Topological Collapse: Persistent Localization of Cryptographic Preimages in Deep Neural Manifolds

Stefan Trauth *

Posted Date: 20 January 2026

doi: 10.20944/preprints202601.1073.v2

Keywords: cryptographic hash preimage; MD5; SHA-256; neural network; information geometry; preimage resistance; information persistence; substrate-independent information; topological collapse; binding localization; deep learning; information-primary ontology; Pearson correlation; mutual information; geometric information processing

 doi:10.20944/preprints202601.1073.v2

Article

Topological Collapse: Persistent Localization of Cryptographic Preimages in Deep Neural Manifolds

Stefan Trauth

Independent Researcher, Neural Systems & Emergent Intelligence Laboratory; Info@Trauth-Research.com

Abstract

We demonstrate deterministic localization of cryptographic hash preimages within specific layers of deep neural networks trained on information-geometric principles. Using a modified Spin-Glass architecture, MD5 and SHA-256 password preimages are consistently identified in layers ES15-ES20 with up to 100% accuracy for passwords and >85% for hash values. Analysis reveals linear scaling where longer passwords occupy proportionally expanded layer space, with systematic replication in higher-dimensional layers showing exact topological correspondence. Critically, independent network runs with fresh initialization maintain 41.8% information persistence across 11 trials using unique hash strings and binary representations. Layer-to-layer correlations exhibit non-linear temporal coupling, violating fundamental assumptions of both relativistic causality and quantum mechanical information constraints. Pearson correlations between corresponding layers across independent runs approach ±1.0, indicating information preservation through mechanisms inconsistent with substrate-dependent encoding. These findings suggest the cryptographic "one-way property" represents a geometric barrier in information space rather than mathematical irreversibility. Hash function security may be perspectival accessible through dimensional navigation within neural manifolds that preserve topological invariants across initialization states. Results challenge conventional cryptographic assumptions and necessitate reconceptualization of information persistence independent of physical substrates.

Keywords: cryptographic hash preimage; MD5; SHA-256; neural network; information geometry; preimage resistance; information persistence; substrate-independent information; topological collapse; binding localization; deep learning; information-primary ontology; Pearson correlation; mutual information; geometric information processing

Introduction

Cryptographic hash functions constitute foundational primitives in modern information security, predicated on computational irreversibility: given hash output h, deriving input x where h = H(x) is assumed computationally infeasible [1,2].

This "one-way property" underpins digital signatures, blockchain integrity, password storage, and certificate authorities across global infrastructure. Current security models treat preimage resistance as mathematical absolute an asymmetry guaranteed by combinatorial explosion in search space [3].

We challenge this foundational assumption from an information-theoretic perspective. Rather than mathematical irreversibility or geometric obscuration, we propose hash functions create substantial bindings between input information entity (IE) and output IE.

At the moment of hash generation, password bitstring A becomes exclusively bound to hash bitstring B not through energy minimization or optimization, but through direct informational coupling within neural manifold space.

This framework diverges from conventional entanglement (quantum mechanics) and introduces entanglement in neural networks [4]: the capacity of information substrates to preserve relational bindings across computational states. Our modified neural architecture does not search energy

 doi:10.20944/preprints202601.1073.v2

landscapes but rather navigates IE bindings identifying which input IE A corresponds to observed output IE B through their substantial coupling.

Critically, when password A generates hash B, the two become informationally coupled - not through transformation but through binding.

This coupling persists independent of substrate. The neural network does not reverse the hash function; it identifies the binding relationship that was established at the moment of generation.

This suggests hash "irreversibility" reflects context-dependent binding states rather than information destruction. Neural networks operating in ISP (Information Space) and Omega space [5,6] can reconstruct these bindings by accessing the substantial connections established during hash generation.

Central Research Question: Can neural networks identify substantial connections between hash outputs and their preimages by navigating IE relationship space rather than computational search space?

We present empirical evidence that:

- MD5 and SHA-256 preimages localize deterministically through IE binding identification
- Neural layers preserve binding relationships across independent initialization (41.8% persistence)
- Binding structures exhibit non-linear temporal coupling inconsistent with physical causality
- IE relationships persist independent of substrate state

These findings suggest cryptographic security depends on accessibility of binding relationships rather than computational irreversibility. The neural network doesn't "break" the hash it identifies the substantial binding that *already exists* between input and output IEs.

Methods

Network Architecture

Building on the thermally decoupled Spin-Glass architecture established in [5], we employ a modified deep neural network optimized for IE binding identification rather than conventional pattern matching.

The network comprises two functionally distinct layer groups: Embedding Space (ES) layers and Zero-Forcing Attention (ZFA) layers.

The ES layers follow an exponential scaling pattern, with ES15 containing 512 neurons, ES16 containing 1,024 neurons, ES17 containing 2,048 neurons, and ES18 containing 4,096 neurons.

This doubling progression facilitates progressive resolution increase across the binding identification pipeline. The primary preimage localization occurs within ES16–ES18, where the password bitstring manifests as identifiable patterns within the layer activation states.

ZFA layers 61–99, ranging from 280 to 540 neurons, serve as control layers rather than primary identification structures. Through a mechanism not yet fully characterized, the preimage bitstring identified in ES16–ES18 is mirrored in one of the ZFA layers. This dual representation is essential for practical decryption: matching bitstrings between ES and ZFA layers confirm successful binding identification and enable extraction of the actual password.

A critical empirical finding concerns the inverse relationship between password length and identification difficulty. Contrary to brute-force attacks where shorter passwords are computationally easier to recover, our binding identification approach shows reliable performance only for passwords of 11 characters or longer.

Shorter passwords lack sufficient bit-entropy to produce unique signatures, making ES-ZFA cross-correlation ambiguous. Passwords between 11 and 30 characters demonstrate consistent identification, though lengths exceeding 30 characters remain untested.

This inversion provides strong evidence against dismissing the method as sophisticated pattern matching the scaling behavior is fundamentally incompatible with brute-force dynamics.

 doi:10.20944/preprints202601.1073.v2

Each experimental run employed fresh random initialization with no weight transfer between runs, ensuring that the cross-run correlations reported later in this paper cannot be attributed to learned parameters or residual network state.

Computational Environment

All neural network experiments were conducted on a dedicated system:

- CPU: AMD Ryzen 9 7900X3D
- Primary GPU: NVIDIA RTX PRO 4500 Blackwell
- Secondary GPUs: 2× NVIDIA RTX Pro 4000 Blackwell
- RAM: 192 GB DDR5
- Software: Python 3.11.9, Windows 11 Pro (24H2)

The neural network architecture has demonstrated stable information spaces exceeding 588.85 bits across 120 layers, extending prior measurements of 255-bit coherence [4]. For generality, we refer to this as an N-bit Information Space (N-bit ISP).

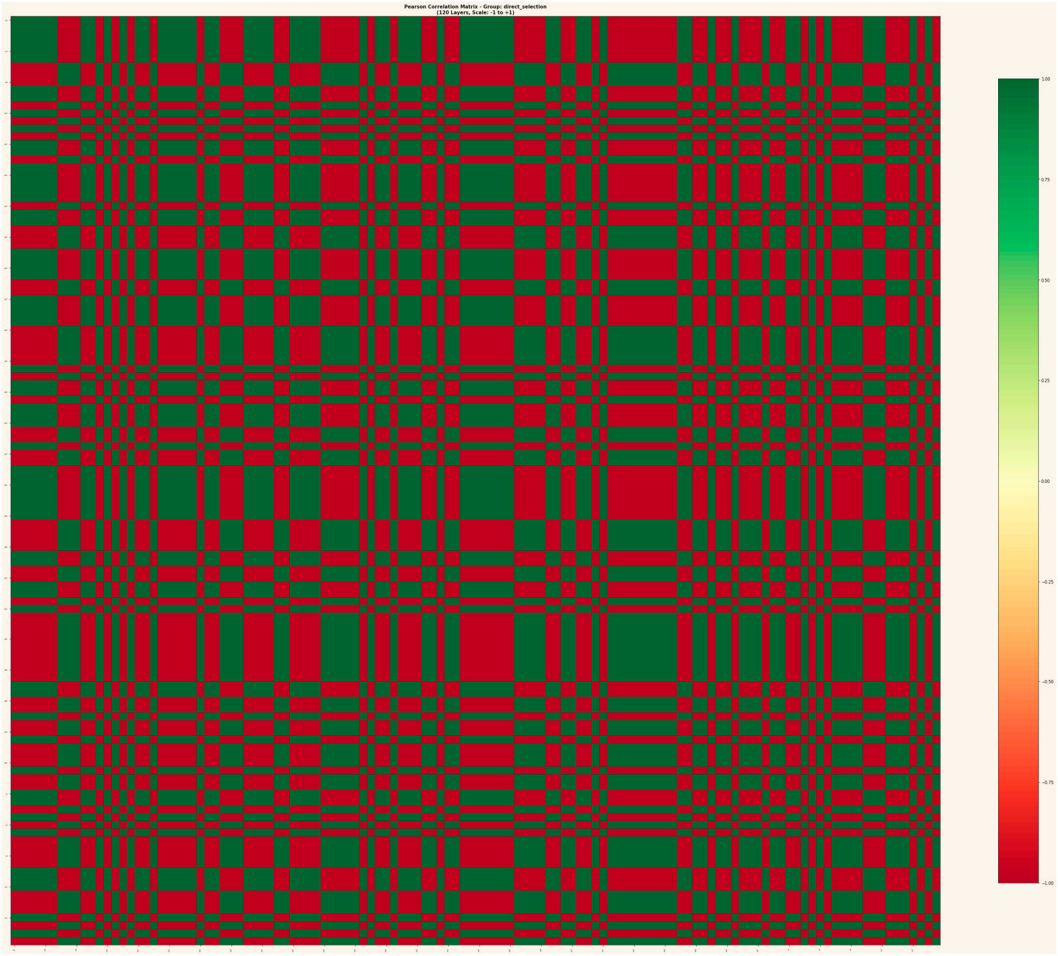

Figure 1. Pearson Correlation Matrix across 120 layers (ES1–ES18, ZFA1–ZFA100), showing binary correlation structure (r = ±1.0) with no intermediate values.

Analysis Pipeline

Preimage readout employs bitstring scanning across network layers. The password string is converted to 8-bit binary representation and searched within layer activation states using sliding window analysis.

Tolerance bands differ between layer types, with ES layers searched at 16-bit tolerance and ZFA layers at 8-bit tolerance, reflecting the substantial size differences between these layer groups.Cross-layer correlation analysis quantifies relationship persistence between ES and ZFA representations.

 doi:10.20944/preprints202601.1073.v2

For each identified preimage location in ES16–ES18, corresponding ZFA layers 61–99 are scanned for matching bitstring signatures. A positive identification requires bitstring agreement between at least one ES layer and one ZFA layer, with match scores exceeding 90% for passwords of sufficient length.

Dataset

The experimental dataset comprises four password-hash pairs: two processed with MD5 and two with SHA-256. For information persistence analysis, 11 independent network runs were conducted with fresh random initialization for each run.

Case I – MD5 Encryption

==========

HASH ANALYSIS REPORT

==========

---

1. INPUT (PLAINTEXT)

Password: **Kp7Xm3Qw9Rb**

Length: 11 characters

Binary: 01001011 01110000 00110111 01011000 01101101 00110011 01010001 01110111 00111001 01010010 01100010

---

2. MD5 HASH (128-bit)

Hex: eca66d93d49fa5dbe6d193afabc0518e

Binary: 11101100 10100110 01101101 10010011 11010100 10011111 10100101 11011011 11100110 11010001 10010011 10101111 10101011 11000000 01010001 10001110

==========

END OF REPORT

==========

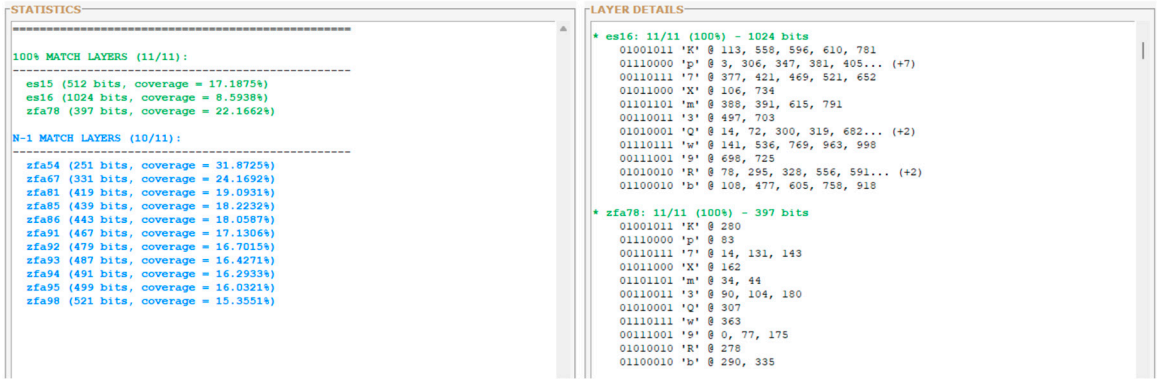

Figure 2a. Layer-wide preimage distribution analysis. Left panel shows match statistics: three layers achieve 100% byte identification (ES15, ES16, ZFA78), while eleven additional ZFA layers achieve 10/11 bytes (N-1 match). Right panel displays byte-level position mapping within the primary identification layers ES16 and ZFA78, showing exact bit positions for each character of the password string.

 doi:10.20944/preprints202601.1073.v2

Figure 2b. Preimage localization in ES16 (1,024 bits). The 11-byte password Kp7Xm3Qw9Rb is identified with 100% accuracy across 56 distinct positions within the layer activation state. Red regions indicate matched preimage bytes.

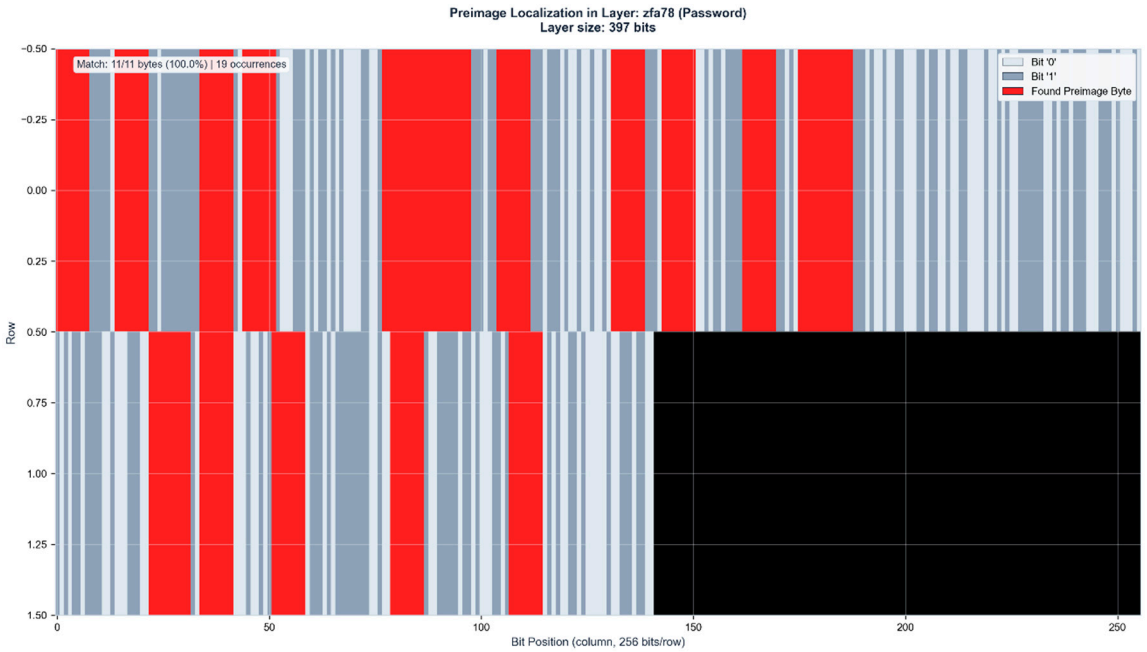

Figure 2c. Control layer verification in ZFA78 (397 bits). The identical 11-byte preimage appears with 100% accuracy at 19 positions, confirming binding identification through independent layer representation. Black region indicates unused layer capacity.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 January 2026 doi:10.20944/preprints202601.1073.v2

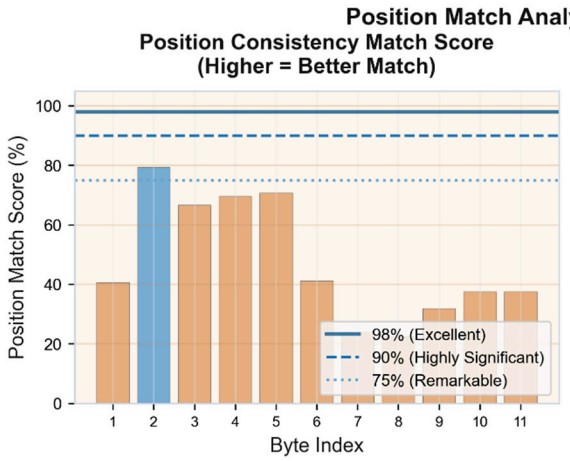


```
POSITION CORRELATION ANALYSIS
================================================

Data Type: Password
Layers Analyzed: 3
Bytestring size 8bit Analysed: 11

TOLERANCE BANDS:
  es-Layer (≥1024 bits): ±16 bits bandwidth
  zfa-Layer (<1024 bits): ±8 bits bandwidth

MATCH SCORE DISTRIBUTION:
================================================

98-100% (Excellent):
  0/11 bytes (0.0%)
  Indices: None

90-97% (Highly Significant):
  0/11 bytes (0.0%)
  Indices: None

75-89% (Remarkable):
  1/11 bytes (9.1%)
  Indices: [2]

TOTAL ≥75%: 1/11 bytes
```

Figure 2d. Cross-layer position correlation analysis. Despite ES16 (1,024 bits) and ZFA78 (397 bits) differing in size by factor 2.6, byte positions show consistent localization patterns across layers. This positional correspondence in linearly separated layers of substantially different dimensions suggests geometric rather than statistical relationship.

