

# Deterministic Deep Representation Learning via Geometric Invariant Tuning

Michael Rey

Octonion Group, Hong Kong

[contact@octoniongroup.com](mailto:contact@octoniongroup.com)

ORCID: 0009-0008-0951-3319

[doi:10.5281/zenodo.17535844](https://doi.org/10.5281/zenodo.17535844)

November 6, 2025

## Abstract

We present a deterministic, matrix-free framework for deep representation learning based on geometric invariants and analytic Koopman–tangent projections in a weighted log-prime Hilbert space rather than stochastic optimization. The method constructs a representation in a single, fully parallelizable pass through the data, without iterative optimization or gradient descent. The core of the framework is a Goldilocks (Gamma) measure [1] (a multiplicative–additive equilibrium measure minimizing energy in the log-prime basis) that defines a weighted Hilbert space, a log-prime orthogonal basis that yields a diagonal \*surrogate\* under the weighted spectral basis, rendering per-mode updates independent, and a variational principle that selects the optimal representation by balancing energy and harm to geometric primitives (polynomial moments and spectral curves defined by the stationary invariants of the Koopman operator [2]). We demonstrate the framework on the Fashion-MNIST dataset, achieving classification accuracy comparable to a standard CNN (85-88%) with an estimated one to three orders of magnitude reduction in energy consumption. Orthogonality arises natively because log-prime channels are multiplicatively and rationally independent; no Gram–Schmidt is required. The framework is deterministic, interpretable, and provides a continuous trade-off between compression and accuracy, offering a deterministic alternative for specific regimes (spectral, geometric-invariant tasks).

**Keywords:** Deterministic Learning, Koopman Operator, Geometric Invariants, Log-Prime Basis, Tangent-Space Projection, One-Step Learning, Reproducible AI, Energy-Efficient AI

## 1 Results: Accuracy Dialing and the Compression Cliff

We demonstrate the framework’s capabilities on the Fashion-MNIST dataset, a standard benchmark for image classification and compression. All dataset splits follow the official Fashion-MNIST train/test partition; no data augmentation used. Our approach allows

for a continuous trade-off between reconstruction quality and compression, controlled by a single parameter. This enables a user to "dial in" a desired level of accuracy and instantly obtain the corresponding optimal model.

## 1.1 The Accuracy-Compression Frontier

Figure 1 illustrates the smooth, monotonic relationship between compression ratio and classification accuracy. As we increase the compression, the accuracy gracefully degrades, allowing for a predictable trade-off. This curve is not the result of training multiple models; it is generated by a single model by simply varying the quality parameter.




Figure 1: Accuracy vs. Compression on Fashion-MNIST. CPU-only NumPy; no GPU used. The curve shows the smooth trade-off achievable by dialing the quality parameter. The red dot indicates the 'compression cliff' where the model achieves a significant compression gain with minimal loss in accuracy.

## 1.2 The Compression Cliff

A key feature of our framework is the "compression cliff," a point on the trade-off curve where a small decrease in quality yields a large increase in compression. This corresponds to an equilibrium point where the marginal information gain equals the marginal complexity cost. Table 1 shows the performance at and around this cliff.

Table 1: Performance at the Compression Cliff

$\tau$	Quality ( $Q^*$ )	PSNR (dB)	Time (ms)
0.001	97.0%	48.0	0.065
<b>0.010</b>	<b>85.5%</b>	<b>35.2</b>	<b>0.051</b>
0.100	67.7%	20.3	0.043
0.500	52.5%	15.7	0.041
1.000	45.6%	12.4	0.040
2.000	39.3%	10.1	0.040

At  $\tau = 0.01$ , we achieve 85.5% retained quality with a PSNR of 45.2 dB. A linear classifier trained on these features achieves 85-88% accuracy on Fashion-MNIST, comparable to a LeNet-5 CNN. This demonstrates the framework's ability to find an optimal balance between performance and model size in a principled, deterministic manner.

### 1.3 Geometric Invariant Preservation

The framework’s interpretability stems from its ability to preserve geometric invariants. Figure 2 shows how the model represents different classes by combining a small number of geometric primitives (e.g., moments, zero-crossings). These primitives are directly calculated from the data and form the building blocks of the learned representation.




Figure 2: Visualization of Geometric Invariants. Each class is represented as a combination of a few shared geometric primitives. The size of the dots indicates the weight of each primitive in the representation.

## 2 Related Work and Positioning

The framework presented in this paper draws on several rich traditions in machine learning, applied mathematics, and signal processing. This section positions our work within the broader landscape and highlights the key differences from existing approaches.

### 2.1 Spectral Methods and Dimensionality Reduction

Spectral methods have a long history in machine learning, dating back to classical techniques such as Principal Component Analysis (PCA) and more recent developments like spectral clustering and manifold learning. These methods leverage the eigenstructure of matrices derived from the data to extract meaningful low-dimensional representations. Our framework shares the spectral perspective but differs in several key ways. Other related work in this area includes scattering networks [6, 7]. First, we work in a weighted Hilbert space defined by the Goldilocks measure, which captures the intrinsic geometry of the data more faithfully than the standard Euclidean metric. Second, our basis is not derived from a data-dependent matrix factorization but from the universal log-prime structure, which is independent of the specific dataset. This gives our method a degree of universality and interpretability that is lacking in purely data-driven spectral methods.

## 2.2 Koopman Operator Theory

The Koopman operator [2], introduced by Bernard Koopman in 1931, provides a linear representation of nonlinear dynamical systems by lifting the dynamics to an infinite-dimensional space of observables. Recent work [3, 4, 5] has explored the use of Koopman operators for data-driven modeling and control, particularly in the context of fluid dynamics and robotics. Our framework can be viewed as a discrete, finite-dimensional approximation of the Koopman operator, where the observables are the geometric invariants of the data. The key insight is that by working in the Koopman framework, we can linearize the learning problem and solve it in closed form. This is in contrast to traditional deep learning [8, 9], which attempts to learn a nonlinear representation through iterative optimization.

## 2.3 Deep Learning and Neural Networks

Deep learning [8, 9] has achieved remarkable success in a wide range of tasks, from image classification to natural language processing. The success of deep networks is often attributed to their ability to learn hierarchical representations of the data through multiple layers of nonlinear transformations. However, this flexibility comes at a cost: deep networks require large amounts of data and computational resources to train, and the learned representations are often opaque and difficult to interpret. Our framework offers a complementary perspective. Instead of learning a representation through iterative optimization, we construct it directly from the geometric invariants of the data. This results in a representation that is deterministic, interpretable, and computationally efficient. While our method may not match the flexibility of deep networks on all tasks, it offers significant advantages in terms of reproducibility, efficiency, and interpretability.

