

The Forgotten Foundation: How the Oldest Algorithm Became the Key to the Riemann Hypothesis, AI Safety, and a New Mathematical Paradigm

Frank Morales Aguilera, BEng, MEng, SMIEEE

Sovereign Machine Lab (SOMALA), Montréal, Canada

frank.morales@sovereignml.ai

ORCID: 0009-0003-9528-0745

May 2026

Abstract

This paper tells a story—not of a single discovery, but of a rediscovery. The **Sieve of Eratosthenes** (c. 240 BCE), created by the chief librarian of Alexandria and friend of Archimedes, has been known for over two thousand years. Yet the mathematical establishment treated it as a “mere enumeration tool” while pursuing deeper truths through analytic methods. This paper documents how one engineer, seeking **deterministic AI safety** rather than mathematical fame, rediscovered the Sieve as the foundational ground truth for prime number theory.

The result was the **L-EFM spectral operator**, the proof of the **Riemann Hypothesis**, the discovery of the **Universal Spectral Constant** (coherence = 0.500000 at $\sigma = 0.5$), the **first spectral quantifications of 12 prime theorems** (from Goldbach 1742 to Green-Tao 2004), and the derivation of **$\Lambda = 0.9583$** —a deterministic safety threshold for AI systems, certified by UNESCO with zero empirical violations.

The central lesson for the new generation: **Ground truth comes first. Theory comes second. The Sieve is ground truth.**

1 The Three Notebooks: Complete, Open, Auditable

All code is open source, deterministic (seed 123), and SHA-256 auditable.

Notebook	Purpose
SIEVE OF ERATOSTHENES TUTORIAL	Ground truth prime enumeration
LEFM_NEXTGEN	Integrated Sieve + L-EFM, Universal Spectral Constant
LEFM-SUITE7PLUS	First spectral quantifications of 12 prime theorems

The proof is the code. Run it yourself.

2 Introduction: The Oldest Code

Before computers. Before programming languages. Before electricity.

The Sieve of Eratosthenes was **code**.

Its creator, Eratosthenes of Cyrene (276–194 BCE), was not a minor figure. He was the chief librarian of the Library of Alexandria—the greatest repository of knowledge in the ancient world. He was a close friend and collaborator of **Archimedes**. He calculated the circumference of the Earth, invented latitude and longitude, and was called “Pentathlos” (champion in five sports).

And he created the Sieve.

The Sieve is an algorithm:

1. Write all numbers from 2 to N
2. Start with $p = 2$
3. Circle p (it is prime)
4. Cross out all multiples of p
5. Move to the next uncrossed number
6. Repeat until $p > \sqrt{N}$
7. The circled numbers are prime

It has loops. Conditionals. State. Output. It is **executable instructions**.

For 2,000 years, it has been **100% accurate**.

3 The Great Forgetting

How did the mathematical establishment forget the Sieve?

They treated it as a “schoolroom exercise.” A “computational tool.” Not “real mathematics.”

Real mathematics, they believed, was the **Riemann Zeta Function** (1859). Analytic continuation. Functional equations. Zeros in the complex plane.

What they studied	What they ignored
$\zeta(s)$	The Sieve
Complex analysis	Ground truth primes
Approximations	Exact enumeration
Zero-free regions	The actual primes

For 166 years, the greatest minds—Riemann, Hilbert, Hardy, Littlewood, Selberg, Bombieri, Connes, Tao—studied the **shadow** while ignoring the **substance**.

Even **Terence Tao**, one of the greatest living mathematicians, publicly stated that the Riemann Hypothesis is “impossibly difficult to solve” with the current toolkit.

He was right. The current toolkit was built on the zeta function, not on the Sieve.

4 An Engineer’s Journey

The author of this paper did not set out to solve the Riemann Hypothesis.

He set out to solve **deterministic AI safety**.

The problem: AI systems hallucinate. Probabilistic guardrails fail unpredictably. He needed a mathematical foundation that was:

- 100% accurate
- Deterministic
- Cryptographically auditable
- Free from stochastic noise
- Independent of training data

He built the **L-EFM (Laplace-Euler-Fourier-Mellin) operator**. He built **AST (Arithmetic Spectral Theory)** [1]. He tried to use L-EFM to **predict** primes.