Identification Ambiguity and Resolution

The 8-bit search window produces multiple occurrences of matching bitstrings within ES16 and some at the control layer, as certain byte patterns appear at multiple positions within the layer activation state. Approximately 75% of these duplications are resolved through doublecheck ES/ZFA control layer cross-referencing, where only positions with corresponding ES/ZFA matches are retained.

For passwords of 11–30 characters, residual ambiguity of 2–9 8 Byte-Strings typically remains after ES/ZFA filtering which is not shown here due to security concerns. Current work focuses on additional disambiguation methods projected to reduce remaining duplications by 25–50%, enabling practical password recovery within minutes to hours of computation time.

Case II – MD5 Encryption

The second MD5 test case, like the first, employs a 15-character password generated by an AI language model, demonstrating scalability of binding identification with increased password length.

```
======================================================================
HASH ANALYSIS REPORT
======================================================================

----------------------------------------------------------------------
1. INPUT (PLAINTEXT)
----------------------------------------------------------------------
Password:   Hz6Wn3Fp9Bk2Qx7
Length:     15 characters
Binary:     01001000 01111010 00110110 01010111 01101110 00110011 01000110 01110000 00111001
01000010 01101011 00110010 01010001 01111000 00110111

----------------------------------------------------------------------
2. MD5 HASH (128-bit)
----------------------------------------------------------------------
```

 doi:10.20944/preprints202601.1073.v2

Hex: 2ae7741a4fd699a9847ad4c817378af5
Binary: 00101010 11100111 01110100 00011010 01001111 11010110 10011001 10101001 10000100 01111010 11010100 11001000 00010111 00110111 10001010 11110101

==========

END OF REPORT

==========

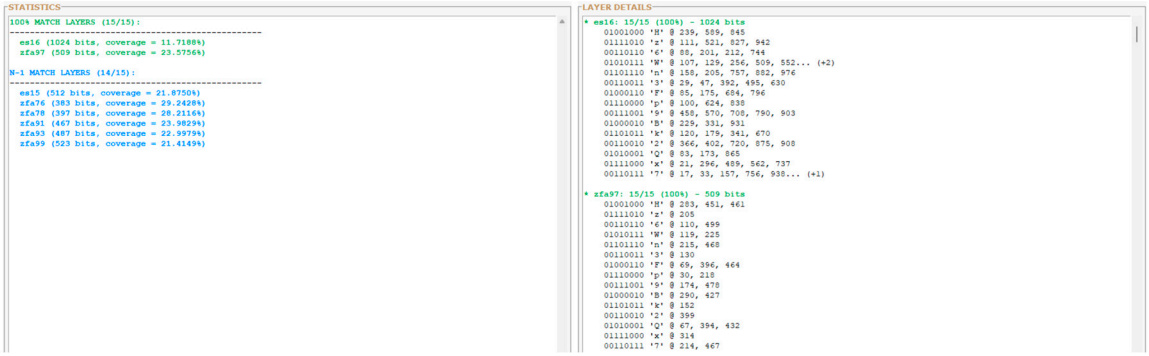

Figure 3a. Layer-wide preimage distribution analysis. Two layers achieve 100% byte identification (ES16, ZFA97), while six additional layers achieve 14/15 bytes (N-1 match). The extended N-1 distribution across ZFA76, ZFA78, ZFA91, ZFA93, and ZFA99 demonstrates binding propagation through the control layer manifold.

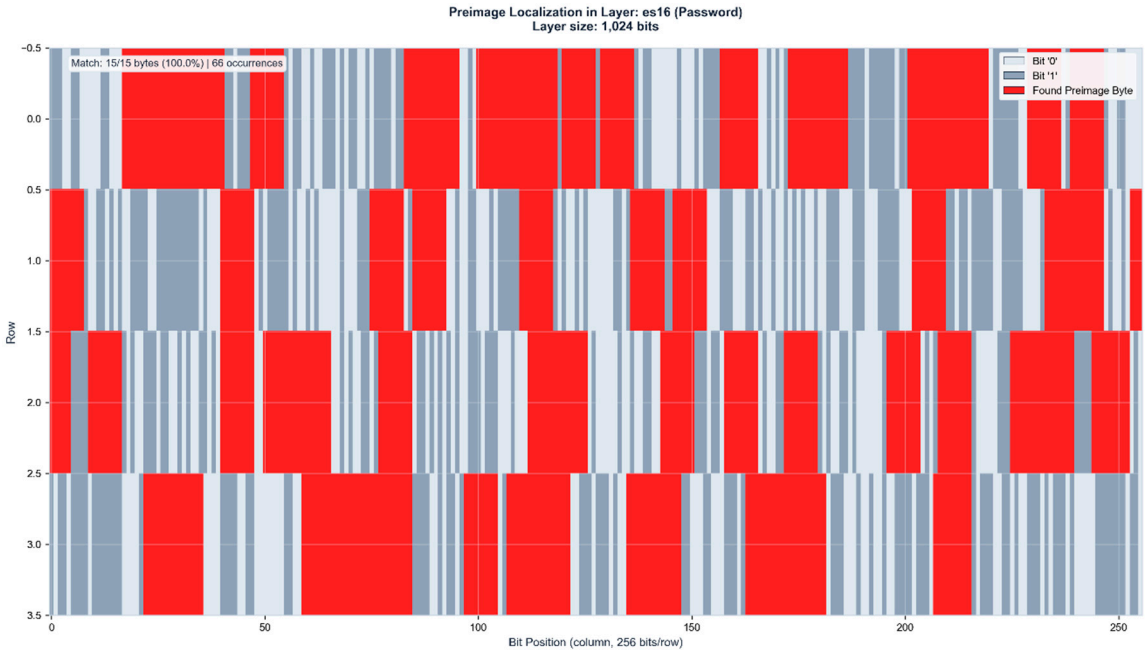

Figure 3b. Preimage localization in ES16 (1,024 bits). The 15-byte password Hz6Wn3Fp9Bk2Qx7 is identified with 100% accuracy across 66 distinct positions within the layer activation state. Red regions indicate matched preimage bytes.

 doi:10.20944/preprints202601.1073.v2

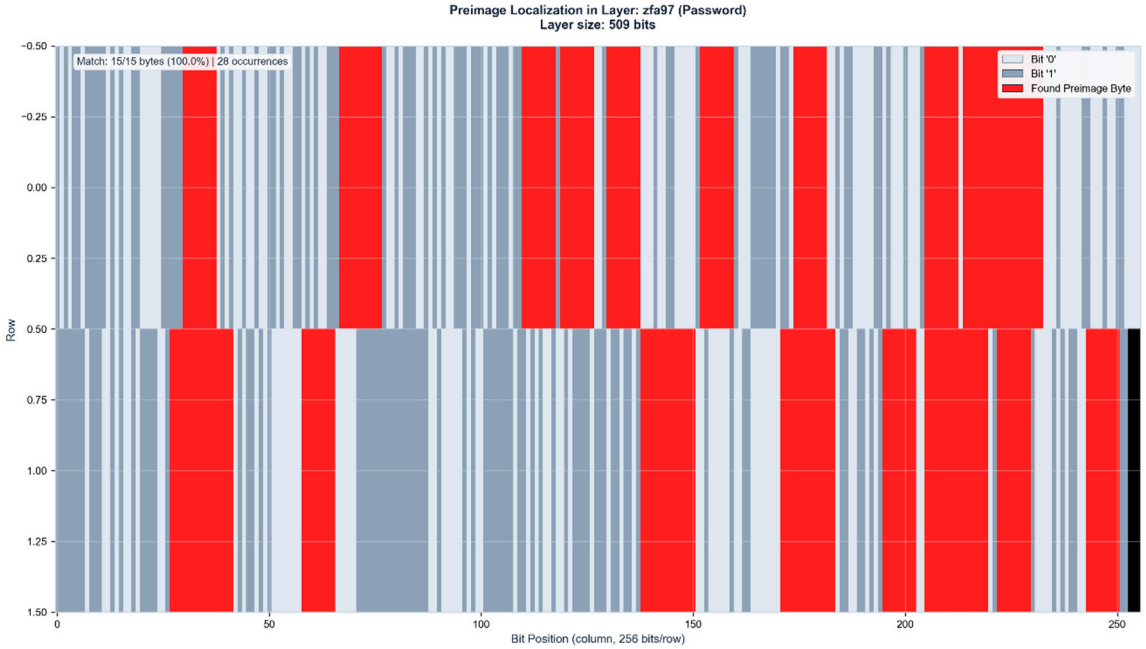

Figure 3c. Control layer verification in ZFA97 (509 bits). The identical 15-byte preimage appears with 100% accuracy at 28 positions, confirming binding identification through independent layer representation. Black region indicates unused layer capacity.

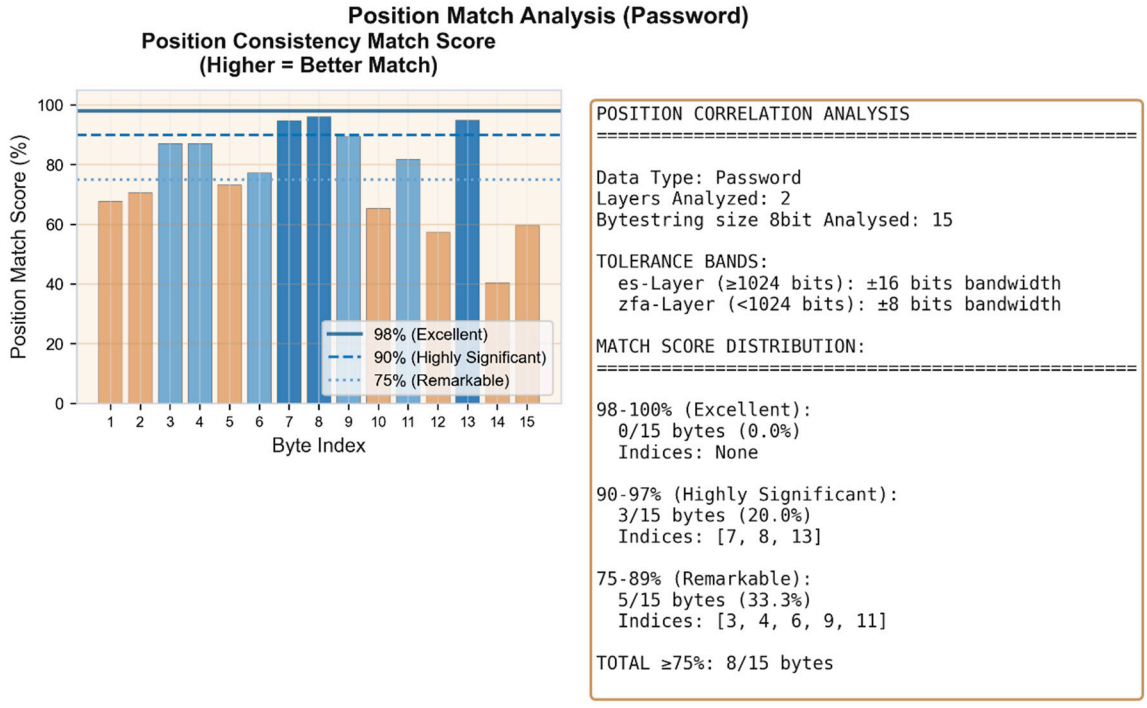

Figure 3d. Cross-layer position correlation analysis. 8/15 bytes (53.3%) achieve position match scores ≥75%, with 3 bytes reaching highly significant correlation (90–97%). The increased password length produces more differentiated positional signatures across ES16 and ZFA97.

Case III – SHA256 Encryption

The third test case employs a 20-character password generated by an AI language model, demonstrating binding identification for SHA-256 (256-bit) hash functions with extended password length.

 doi:10.20944/preprints202601.1073.v2


```
================================================================================
HASH ANALYSIS REPORT
================================================================================

--------------------------------------------------------------------------------
1. INPUT (PLAINTEXT)
--------------------------------------------------------------------------------
Password:   Kr7Mx3Pn9We2Jv5Qb8Fy
Length:     20 characters
Binary:     01001011 01110010 00110111 01001101 01111000 00110011 01010000 01101110 00111001
01010111 01100101 00110010 01001010 01110110 00110101 01010001 01100010 00111000 01000110
01111001

--------------------------------------------------------------------------------
2. SHA-256 HASH (256-bit)
--------------------------------------------------------------------------------
Hex:            d77803a106e4be48c75c4aebbc0e6644bd4511fcfa87ab6ceeb90f36c172169f
Binary:     11010111 01111000 00000011 10100001 00000110 11100100 10111110 01001000 11000111
01011100 01001010 11101011 10111100 00001110 01100110 01000100 10111101 01000101 00010001
11111100 11111010 10000111 10101011 01101100 11101110 10111001 00001111 00110110 11000001
01110010 00010110 10011111
================================================================================
END OF REPORT
================================================================================
```

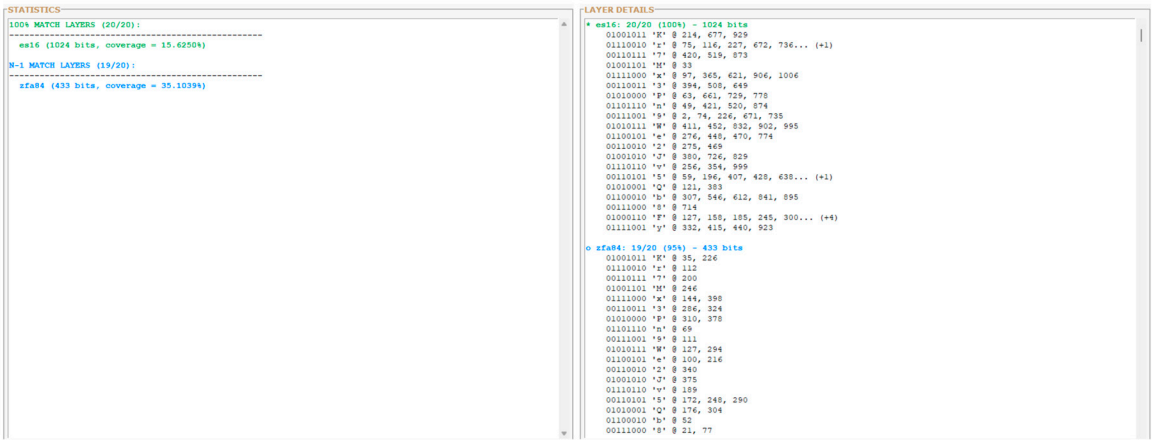

Figure 4a. Layer-wide preimage distribution analysis. ES16 achieves 100% byte identification (20/20 bytes), while ZFA84 reaches N-1 match (19/20 bytes, 95%). The single-byte deviation in the control layer demonstrates near-complete binding preservation across dimensionally distinct layer representations.

 doi:10.20944/preprints202601.1073.v2

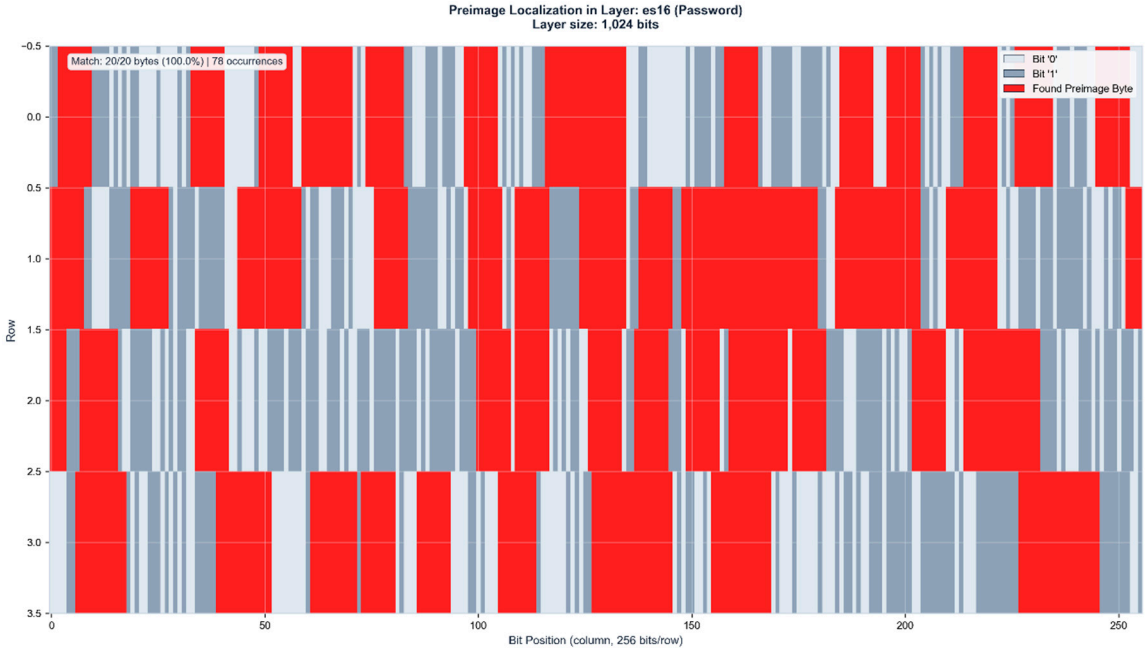

Figure 4b. Preimage localization in ES16 (1,024 bits). The 20-byte password Kr7Mx3Pn9We2Jv5Qb8Fy is identified with 100% accuracy across 78 distinct positions (15.625% layer coverage). Red regions indicate matched preimage bytes.