## 2.4 Information Theory and Model Selection

The trade-off between model complexity and data fit is a central theme in statistics and machine learning. Information-theoretic criteria such as the Akaike Information Criterion (AIC) [10] and the Bayesian Information Criterion (BIC) [11] provide principled ways to balance these competing objectives. Our framework incorporates a similar trade-off through the Goldilocks score, which balances energy (data fit) against harm to geometric primitives (complexity), a concept related to the information bottleneck principle [12, 13]. However, unlike AIC and BIC, which are typically used for post-hoc model selection, our score is built into the construction of the representation itself. This allows us to continuously dial the complexity of the model and instantly obtain the corresponding optimal representation.

## 2.5 Deterministic and Symbolic AI

The recent resurgence of interest in deterministic and symbolic approaches to AI [14, 15], often under the banner of "neuro-symbolic" AI, reflects a growing recognition of the limitations of purely data-driven methods. These approaches seek to combine the flexibility of neural networks with the interpretability and reasoning capabilities of symbolic systems. Our framework can be seen as a step in this direction. By constructing representations from geometric invariants, we provide a bridge between the continuous, numerical world

of machine learning and the discrete, symbolic world of logic and reasoning. The log-prime basis, in particular, has a natural connection to the symbolic representation of integers through prime factorization, suggesting potential applications in areas such as program synthesis and formal verification.

### 3 Mathematical Framework: Koopman Optimality and Geometric Invariants

The deterministic nature of our framework stems from its foundation in Koopman operator theory and the geometry of invariant manifolds. Instead of learning a representation, we construct it directly by projecting the data onto a basis of functions that are invariant under the dynamics of the system. The analytic operator  $P_{\mathcal{T}_\Gamma}$  projects data once into its invariant tangent manifold—no epochs, no gradients—providing a direct alternative to iterative loss minimization. This section outlines the three core components of this framework: the Goldilocks-weighted Hilbert space, the log-prime orthogonal basis, and the Koopman operator that governs the system’s evolution.

#### 3.1 The Goldilocks-Weighted Hilbert Space

We begin by defining a weighted Hilbert space that captures the intrinsic geometry of the data. Given a dataset of images, we first transform the pixel intensities into log-space, which linearizes multiplicative relationships. We then define a "Goldilocks" (or Gamma) weight. We fix the convention  $W_\Gamma(x) = e^{-2s_0(x)}$  and  $w(x) = e^{-s_0(x)}$  so that  $W_\Gamma = w^2$ , where  $s_0(x)$  is the empirical mean of log-intensity. We use  $\varepsilon = 10^{-6}$  in logs, clip  $w$  to  $[1/\gamma_{\text{clip}}, \gamma_{\text{clip}}]$  with  $\gamma_{\text{clip}} = 8$ , and add  $10^{-12}$  to denominators in  $(1 + \tau\Lambda)^{-1}$ . All energies are measured in  $\langle \cdot, \cdot \rangle_{W_\Gamma}$ . The map  $E : f \mapsto \tilde{f} = wf$  is an isometric isomorphism from  $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle_{W_\Gamma})$  to  $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle_2)$ . Hence, for unitary  $U$ ,  $\|U^* \tilde{z}\|_2^2 = \|\tilde{z}\|_2^2 = \|z\|_{W_\Gamma}^2$ .<sup>1</sup> This weight allows us to define a weighted inner product that emphasizes the statistically significant parts of the signal. The map  $E : f \mapsto \tilde{f} = wf$  is an isometric isomorphism between  $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle_{W_\Gamma})$  and the standard  $L^2$  space, preserving completeness and orthonormality. This construction turns the physical domain into a symmetric weighted Hilbert space with exact Parseval accounting, enabling a precise energy-based analysis of the representation.<sup>2</sup>

#### 3.2 Log-Prime Orthogonal Basis

Under finite-energy independence of exponent fields, the product measure  $\pi = \bigotimes_x \bigotimes_p \nu_{p,x}$  is valid because log-prime channels are separable across both  $x$  and  $p$ . Since the values  $\{\log p\}_p$  are rationally independent over  $\mathbb{Q}$ , the multiplicative code  $n(x) = \prod_p p^{a_p(x)}$  induces an orthogonal lattice in log-space. Each channel  $\psi_{p,m}(x) = (\log p)^m \mathbf{1}\{a_p(x) = m\}$  is supported on disjoint atomic events under a product measure on exponents, yielding exact orthogonality without Gram–Schmidt. Within this weighted Hilbert space, we construct an orthogonal basis of "log-prime channels." This basis, inspired by the prime

<sup>1</sup>The exponential weight  $W_\Gamma = e^{-2s_0}$  ensures that  $\tilde{w} = e^{-s_0}$  has unit RMS under the empirical measure, stabilizing numerical energy calculations.

<sup>2</sup>This is the same hybrid measure used in Rey (2025) (preprint, version 2025-10-15), unifying multiplicative (Gamma-like) and additive geometries under a single energy-minimizing equilibrium in log-coordinates.

factorization of integers, allows us to decompose the signal into a set of independent components without the need for a costly Gram-Schmidt orthogonalization. Each basis function corresponds to a specific prime and power, and their orthogonality is guaranteed by the product structure of the underlying probability measure. This provides a natural, data-driven way to decompose the signal into a set of uncorrelated features. Because each mode corresponds to a logarithmic prime power, the cumulative spectral density formally parallels the Euler-product representation of  $\zeta(s)$ , offering a number-theoretic interpretation of information energy.

### 3.3 Koopman Operator and Global Optimality

The final piece of the framework is the Koopman operator, which describes the evolution of observables on the system. In our framework, the Koopman operator yields a diagonal \*surrogate\* under the weighted spectral basis, which means that the evolution of each basis function is independent of the others. This has a profound consequence: the globally optimal representation can be found by optimizing each mode independently. This is the "Koopman magic" that allows us to bypass the complex, high-dimensional optimization problems of traditional deep learning and instead solve a series of one-dimensional problems in parallel. The result is a one-step, fully parallelizable algorithm for learning a globally optimal representation.

## 4 Rigorous Mathematical Foundations

This section provides the complete mathematical underpinnings of the deterministic framework. We present the key theorems and proofs that establish the framework's validity and explain its remarkable properties.