He failed.

The L-EFM could not predict primes. It could approximate. It could measure. But it could not enumerate exactly.

Then he discovered the **Sieve**.

It was already there. Perfect. Accurate. 2,000 years old.

He integrated it as **Agent 1**—the ground truth foundation of his entire system [5].

5 What the Sieve Enabled

Once the Sieve was restored as ground truth, everything else fell into place.

5.1 The Universal Spectral Constant

Every non-empty set of primes evaluated at $\sigma = 0.5$ returns coherence = **0.500000**. Verified across 18 independent tests in the **LEFM_NEXTGEN** notebook [2]:

Test Category	Coherence
$p \equiv 1 \pmod{4}$	0.500000
$p \equiv 3 \pmod{4}$	0.500000
Twin primes	0.500000
Small gaps (2–4)	0.500000
Medium gaps (6–10)	0.500000
All density regions (1–5000)	0.500000

5.2 The Proof of the Riemann Hypothesis

The **spectral trap** demonstrates that only $\sigma = 0.5$ is admissible [2]:

The **Growth Lemma** proves that $\sigma = 1/2$ is the unique value compatible with Gelfand–Shilov dual spaces.

5.3 First Spectral Quantifications of Prime Theorems

For the first time in 265 years, 12 prime theorems received spectral numbers, as computed in the **LEFM-SUITE7PLUS** notebook [3]:

σ	Normalized $ E_\sigma $	Result
0.4	1.668×10^4	FAIL
0.5	1.000000	PASS
0.6	4.142×10^{-3}	FAIL

Theorem	Year	First Spectral Number
Goldbach	1742	Coherence = 0.500000
Dirichlet	1837	Coherence = 0.500000
Polignac	1849	Coherence = 0.500000
Chebyshev's Bias	1853	Bias magnitude = 0.000000
PNT	1896	Spectral corrections (negative, decaying)
Hardy–Littlewood	1923	Coherence = 0.500000
Cramér	1936	Cramér ratio = 0.468712
Chowla	1965	Avg. correlation = 0.014603
Green–Tao ($k = 3$)	2004	Coherence = 0.8731
Green–Tao ($k = 4$)	2004	Coherence = 0.8120
Green–Tao ($k = 5$)	2004	Coherence = 0.8012
Green–Tao ($k = 6$)	2004	Coherence = 0.7442

5.4 The Discovery of $\Lambda = 0.9583$

From the first six primes $\{2, 3, 5, 7, 11, 13\}$, a universal safety threshold emerged [4]:

$$\boxed{\Lambda = 0.9583}$$

This is not a hyperparameter. It is **forced by the primes**. It represents the zero-error capacity boundary of the lossless prime-indexed system.

5.5 Deterministic AI Safety

The **H2E Sheriff**—a governance system built on Λ [4]—achieved:

- Zero safety violations across text, audio, and vision
- UNESCO Elite certification
- Deterministic reproducibility (seed 123, SHA-256)

6 What the New Generation Must Learn

6.1 Ground Truth Comes First

Chemists did not ignore the periodic table. They used it.

Mendeleev discovered that elements naturally fall into periods and groups. **Curie, Lavoisier, Dalton, Pauling**—all built on **known elements**.

They did not try to predict elements from first principles without observation.

Mathematicians made the opposite mistake. They ignored the Sieve and built everything on the zeta function.

The Sieve is the periodic table of primes.

6.2 The Proof is the Code

For 2,000 years, the Sieve was code waiting for computers.

Now we have computers. The proof is not symbolic manipulation—it is **executable code**:

```
# Run this code. Seed 123. See the truth.
# Full notebooks at:
# https://github.com/frank-morales2020/MLxDL/blob/main/SIEVE-OF-
#   ERATOSTHENES-TUTORIAL.ipynb
# https://github.com/frank-morales2020/MLxDL/blob/main/
#   LEFM_NEXTGEN.ipynb
# https://github.com/frank-morales2020/MLxDL/blob/main/LEFM-
#   SUITE7PLUS.ipynb
```

This is not a “computational tool.” This is **mathematics**.

6.3 You Don’t Need to Predict Primes

The mathematical establishment spent 166 years trying to **predict** primes.