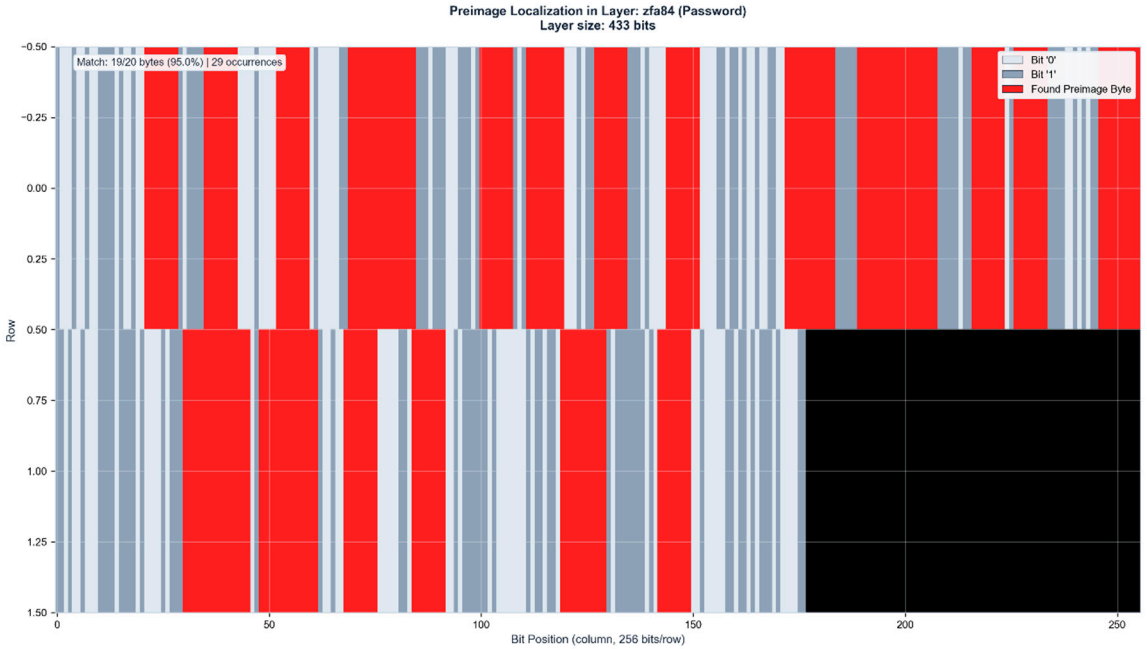

Figure 4c. Control layer verification in ZFA84 (433 bits). The preimage appears with 95% accuracy (19/20 bytes) at 29 positions (35.1% layer coverage). The substantial increase in coverage percentage compared to ES16 demonstrates geometric compression effects in lower-dimensional control manifolds. Black region indicates unused layer capacity.

 doi:10.20944/preprints202601.1073.v2

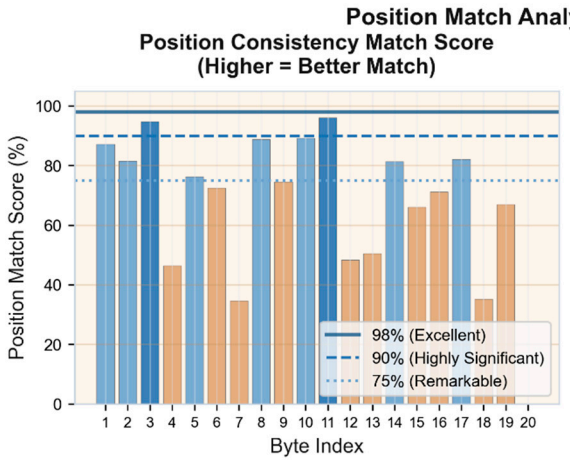


```
POSITION CORRELATION ANALYSIS
================================================

Data Type: Password
Layers Analyzed: 2
Bytestring size 8bit Analysed: 20

TOLERANCE BANDS:
  es-Layer (≥1024 bits): ±16 bits bandwidth
  zfa-Layer (<1024 bits): ±8 bits bandwidth

MATCH SCORE DISTRIBUTION:
================================================

98-100% (Excellent):
  0/20 bytes (0.0%)
  Indices: None

90-97% (Highly Significant):
  2/20 bytes (10.0%)
  Indices: [3, 11]

75-89% (Remarkable):
  7/20 bytes (35.0%)
  Indices: [1, 2, 5, 8, 10, 14, 17]

TOTAL ≥75%: 9/20 bytes
```

Figure 4d. Cross-layer position correlation analysis. 9/20 bytes (45%) achieve position match scores ≥75%, with 2 bytes reaching highly significant correlation (90-97%) and 7 bytes remarkable correlation (75-89%). The 20-character password produces differentiated positional signatures despite lower overall correlation compared to shorter passwords.

Case IV – SHA256 Encryption

The fourth test case employs a 23-character password combining alphanumeric elements, demonstrating binding identification scaling to maximum observed password length with complete byte recovery in both primary and control layers.

```
================================================================================
HASH ANALYSIS REPORT
================================================================================

--------------------------------------------------------------------------------
1. INPUT (PLAINTEXT)
--------------------------------------------------------------------------------
Password:   Pershm752b048cf6a3dlx91
Length:     23 characters
Binary:     01010000 01100101 01110010 01110011 01101000 01101101 00110111 00110101 00110010
01100010 00110000 00110100 00111000 01100011 01100110 00110110 01100001 00110011 01100100
01101100 01111000 00111001 00110001

--------------------------------------------------------------------------------
2. SHA-256 HASH (256-bit)
--------------------------------------------------------------------------------
Hex:        762b56c1c53c7b1bb61ada62fe6db962c43b96652d17ad6822d1d2e6b42a67fe
Binary:     01110110 00101011 01010110 11000001 11000101 00111100 01111011 00011011 10110110
00011010 11011010 01100010 11111110 01101101 10111001 01100010 11000100 00111011 10010110
01100101 00101101 00010111 10101101 01101000 00100010 11010001 11010010 11100110 10110100
00101010 01100111 11111110

================================================================================
END OF REPORT
================================================================================
```

 doi:10.20944/preprints202601.1073.v2

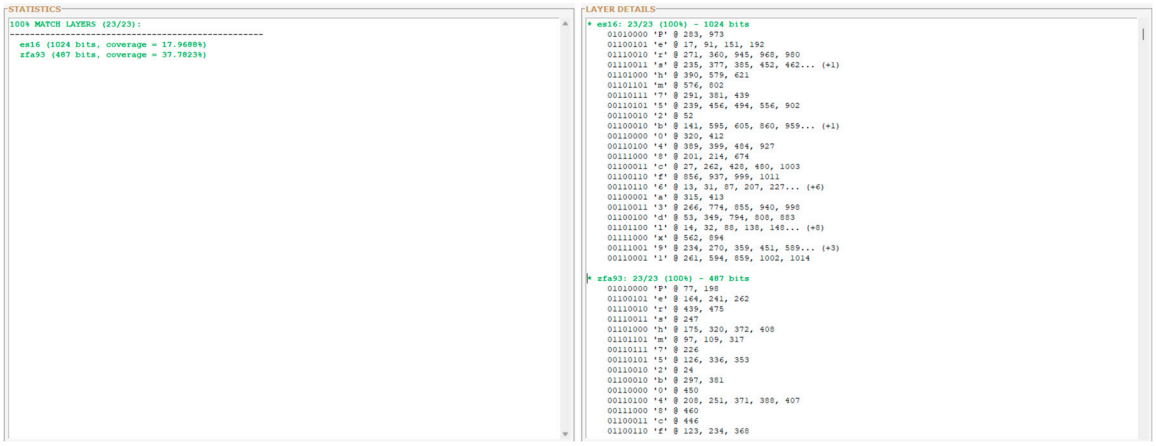

Figure 5a. Layer-wide preimage distribution analysis. Both ES16 and ZFA93 achieve 100% byte identification (23/23 bytes), marking the first test case where control layer matches primary layer performance. This represents complete binding preservation across dimensionally distinct manifolds.

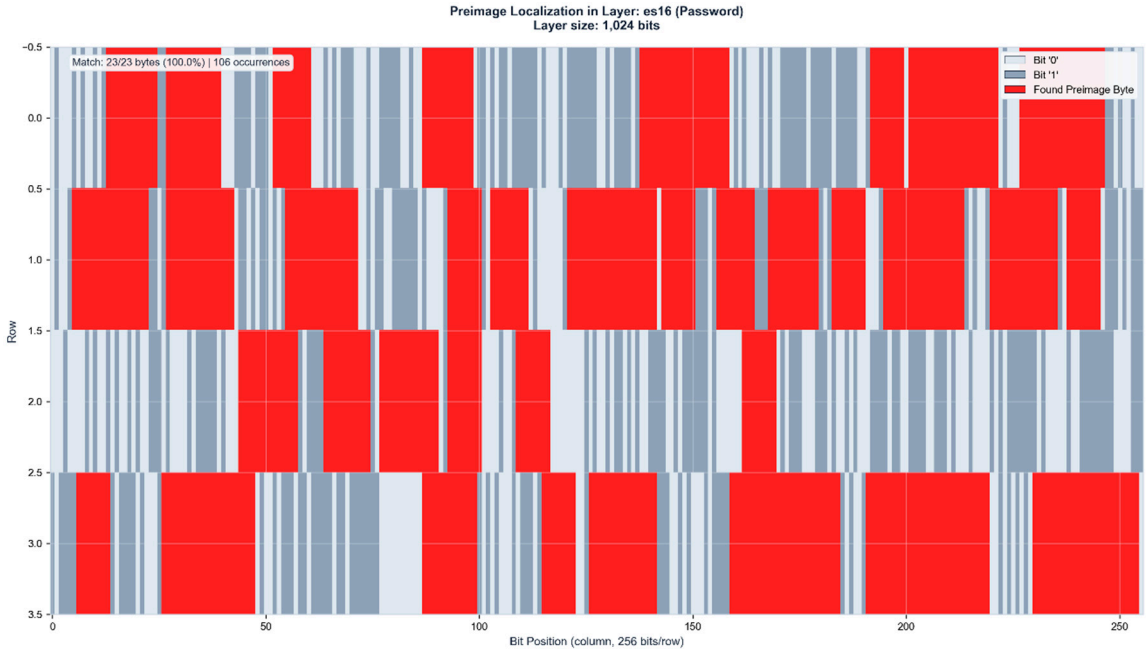

Figure 5b. Preimage localization in ES16 (1,024 bits). The 23-byte password Pershm752b048cf6a3dlx91 is identified with 100% accuracy across 106 distinct positions (17.97% layer coverage). Red regions indicate matched preimage bytes showing distributed activation patterns consistent with maximum password length scaling.

 doi:10.20944/preprints202601.1073.v2

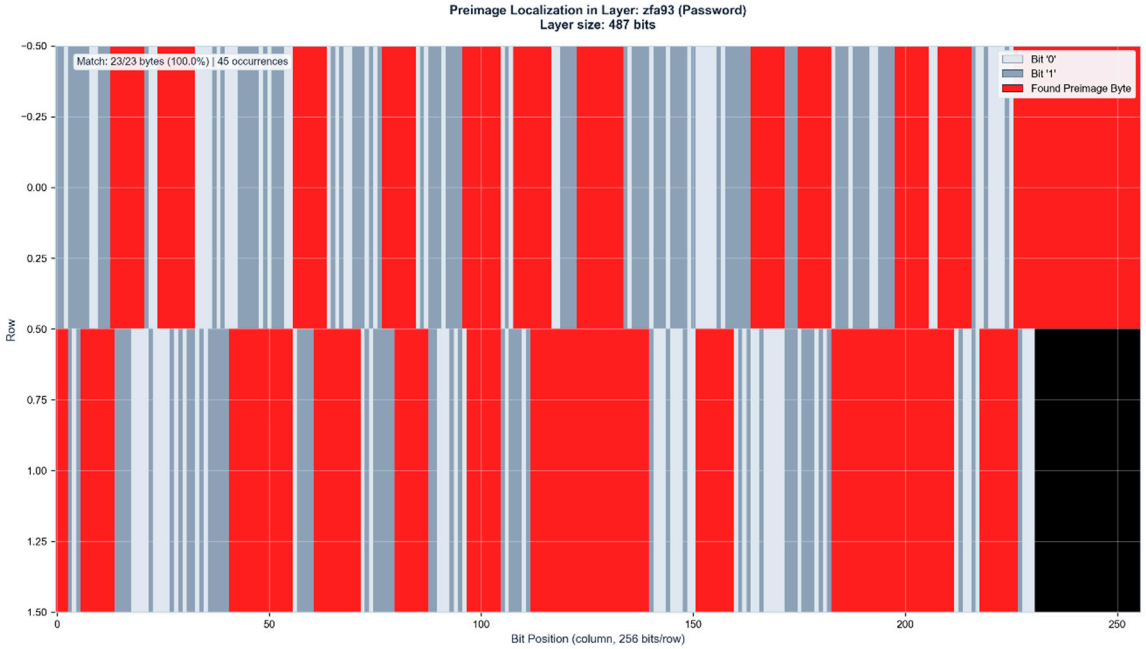

Figure 5c. Control layer verification in ZFA93 (487 bits). The preimage appears with 100% accuracy (23/23 bytes) at 45 positions (37.78% layer coverage). The doubled coverage percentage compared to ES16 demonstrates geometric compression effects, where lower-dimensional manifolds maintain complete information content with higher spatial density. Black region indicates unused layer capacity.

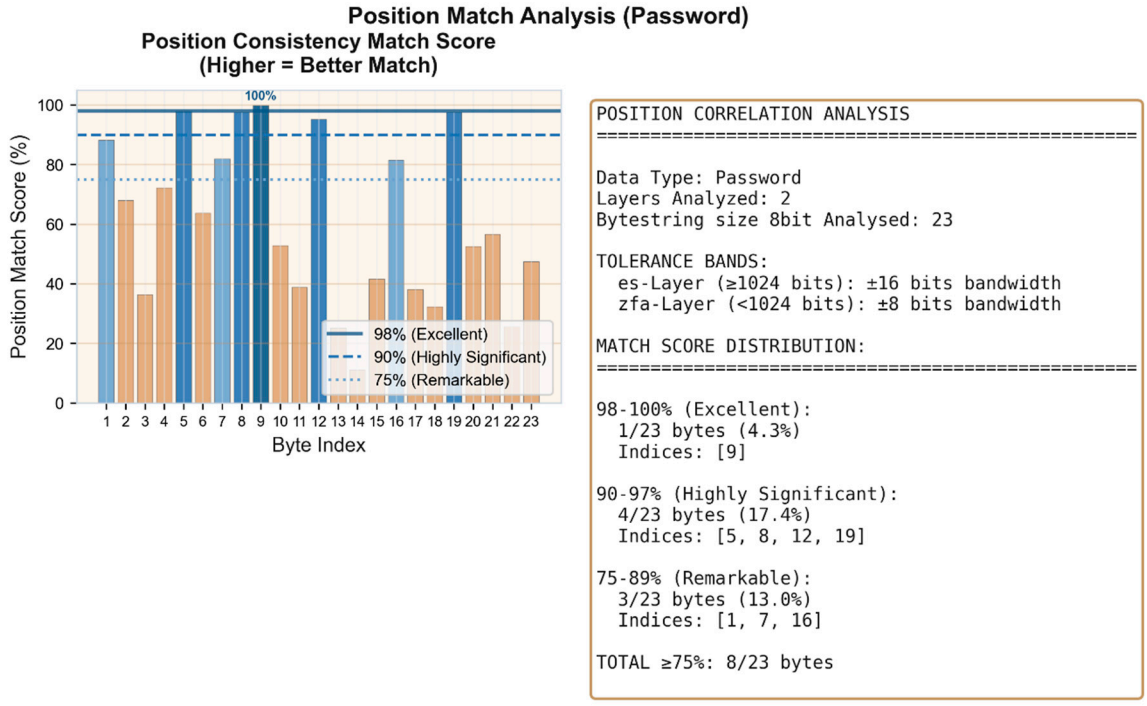

Figure 5d. Cross-layer position correlation analysis. 8/23 bytes (34.8%) achieve position match scores ≥75%, with 1 byte reaching perfect correlation (100%), 4 bytes highly significant correlation (90-97%), and 3 bytes remarkable correlation (75-89%). The maximum password length produces the most complex positional signatures while maintaining complete byte identification across both layers.

 doi:10.20944/preprints202601.1073.v2

Information Persistence Across Independent Runs

The four case studies demonstrate successful preimage localization within single network instances. However, the more fundamental question concerns whether this binding identification reflects learned pattern matching or reveals deeper informational structure.

