fig. 4a). Both surrogates carry no information on prime powers. The GUE result is the sharper one: the zeros and GUE eigenvalues share their two-point statistics, so the averaged form factor of GUE reproduces the ramp of Section 3.2, but the

positions of the individual zeros are what encode the primes, and a GUE spectrum has no echoes at log *n*. A statistic that depends on the zeros only through their pair correlation cannot read primes.

Table 4
Attenuation of the echo at $n = 7$ under independent Gaussian errors of size δ in the ordinates ($N = 1000$, mean of 60 realisations).

δ	measured $\Lambda(7)/\Lambda(7)$	Eq. (8): $e^{-(\delta \log 7)^2/2}$	|difference|
0.00	**0.997**	**1.000**	**0.003**
0.05	0.992	0.995	0.003
0.10	0.976	0.981	0.005
0.20	0.938	0.927	0.011
0.30	0.835	0.843	0.008
0.50	0.598	0.623	0.025

The measured attenuation follows the Debye–Waller factor of Proposition 2 to within 0.03 in absolute value at every δ (Fig. 4b), the residual being of the size expected from 60 realisations.

5.6 Short-time law, plateau and ramp

Table 5
Short-time law (3), Mandelstam–Tamm time and long-time mean of C_N (average over $u \in [0, 3000]$, 60 001 equispaced points).

N	γ_N	σ	π/(2σ)	$C_N(10^{-3})$	$1 - \sigma^2 10^{-6}$	mean of C_N	$1/N$
100	236.52	60.94	0.02577	0.9963	0.9963	0.01000	0.01000
200	396.38	104.68	0.01501	0.9891	0.9890	0.00501	0.00500
500	811.18	218.48	0.00719	0.9532	0.9523	0.00203	0.00200
1000	1419.42	386.46	0.00406	0.8594	0.8507	0.00101	0.00100

The short-time law (3) is accurate while σu is small and departs as σu approaches 1 (0.8594 against 0.8507 at $N = 1000$, $u = 10^{-3}$). The long-time mean equals $1/N$ within 1.4% (Proposition 1d). The Mandelstam–Tamm time π/(2σ) shrinks from 0.026 to 0.0041 between $N = 100$ and 1000, so a larger set of zeros loses its aligned state faster.

The mean of $NC_N(u)$ in windows of u (Fig. 5) rises with u and saturates at 1 near $u \approx 5$, in agreement with the Heisenberg time $L_1 = 4.42$ of Heuristic 1. The agreement with the straight line u/L_1 is approximate: individual windows deviate by between 1% and 28% (the largest being 0.24 measured against 0.33 predicted for $1.2 \le u < 1.8$), so (7) is confirmed as a trend, not as an exact law.

6. Discussion

What is classical and what is added. The link between sums over zeros and prime powers (Landau, Gonek), the interpretation of the coherence as a survival probability or form factor, and the diagonal explanation of its ramp are classical [2, 3, 6, 7, 9]. The additions of this paper are quantitative and empirical: a calibrated finite-N estimator, the noise law (10), the readout limit (11) with its interpretation as a cost of about $\log(T/2\pi e)$ zeros per prime, the surrogate test, and the exact Debye–Waller factor (8). To the author's knowledge a systematic measurement of this kind has not been published, but a full literature search should be completed before submission.

Interpretation of the cost law. Eq. (11) says that the number of primes that can be read grows in proportion to the height T of the zeros, one prime for each interval of length $2\pi/0.97$, while the number of zeros grows as $T \log T$. The information per zero therefore decreases slowly with height. This gives a concrete, quantitative sense in which the zeros are an efficient but not a compressive encoding of the primes.

Possible physical realisation. A field with line frequencies $\nu_n = \nu_0\gamma_n$ has intensity $I(t) = N^2 C_N(u)\cdot$(line power) with $u = 2\pi\nu_0 t$, so a photodiode records the survival probability directly, and the square-law detector output contains a beat note at every difference $\nu_0(\gamma_n - \gamma_m)$. For $N = 100$ and $\nu_0 = 1$ MHz the beat notes extend to 222 MHz, and the echo of $n = 2$ is expected at $t \approx 110$ ns with relative height 0.030 (computed value for the first 100 zeros). Because the physics of the intensity is ordinary multimode interference, such an experiment would characterise the source, not test any new law of light. The coherence defined here also appeared in earlier preprints of the author under a physical interpretation [18]; the results above are purely mathematical and numerical and do not use that interpretation.