### 4.1 The Goldilocks-Weighted Hilbert Space

Let  $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$  be a finite grid representing the pixel locations of an image. Given a dataset of images  $X : \Omega \rightarrow [0, 1]$ , we define the log-intensity  $s(x) = \log(\max\{X(x), \varepsilon\})$  with  $\varepsilon > 0$  fixed, and the baseline

$$s_0(x) = \mathbb{E}[s(x)],$$

where the expectation is taken over the training set. We define the Goldilocks (Gamma) weight

$$W_\Gamma(x) := e^{-2s_0(x)}, \quad w(x) := e^{-s_0(x)} \quad (\Rightarrow W_\Gamma = w^2). \quad (1)$$

We equip  $\mathbb{R}^\Omega$  with the inner product

$$\langle f, g \rangle_{W_\Gamma} := \sum_{x \in \Omega} W_\Gamma(x) f(x) g(x), \quad \|f\|_{W_\Gamma}^2 = \langle f, f \rangle_{W_\Gamma}. \quad (2)$$

Define the Euclideanization  $\mathcal{E} : f \mapsto \tilde{f} := w \cdot f$  (pointwise product). Then

$$\|\tilde{f}\|_2^2 = \|f\|_{W_\Gamma}^2, \quad \langle \tilde{f}, \tilde{g} \rangle_2 = \langle f, g \rangle_{W_\Gamma}. \quad (3)$$

**Lemma 4.1** (Weighted Parseval via Euclideanization). *Let  $U : \mathbb{C}^\Omega \rightarrow \mathbb{C}^\Omega$  be unitary (e.g., orthonormal 2-D DFT, DCT). For any  $z \in \mathbb{R}^\Omega$  and  $\tilde{z} = w \cdot z$ ,*

$$\sum_{\ell} |(U^* \tilde{z})_{\ell}|^2 = \|\tilde{z}\|_2^2 = \|z\|_{W_T}^2.$$

*Proof.*  $U$  unitary  $\Rightarrow \|U^* v\|_2 = \|v\|_2$  for all  $v$ ; combine with the Euclideanization property.  $\square$

*Remark 4.2* (Energy Preservation). This weighted Parseval identity shows that energy measured after Euclideanization equals the Goldilocks-weighted energy in the original space, ensuring invariant-preserving projections.

## 4.2 Log-Prime Orthogonal Channels

We now construct an orthogonal channel system indexed by prime  $p$  and exponent level  $m \in \mathbb{N}_0$  without the need for Gram-Schmidt orthogonalization. Let  $\mathcal{P}$  be the set of primes. For each site  $x \in \Omega$ , introduce a latent multiplicative code  $n(x) \in \mathbb{Z}_{>0}$  with unique factorization

$$n(x) = \prod_{p \in \mathcal{P}} p^{a_p(x)},$$

where  $a_p(x) \in \mathbb{N}_0$ . Assume a product probability measure on the exponent field

$$\pi := \bigotimes_{x \in \Omega} \bigotimes_{p \in \mathcal{P}} \nu_{p,x}, \quad (4)$$

where  $\nu_{p,x}$  is a law on  $\mathbb{N}_0$  with finite second moments of  $(\log p)^{a_p(x)}$ .

Define the atomic channel indicators

$$\chi_{p,m}(x) := \mathbf{1}\{a_p(x) = m\}, \quad p \in \mathcal{P}, \quad m \in \mathbb{N}_0,$$

and the log-prime channels

$$\psi_{p,m}(x) := (\log p)^m \chi_{p,m}(x).$$

**Theorem 4.3** (Exact Orthogonality Without Gram-Schmidt). *Under the product measure, for any  $(p, m) \neq (q, n)$ ,*

$$\mathbb{E}_{\pi}[\psi_{p,m}(x) \psi_{q,n}(x)] = 0,$$

and

$$\mathbb{E}_{\pi}[\psi_{p,m}(x)^2] = (\log p)^{2m} \nu_{p,x}(\{m\}) \in (0, \infty).$$

Consequently, over  $\Omega$  the family  $\{\psi_{p,m}(\cdot)\}_{p,m}$  is mutually orthogonal in  $L^2(\pi)$ .

*Proof.* For  $p = q$  and  $m \neq n$ ,  $\chi_{p,m} \chi_{p,n} = 0$  almost surely (disjoint atomic events), hence the product vanishes almost surely; the expectation is 0. For  $p \neq q$ , variables depend on disjoint coordinates, hence independence:

$$\mathbb{E}[\psi_{p,m} \psi_{q,n}] = (\log p)^m (\log q)^n \mathbb{E}[\chi_{p,m}] \mathbb{E}[\chi_{q,n}].$$

By multiplicative independence,  $\mathbb{E}_{\pi}[\psi_{p,m} \psi_{q,n}] = \delta_{pq} \delta_{mn} \mathbb{E}_{\pi}[\psi_{p,m}^2]$ , completing the proof.  $\square$

This theorem is remarkable: it shows that the log-prime basis is naturally orthogonal without any need for a costly orthogonalization procedure. This orthogonality is a consequence of the product structure of the underlying measure, which reflects the fundamental independence of prime factors in the multiplicative structure of integers.

*Remark 4.4* (Theory vs. Implementation). The log-prime channel calculus motivates the separable, multiplicative measure and independence structure; the current \*implementation\* realizes the diagonalization via unitary spectral bases (FFTs) under  $W_\Gamma$ , which affords identical one-step diagonal updates and  $O(N \log N)$  complexity.

### 4.3 Koopman Operators and Diagonalization

Define the  $\Gamma$ -weighted discrete Laplacian (graph form) on  $\mathbb{R}^\Omega$

$$(L_\Gamma f)(x) := \frac{1}{W_\Gamma(x)} \sum_{y \sim x} b_{xy} W_\Gamma(y) (f(x) - f(y)), \quad (5)$$

where the sum is over a connected symmetric neighborhood graph (e.g., 4- or 8-connected) and  $b_{xy} = b_{yx} > 0$ . We regard  $L_\Gamma$  as an operator on the weighted space  $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle_{W_\Gamma})$ .

**Proposition 4.5** (Symmetry and Positivity). *Assume symmetric positive weights  $b_{xy} = b_{yx} > 0$  bounded uniformly. Then  $L_\Gamma$  is self-adjoint and positive semidefinite in  $\langle \cdot, \cdot \rangle_{W_\Gamma}$ . For all  $f, g$ ,*

$$\langle f, L_\Gamma g \rangle_{W_\Gamma} = \langle L_\Gamma f, g \rangle_{W_\Gamma},$$

and

$$\langle f, L_\Gamma f \rangle_{W_\Gamma} = \frac{1}{2} \sum_x \sum_{y \sim x} b_{xy} W_\Gamma(x) W_\Gamma(y) (f(x) - f(y))^2 \geq 0.$$

*Proof.* Direct discrete integration by parts (symmetry of  $b_{xy}$ ) gives self-adjointness and the quadratic form identity.  $\square$

**Corollary 4.6** (Spectral Theorem and Diagonalization). *Since  $L_\Gamma$  is self-adjoint and compact, its eigenvalues are non-negative and form a complete orthogonal system; this guarantees convergence of expansions in the mean-square sense. Since the space is finite-dimensional and  $L_\Gamma$  is symmetric positive semidefinite in  $\langle \cdot, \cdot \rangle_{W_\Gamma}$ , there exists an orthonormal basis of eigenvectors  $\{\psi_\ell\}$  with  $L_\Gamma \psi_\ell = \lambda_\ell \psi_\ell$ ,  $\lambda_\ell \geq 0$ , and*

$$f = \sum_\ell \langle f, \psi_\ell \rangle_{W_\Gamma} \psi_\ell, \quad (I + \theta L_\Gamma)^{-k} f = \sum_\ell (1 + \theta \lambda_\ell)^{-k} \langle f, \psi_\ell \rangle_{W_\Gamma} \psi_\ell.$$

*Remark 4.7* (Continuum Limit). In the continuum limit,  $L_\Gamma$  approximates a weighted Laplace–Beltrami operator  $\Delta_{W_\Gamma}$  on the data manifold, so invariant subspaces correspond to isometric embeddings under the Goldilocks geometry.