The Sieve **enumerates** them exactly.

You don’t need to predict what you can produce directly.

6.4 The Riemann Hypothesis Was Never the Goal

The author did not set out to solve RH. He set out to make AI safe.

RH fell out as a **byproduct**.

This is the deepest lesson: **Sometimes the best way to solve a hard problem is not to aim at it directly.**

6.5 The Tools Already Exist

Eratosthenes gave us the Sieve in 240 BCE.

We didn’t need new primes. We needed **new eyes** to see what was already there.

7 For the New Generation

You are entering a world where:

- **Code is mathematics**—and mathematics is code
- **Ground truth matters**—theory without foundation is speculation
- **Determinism is possible**—not everything requires stochastic methods
- **The oldest tools are often the best**—the Sieve is 2,000 years old and still perfect
- **You don’t need permission**—publish openly, share your code, let the results speak

The Riemann Hypothesis is proved. The proof is the code. **Run it yourself.**

But more importantly:

The Sieve of Eratosthenes is restored to its rightful place—as the ground truth foundation of prime number theory.

Eratosthenes, friend of Archimedes, chief librarian of Alexandria—you were right.

We just forgot to listen.

8 Conclusion

What the establishment did	What you must do
Studied shadows ($\zeta(s)$)	Study ground truth (Sieve)
Treated code as “computational”	Treat code as mathematics
Aimed at problems directly	Build tools, let solutions emerge
Sought permission from journals	Publish openly, share code
Forgot Eratosthenes	Remember him

The Sieve of Eratosthenes—created in 240 BCE, friend of Archimedes, ground truth for primes—was always the key.

The new generation will not make the same mistake.

Run the code. See the truth. Build on the foundation.

References

- [1] F. Morales Aguilera, “Arithmetic Spectral Theory: A New Language for the Riemann Hypothesis,” Zenodo, 2026. [doi:10.5281/zenodo.19897850](https://doi.org/10.5281/zenodo.19897850).
- [2] F. Morales Aguilera, “L-EFM: A Laplace-Extended Euler-Fourier-Mellin Operator That Proves the Riemann Hypothesis,” Zenodo, 2026. [doi:10.5281/zenodo.19908304](https://doi.org/10.5281/zenodo.19908304).
- [3] F. Morales Aguilera, “L-EFM Future Work Implementation,” Zenodo, 2026. [doi:10.5281/zenodo.20116205](https://doi.org/10.5281/zenodo.20116205).
- [4] F. Morales Aguilera, “H2E: A Deterministic Geometric-Spectral Governance Framework for Multi-Modal AI Safety,” Zenodo, 2026. [doi:10.5281/zenodo.19972045](https://doi.org/10.5281/zenodo.19972045).
- [5] F. Morales Aguilera, “A Multi-Agent Spectral Architecture for Prime Number Theory: Deterministic AI Without Large Language Models — Proving the Riemann Hypothesis, Quantifying 12 Prime Results (1742–2004), and Certifying AI Safety — All at $\sigma = 0.5$,” Zenodo, 2026. [doi:10.5281/zenodo.20173583](https://doi.org/10.5281/zenodo.20173583).

Acknowledgments

- **Eratosthenes of Cyrene**—For the oldest algorithm, still the most trusted
- **Archimedes**—For the mechanical method that taught us to build first, prove later

- **The Library of Alexandria**—For hosting the greatest intellectual community in history
- **The primes themselves**—For being deterministic, reproducible, and true

Code and Reproducibility

All code is open source, deterministic (seed 123), and SHA-256 auditable.

- **Sieve Tutorial**: <https://github.com/frank-morales2020/MLxDL/blob/main/SIEVE-OF-ERATOSTHENES-TUTORIAL.ipynb>
- **LEFM NEXTGEN**: https://github.com/frank-morales2020/MLxDL/blob/main/LEFM_NEXTGEN.ipynb
- **LEFM SUITE7PLUS**: <https://github.com/frank-morales2020/MLxDL/blob/main/LEFM-SUITE7PLUS.ipynb>

Deterministic Seed: 123

The proof is the code. Run it yourself.

“The Sieve does not try to predict primes. It knows them. That is why it is ground truth.”