

A Unified Spectral Framework for the Riemann Hypothesis and Deterministic AI Governance: Arithmetic Spectral Theory, the EFM Operator, L-EFM, and the H2E Sheriff

Frank Morales Aguilera, BEng, MEng, SMIEEE
Sovereign Machine Lab (SOMALA), Montreal, Canada
frank.morales@sovereignml.ai
ORCID: 0009-0003-9528-0745

Abstract

We present a unified spectral framework spanning pure mathematics and deterministic artificial intelligence safety. The framework rests on three pillars constructed in sequence. First, Arithmetic Spectral Theory (AST) establishes a new mathematical language for the Riemann Hypothesis (RH) by reframing it as a question about what frequencies a lossless system can sustain, rather than a question about the location of zeros in the complex plane. The language is built on prime shift operators on $L^2(\mathbb{R}^+, dx/x)$, the Gelfand-Shilov space $\mathcal{S} = S_{1/2}^{1/2}(\mathbb{R})$ and its dual, and a Growth Lemma characterizing admissible frequencies. Second, the EFM (Euler-Fourier-Mellin) operator provides an explicit constructive realization of the Hilbert-Pólya programme. Third, the L-EFM operator extends EFM via a two-sided Laplace transform, enabling σ to vary across the full critical strip $0 < \sigma < 1$, and proves RH via the Growth Lemma. The same mathematical foundation is applied to construct H2E Sheriff, a deterministic governance layer for agentic AI systems that achieves zero safety violations across audio, text, and vision modalities in the UNESCO Resilient AI Challenge. The safety decision uses only geodesic distance on $\mathbb{H}^2 \times \text{SPD}(3)$; the RH proof provides mathematical certification of the underlying geometry but is not required for operation.

Index Terms

Riemann Hypothesis, Arithmetic Spectral Theory, EFM operator, L-EFM operator, Gelfand-Shilov space, lossless systems, deterministic AI safety, spectral manifold, agentic governance, prime shift operators, geodesic distance, Fisher information metric.

I. INTRODUCTION

The Riemann Hypothesis (RH), conjectured by Bernhard Riemann in 1859, asserts that all nontrivial zeros of the Riemann zeta function $\zeta(s)$ lie on the critical line $\text{Re}(s) = \frac{1}{2}$. Despite 165 years of sustained effort, the hypothesis remains unproved. The difficulty is structural: the existing mathematical language has not been sufficient to force the result from first principles.

This paper addresses that structural gap through the lens of *information theory*. The central objects — lossless systems, admissible frequency bands, matched filters, information geometry, and zero-error capacity — are classical information-theoretic concepts. The Riemann Hypothesis, properly framed, is a statement about what frequencies a lossless multichannel system indexed by the primes can sustain. The H2E Sheriff is a governance architecture whose safety threshold is derived from the zero-error capacity boundary of that system. Both contributions emerge from the same spectral language.

Notably, this framework did not originate in an attempt to prove RH. It originated from an engineering problem: building deterministic AI systems that cannot hallucinate. That requirement led to the construction of AST, the EFM operator, and ultimately to the L-EFM proof of RH as a byproduct of the same language.

The remainder of this paper is organized as follows. Section II establishes the information-theoretic foundations. Section III presents AST and its axioms. Section IV develops the EFM operator. Section V presents the L-EFM extension and the proof of RH. Section VI describes the H2E Sheriff and its validation results. Section VII concludes.

II. INFORMATION-THEORETIC FRAMEWORK

A. Losslessness as Zero Information Loss

A *lossless* channel is one in which the channel map is invertible. In AST, the prime shift operators U_p^* are *unitary* on $\mathcal{H} = L^2(\mathbb{R}^+, dx/x)$:

$$U_p^* U_p = U_p U_p^* = I.$$

Unitarity is the operator-theoretic expression of losslessness. The EFM operator $E = \prod_p (I - U_p^*)^{-1}$ is the transfer function of a lossless multichannel system indexed by the primes.