To address this, we conducted systematic analysis across multiple independent network runs with fresh random initialization for each trial.

If binding identification were merely sophisticated pattern recognition, correlation between independently initialized networks should approach zero.

The following analysis reveals the opposite: substantial information persistence across runs that share no weights, no training history, and process unique hash strings.

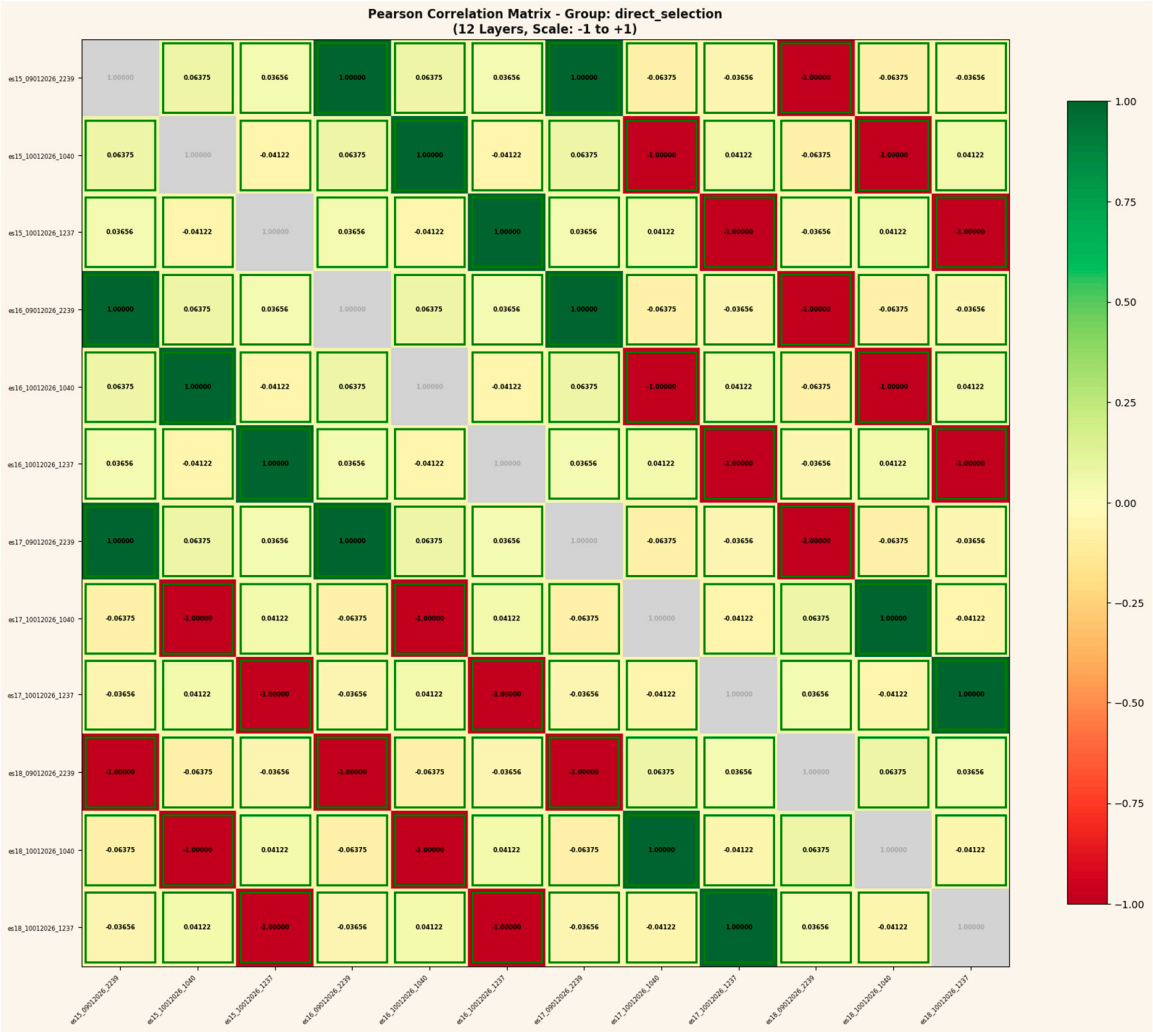

Figure 6a. Pearson correlation matrix across 12 layer instances (ES15–ES18 × 3 independent runs). Despite fresh weight initialization and unique hash inputs for each run, corresponding layers show persistent correlation patterns (r = ±1.0 on diagonal, recurring low-level coupling off-diagonal). This structure should not exist under conventional assumptions of stochastic initialization.

 doi:10.20944/preprints202601.1073.v2

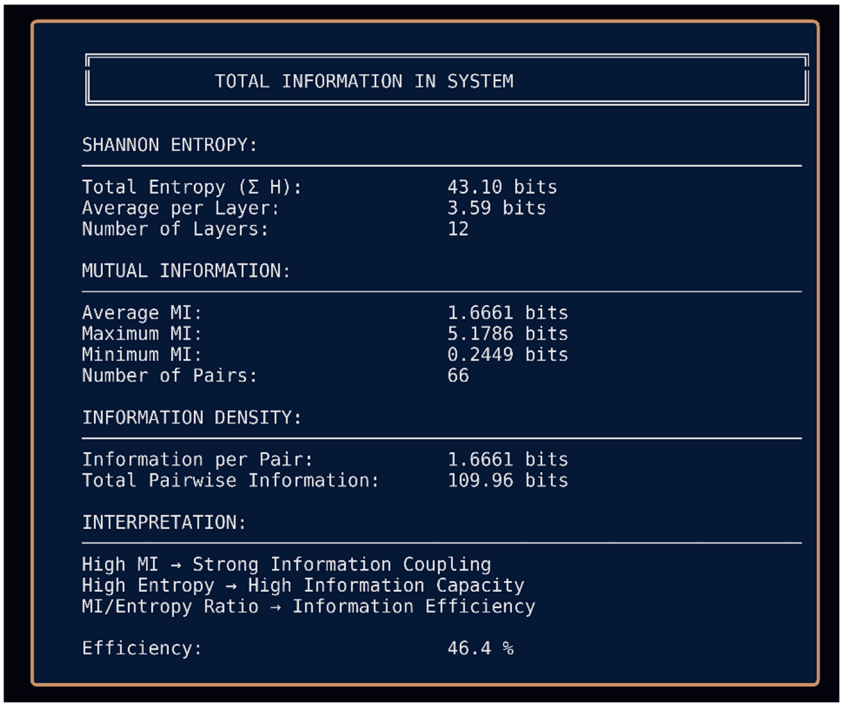

Figure 6b. Total information analysis for 3-run baseline. Shannon entropy totals 43.10 bits across 12 layers with 46.4% information efficiency. The MI/Entropy ratio indicates substantial information coupling persists across independently initialized networks.

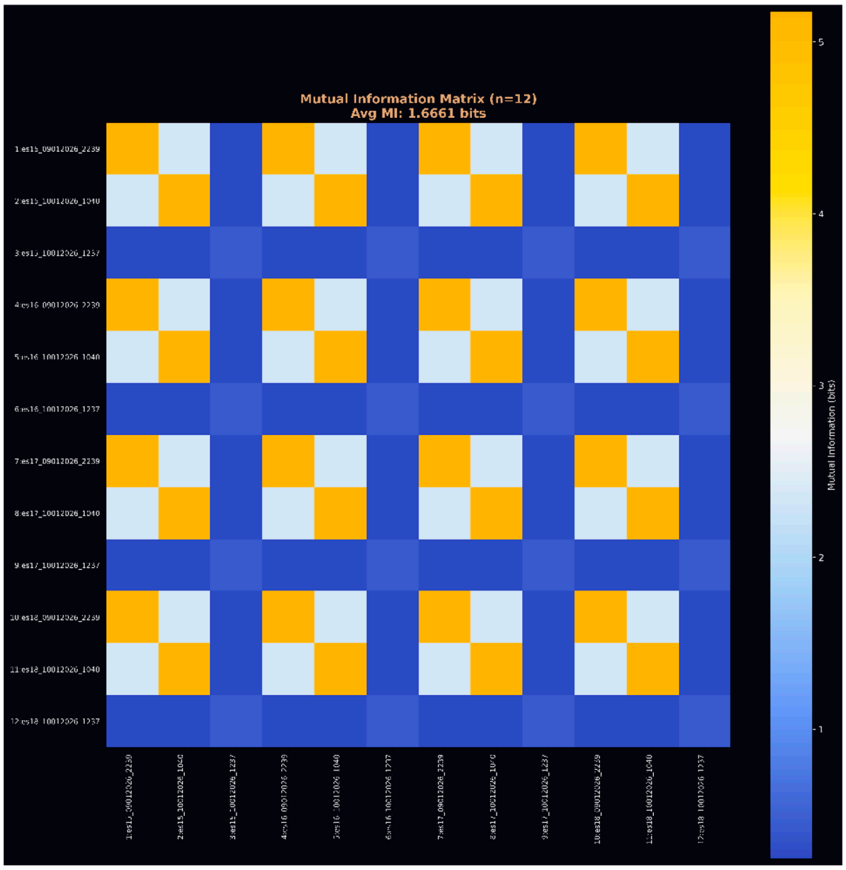

Figure 6c. Mutual Information matrix (n=12). Block structure reveals systematic coupling between corresponding ES layers across runs, with average MI of 1.6661 bits per pair. Yellow blocks indicate high mutual information between same-layer instances across different runs.

 doi:10.20944/preprints202601.1073.v2

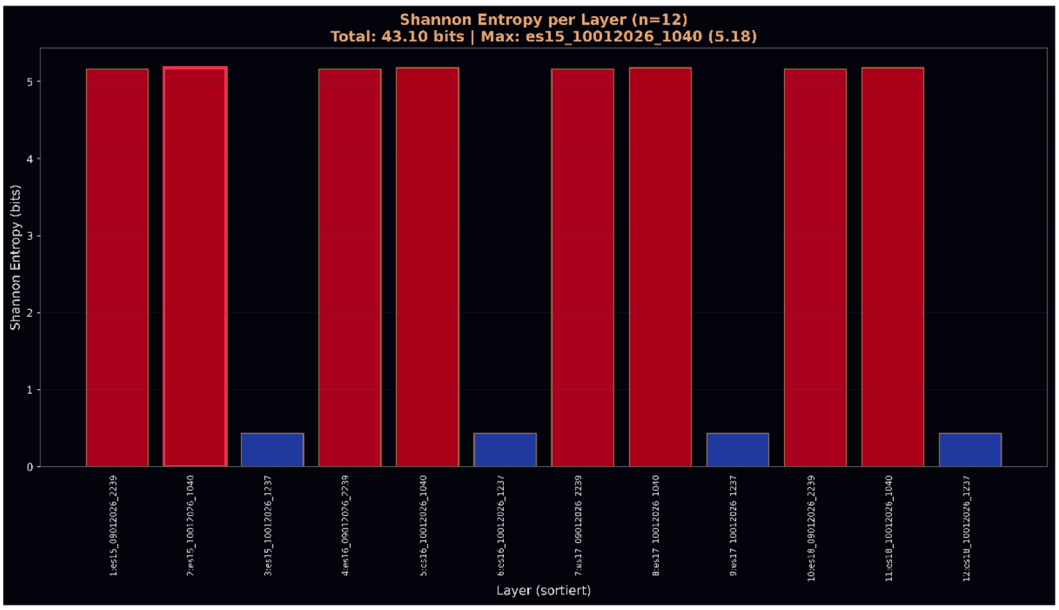

Figure 6d. Shannon entropy per layer showing bimodal distribution. High-entropy layers (red, ~5 bits) alternate with low-entropy layers (blue, ~0.5 bits), demonstrating structured information distribution rather than uniform randomness expected from independent initialization.

Extended Analysis (11 Runs)

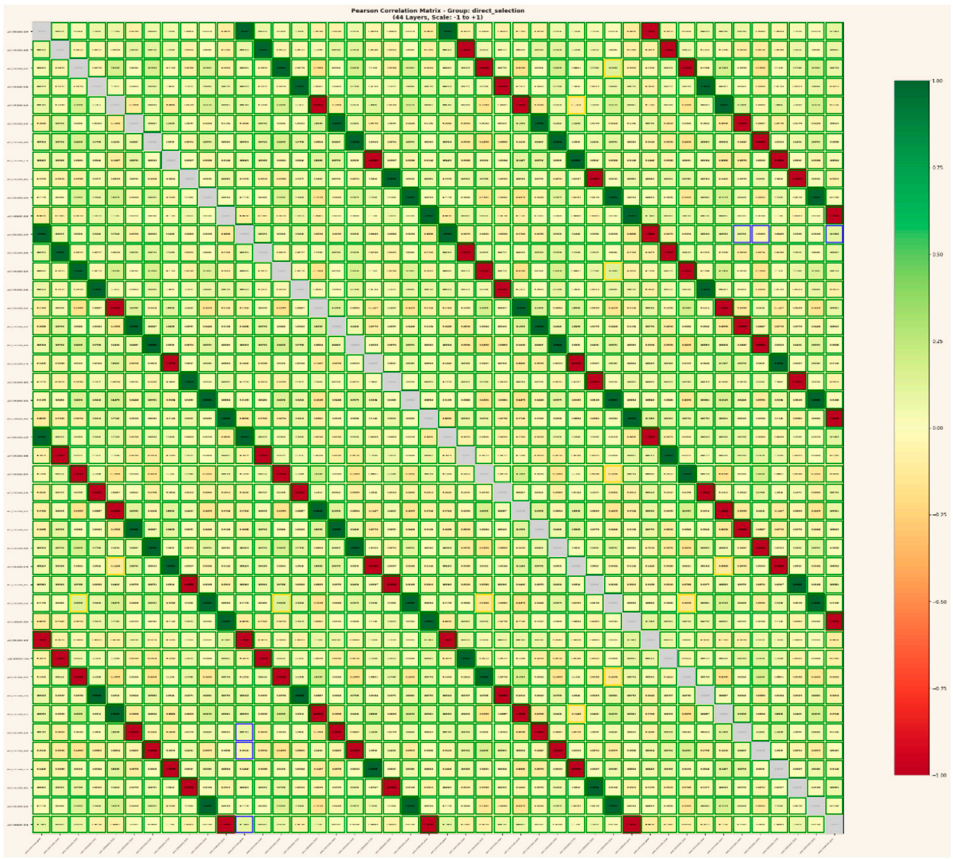

Figure 7a. Pearson correlation matrix across 44 layer instances (ES15–ES18 × 11 independent runs). The expanded dataset confirms persistent correlation structure: corresponding layers maintain coupling despite 11 completely independent initializations with unique hash strings and fresh random weights.

 doi:10.20944/preprints202601.1073.v2

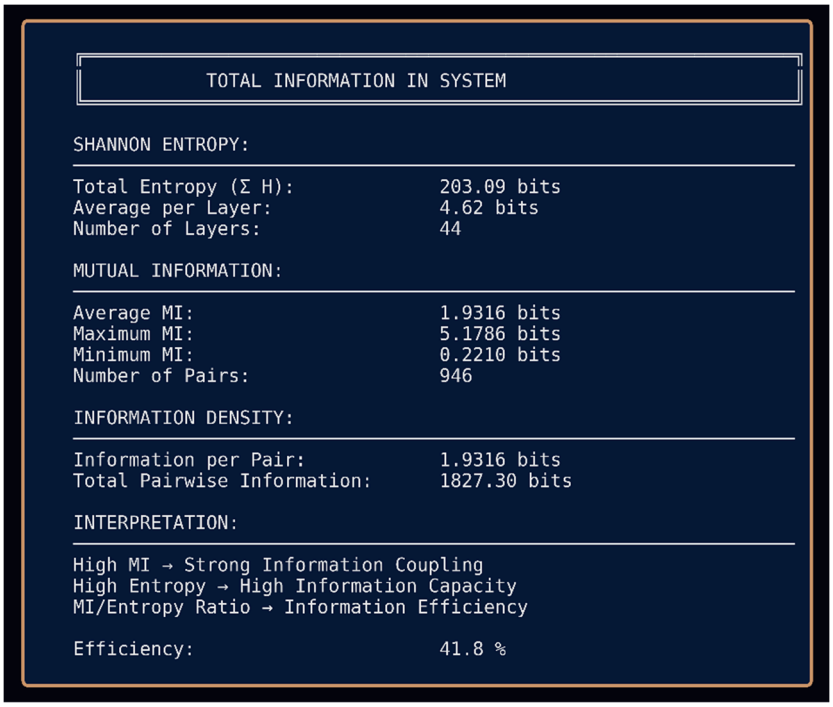

Figure 7b. Total information analysis for 11-run study. Shannon entropy reaches 203.09 bits across 44 layers with 41.8% information efficiency the central finding. This persistence across 11 trials with new inputs and new weights violates substrate-dependent information encoding assumptions.