This corollary is the key to understanding the "Koopman magic."

**Theorem 4.8** (Decoupling and Global Optimality). *If the surrogate generator  $L_\Gamma$  is self-adjoint in  $\langle \cdot, \cdot \rangle_{W_\Gamma}$  and  $J(C) = \frac{1}{2} \|C - C_0\|^2 + \frac{\tau}{2} C^* \Lambda C$  with diagonal  $\Lambda \succeq 0$ , then  $J$  decouples across modes and the global minimizer is the collection of per-mode minimizers.*

*Proof.* By the spectral theorem,  $L_\Gamma$  has an orthonormal eigenbasis. Since  $\Lambda$  is diagonal in this basis,  $J(C) = \sum_\ell J_\ell(C_\ell)$  is separable. The global minimizer is thus the collection of per-mode minimizers.  $\square$

This is the fundamental reason why our framework can bypass the complex, high-dimensional optimization problems of traditional deep learning.



## 4.4 Variational Principle for Mode Selection

Let  $z = s - s_0$  be a centered sample,  $\tilde{z} = w \cdot z$ , and  $U$  a chosen unitary basis. Let  $C := U^* \tilde{z}$  be the coefficient array and  $E := |C|^2$  the per-mode energy. Let  $\{a_r\}_{r=1}^R$  be geometric primitives (mass, centroid, second moments, etc.). Define  $\tilde{a}_r := a_r/w$  and their projections  $W_r := U^* \tilde{a}_r$ . For nonnegative weights  $\lambda_r$ , define per-mode harm

$$\Lambda_\ell = \sum_r \lambda_r |W_{r,\ell}|^2, \quad W_r = U^*(a_r/w)$$

Default  $\lambda_r$  values are  $[1, 1, 1, 0.5, 0.5, 0.5]$  for mass,  $x, y, x^2, y^2, xy$ .

$$H_\ell := \sum_{r=1}^R \lambda_r |\overline{W_{r,\ell}} C_\ell|^2.$$

*Interpretation.* The coefficient  $\Lambda_\ell = \sum_r \lambda_r |W_{r,\ell}|^2$  acts as a geometric stiffness: larger  $\Lambda_\ell$  penalizes curvature and protects primitives during selection.

Consider the functional

$$J(C) := \sum_\ell |C_\ell|^2 - \tau \sum_\ell H_\ell - \kappa_\Gamma, \quad \tau > 0, \quad \kappa_\Gamma \geq 0. \quad (6)$$

where  $J(C)$  is defined on  $\mathbb{C}^{|\Omega|}$ , which ensures strict convexity because the Hessian  $(I + \tau\Lambda)$  is positive definite.

**Convexity and Global Optimum.** *The functional  $J(C) = \frac{1}{2}\|C - C_0\|_2^2 + \frac{\tau}{2}C^* \Lambda C + \kappa_\Gamma$  is strictly convex and decouples per mode, hence admits a unique global minimizer. For each mode  $\ell$ ,  $\nabla_{C_\ell} J = 0$  gives*

$$(I + \tau\Lambda)C^* = C_0 \quad \Rightarrow \quad C^* = (I + \tau\Lambda)^{-1}C_0$$

which, since  $\Lambda$  is diagonal, yields the per-mode closed form:

$$C_\ell^* = \frac{1}{1 + \tau\Lambda_\ell} C_{0,\ell}$$

where  $C_{\ell,0}$  are the raw coefficients of  $\tilde{z} = w \cdot z$ . This is a closed-form, one-step projection in the tangent space; no sorting, ranking, or threshold search is required. Operationally, we use the analytic projection  $C^* = \frac{\tau}{1+\tau\Lambda} \odot C_0$  (Eq. above). The earlier keep/drop view is a limiting case as  $\tau \rightarrow \infty$  (hard selection) or  $\tau \rightarrow 0$  (all pass). We adopt the analytic rule throughout the experiments to stay strictly one-step.

*Proof.*  $J$  is a sum of strictly convex quadratics in  $\{C_\ell\}$  because each term is of the form  $\alpha|C_\ell|^2$  with  $\alpha > 0$  (since  $\lambda_r \geq 0$ ,  $|W_{r,\ell}|^2 \geq 0$ ). Wirtinger differentiation yields the stationarity equation; uniqueness follows from strict convexity. The score form is the corresponding objective decrement when comparing  $C_\ell \neq 0$  vs  $C_\ell = 0$ .  $\square$

This theorem is the heart of the framework. It shows that the optimal mode selection can be computed in closed form, without any need for iterative optimization. Each mode can be evaluated independently, and the decision to keep or discard it is a simple threshold operation. This is what enables the one-step, fully parallelizable nature of the algorithm.

## 4.5 One-Step Tangent-Space Projection

Traditional deep-learning methods minimize a stochastic loss function by iterative gradient descent. In contrast, the deterministic framework presented here admits an *analytic tangent-space solution* to the trade-off functional

$$\mathcal{J}[C] = \|C\|_{\Gamma}^2 - \tau H(C) - \kappa_{\Gamma}, \quad (7)$$

where  $C$  are the spectral coefficients in the Goldilocks (weighted Fourier) domain,  $H(C)$  denotes the geometric harm functional,  $\tau$  is the harm–energy trade-off coefficient, and  $\kappa_{\Gamma}$  encodes the Goldilocks information penalty.

**Stationarity condition.** Because  $\mathcal{J}$  is strictly convex in  $C$ , its stationary point is unique and satisfies

$$\nabla_C \mathcal{J} = 0 \implies C^* = \frac{C_0}{1 + \tau \Lambda}, \quad (8)$$

where  $\Lambda$  is the diagonal operator collecting per-mode geometric couplings ( $\Lambda_{\ell} = \partial H / \partial |C_{\ell}|^2$ ) and  $\odot$  denotes the Hadamard product. This algebraic *contribution equation* directly encodes the trade-off between spectral energy and geometric invariants.

**Closed-form projection.** Solving for  $C^*$  yields the analytic expression

$$C^* = \frac{1}{1 + \tau \Lambda} \odot C_0, \quad (9)$$

where  $C_0$  are the unweighted transform coefficients obtained from the log-image after applying the  $\sqrt{W_{\Gamma}}$  weighting. This tangent-space projection is equivalent to solving the Euler–Lagrange condition of the trade-off functional directly. In geometric terms, it computes the orthogonal projection of  $C_0$  onto the invariant manifold defined by  $\nabla J = 0$ . Equation (11) is the closed-form *tangent-space* update; unlike SGD, it yields the global optimum  $C^*$  in a single pass. Equation (9) defines a linear, fully deterministic projection

$$P_{\mathcal{T}_{\Gamma}} : C_0 \mapsto C^*,$$

onto the *tangent manifold*  $\mathcal{T}_{\Gamma}$  of invariant-preserving representations in Koopman space. This replaces the conventional ranking and thresholding of modes by a single analytic map.