7. Limitations

(i) All laws are empirical and are measured for $N \le 1000$ ($T \le 1419$); the exponent in (10) and the proportionality in (11) have not been tested at larger height, where tables of zeros at much greater height would allow a direct check. (ii) The bootstrap intervals treat the 64 noise points as independent, which they are not, so the intervals are indicative. (iii) The readout limits depend on the choice of AUC level (0.90 or 0.95), on the window ratio 1.5 and on the smoothing; the 0.95 limit is noticeably less stable. (iv) The estimator uses uniform weights and the real part only; other estimators, such as sparse recovery that exploits the non-negativity and sparsity of $\Lambda(n)/\sqrt{n}$, were not tested and might exceed the limit (11). (v) The surrogate test uses 20 realisations of each ensemble. (vi) Landau's formula requires the Riemann Hypothesis only to write $x^{\rho} = \sqrt{x}\, e^{i\gamma \log *x*}$; for the zeros used here this holds numerically. (vii) The derivation of (10) from the uniform Landau formula is not carried out; the agreement of scaling is an observation. (viii) The extrapolation to $X = 10^4$ is outside the tested range.

8. Conclusions

The spectral coherence of the first N Riemann zeros carries the primes as echoes at $u = \log n$, and the readout from a finite set of zeros has simple quantitative behaviour. The error of the estimator grows linearly in $n/(T/2\pi)$ (constant 0.67), prime powers can be told apart from other integers (AUC ≥ 0.90) up to $X/\log X \approx 0.97 \cdot T/2\pi$, which is about 4.6 zeros per readable prime at $N = 1000$, and the effect disappears for surrogates with the same density and pair statistics. Errors in the ordinates act through an exact Debye–Waller factor. A smooth taper does not improve the readout. The natural next steps are a test of the laws at larger height, an analytic derivation of the constant κ from the uniform Landau formula, and a test of sparse-recovery estimators.

Declarations

Ethical Approval:

Not applicable. This study does not involve human participants or animals.

Consent to Publish:

Not applicable.

Consent to Participate:

Not applicable.

Competing Interests:

The author declares no competing financial or non-financial interests.

Funding:

This research received no external funding.

Data and code availability:

The Python code, scripts, and data used to generate all simulations and figures in this study are available from the corresponding author upon reasonable request. The ordinates of the first 1000 zeros (zeros_first_1000.csv) and the script zeta_echo_reproduce.py, which regenerates results.json and all figures, accompany this paper. The script requires numpy, scipy, sympy, pandas and matplotlib.

Author Contributions:

Eng Ebrahem is the sole author of this work and is responsible for all aspects of the research, including conceptualization, methodology, software development, validation, formal analysis, investigation, data

curation, writing—original draft preparation, writing—review and editing, and visualization.

References

1. Riemann B (1859) Über die Anzahl der Primzahlen unter einer gegebenen Grösse. Monatsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin, 671–680
2. Landau E (1911) Über die Nullstellen der Zetafunktion. Math Ann 71:548–564
3. Gonek SM (1993) An explicit formula of Landau and its applications to the theory of the zeta-function. Contemp Math 143:395–413
4. Montgomery HL (1973) The pair correlation of zeros of the zeta function. In *Analytic Number Theory*, Proc. Sympos. Pure Math. 24, 181–193. American Mathematical Society
5. Odlyzko AM (1987) On the distribution of spacings between zeros of the zeta function. Math Comput 48:273–308
6. Berry MV (1985) Semiclassical theory of spectral rigidity. *Proceedings of the Royal Society of London A*, 400, 229–251
7. Bogomolny EB, Keating JP (1996) Gutzwiller's trace formula and spectral statistics: beyond the diagonal approximation. Phys Rev Lett 77:1472–1475
8. Berry MV, Keating JP (1999) The Riemann zeros and eigenvalue asymptotics. SIAM Rev 41:236–266
9. Schumayer D, Hutchinson DAW (2011) Colloquium: Physics of the Riemann hypothesis. Rev Mod Phys 83:307–330
10. Mandelstam L, Tamm I (1991) The Uncertainty Relation Between Energy and Time in Non-relativistic Quantum Mechanics. In: Bolotovskii BM, Frenkel VY, Peierls R (eds) Selected Papers. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-74626-0_8
11. Titchmarsh EC (1986) *The Theory of the Riemann Zeta-Function*(2nd ed., revised by D. R. Heath-Brown). Oxford University Press
12. Edwards HM (1974) Riemann's Zeta Function. Academic
13. Mehta ML (2004) Random Matrices, 3rd edn. Elsevier
14. Haake F (2010) Quantum Signatures of Chaos, 3rd edn. Springer
15. Bohigas O, Giannoni M-J, Schmit C (1984) Characterization of chaotic quantum spectra and universality of level fluctuation laws. Phys Rev Lett 52:1–4
16. Mann HB, Whitney DR (1947) On a test of whether one of two random variables is stochastically larger than the other. Annals Math Stat 18:50–60
17. Platt D, Trudgian T (2021) The Riemann hypothesis is true up to $3 \cdot 10^{12}$. Bull Lond Math Soc 53:792–797
18. Elsayed EE (2026) The Lunet: A New Spectral Entity of Light Resolving Wave-Particle Duality. 10.5281/zenodo.22424449. Zenodo