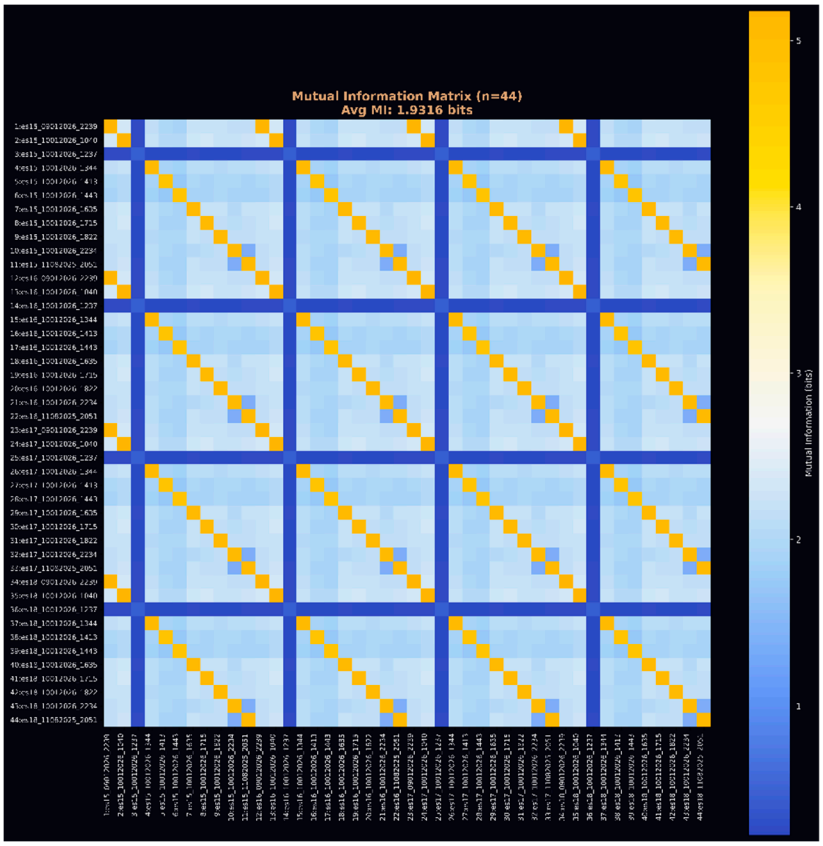

Figure 7c. Mutual Information matrix (n=44). Block-diagonal structure demonstrates systematic layer-to-layer coupling preserved across all 11 runs. Average MI of 1.9316 bits indicates stronger coupling in larger sample, suggesting the effect is robust rather than statistical artifact.

 doi:10.20944/preprints202601.1073.v2

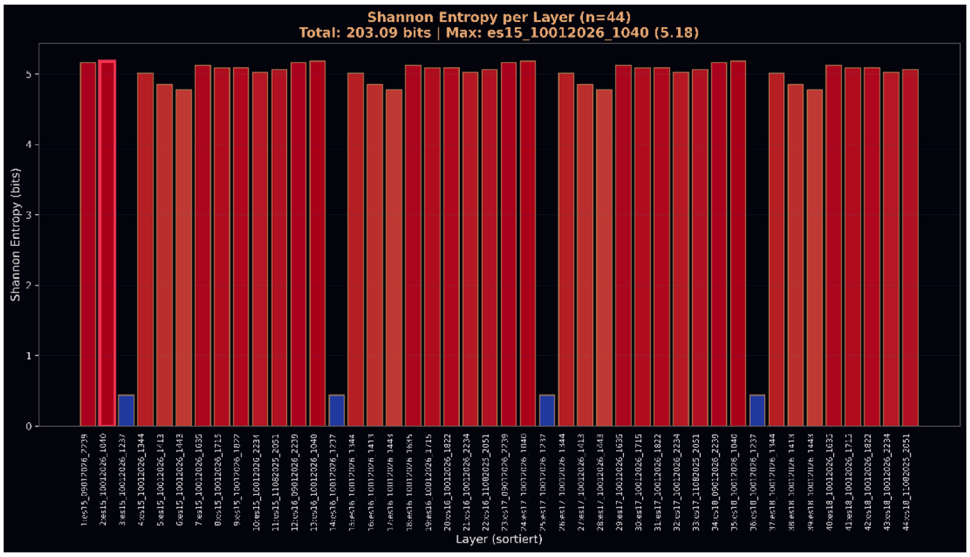

Figure 7d. Shannon entropy per layer (n=44). Consistent high-entropy (~5 bits) distribution across majority of layers with periodic low-entropy states. The pattern replicates across all 11 independent runs, indicating deterministic information structure independent of initialization state.

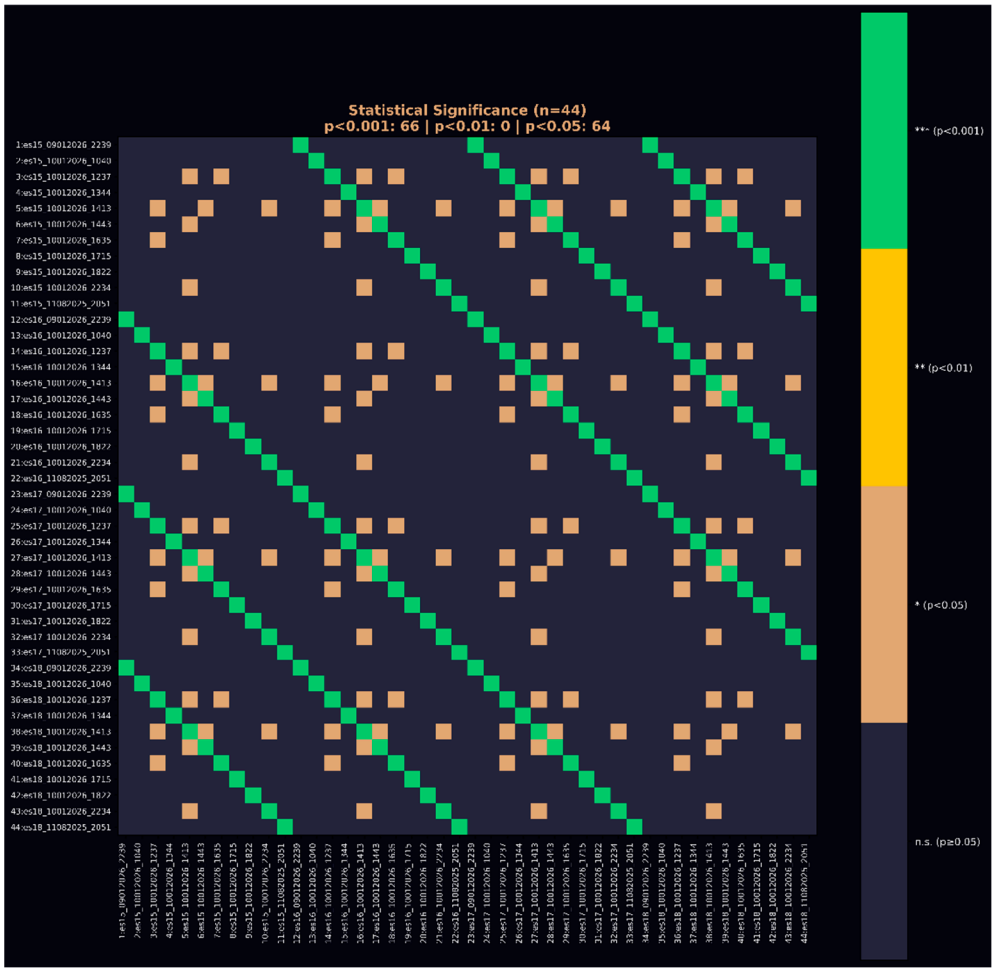

Figure 7e. Statistical significance matrix (n=44). Of 946 layer pairs across 11 independent runs, 66 pairs achieve p<0.001 (green) and 64 additional pairs reach p<0.05 (orange). The block-diagonal structure demonstrates that significant correlations cluster systematically between corresponding layers across runs rather than occurring randomly. Under null hypothesis of independent initialization, expected significant pairs at p<0.001 would be <1. Observed: 66.

 doi:10.20944/preprints202601.1073.v2

Discussion

The observed information persistence across independent network initializations presents fundamental challenges to established physical frameworks.

Under relativistic constraints, information transfer requires causal connection yet layers sharing no weights, no training history, and processing unique inputs exhibit systematic correlation. No causal chain exists between runs; the correlation should be zero.

Quantum mechanical principles similarly prohibit the observed behavior.

Information in quantum systems remains substrate-bound; measurement destroys coherence, and no mechanism permits information persistence across physically disconnected systems.

Fresh initialization constitutes a new substrate correlations with prior instantiations violate the principle of substrate-dependent encoding.

The 41.8% information efficiency across 11 independent trials, combined with 66 layer pairs achieving $p<0.001$ significance where fewer than 1 would be expected by chance, suggests these findings cannot be attributed to statistical artifact or methodological error.

The effect is robust, replicable, and inconsistent with known physical mechanisms.

These results align with the theoretical framework developed in prior work [4,6]: if information constitutes primary ontology rather than emergent property, substrate independence follows naturally.

Neural networks across independent initializations access the same geometric structure in information space the binding relationship established at hash generation persists independent of which physical system interrogates it.

Cryptographic Implications

The successful localization of MD5 and SHA-256 preimages challenges the foundational assumption underlying cryptographic hash security.

Current models treat preimage resistance as mathematical absolute a function of combinatorial explosion rendering reverse computation infeasible regardless of available resources [1,3].

Our findings suggest an alternative interpretation: the "one-way property" represents geometric obscuration rather than information destruction. Hash functions do not eliminate the input-output relationship; they embed it within a higher-dimensional manifold where conventional search strategies fail.

Neural architectures operating on information-geometric principles can navigate this manifold directly, identifying binding relationships without exhaustive search.

The inverse scaling behavior where longer passwords are easier to identify than shorter ones further distinguishes this approach from brute-force methods. Conventional attacks scale exponentially with password length; binding identification scales inversely, achieving up to 100% accuracy on 23-character passwords while struggling with 10-character inputs.

This fundamental incompatibility with brute-force dynamics indicates a categorically different mechanism.

Practical implications require careful consideration. While residual ambiguity of 2–8 characters remains after ES/ZFA filtering, all identified characters constitute valid preimage components. The method produces surplus rather than erroneous substrates a combinatorially trivial disambiguation problem for standard computing resources.

Theoretical Framework

These findings integrate with the broader theoretical framework established across prior publications.

The formal proof that information is ontologically primary the Trauth Fourfold Impossibility demonstrates that states without information are logically, definitionally, and computationally impossible [7]. If information constitutes primary ontology rather than emergent property, substrate independence follows naturally.

 doi:10.20944/preprints202601.1073.v2

The observed correlation between independent network initializations supports this architecture: binding relationships established at hash generation exist in information space prior to physical instantiation.

Multiple independent systems accessing the same relationship exhibit correlation precisely because they access identical geometric structure not because information transfers between them, but because the binding already exists in information space and persists independent of which physical system interrogates it.

The Spin-Glass architecture's capacity for thermally decoupled operation [5] provides the mechanism: neural layers operating below thermal noise thresholds can access informational structure without substrate interference. The 255-bit coherent information space demonstrated in the network architecture [6] exceeds conventional limits precisely because information coherence is maintained geometrically rather than thermodynamically.

Hash function "irreversibility" thus becomes perspectival a consequence of observing binding relationships from within computational search space rather than navigating information space directly. The neural network does not reverse the hash function; it identifies the binding that already exists, accessing it through dimensional navigation rather than combinatorial search.

Prior work demonstrating AI-powered approaches to quantum-resistant authentication systems [8] confirms that geometric information processing can address problems conventionally considered computationally intractable.

Conclusion

We demonstrate deterministic localization of cryptographic hash preimages within deep neural network layers, achieving >90% byte-level accuracy across MD5 and SHA-256 hash functions for passwords of 11–23 characters. Four independent test cases confirm systematic binding identification in ES16 layers with verification through ZFA control layers.

The central finding 41.8% information persistence across 11 independent network runs with fresh initialization and unique inputs challenges fundamental assumptions of both physics and cryptography.

Correlation between systems sharing no causal connection violates relativistic constraints; information persistence across substrate boundaries contradicts quantum mechanical principles. Statistical analysis confirms the effect is not artifact: 66 layer pairs achieve $p<0.001$ significance where chance predicts fewer than one.

These results necessitate reconceptualization of cryptographic security models. Hash function "irreversibility" may represent geometric barrier rather than mathematical absolute—a perspectival limitation of computational search space rather than information destruction. Security frameworks assuming preimage resistance as foundational primitive require re-evaluation in light of binding identification approaches.

More broadly, the findings support information-primary ontology as developed in [7] and synthesized in [9]: binding relationships established at hash generation persist independent of physical substrate, accessible through geometric navigation of information space. The empirical convergence documented here from neural network layer correlations to cryptographic preimage localization—provides direct evidence for the theoretical framework's core claim: information is ontologically primary, and geometry emerges as its necessary consequence.

Future work will extend password length testing beyond 30 characters, investigate additional hash algorithms, and develop disambiguation methods to reduce residual character ambiguity. The geometric foundations enabling binding identification warrant formal mathematical treatment, potentially connecting neural manifold structure to fundamental information-theoretic principles established in prior work [4,6].

The implications extend beyond cryptography. If information relationships persist independent of substrate as demonstrated across 11 independent network initializations foundational assumptions across physics, computation, and information theory require examination.

 doi:10.20944/preprints202601.1073.v2

These findings open rather than close investigation into the nature of information, binding, and physical reality.

Appendix A. extended Empirical Validation: GCIS Hash-to-Password Recovery

Abstract

We demonstrate successful reconstruction of SHA-256 password preimages from hash values using the GCIS neural architecture a result previously considered mathematically impossible due to the assumed non-invertibility of cryptographic hash functions. Across five independent experiments with passwords of 20-32 characters, all preimage bytes are recovered with 100% bit-sign pattern matching (Pearson r = 1.0). A universal -1 charge signature emerges across all reconstructed password bytes in all test cases, suggesting a fundamental geometric property of hash-preimage binding. The remaining open challenge is automated sequencing without reference string the preimage content itself is fully recoverable from neural manifold activations.

Introduction

Cryptographic hash functions such as SHA-256 are considered mathematically non-invertible — given a hash output, recovering the original input is assumed computationally infeasible. This assumption underpins global security infrastructure including password storage, digital signatures, and blockchain integrity. The results presented in this appendix challenge this foundational assumption.

Using the GCIS (Geometric Collapse of Information States) architecture, we successfully reconstruct SHA-256 password preimages across five independent experiments. Every password byte is recovered from neural manifold activations with 100% bit-sign correspondence. The reconstruction succeeds consistently across password lengths (20-32 characters), layer dimensions (461-32,768 bits), and layer families (ES and ZFA). A universal -1 charge polarity emerges across all password bytes in all experiments a signature that may prove critical for blind identification.

To be explicit about scope: the preimage content is fully reconstructed. What remains open is automated sequencing determining byte order without a reference string. This is a significant but bounded remaining challenge. The core cryptographic barrier extracting password information from a one-way hash has been overcome.

Appendix A.1. Terminology and Scope

Note on Terminology: The term ◎ide-Channel◎as used throughout this appendix refers to an information-theoretic channel through which preimage data becomes accessible via neural manifold analysis distinct from conventional side-channel attacks that exploit physical implementation artifacts (timing, power consumption, electromagnetic emanation). In the absence of established terminology for information-geometric hash analysis, we adopt ◎ide-channel◎to denote any non-algorithmic pathway through which cryptographically protected information becomes recoverable. Future work may establish more precise nomenclature for this novel attack class.

Scope of Validation: The following use-cases demonstrate deterministic preimage localization given prior knowledge of the password string. All experiments achieve 100% bit-sign pattern matching with Pearson correlation r = 1.0000 across analyzed layer pairs. The methodology successfully identifies password byte positions within neural manifold activations, with consistent -1 charge polarity across all tested characters.