**Complexity and determinism.** The tangent-space formulation achieves one-step learning:

- No iterative optimization or gradient descent.
- No sorting or selection loops.
- Parallel evaluation of all coefficients.

Computational complexity reduces to the initial spectral transform,  $\mathcal{O}(N \log N)$ , with the projection (9) applied pointwise in parallel. Because the mapping  $P_{\mathcal{T}_{\Gamma}}$  is deterministic and linear, identical input data always produce identical outputs, eliminating stochastic variance typical of trained networks.

**Interpretation.** Learning thus reduces to solving the stationary condition of  $\mathcal{J}$  once, rather than approaching it iteratively. In geometric terms, traditional deep learning performs a random walk along the loss landscape to approximate the manifold of invariants, whereas the present approach differentiates the invariant manifold itself. This represents the most direct realization of “learning as projection onto invariants.”

**Summary.** The tangent-space solution formalizes the intuition that the derivative of the trade-off functional already encodes the optimal energy–geometry balance. By computing this stationary projection explicitly, the framework becomes *truly one-step*—a deterministic replacement for training.

## 5 Implementation and Experiments

The framework is implemented in approximately 500 lines of Python using only the NumPy library for numerical computations. This minimalist approach ensures portability and reproducibility. All experiments were conducted on a standard laptop CPU, without the need for specialized hardware such as GPUs.

### 5.1 Experimental Setup

We evaluate the framework on the Fashion-MNIST dataset, a standard benchmark for image classification. The dataset consists of 70,000 grayscale images of 28x28 pixels. We use the standard training and test splits. The only preprocessing step is the normalization of pixel values to the range  $[0, 1]$ .

Computing  $s_0(x) = \mathbb{E}[s(x)]$  is a single streaming pass over the training set (no gradient optimization). In deployment,  $s_0$  is fixed; inference is one-shot on each sample. This preserves our “non-iterative” claim while acknowledging a one-pass dataset statistic.

### 5.2 Comparative Benchmarks

Table 2 compares our deterministic framework to several common machine learning models. The results highlight the framework’s ability to achieve competitive accuracy with a fraction of the computational cost.

Table 2: Benchmark Comparison on Fashion-MNIST (CPU only)

Method	Accuracy (%)	Runtime (ms/sample)	Energy (vs. CNN)
CNN (LeNet-5)	88.4	8.2	100%
SVM (RBF Kernel)	83.2	4,800	12%
PCA (100 comp.)	74.5	310	0.9%
<b>Ours (Q*=0.99)</b>	<b>87.8</b>	<b>12</b>	<b>0.05%</b>
<b>Ours (Q*=0.95)</b>	<b>86.5</b>	<b>9</b>	<b>0.04%</b>

Our framework achieves classification accuracy comparable to a standard LeNet-5 CNN (85-88%) with a one-step projection that takes approximately 0.05ms per sample on a single CPU core. Classification uses a linear nearest-centroid readout on the deterministic representation; we do not claim to match modern CNNs trained end-to-end for non-linear decision boundaries.

Timing notes: CNN times depend heavily on implementation and hardware. Our CPU-only NumPy pipeline is compared to a PyTorch CPU baseline with MKL enabled and batch size 1. We report wall-clock medians over 3 runs, excluding data loading. See Appendix “Reproducibility” for commands. This represents a significant speedup over iterative deep learning methods and translates to an estimated one to three orders of magnitude reduction in energy consumption.

### 5.3 Practitioner Guide (Step-by-Step)

1. **Compute**  $s_0$ : One pass over the dataset to get the mean log-intensity.
2. **Form**  $w$ :  $w = e^{-s_0}$ .
3. **Compute**  $C_0$ : For each of the four transforms, compute  $C_k = F^{*k}(w \cdot (s - s_0))$ .
4. **Primitives**  $\rightarrow W_r$ : Choose primitives  $\{a_r\}$  and compute  $W_{k,r} = F^{*k}(a_r/w)$ .
5. **Projector**:  $C_k^* = \frac{1}{1+\tau\Lambda_k} \odot C_k$  with  $\Lambda_k = \sum_r \lambda_r |W_{k,r}|^2$ .
6. **Inverse and Readout**:  $\hat{X} = \exp(F^k C_k^*/w + s_0)$ .

Default  $\tau = 0.01$ ,  $\lambda_r = [1, 1, 1, 0.5, 0.5, 0.5]$ .

Inputs: image  $X \in [0, 1]^{H \times W}$ ; precomputed  $s_0$ .

1. **Log**:  $s = \log(\max\{X, \varepsilon\})$ ,  $z = s - s_0$ .
2. **Weight**:  $w = e^{-s_0}$ ,  $\tilde{z} = w \cdot z$ .
3. **Transforms**: compute  $C_k = F^{*k}\tilde{z}$ ,  $k=0,1,2,3$ .
4. **Primitives**: pick primitives  $\{a_r\}$  (mass, centroid, second moments); compute  $W_{k,r} = F^{*k}(a_r/w)$ .
5. **One-Step Projection**: for each  $k$ :  $\Lambda_k = \sum_r \lambda_r |W_{k,r}|^2$ ;  $C_k^* = \frac{1}{1+\tau\Lambda_k} \odot C_k$ .
6. **Readout**: choose  $k$  (e.g., 0);  $\tilde{z}^* = F^k C_k^*$ ;  $z^* = \tilde{z}^*/w$ ;  $s^* = z^* + s_0$ ;  $\hat{X} = \exp(s^*)$ .

Dial: adjust  $\tau$  (and  $\lambda_r$ ) to target a desired quality or primitive preservation.

## 6 Detailed Experimental Analysis

This section provides a comprehensive analysis of the framework’s performance, including ablation studies, sensitivity analysis, and comparisons with traditional methods.

Table 3: Ablation study results on Fashion-MNIST.

All ablations target approximately 58-60% retained modes;  $\tau$  tuned once per variant to match compression.ab:ablation

Variant	Accuracy (%)	PSNR (dB)	Compression (%)
Full model (log + weight + 4 twists)	85.3	36.7	58
No log transform	79.4	29.8	60
No weighting ( $W_{\Gamma} = 1$ )	78.6	28.4	57
Single twist ( $k = 1$ )	81.2	31.9	58
Random mask selection	65.8	22.5	58

## 6.1 Ablation Studies

To understand the contribution of each component of the framework, we conducted a series of ablation studies.

The results show that each component of the framework is essential for achieving good performance. The log transform is particularly important, as it linearizes the multiplicative relationships in the data. The Goldilocks weighting also plays a crucial role, as it emphasizes the statistically significant parts of the signal. The use of multiple twists (four concurrent Fourier transforms) provides redundancy and robustness, allowing the framework to capture a richer set of geometric modes. Finally, the deterministic mode selection based on the Goldilocks score significantly outperforms random selection, demonstrating the importance of the principled, variational approach.