Figures

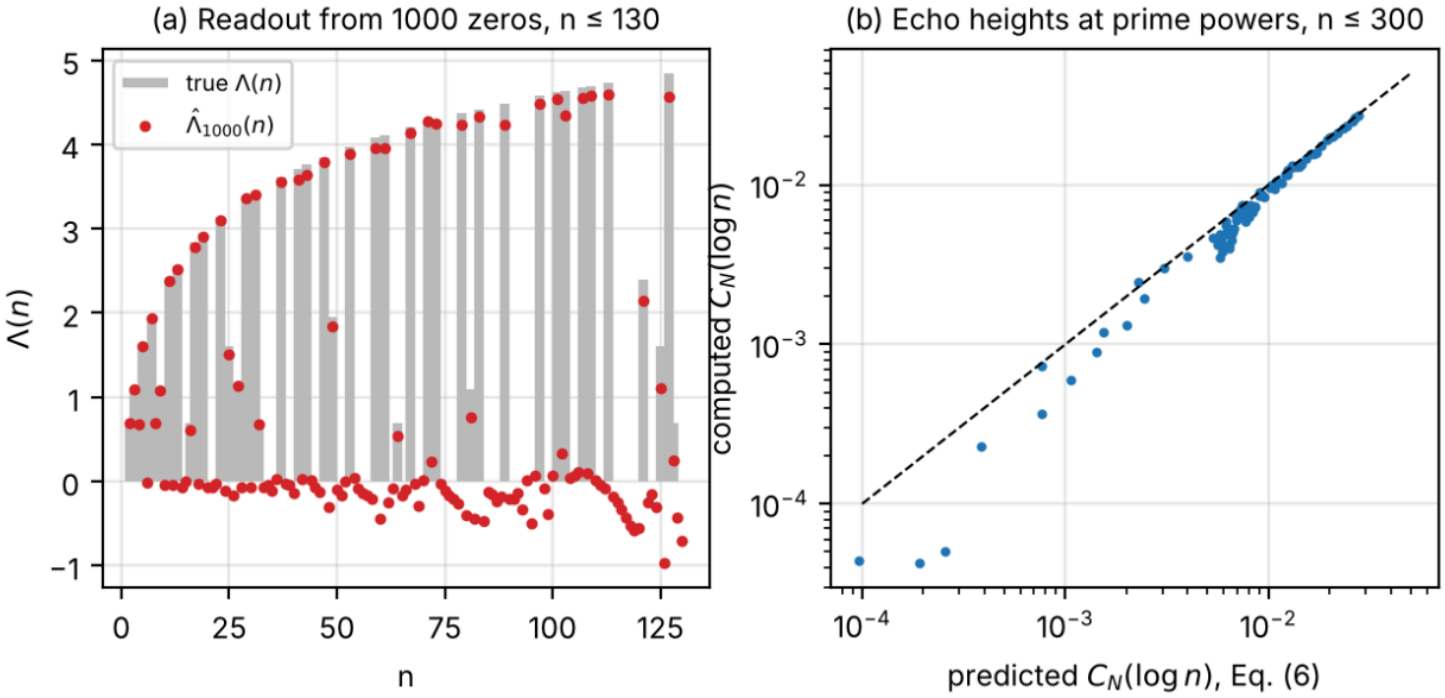

Figure 1

(a) The von Mangoldt function $\Lambda(n)$ (grey bars) and its readout $\hat{\Lambda}_{1000}(n)$ (red points) from the first 1000 zeros, $n \le 130$. (b) Computed echo heights $C_{1000}(\log n)$ at prime powers $n \le 300$ against the prediction of Eq. (6); the dashed line is equality.