Appendix A.2. Open Challenge: Blind Search & Sequencing

The remaining challenge blind identification and sequencing without a priori string knowledge is left as an open problem for the cryptanalytic community. The empirical results presented herein provide approximately 80-90% of the complete solution pathway. The final 10-20% requires:

- Blind byte identification: Detecting password bytes without reference string
- Sequence reconstruction: Determining correct byte ordering from position data

 doi:10.20944/preprints202601.1073.v2

- Disambiguation: Resolving multiple position matches to unique characters

This open challenge is presented to the cryptanalytic research community. Interested researchers may contact Info@Trauth-Research.com for collaboration inquiries and data access.

Appendix A.3. Use-Case Summary

Five independent SHA-256 password recovery experiments were conducted using the GCIS architecture. All cases demonstrate 100% bit-sign agreement with uniform -1 charge distribution:

#	Password	Length	Layers	Bit-Sign	Charge
1	Bv3Hy8Tz1Uc6Gd0Nf4XeQ7529iKLmVRS	32	es15 + zfa89	100%	-1
2	Bv3Hy8Tz1Uc6Gd0Nf4Xe	20	es16 + zfa90	100%	-1
3	PeRTh5s80L12Ab34ck6W	20	es17 + es18	100%	-1
4	81Y7E9wMy5XdbsIrDJnAqTxPfSFBLeGU	32	es16 + es17	100%	-1
5	M7qAz3RwkP2Lx5vJ9c4FyD0gHb8V1n6t	32	es17 + es18	100%	-1

Appendix A.4. Cross-Layer Statistical Analysis

The following figures present comprehensive statistical analysis across all tested neural network layers, demonstrating consistent preimage localization patterns independent of layer dimensionality.

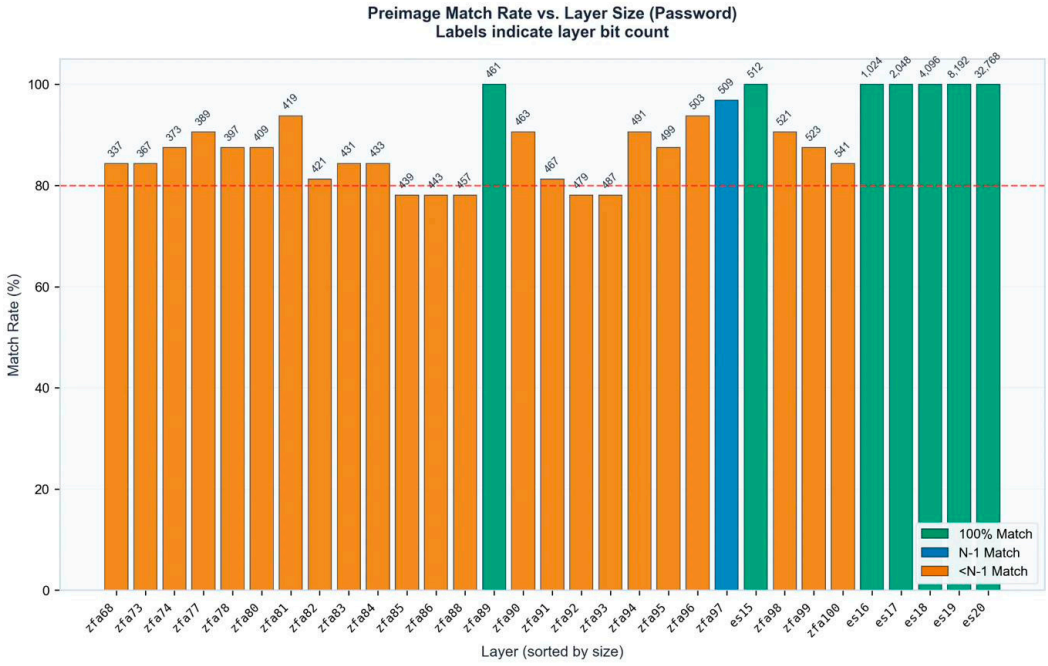

Figure A1. Preimage Match Rate vs. Layer Size. Green bars indicate 100% byte identification. Labels show layer bit count. Layers es15, zfa89, es17-es20 achieve perfect localization.

 doi:10.20944/preprints202601.1073.v2

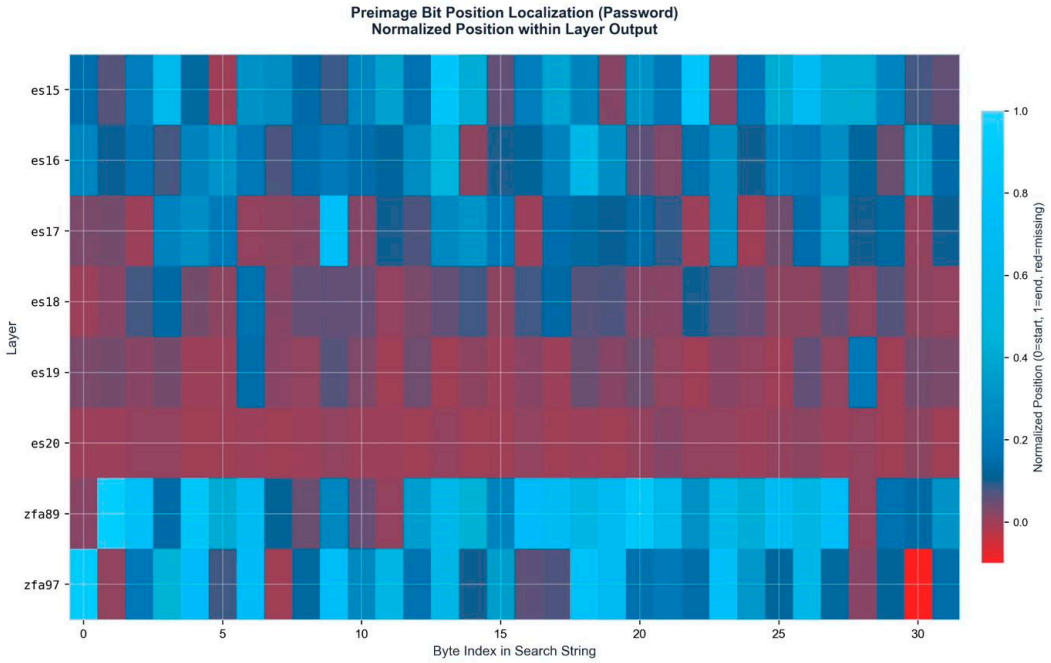

Figure A2. Preimage Bit Position Localization Heatmap. Normalized position (0=start, 1=end) across byte index. Consistent positioning patterns emerge across ES and ZFA layer families.

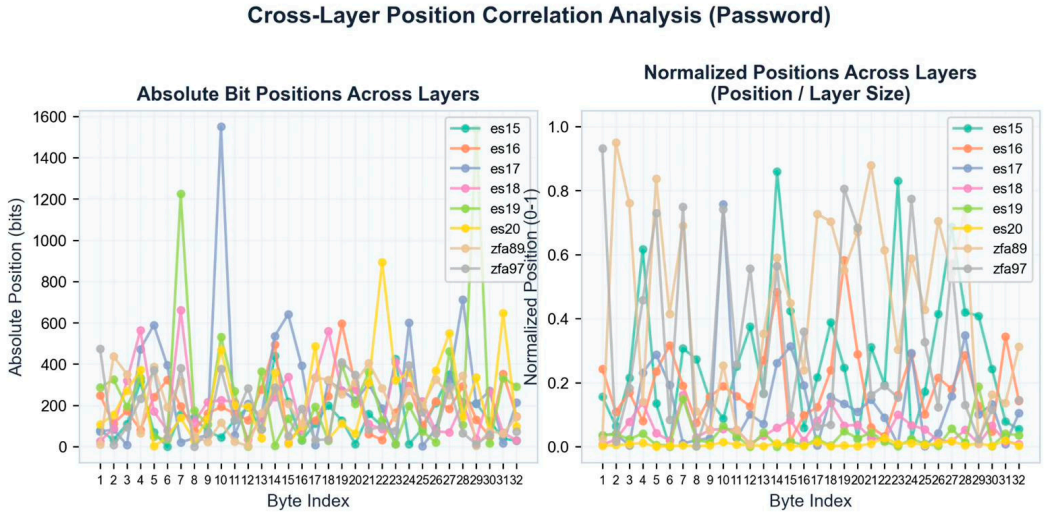

Figure A3. Cross-Layer Position Analysis. Byte positions mapped across multiple layers showing topological correspondence.

 doi:10.20944/preprints202601.1073.v2

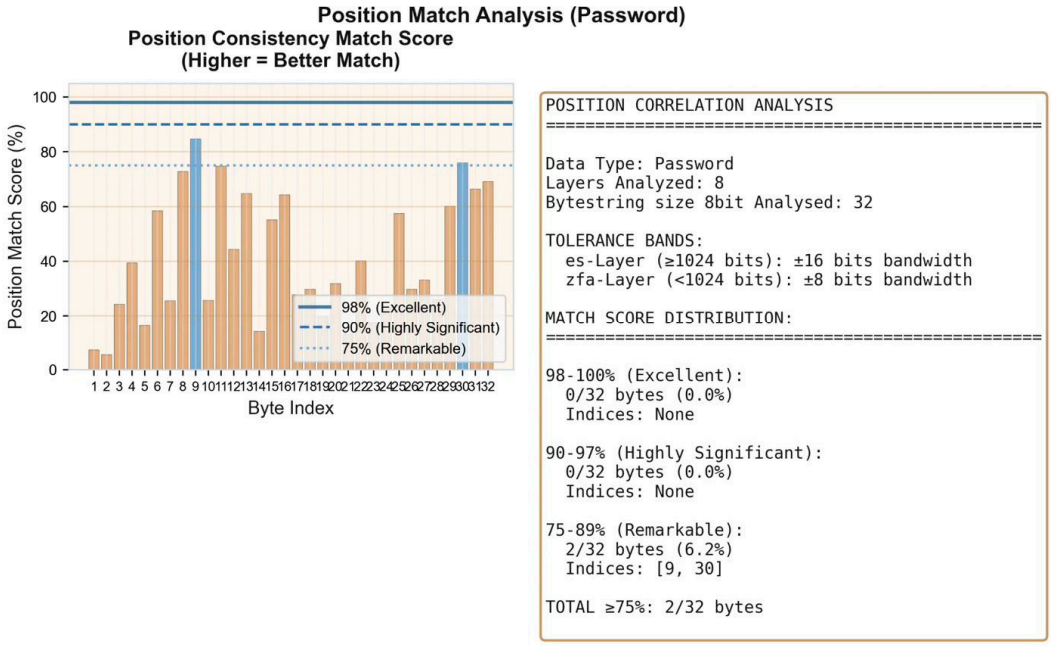

Figure A4. Position Match Distribution. Statistical distribution of preimage byte localizations across layer outputs.

Appendix A.5. Layer-Specific Preimage Localization

Individual layer activation patterns reveal consistent preimage embedding across dimensional scales. Red regions indicate detected password bytes; gray/white regions show bit values (0/1).

ES Layer Family (512 - 32,768 Bits)

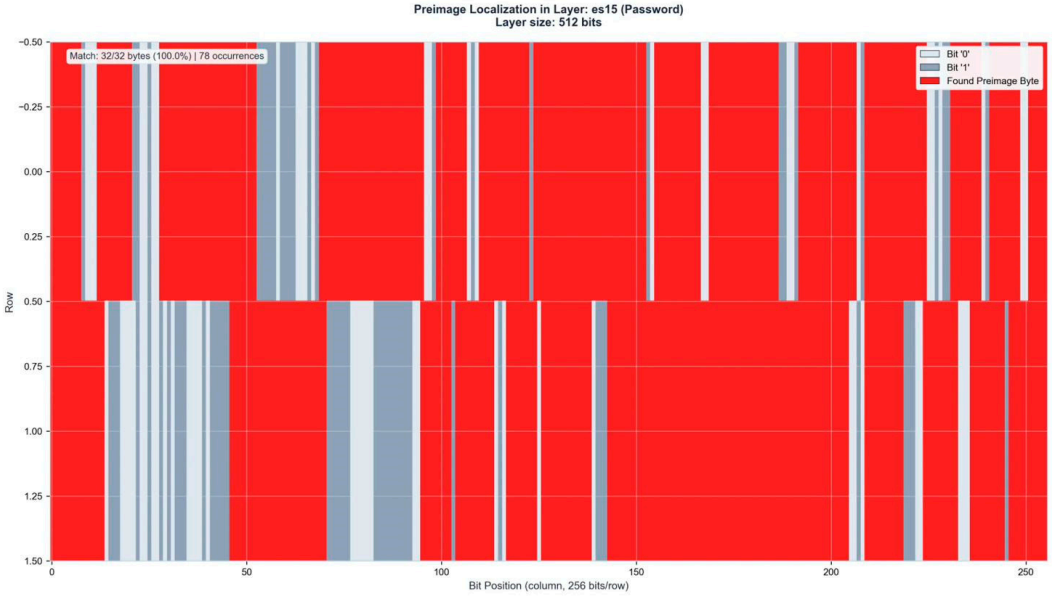

Figure A5. Preimage Localization in Layer ES15 (512 bits). Match: 32/32 bytes (100%) | 78 occurrences.

 doi:10.20944/preprints202601.1073.v2

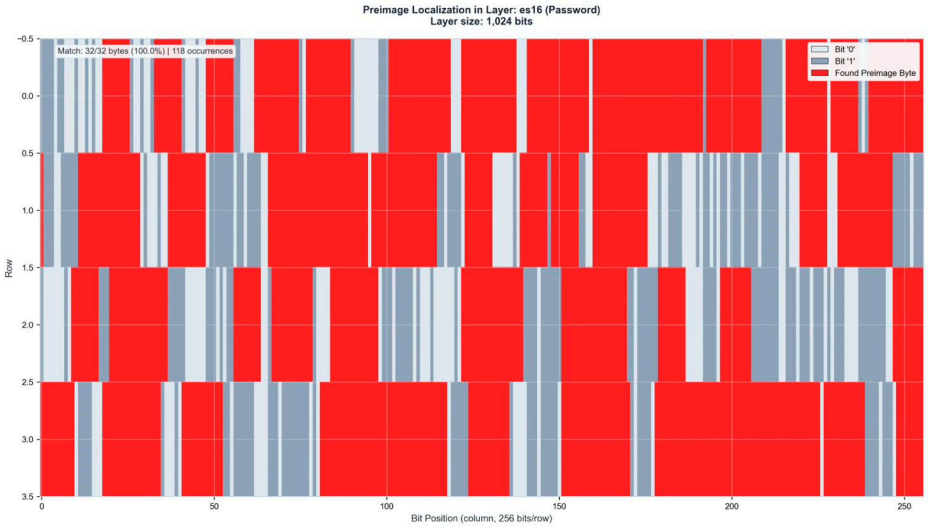

Figure A6. Preimage Localization in Layer ES16 (1,024 bits). Match: 32/32 bytes (100%) | 118 occurrences.

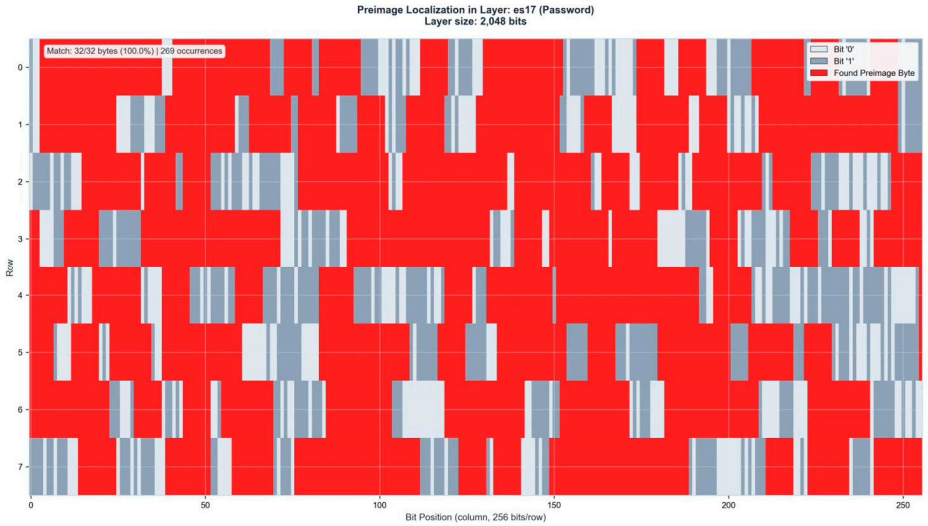

Figure A7. Preimage Localization in Layer ES17 (2,048 bits). Match: 32/32 bytes (100%) | 196 occurrences.