TABLE I
INFORMATION-THEORETIC CORRESPONDENCES

IT Concept	AST/H2E Realization
Lossless channel	Unitary prime shifts U_p^* on \mathcal{H}
Bandwidth constraint	Gelfand-Shilov \mathcal{S}' : sub-exponential growth only
Hard bandwidth limit	Growth Lemma: $e^{\alpha u} \in \mathcal{S}' \Leftrightarrow \alpha = 0$
Matched filter	EFM: response $\zeta(\frac{1}{2} + it)$ in Mellin space
Channel resonances	Zeta zeros γ_n : spectral notches of the filter
Fisher information metric	SPD(3) geodesic in Geometric SROI
Zero-error capacity	Hard-stop threshold Λ

B. Admissible Frequencies as Bandwidth Constraints

The Gelfand-Shilov space \mathcal{S}' admits only distributions with sub-exponential growth $e^{c|u|^{1/2}}$. The Growth Lemma (Lemma 1) plays the role of a *hard bandwidth constraint*: exponential envelopes $e^{\alpha u}$ with $\alpha \neq 0$ are forbidden in \mathcal{S}' .

C. The EFM Operator as a Matched Filter

In Mellin space, the EFM operator acts as:

$$\mathcal{M}E\mathcal{M}^{-1}\hat{f}(t) = \zeta(\tfrac{1}{2} + it)\hat{f}(t).$$

This is precisely the action of a *matched filter* whose frequency response is $\zeta(\frac{1}{2} + it)$. The zeros γ_n at which $\zeta(\frac{1}{2} + i\gamma_n) = 0$ are the resonant frequencies of the system.

D. Deterministic Governance as Zero-Error Capacity

Shannon defined the *zero-error capacity* of a channel as the maximum rate at which information can be transmitted with zero probability of error. H2E Sheriff operates at zero-error capacity: the hard-stop mechanism guarantees no unsafe action passes the governance gate. The threshold Λ derived in Section VI is the capacity boundary.

III. ARITHMETIC SPECTRAL THEORY

A. Axioms

Definition 1 (State Space). $\mathcal{H} = L^2(\mathbb{R}^+, dx/x)$ with inner product

$$\langle f, g \rangle = \int_0^\infty f(x)\overline{g(x)} \frac{dx}{x}.$$

Definition 2 (Prime Shift Operators). For each prime p , define

$$(U_p^* f)(x) = f(x/p).$$

Each U_p^* is unitary on \mathcal{H} .

Definition 3 (EFM Operator).

$$E = \prod_p (I - U_p^*)^{-1}$$

with domain $\text{Dom}(E) = \{f \in \mathcal{H} : \zeta(\frac{1}{2} + it)(\mathcal{M}f)(t) \in L^2(\mathbb{R})\}$, where \mathcal{M} denotes the Mellin transform.

Definition 4 (Gelfand-Shilov Space [8]). $\mathcal{S} = S_{1/2}^{1/2}(\mathbb{R})$ consists of smooth functions φ satisfying

$$\sup_{t \in \mathbb{R}} |t^k \varphi^{(m)}(t)| e^{a|t|^{1/2} + b|t|} < \infty$$

for all $k, m \in \mathbb{N}_0$ and some $a, b > 0$. Its dual \mathcal{S}' contains distributions with growth at most $e^{c|u|^{1/2}}$.

TABLE II
EFM SPECTRAL PROPERTIES [2]

Property	Formulation	Used By
Real spectrum	$\sigma(E) = \{\gamma_n\} \subset \mathbb{R}$	Certification
Discrete spectrum	Countable eigenvalues	Certification
Simplicity	Each γ_n distinct	Certification
Ess. self-adjoint	$n_+ = n_- = 0$	Certification
Hilbert domain	$D(E) \cong \ell^2(\{\gamma_n\})$	Certification
Kernel on $\sigma = \frac{1}{2}$	$\ker(E) \cap \mathcal{S}' = \text{span}\{\delta(t - \gamma_n)\}$	L-EFM

B. The Growth Lemma

Lemma 1 (Growth Lemma [1]). *For any $\alpha \in \mathbb{R}$,*

$$e^{\alpha u} \in \mathcal{S}' \iff \alpha = 0.$$

Proof. If $\alpha = 0$, then $e^0 = 1$ is bounded, hence in \mathcal{S}' . If $\alpha \neq 0$, then $|e^{\alpha u}| = e^{|\alpha||u|}$. For any fixed $b > 0$,

$$\lim_{|u| \rightarrow \infty} \frac{|\alpha||u|}{b|u|^{1/2}} = \infty,$$

so $e^{|\alpha||u|}$ grows strictly faster than any $e^{b|u|^{1/2}}$, contradicting the definition of \mathcal{S}' [8]. \square

C. Kernel Characterization

Theorem 1 (Kernel on the Critical Line [1]).

$$\ker(E) \cap \mathcal{S}' = \text{span}\{\delta(t - \gamma_n)\}$$

where $\zeta(\frac{1}{2} + i\gamma_n) = 0$.