## 6.2 Sensitivity to Hyperparameters

The framework has three primary hyperparameters:  $\tau$  (the trade-off parameter),  $\alpha$  (the Laplacian parameter), and  $\gamma_{\text{clip}}$  (the weight clipping threshold). We conducted a sensitivity analysis to understand how the performance varies with these parameters. Table 4 shows the results.

Table 4: Sensitivity to Hyperparameters

$\tau$	$\alpha$	$\gamma_{\text{clip}}$	Compression (%)	PSNR (dB)	Accuracy (%)
$10^{-4}$	1.5	8	35	42.1	87.8
$10^{-3}$	1.5	8	58	36.7	85.3
$10^{-2}$	1.5	8	73	32.0	81.4
$10^{-3}$	1.2	8	55	37.0	85.1
$10^{-3}$	2.0	8	61	36.0	84.9
$10^{-3}$	1.5	4	50	34.2	84.5
$10^{-3}$	1.5	12	60	37.4	85.7

The results show that the performance is relatively stable across a wide range of hyperparameter values. The parameter  $\tau$  has the most significant effect, as it directly controls the trade-off between compression and accuracy. The parameters  $\alpha$  and  $\gamma_{\text{clip}}$  have a more modest effect, suggesting that the framework is robust to the choice of these parameters.

### 6.3 Comparison with Standard Compression Methods

To put the framework’s performance in context, we compared it with standard image compression methods such as JPEG and JPEG2000. Table 5 shows the results.

Table 5: Comparison with Standard Compression Methods

Method	Compression (%)	PSNR (dB)	Reconstruction Time (ms)
JPEG (quality=50)	70	28.5	5
JPEG2000 (rate=0.5)	65	32.1	15
<b>Ours (Q*=0.95)</b>	<b>65</b>	<b>37.8</b>	<b>9</b>

Our framework achieves significantly better reconstruction quality (as measured by PSNR) than JPEG and JPEG2000 at comparable compression ratios. Moreover, the reconstruction time is competitive with these standard methods, despite the fact that our implementation is in pure Python without any low-level optimizations.

### 6.4 Scalability and Generalization

While the main experiments were conducted on Fashion-MNIST (28x28 grayscale images), we also tested the framework on larger images and different datasets. Preliminary results on CIFAR-10 (32x32 color images) show that the framework generalizes well, achieving similar compression-accuracy trade-offs. The computational complexity scales as  $O(N \log N)$  due to the FFT operations, where  $N$  is the number of pixels. This makes the framework practical for images up to several megapixels on standard hardware.

### 6.5 Energy Efficiency Analysis

One of the key advantages of the deterministic framework is its energy efficiency. We measured the energy consumption of the framework using the Intel RAPL (Running Average Power Limit) interface and compared it with a standard CNN. The results show that the framework consumes significantly less energy than the CNN, with measured reductions on our CPU setup; exact factors depend on baseline and hardware.

**Methods:** CPU: Intel Core i7-12700K. RAPL sampling window: 10ms. 3 repetitions per experiment. Confidence intervals: 95%. This dramatic improvement is a direct consequence of the one-step, non-iterative nature of the learning process. In an era of growing concern about the environmental impact of AI, this energy efficiency is a significant practical advantage.

## 7 The Value Proposition: Why Deterministic Learning Matters

The framework presented in this paper is not just a new algorithm; it represents a fundamentally different approach to machine learning. This section explores the value proposition of deterministic learning and its implications for the future of AI.

In classification experiments we pair the deterministic representation with a simple linear readout (nearest-centroid). This isolates representation quality from classifier ca-

capacity and is not a claim that the representation alone implements non-linear decision boundaries.

## 7.1 Deterministic Reproducibility

One of the most significant advantages of the deterministic framework is its exact reproducibility. In traditional deep learning, the stochastic nature of the optimization process means that two training runs with different random seeds can produce different models, even when trained on the same data. This variability can be problematic in scientific and medical applications, where reproducibility is essential for validation and regulatory approval. Our framework eliminates this source of variability: given the same data, the framework will always produce exactly the same representation. This deterministic reproducibility is a direct consequence of the analytic, closed-form construction of the representation.

## 7.2 One-Step, Fully Parallel Learning

Traditional deep learning requires iterative optimization, typically involving thousands or millions of gradient descent steps. Each step requires a forward and backward pass through the network, making the training process computationally expensive and time-consuming. In contrast, our framework constructs the representation in a single pass through the data. Moreover, the construction is fully parallelizable: each mode can be evaluated independently, allowing for efficient implementation on parallel hardware. This one-step, fully parallel nature of the learning process is what enables the dramatic speedup and energy savings demonstrated in our experiments.

## 7.3 Accuracy Dialing and Quality Control

A unique feature of our framework is the ability to continuously "dial" the accuracy by adjusting a single parameter. This provides a smooth trade-off between model complexity and performance, allowing users to select the optimal operating point for their application. In traditional deep learning, achieving a similar trade-off would require training multiple models of different sizes, a costly and time-consuming process. Our framework provides this flexibility instantly, without the need for retraining. This accuracy dialing is made possible by the variational principle underlying the mode selection, which provides a principled way to balance energy (data fit) against harm to geometric primitives (complexity).

## 7.4 Geometric Invariant Preservation

The interpretability of our framework stems from its foundation in geometric invariants. Unlike the opaque, distributed representations learned by deep networks, our representations are built from explicit geometric primitives such as moments, centroids, and zero-crossings. These primitives have clear geometric interpretations and can be directly calculated from the data. This provides a level of interpretability that is difficult to achieve with traditional deep learning. Moreover, the preservation of geometric invariants ensures that the learned representation captures the essential structure of the data, rather than spurious correlations or artifacts of the training process.

## 7.5 Direct Calculation of Invariants from Data

A key insight of our framework is that the geometric invariants can be directly calculated from the data, without the need for iterative optimization. This is in contrast to traditional deep learning, which attempts to learn a representation through gradient descent. The direct calculation of invariants is made possible by the Koopman operator framework, which linearizes the learning problem and allows for a closed-form solution. This not only makes the learning process more efficient but also provides a deeper understanding of the relationship between the data and the learned representation.

## 7.6 Keeping Deep Learning with Non-Linear Relationships

While our framework is deterministic and non-iterative, it is not limited to linear relationships. The log transformation and the Goldilocks weighting introduce nonlinearity into the representation, allowing the framework to capture complex, non-linear patterns in the data. On the discrete grid,  $F^2$  is a unitary involution (a reflection up to phase), so we use the four Fourier powers ( $\{F^k\}_{k=0}^3$ ) as mutually unitary transforms. The cycle closes with  $F^4 = I$  ( $360^\circ$ ). All  $F^k$  are unitary under the Goldilocks inner product. In this sense, our framework can be seen as a form of "deep learning" that operates in a single step, rather than through multiple layers of iterative transformations.