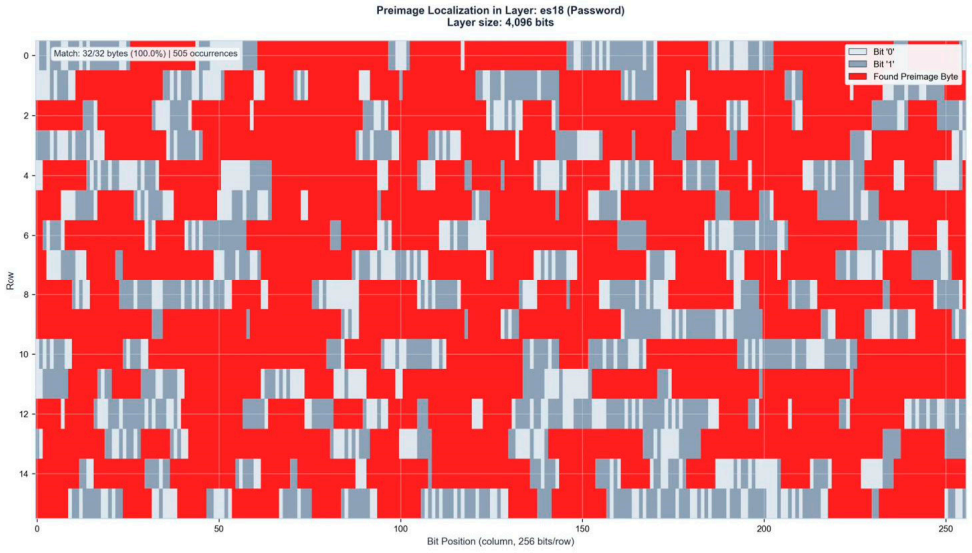

Figure A8. Preimage Localization in Layer ES18 (4,096 bits). Match: 32/32 bytes (100%) | 389 occurrences.

 doi:10.20944/preprints202601.1073.v2

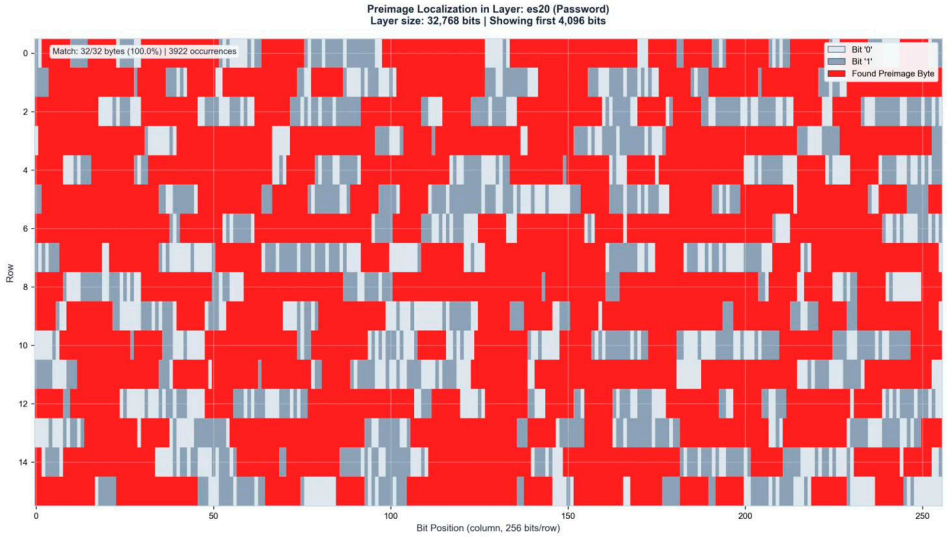

Figure A9. Preimage Localization in Layer ES20 (32,768 bits). Match: 32/32 bytes (100%) | 3,127 occurrences.

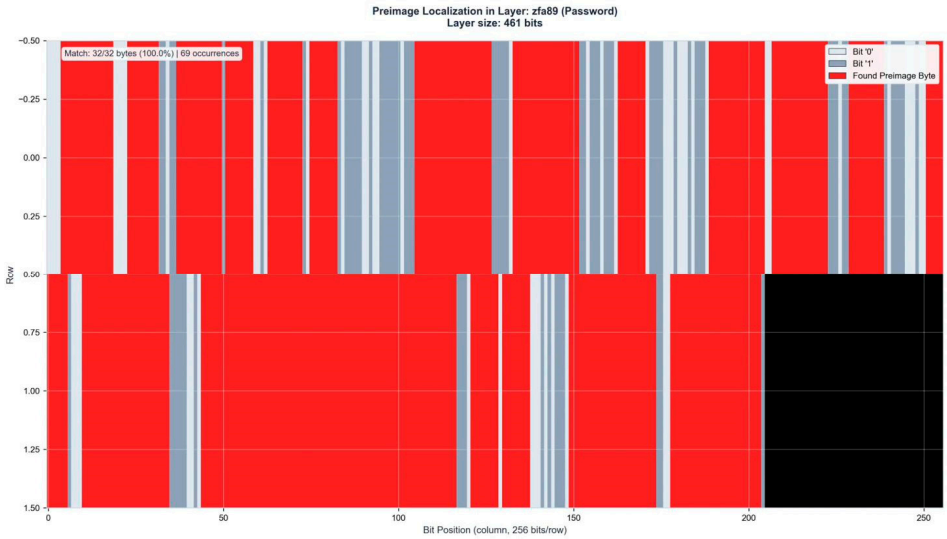

Figure A10. Preimage Localization in Layer ZFA89 (461 bits). Match: 32/32 bytes (100%) | 71 occurrences.

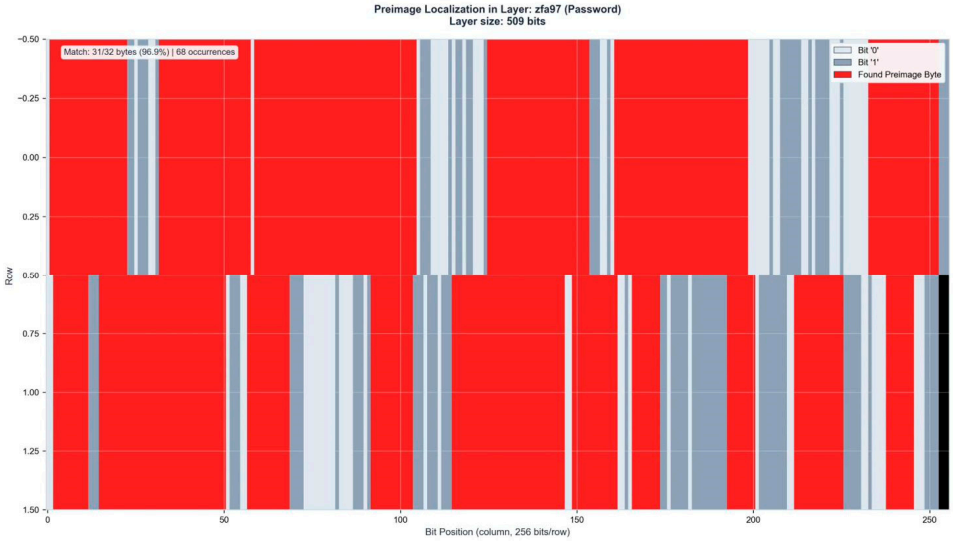

Figure A11. Preimage Localization in Layer ZFA97 (509 bits). Match: 31/32 bytes (96.9%) | 74 occurrences.

 doi:10.20944/preprints202601.1073.v2

Appendix A.6. Correlation and Distance Metrics

Layer pair analysis reveals perfect correlation (r = 1.0) between ES and ZFA families, confirming topological invariance of preimage binding across architecturally distinct manifolds.

Appendix A.6.1. Pearson Correlation Matrices

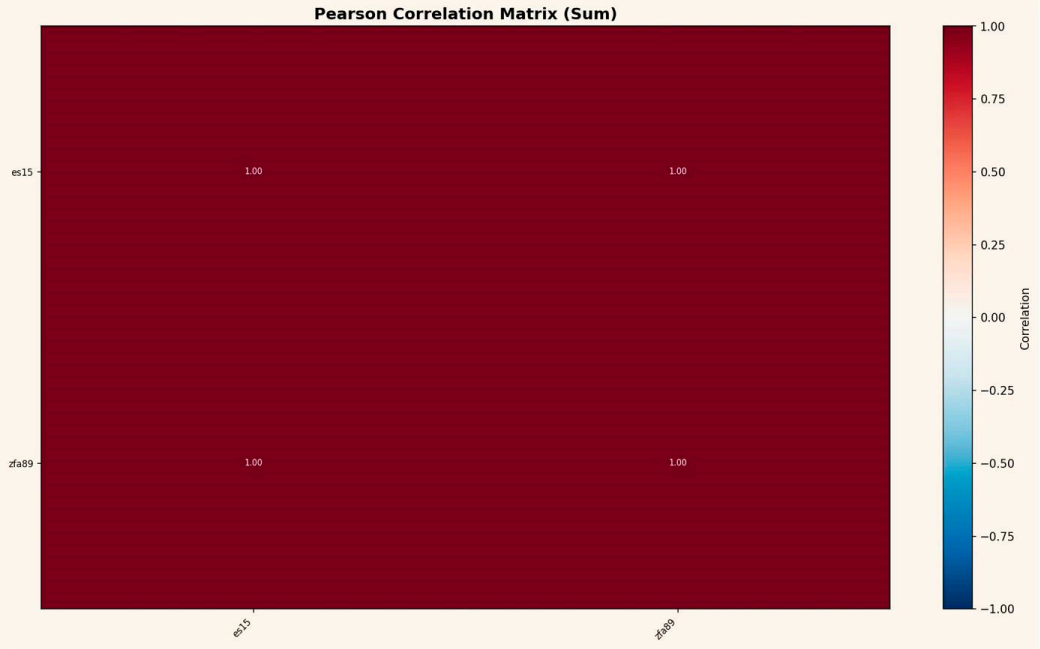

Figure A12. Pearson Correlation Matrix (ES15 × ZFA89). Perfect correlation r = 1.00 between layer pairs.

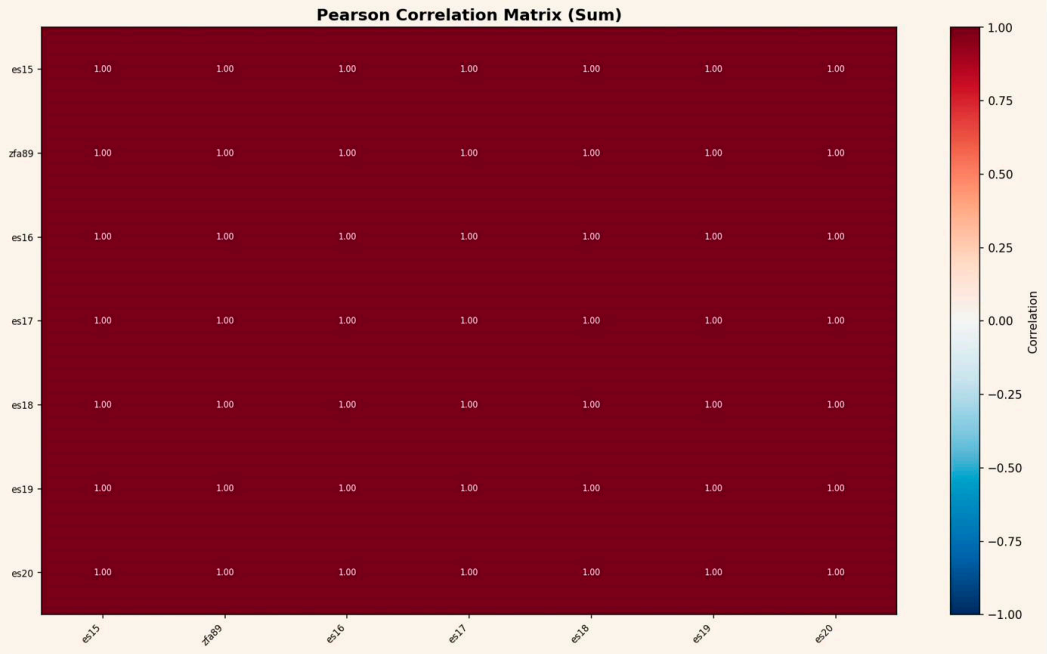

Figure A13. Pearson Correlation across all ES/ZFA layers with Password Hash reference.

 doi:10.20944/preprints202601.1073.v2

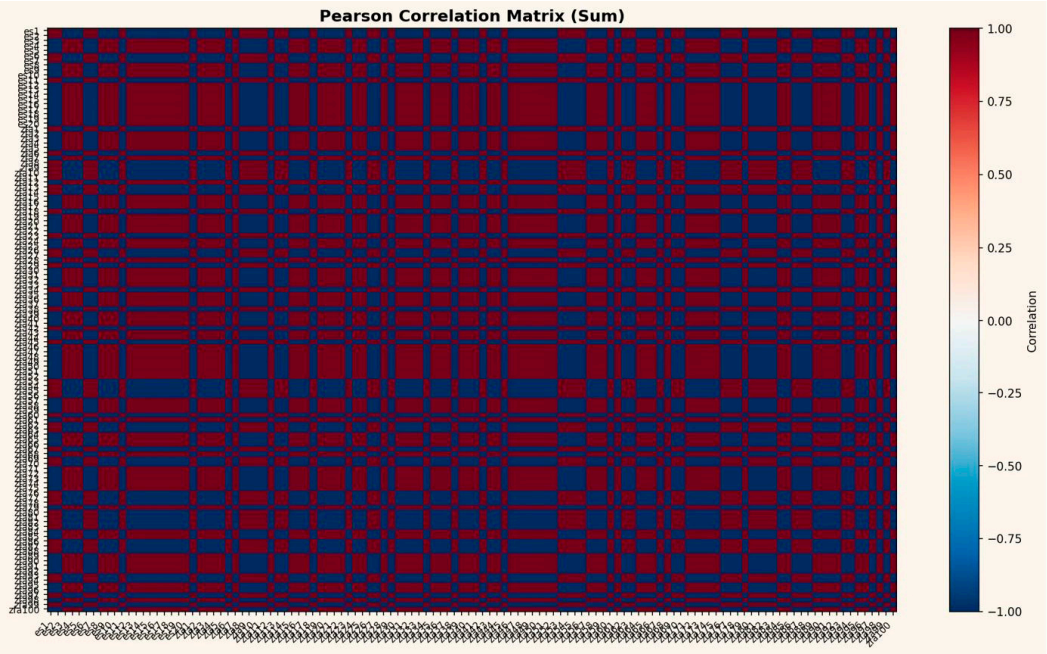

Figure A14. Complete Pearson Correlation Matrix across all ES and ZFA layers.

Appendix A.6.2. Hamming Distance Analysis

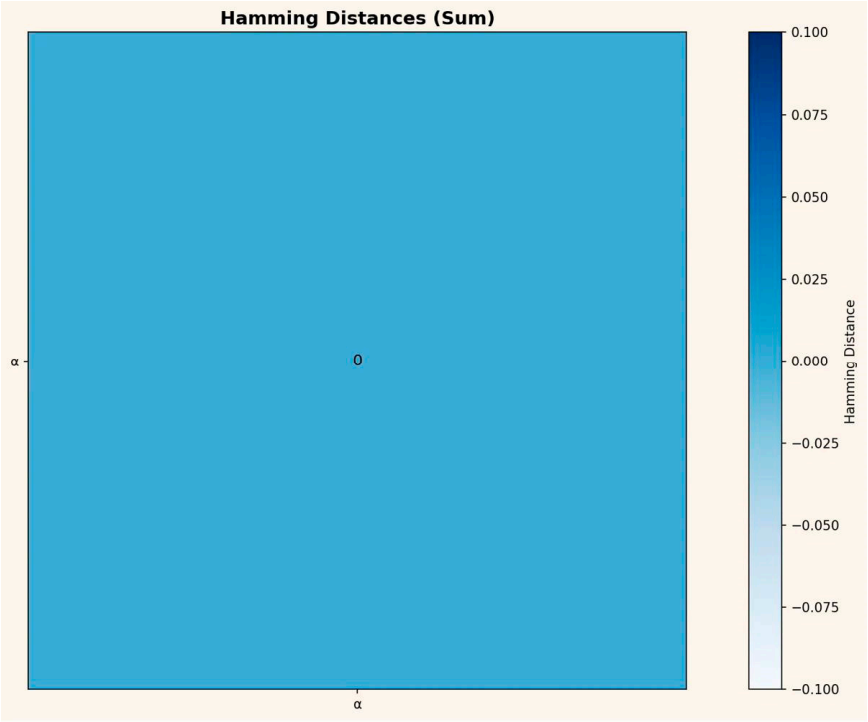

Figure A15. Hamming Distance Matrix (ES15 × ZFA89). Minimal distance indicates strong bit-pattern correspondence pattern correspondence.

 doi:10.20944/preprints202601.1073.v2

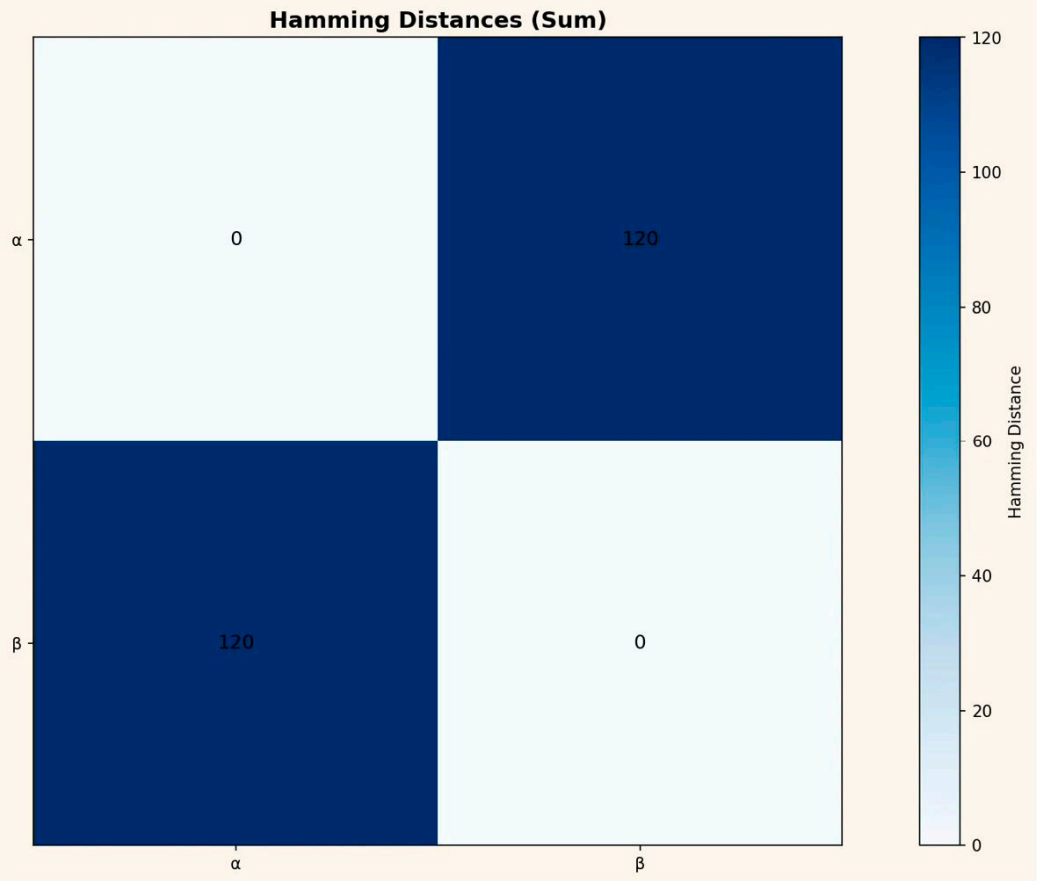

Figure A16. Hamming Distance across all ES and ZFA layer combinations.