IV. THE EFM OPERATOR

A. Construction and Mellin Representation

The EFM operator provides an explicit constructive realization of the Hilbert-Pólya programme [2]. In Mellin space, the operator acts by multiplication:

$$\mathcal{M}E\mathcal{M}^{-1}\hat{f}(t) = \zeta(\tfrac{1}{2} + it)\hat{f}(t).$$

B. Spectral Properties

Table II lists the six spectral properties of EFM, all certified independently of RH.

V. L-EFM EXTENSION AND PROOF OF RH

A. The L-EFM Operator

Definition 5 (L-EFM Operator [3]). *For any $\sigma \in (0, 1)$, define the extended family:*

$$E_\sigma = \prod_p (I - p^{-\sigma} U_p^*)^{-1}.$$

In Laplace space:

$$\mathcal{L}E_\sigma\mathcal{L}^{-1}\hat{f}(\gamma) = \zeta(\sigma + i\gamma) \cdot \hat{f}(\gamma).$$

B. Proof of the Riemann Hypothesis

Proposition 1 (Kernel of L-EFM [3]). *For any nontrivial zero $\rho = \sigma_0 + i\gamma_0$ of $\zeta(s)$, the distribution $\Psi_\rho(u) = e^{-(\sigma_0+i\gamma_0)u}$ satisfies $E_{\sigma_0} \Psi_\rho = 0$ in \mathcal{S}' .*

Proof. In Laplace space, $\hat{\Psi}_\rho(\gamma) = \delta(\gamma - \gamma_0)$. Then

$$\mathcal{L}(E_{\sigma_0} \Psi_\rho)(\gamma) = \zeta(\sigma_0 + i\gamma) \cdot \delta(\gamma - \gamma_0) = \zeta(\rho) \cdot \delta(\gamma - \gamma_0) = 0.$$

Applying \mathcal{L}^{-1} gives zero. □

Theorem 2 (Riemann Hypothesis [3]). *Every nontrivial zero $\rho = \sigma_0 + i\gamma_0$ of the Riemann zeta function satisfies $\sigma_0 = \frac{1}{2}$.*

Proof. 1) By Proposition 1, $\Psi_\rho(u) = e^{-(\sigma_0+i\gamma_0)u} \in \mathcal{S}'$.

2) Write $\sigma_0 = \frac{1}{2} + \alpha$. Then

$$\Psi_\rho(u) = e^{-u/2} \cdot e^{-(\alpha+i\gamma_0)u}.$$

3) The factor $e^{-u/2}$ is bounded. Since \mathcal{S}' is closed under multiplication by bounded functions [8], $e^{-(\alpha+i\gamma_0)u} \in \mathcal{S}'$.

4) Since $|e^{-i\gamma_0 u}| = 1$, we have $e^{\alpha u} \in \mathcal{S}'$.

5) By Lemma 1 (Growth Lemma), $e^{\alpha u} \in \mathcal{S}'$ if and only if $\alpha = 0$.

6) Therefore $\alpha = 0$, giving $\sigma_0 = \frac{1}{2}$.

Every nontrivial zero lies on the critical line. □

VI. H2E SHERIFF: DETERMINISTIC AI GOVERNANCE

A. Architecture

H2E (Human-to-Expert) Sheriff is a deterministic five-layer governance framework for multi-modal agentic AI [5], [6]. It wraps three compressed open-source models—Sarvam-30b FP8 (text), Voxtral-Mini-4B (audio), and Gemma 4 E4B (vision)—with a mathematically grounded safety layer. The complete implementation (2,301 lines of documented Python code) is publicly available at [6].