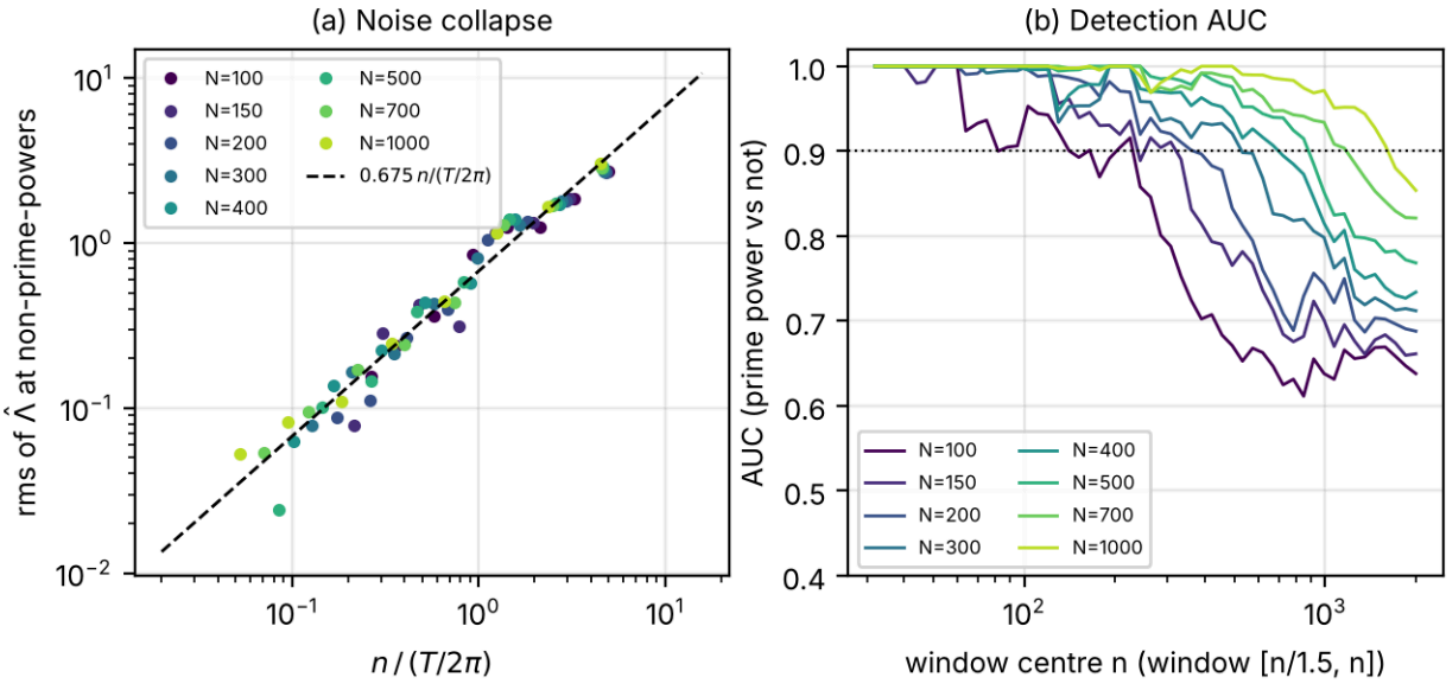

Figure 2

(a) Root-mean-square readout noise at non-prime-powers against $n/(T/2\pi)$ for eight values of N; the dashed line is Eq. (10). (b) Detection AUC of prime powers against other integers in windows $(c/1.5, c]$; the dotted line marks AUC = 0.90.