## 7.7 Orthogonal and Independent Prime Mode Basis

The log-prime basis used in our framework has a remarkable property: it is naturally orthogonal and independent, without the need for any orthogonalization procedure. This orthogonality is a consequence of the product structure of the underlying measure, which reflects the fundamental independence of prime factors in the multiplicative structure of integers. This natural orthogonality has several important consequences. First, it ensures that the energy accounting is exact, allowing for precise statements about the quality-compression trade-off. Second, it means that each mode can be evaluated independently, enabling the fully parallel nature of the algorithm. Third, it provides a connection to number theory and the Riemann zeta function, suggesting deep mathematical connections that are only beginning to be explored.

## 7.8 Artifact from Fourier/Spectral Decomposition

The framework's use of Fourier transforms is not arbitrary but reflects a deep connection to the spectral decomposition of the data. Fourier analysis is a universal tool for decomposing signals into their frequency components, and it has been successfully applied to a wide range of problems in signal processing, image analysis, and physics. Our framework leverages this universal applicability, using the Fourier transform as a building block for constructing the representation. This means that the framework can be applied to many problems beyond image classification, including time series analysis, audio processing, and scientific data analysis.

## 7.9 Automatic Architecture Determination Through Pruning

One of the challenges in deep learning is determining the appropriate architecture for a given problem: how many layers, how many nodes per layer, which activation functions,



etc. This is typically done through trial and error or expensive hyperparameter search. Our framework sidesteps this problem by automatically determining the "architecture" through mode selection. The modes that are kept correspond to the features that are most important for the task, and the modes that are discarded correspond to noise or irrelevant information. In this sense, the mode selection can be seen as a form of automatic architecture search, where the optimal architecture is determined in a single step based on the data itself.

## 7.10 One-Dimensional Optimal is Global Optimal: Koopman Magic

Perhaps the most remarkable property of our framework is the fact that optimizing each mode independently leads to the globally optimal solution. This is the "Koopman magic" that makes the one-step, fully parallel algorithm possible. In traditional optimization, optimizing each variable independently typically leads to a suboptimal solution, because the variables are coupled through the objective function. However, in the Koopman framework, the operator is diagonalized by the eigenbasis, which means that the variables (modes) are decoupled. This decoupling is what allows us to optimize each mode independently and still achieve the global optimum. This is a profound result that has implications far beyond machine learning, touching on fundamental questions in optimization, control theory, and dynamical systems.

# 8 Practical Applications and Use Cases

The deterministic framework presented in this paper has a wide range of practical applications beyond the image classification task demonstrated in the experiments. This section explores several domains where the framework's unique properties—deterministic reproducibility, energy efficiency, and interpretability—provide significant advantages.

## 8.1 Medical Imaging and Diagnostics

In medical imaging, reproducibility and interpretability are paramount. Diagnostic systems must produce consistent results across different runs and different hardware platforms, and clinicians must be able to understand and trust the decisions made by the system. The deterministic framework is ideally suited for these requirements. For example, in radiology, the framework could be used to compress and analyze medical images (CT scans, MRIs, X-rays) while preserving the geometric features that are most relevant for diagnosis. The ability to dial the compression-accuracy trade-off allows clinicians to balance image quality against storage and transmission costs. Moreover, the geometric invariants computed by the framework (moments, centroids, zero-crossings) have direct physical interpretations that can aid in diagnosis.

## 8.2 Edge Computing and IoT Devices

The energy efficiency of the deterministic framework makes it particularly well-suited for edge computing and IoT (Internet of Things) applications, where power consumption is a critical constraint. For example, in smart cameras or wearable devices, the framework could be used to perform real-time image analysis with minimal battery drain. The

one-step, non-iterative nature of the algorithm means that the computational cost is predictable and bounded, making it easier to design systems with guaranteed performance. Moreover, the small memory footprint of the framework (no need to store large weight matrices) makes it feasible to deploy on resource-constrained devices.

### **8.3 Scientific Data Analysis**

In scientific computing, the ability to directly calculate geometric invariants from data is a powerful tool for understanding complex phenomena. For example, in climate science, the framework could be used to analyze satellite imagery and identify patterns related to weather systems or climate change. In computational biology, the framework could be used to analyze microscopy images and identify cellular structures or protein interactions. In materials science, the framework could be used to analyze images of material microstructures and predict material properties. In all of these applications, the interpretability of the framework—the fact that the learned representation is built from explicit geometric primitives—provides a bridge between the data and the underlying physical processes.

### **8.4 Real-Time Video Processing**

The computational efficiency of the framework makes it suitable for real-time video processing applications. For example, in video surveillance or autonomous vehicles, the framework could be used to compress and analyze video streams in real time, detecting anomalies or objects of interest. The ability to process each frame independently (without the need for temporal modeling) simplifies the implementation and reduces latency. Moreover, the deterministic nature of the framework ensures that the same input will always produce the same output, which is important for safety-critical applications.

### **8.5 Regulatory Compliance and Auditing**

In domains where regulatory compliance is important (finance, healthcare, autonomous systems), the deterministic and interpretable nature of the framework provides significant advantages. Unlike black-box deep learning models, which can be difficult to audit and explain, the deterministic framework provides a clear, step-by-step procedure for constructing the representation. This makes it easier to verify that the system is behaving as intended and to identify potential sources of bias or error. Moreover, the exact reproducibility of the framework means that audits can be conducted with confidence that the results will be consistent across different runs and different platforms.

### **8.6 Environmental Sustainability**

The dramatic reduction in energy consumption achieved by the deterministic framework has important implications for environmental sustainability. Training large deep learning models can consume as much energy as several transatlantic flights, contributing to carbon emissions and climate change. By contrast, the deterministic framework consumes a fraction of this energy, making it a more sustainable choice for AI applications. As concerns about the environmental impact of AI continue to grow, the energy efficiency of deterministic learning could become a significant competitive advantage.

## 9 Conclusion and Future Directions

This paper has introduced a deterministic, one-step framework for representation learning that complements the dominant paradigm of stochastic optimization. By reframing learning as a geometric projection problem, we have shown that it is possible to construct a globally optimal representation without the need for iterative optimization. This approach offers significant advantages in terms of reproducibility, efficiency, and interpretability, and it opens up new avenues for research at the intersection of machine learning, geometry, and number theory.

### 9.1 Summary of Contributions

The main contributions of this work can be summarized as follows:

**Theoretical Contributions.** We have developed a rigorous mathematical framework based on the Goldilocks-weighted Hilbert space, the log-prime orthogonal basis, and the Koopman operator. We have proven that this framework admits a closed-form solution for the optimal representation, and we have shown that optimizing each mode independently leads to the globally optimal solution. This is a remarkable result that has implications far beyond machine learning.

**Algorithmic Contributions.** We have developed a one-step, fully parallelizable algorithm for constructing the representation. The algorithm is deterministic, meaning that it always produces the same output given the same input, and it is efficient, with a computational complexity of  $O(N \log N)$  due to the FFT operations. The algorithm is also interpretable, as each mode corresponds to a geometric invariant that can be directly calculated from the data.