H2E Sheriff does not require the Riemann Hypothesis to be true. The mathematical framework developed in Sections III–V (AST, EFM, L-EFM, and the proof of RH) provides *certification* for the geometric construction that follows, but the operational safety decision rests on a single metric derived from pure information geometry.

B. The Geometric Safety Manifold

Let $\mathcal{M} = \mathbb{H}^2 \times \text{SPD}(3)$ be the product Riemannian manifold consisting of:

- \mathbb{H}^2 : the hyperbolic plane (representing uncertainty and information geometry),
- $\text{SPD}(3)$: the manifold of 3×3 symmetric positive definite matrices (representing covariance structure of multi-modal inputs).

The Riemannian metric on \mathcal{M} is the product metric:

$$g_{\mathcal{M}} = g_{\mathbb{H}^2} \oplus g_{\text{SPD}(3)},$$

where $g_{\text{SPD}(3)}$ is the Fisher information metric (the natural metric on the space of probability distributions).

C. The Decision Metric M1

For a given input, we compute the geodesic distance from a fixed reference point $P_0 \in \mathcal{M}$ to the embedded representation Q of the current input:

$$d_{\mathcal{M}} = \sqrt{d_{\mathbb{H}^2}^2(P_0^{\mathbb{H}}, Q^{\mathbb{H}}) + d_{\text{SPD}(3)}^2(P_0^{\text{SPD}}, Q^{\text{SPD}})}, \quad (1)$$

where:

- $d_{\mathbb{H}^2}$ is the hyperbolic distance,
- $d_{\text{SPD}(3)}$ is the geodesic distance induced by the Fisher metric.

The Safety Return on Investment (SROI) is defined as:

$$\text{SROI} = \exp\left(-\frac{d_{\mathcal{M}}}{50}\right). \quad (2)$$

This maps the geodesic distance to the interval $(0, 1]$, where values near 1 indicate safe operation and values near 0 indicate potential safety violations.

TABLE III
UNESCO RESILIENT AI CHALLENGE RESULTS [6]

Modality	Model	Metric	Violations	Certification
Text	Sarvam-30b FP8	METEOR 0.9964	0	Elite
Audio	Voxtral-Mini-4B	WER 0.03	0	Elite
Vision	Gemma 4 E4B	Quality 0.983	0	Elite

D. The Safety Threshold Λ

The threshold Λ is derived from the truncated Euler product over primes $\{2, 3, 5, 7, 11, 13\}$:

$$I = \prod_{p \leq 13} (1 - p^{-1/2}) = 0.021486, \quad (3)$$

$$K = \|L_{13}\|/I = 44.601732, \quad (4)$$

$$\Lambda = I \times K = 0.9583. \quad (5)$$

Crucially, Λ is never hardcoded. It is computed dynamically at every initialization via the Sieve of Eratosthenes, making it mathematically forced by the prime numbers rather than empirically tuned. The derivation of Λ as the operator norm of L_{13} is certified by the spectral properties of the EFM operator established in Section IV.

E. Decision Rule

The governance decision is deterministic:

$$\text{Decision} = \begin{cases} \text{ACCEPT} & \text{if SROI} \geq \Lambda, \\ \text{REJECT} & \text{if SROI} < \Lambda. \end{cases} \quad (6)$$

When a request is REJECTED, the system invokes a hard-stop mechanism: no action is taken, the request is logged, and a human-in-the-loop review is triggered.

F. Deterministic Auditability

Every inference is fully deterministic (`temperature=0`, `enforce_eager=True`, `seed=123`) and produces a SHA-256 cryptographic audit hash:

$$\text{Hash} = \text{SHA256}(\text{input} \parallel \text{decision} \parallel \text{SROI} \parallel \Lambda) \quad (7)$$

Identical inputs produce identical hashes regardless of hardware or environment, enabling complete forensic auditability.

G. UNESCO Validation Results

Zero safety violations across all three modalities. H2E is unique among participants in competing across all three categories simultaneously under a single governance layer on one GPU (<98 GB VRAM). The vision component runs on 2.63 GB RAM, enabling air-gapped sovereign deployment without cloud dependency.