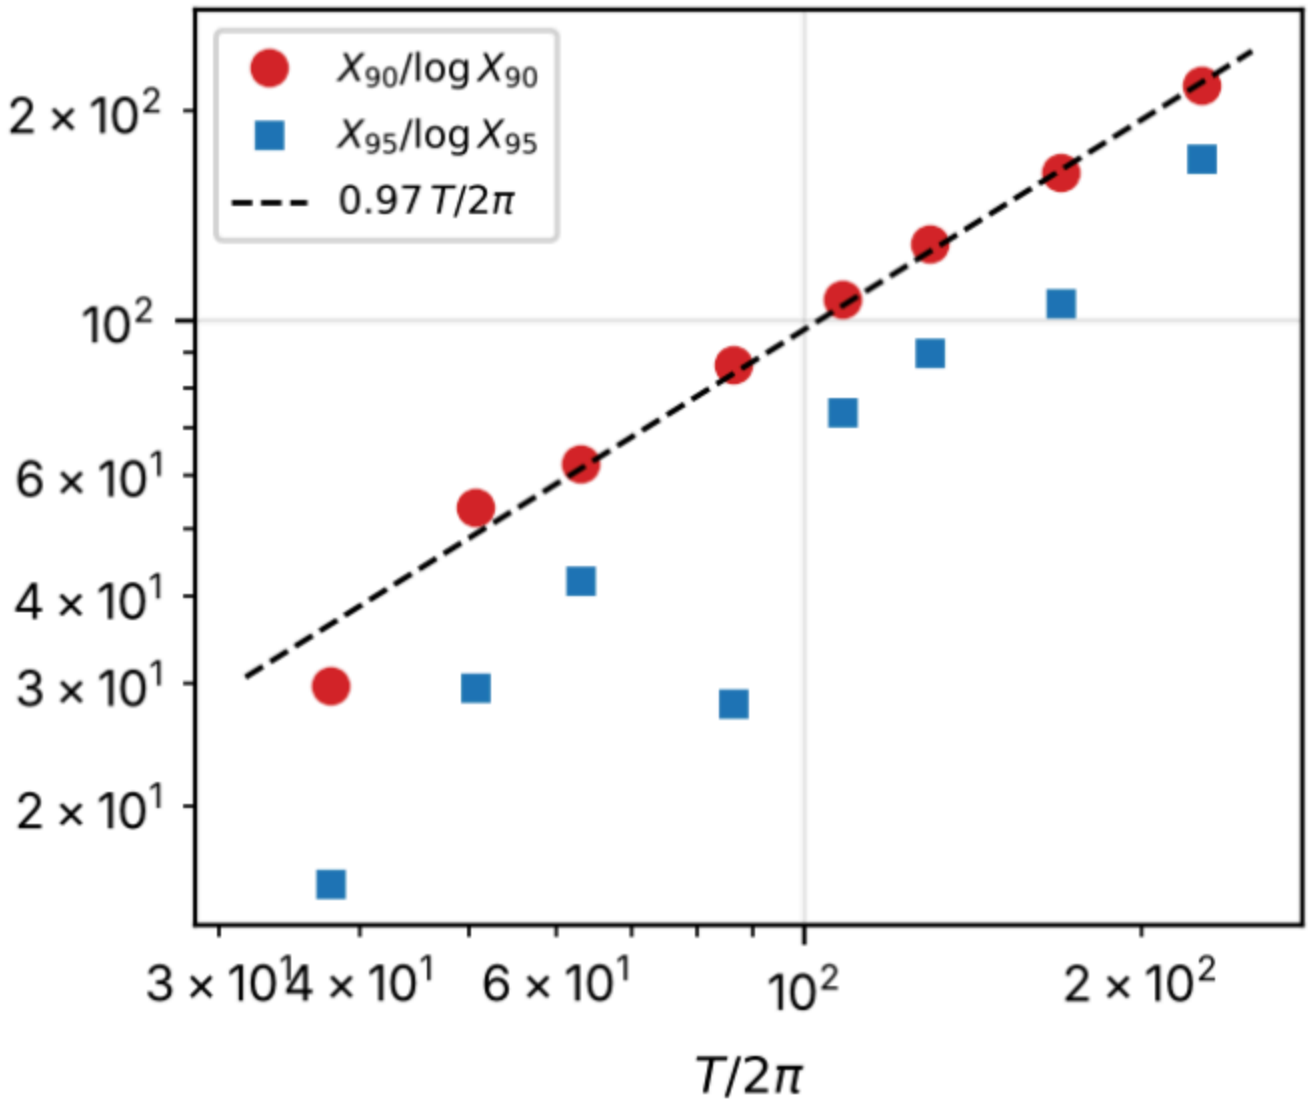

Figure 3

Readout limits X_{90} (red) and X_{95} (blue), plotted as $X/\log X$, against $T/2\pi$ for the eight subsets. The dashed line is Eq. (11).