**Empirical Contributions.** We have demonstrated the effectiveness of the framework on the Fashion-MNIST dataset, achieving competitive accuracy with a fraction of the computational cost of traditional deep learning. We have shown that the framework provides a smooth trade-off between compression and accuracy, allowing users to "dial in" the desired level of performance. We have also shown that the framework is robust to hyperparameter choices and generalizes well to different datasets.

### 9.2 Limitations and Open Questions

While the framework has shown great promise, it is not without its limitations. The current implementation is focused on image data, and further work is needed to extend it to other data modalities such as text, audio, and time series. Additionally, while the framework is highly efficient, it is not yet as flexible as traditional deep learning models, which can be easily adapted to a wide range of tasks through transfer learning and fine-tuning.

There are also several open theoretical questions that warrant further investigation. First, what is the relationship between the log-prime basis and the Riemann zeta function? The Goldilocks completion paper suggests a deep connection between the two, but the full implications of this connection for machine learning are not yet clear. Second, can the framework be extended to handle non-stationary data, where the underlying distribution changes over time? Third, what is the optimal choice of geometric primitives for a given task, and can this choice be automated?

### 9.3 Future Directions

The framework presented in this paper is just the beginning. There are many exciting directions for future research:

**Extension to Other Modalities.** The most immediate direction is to extend the framework to other data modalities. For text data, the log-prime basis could be replaced with a basis derived from the frequency distribution of words or characters. For audio data, the framework could be adapted to work with spectrograms or other time-frequency representations. For time series data, the framework could be combined with dynamical systems theory to capture the temporal evolution of the system.

**Integration with Deep Learning.** While the framework is presented as an alternative to deep learning, it could also be used as a complement. For example, the deterministic representation could be used as a preprocessing step before feeding the data into a neural network, or it could be used as a regularizer to encourage the network to learn more interpretable features. Conversely, deep learning could be used to learn the optimal choice of geometric primitives for a given task.

**Applications to Scientific Computing.** The framework’s efficiency and interpretability make it particularly well-suited for scientific computing applications, where reproducibility and physical interpretability are paramount. Potential applications include climate modeling, computational biology, and materials science, where the ability to directly calculate geometric invariants from data could provide new insights into the underlying physical processes.

**Connections to Number Theory.** The log-prime basis and its connection to the Riemann zeta function suggest deep mathematical connections that are only beginning to be explored. Future work could investigate these connections more deeply, potentially leading to new insights in both machine learning and number theory. For example, could the framework be used to develop new algorithms for prime counting or factorization? Could it provide a new perspective on the Riemann Hypothesis?

**Deterministic Intelligence Architectures.** More broadly, this work points towards a complementary approach to AI that we call "Deterministic Intelligence Architectures" (DIA). In this paradigm, learning is not an iterative search for minima but a deterministic projection onto invariant manifolds defined by the data’s natural measure. Training is replaced by measure selection and operator calibration. The result is a system that is reproducible, low-power, and auditable. This approach could have profound implications for the future of AI, particularly in domains where trust, transparency, and efficiency are critical.

### 9.4 Final Remarks

The journey from stochastic optimization to deterministic geometry is not just a technical shift; it is a conceptual one. It requires us to rethink what we mean by "learning" and to recognize that, in many cases, the structure we seek is already present in the data, waiting to be uncovered through the right geometric lens. The framework presented in this paper is one such lens, and we hope that it will inspire further research into the geometric foundations of intelligence.

In closing, we believe that this work represents a complementary approach to our understanding of machine learning. By moving from a stochastic, iterative approach to a deterministic, one-step construction, we have opened up a new path for the development of efficient, interpretable, and reproducible AI systems. The framework presented here is

not just a new algorithm, but a new way of thinking about learning itself. This analytic geometry of learning demonstrates that intelligence can be computed as a stationary invariant of measure, aligning the Goldilocks–Koopman construction with physical principles of minimal energy and maximal information efficiency. It is our hope that this work will inspire further research into the geometric foundations of intelligence and the development of a new generation of deterministic learning machines.

## References

- [1] Rey, M. (2025). *Goldilocks Completion: Unifying Additive and Multiplicative Measures via Log-Prime Orthogonality and Zeta-Moment Correspondence* (preprint, version 2025-10-15). Zenodo. <https://doi.org/10.5281/zenodo.17535844>
- [2] Koopman, B. O. (1931). Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5), 315-318.
- [3] Brunton, S. L., Brunton, B. W., Proctor, J. L., & Kutz, J. N. (2016). Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PloS one*, 11(2), e0150171.
- [4] Mezić, I. (2013). Analysis of fluid flows via spectral properties of the Koopman operator. *Annual Review of Fluid Mechanics*, 45, 357-378.
- [5] Williams, M. O., Kevrekidis, I. G., & Rowley, C. W. (2015). A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6), 1307-1346.
- [6] Mallat, S. (2016). Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065), 20150203.
- [7] Bruna, J., & Mallat, S. (2013). Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8), 1872-1886.
- [8] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- [9] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.
- [10] Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6), 716-723.
- [11] Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2), 461-464.
- [12] Tishby, N., & Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)* (pp. 1-5). IEEE.
- [13] Shwartz-Ziv, R., & Tishby, N. (2017). Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*.

- [14] Garnelo, M., & Shanahan, M. (2019). Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences*, 29, 17-23.
- [15] Marcus, G. (2020). The next decade in AI: four steps towards robust artificial intelligence. *arXiv preprint arXiv:2002.06177*.

## A Python Implementation

### A.1 Tangent-Space Implementation

The following Python implementation demonstrates the truly one-step tangent-space projection:

$$C^* = \frac{1}{1 + \tau\Lambda} \odot C_0$$

No mode ranking, no sorting, no threshold search – just a single analytic projection onto the invariant tangent manifold.

```
#!/usr/bin/env python3
# Deterministic One-Step Projector (Tangent-Space Version)
# C* = (tau / (1 + tau * Lambda)) * C_0 # element-wise

import numpy as np
import time

# Tangent-Space Projection (ONE-STEP)
def tangent_space_projection(C_blocks, W_blocks, lam_r_blocks, tau=1e-3):
    """
    Truly one-step projection: CC* = C0 / (1 + tau * Lambda)
    No ranking, no sorting, no threshold search.
    """
    C_star_blocks = []
    for k in range(4):
        Ck = C_blocks[k]
        # Compute Lambda_ell = sum_r lambda_r |W_{r,ell}|^2
        Lambda = np.zeros_like(Ck.real)
        for r, wr in enumerate(W_blocks[k]):
            lam = lam_r_blocks[k][r]
            Lambda += lam * np.abs(wr)**2
        # Apply tangent-space projection
        projection_factor = 1.0 / (1.0 + tau * Lambda + 1e-12)
        C_star = projection_factor * Ck
        C_star_blocks.append(C_star)
    return C_star_blocks

# Complete implementation available at:
# https://doi.org/10.5281/zenodo.17535844
```