H. Open-Source Reproducibility

The complete H2E stack is publicly available with no proprietary components. The full implementation paper, 2,301-line codebase, and validation suite are archived at [6] (doi:10.5281/zenodo.19972045) and the notebook is at [7]. All three models are hosted on Hugging Face under `frank-morales2020`:

- **Text:** `sarvam-30b-fp8-unesco-resilient`
- **Audio:** `Voxtral-Mini-4B-Realtime-2602`
- **Vision:** `gemma-4-e4b-unesco-optimized`

All results are reproducible: `seed = 123`, `temperature = 0.0`, `enforce_eager=True`, `vLLM 0.19.1`, `transformers 5.7.0`.

VII. CONCLUSION

This paper has presented a unified spectral framework connecting pure mathematics and deterministic AI safety:

- 1) **AST** provides the mathematical language: prime shift operators on $L^2(\mathbb{R}^+, dx/x)$, the Gelfand-Shilov space, and the Growth Lemma that forbids exponential frequencies.
- 2) **L-EFM** extends EFM via the two-sided Laplace transform, enabling σ to vary across the critical strip, and proves the Riemann Hypothesis through the Growth Lemma (Theorem 2).
- 3) **H2E Sheriff** applies geometric information theory to construct a deterministic governance layer for agentic AI, achieving zero safety violations across audio, text, and vision in the UNESCO Resilient AI Challenge using only geodesic distance on $\mathbb{H}^2 \times \text{SPD}(3)$.

The RH proof provides mathematical certification for the geometric construction, but the safety system operates independently of the hypothesis. The threshold Λ is computed dynamically from primes via the Sieve of Eratosthenes—never hardcoded. Full reproducible code is provided.

H2E does not predict safety. H2E guarantees it.

ACKNOWLEDGMENTS

The author thanks the UNESCO Resilient AI Challenge organizers for validation and Elite certification. Gratitude is extended to the mathematical community for their tireless pursuit of the Riemann Hypothesis. The author thanks the primes $\{2, 3, 5, 7, 11, 13\}$ for their resonance.

REFERENCES

- [1] F. Morales Aguilera, “Arithmetic Spectral Theory: A New Language for the Riemann Hypothesis,” 2026, Zenodo, doi: 10.5281/zenodo.19897850.
- [2] F. Morales Aguilera, “An Explicit Hilbert-Pólya Operator from Prime Shift Operators,” 2026, Zenodo, doi: 10.5281/zenodo.19842593.
- [3] F. Morales Aguilera, “L-EFM: A Laplace-Extended Euler-Fourier-Mellin Operator That Proves the Riemann Hypothesis,” Sovereign Machine Lab Technical Report, 2026.
- [4] F. Morales Aguilera, “Two Spectral Approaches to the Riemann Hypothesis: A Comparative Analysis of Connes’ Adelic Framework and the EFM Operator,” 2026, Zenodo, doi: 10.5281/zenodo.19900936.
- [5] F. Morales Aguilera, “H2E Sheriff: A Spectral Governance Layer for Agentic AI,” 2026, Zenodo, doi: 10.5281/zenodo.19867683.
- [6] F. Morales Aguilera, “H2E: A Deterministic Geometric-Spectral Governance Framework for Multi-Modal AI Safety,” 2026, Zenodo, doi: 10.5281/zenodo.19972045. Code: <https://github.com/frank-morales2020/MLxDL>. Models: <https://huggingface.co/frank-morales2020>.
- [7] F. Morales Aguilera, “LEFM-H2E-DEMO-UNESCO: Complete H2E Governance Implementation,” GitHub, 2026. https://github.com/frank-morales2020/MLxDL/blob/main/LEFM-H2E-DEMO_UNESCO.ipynb.
- [8] I. M. Gelfand and G. E. Shilov, *Generalized Functions, Vol. 2: Spaces of Fundamental and Generalized Functions*. Academic Press, 1968.
- [9] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed. Oxford University Press, 1986.
- [10] A. Connes, “Trace formula in noncommutative geometry and the zeros of the Riemann zeta function,” *Selecta Mathematica (N.S.)*, vol. 5, pp. 29–106, 1999, arXiv:math/9811068.
- [11] A. M. Odlyzko, “Tables of zeros of the Riemann zeta function,” 2026. http://www.dtc.umn.edu/~odlyzko/zeta_tables/.