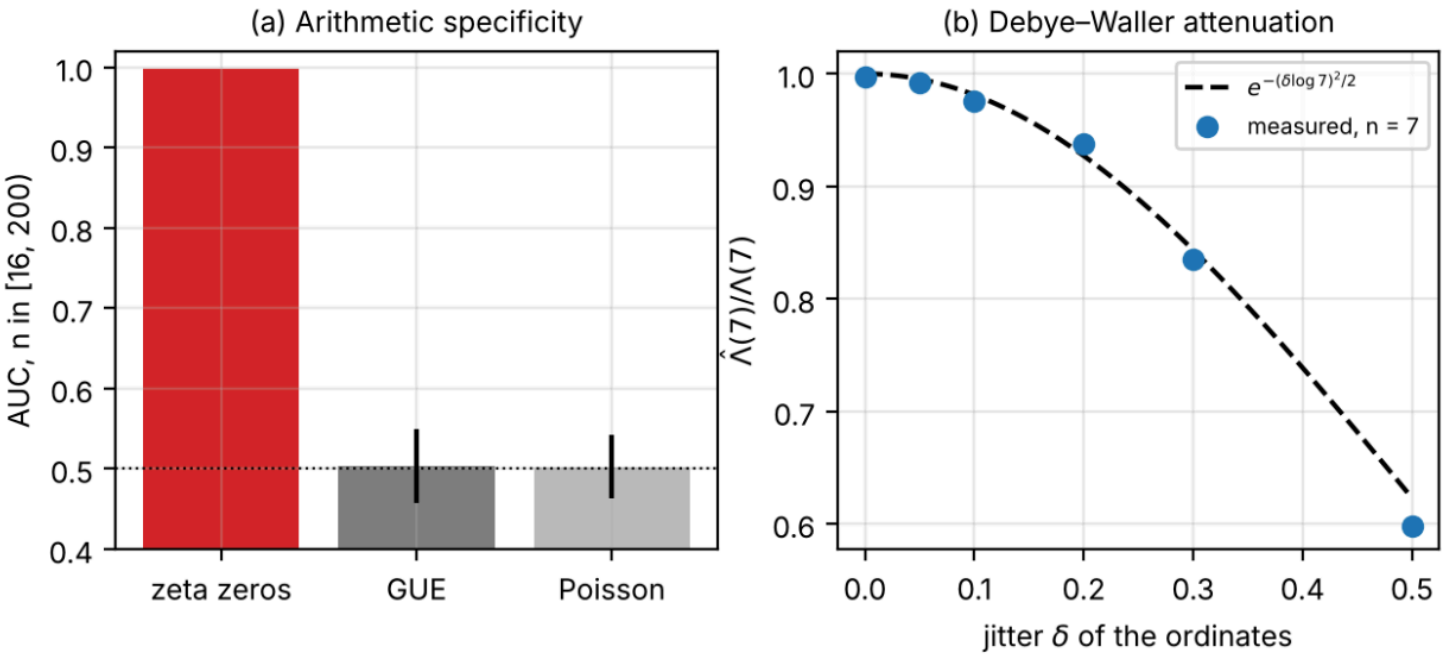

Figure 4

(a) AUC of the readout for the zeta zeros and for GUE and Poisson surrogates of equal size and density (error bars: standard deviation over 20 realisations). (b) Attenuation of $\hat{\Lambda}(7)/\Lambda(7)$ with ordinate jitter δ; the dashed line is Eq. (8).

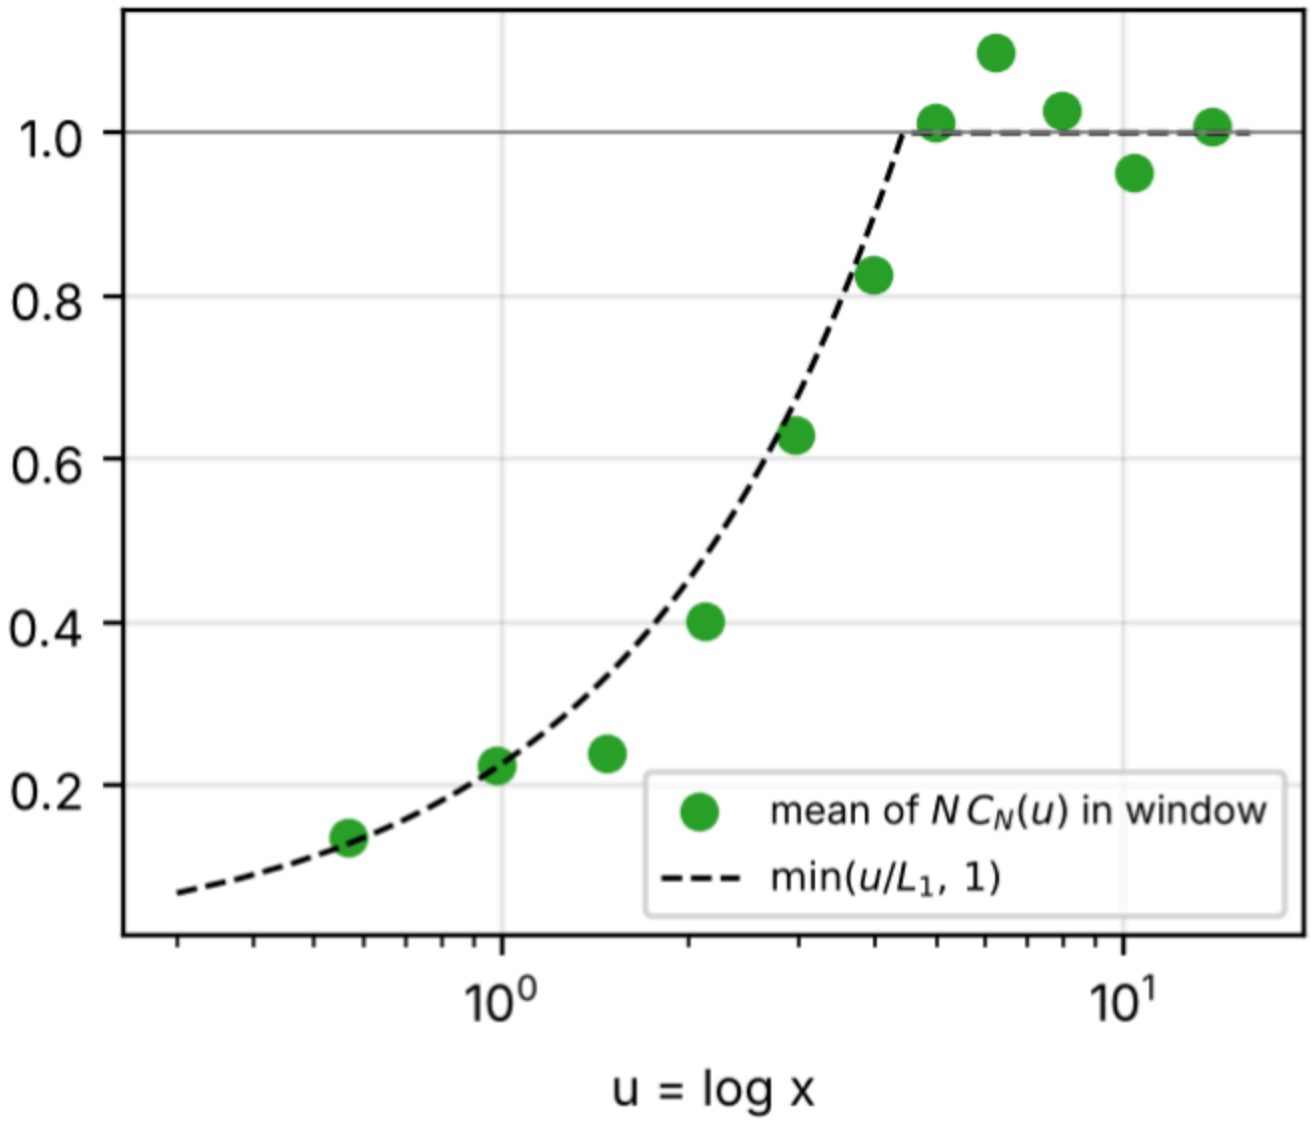

Figure 5

Mean of $NC_N(u)$ in windows of u for N = 1000 (points) and the prediction min(u/L_1, 1) of Eq. (7) (dashed).