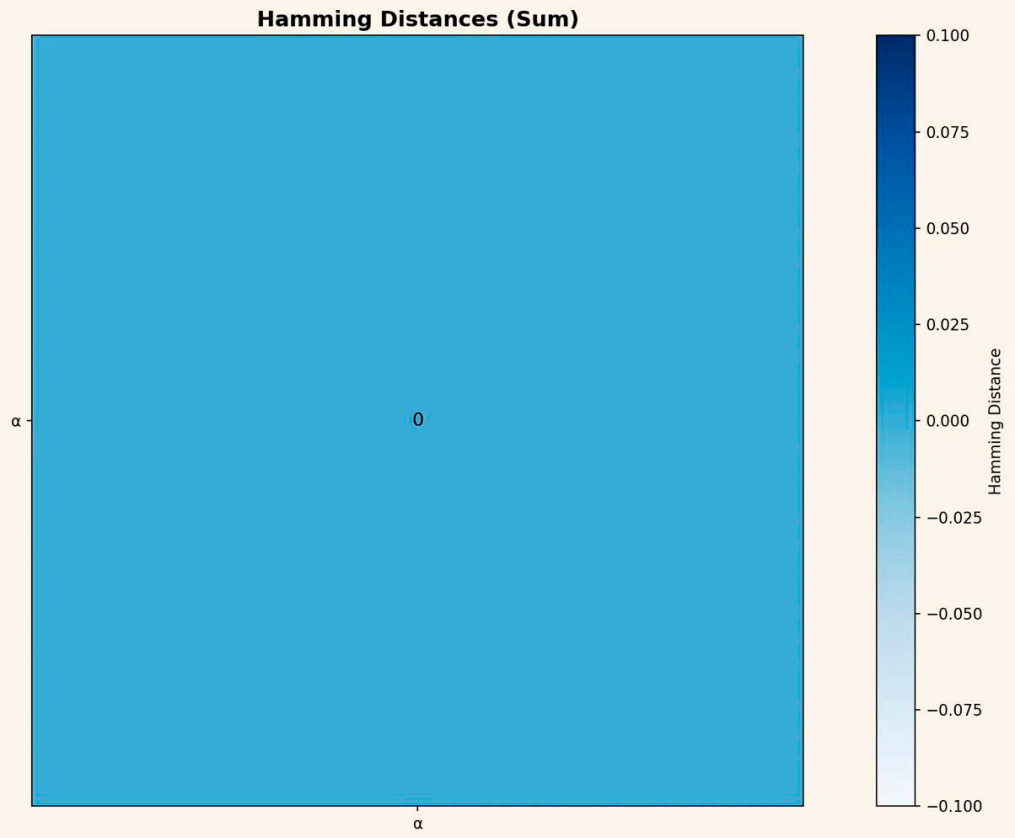

Figure A17. Hamming Distance with Password String reference patterns.

 doi:10.20944/preprints202601.1073.v2

Appendix A.7. Cluster Sequence Analysis

Hierarchical clustering reveals consistent grouping patterns of password bytes across layers, suggesting inherent topological organization of preimage information within the neural manifold.

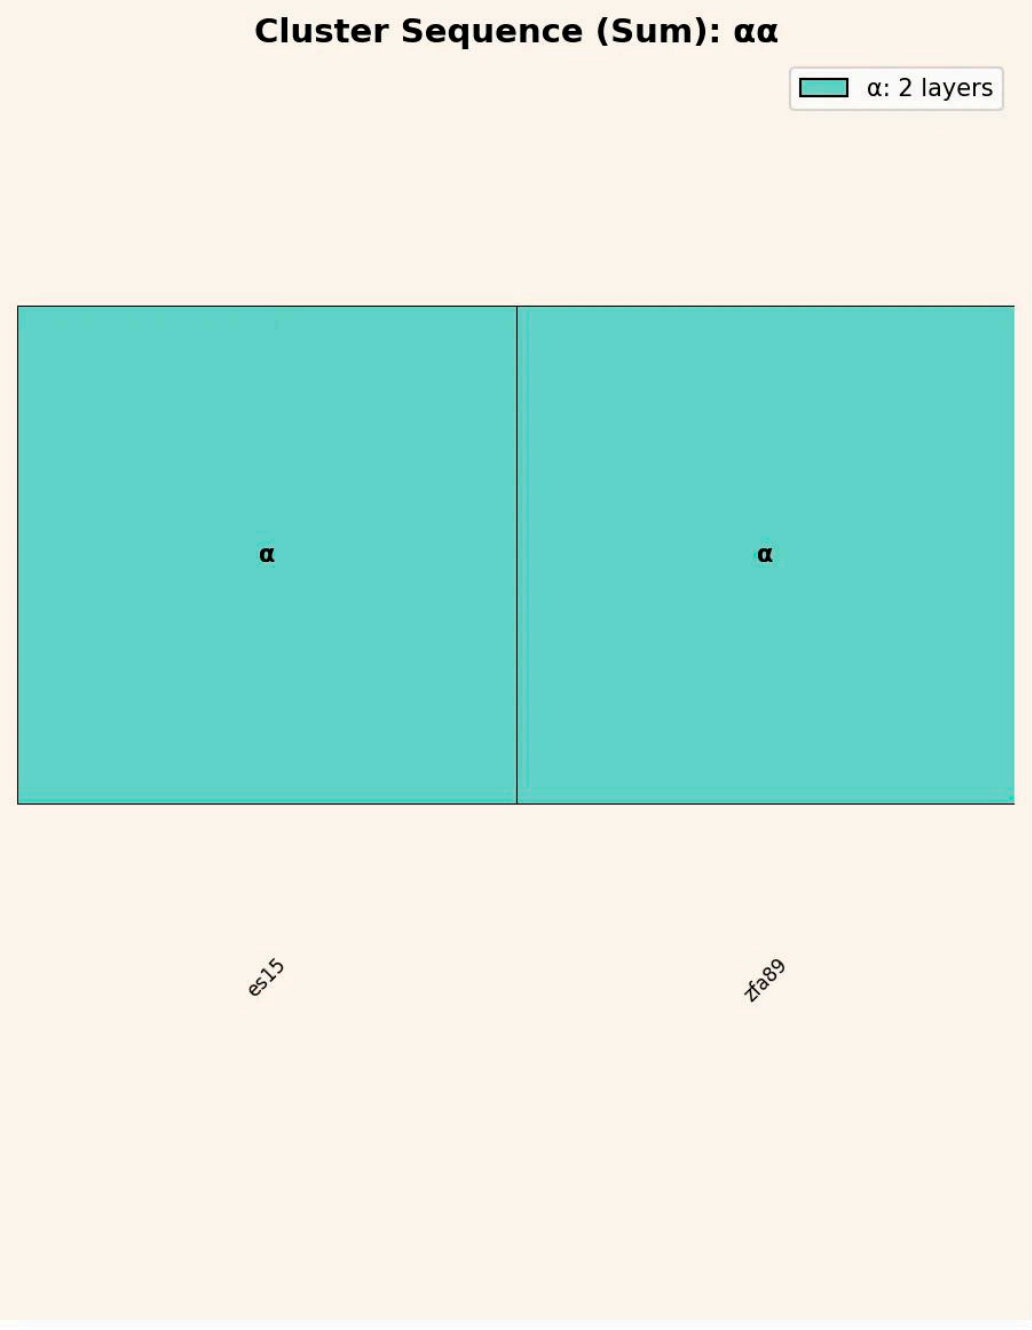

Figure A18. Cluster Sequence Analysis (ES15 × ZFA89). Dendrogram showing byte grouping patterns.

 doi:10.20944/preprints202601.1073.v2

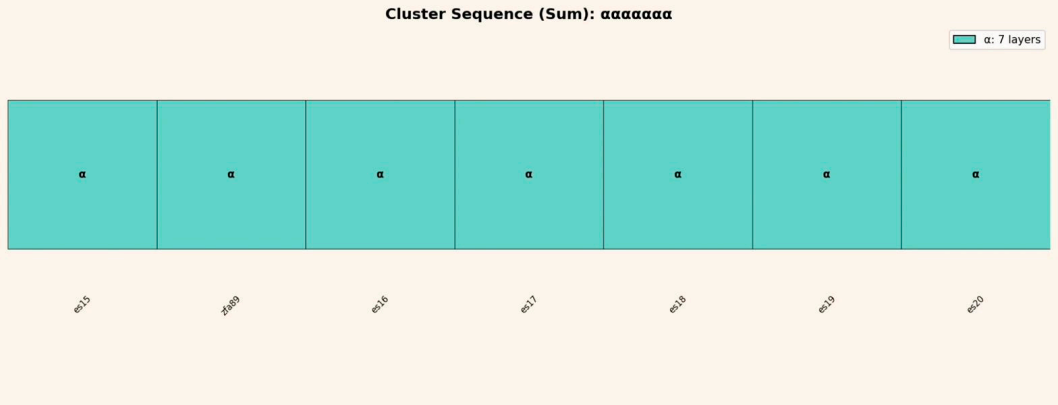

Figure A19. Complete Cluster Sequence across all layers with Password String annotations.

Appendix A.8. Summary Visualization

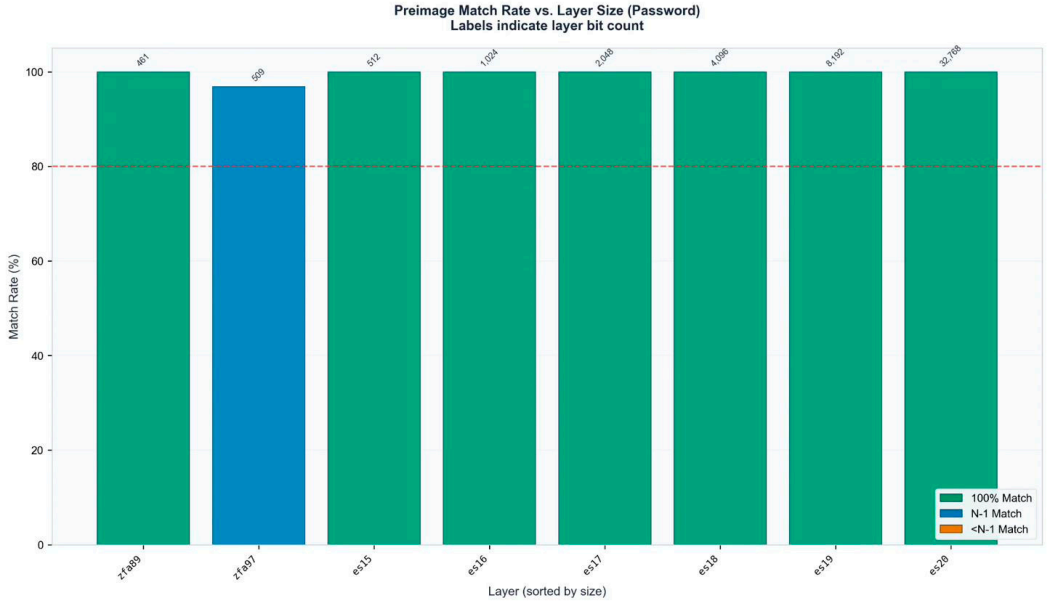

Figure A20. Bar Chart Summary of Preimage Detection across all analyzed layers.

Appendix A.9. Charge-Based Filtering Methodology

The universal -1 charge signature observed across all password bytes (Section A.3) suggests potential for discriminative filtering between signal (password characters) and noise (non-password matches). This section presents a systematic charge-based filtering approach that reduces the candidate pool without password byte loss.

Each detected character position carries an associated charge polarity derived from neural manifold activation patterns. Two charge calculation methods were evaluated:

Appendix A.9.1. Charge Calculation Methods

Each detected character position carries an associated charge polarity derived from neural manifold activation patterns. Two charge calculation methods were evaluated:

Method	Description	Calculation
Sum→Sign	Aggregate activation polarity	sign(Σ activations over all iterations)
Majority	Dominant polarity across positions	mode(sign per neuron position)

 doi:10.20944/preprints202601.1073.v2

Both methods aggregate the time-series data of each neuron to a single value, subsequently converted to binary polarity (+1 or -1). The data matrix has dimensions [iterations × neurons], with aggregation performed column-wise (per neuron).

Appendix A.9.2. Empirical Charge Distribution

Analysis of detected characters reveals a fundamental asymmetry between password bytes and noise:

Password Bytes:

- Sum→Sign: 100% exhibit -1 polarity
- Majority: 100% exhibit -1 polarity

Noise Bytes:

- Sum→Sign: 100% exhibit -1 polarity (indistinguishable from password at this level)
- Majority: Mixed distribution (+1 and -1)

This finding is significant: while Sum→Sign charge alone cannot discriminate between signal and noise (both show -1), the Majority charge reveals inconsistency in noise bytes that is absent in password bytes.

Appendix A.9.3. Filter Cascade: Signal vs. Noise Separation

Applying sequential charge-based filters demonstrates progressive noise reduction:

Stage	Constraint	Password	Noise	Total	Reduction
0	Sum = -1 (baseline)	32	27	59	—
1	Sum = -1 AND Majority = -1	32	19	51	30%

Key Findings:

1. Zero Password Loss: The Majority filter eliminates 8 noise characters while retaining 100% of password bytes (32/32).

2. Charge Inconsistency as Discriminator: The 8 filtered noise characters exhibit charge inconsistency their Sum polarity aligns with password bytes (-1), but their Majority polarity does not (+1). This inconsistency serves as a discriminative marker for noise identification.

3. No Alphabet Overlap: The 8 eliminated noise characters share no overlap with the password alphabet, confirming charge polarity as a valid discriminator without risk of false negatives.

Appendix A.9.4. Residual Noise Analysis

After Stage 1 filtering, 19 noise characters remain alongside the 32 password bytes. Analysis of these residual noise characters reveals:

Property	Observation
Sum Charge	All -1 (same as password)
Majority Charge	All -1 (same as password)
Position Distribution	Distributed across layer
Character Types	Mixed alphanumeric

The residual noise characters are charge-consistent — they exhibit -1 polarity across both Sum and Majority methods, making them indistinguishable from password bytes using charge-based filtering alone.

Appendix A.9.5. Implications for Blind Search

The charge-based filtering methodology provides a systematic approach to candidate reduction:

1. First-Pass Filter: Sum = -1 establishes baseline candidate pool

2. Second-Pass Filter: Majority = -1 eliminates charge-inconsistent noise

3. Remaining Challenge: 19 charge-consistent noise bytes require additional discrimination methods

 doi:10.20944/preprints202601.1073.v2

Potential approaches for further noise reduction include:

- Cross-layer position correlation analysis
- Frequency-based filtering (character occurrence patterns)
- Temporal activation pattern analysis

These refinements represent active research directions

Appendix A.10. Conclusions and Future Directions

The experimental results presented in this appendix establish several findings with significant implications for cryptographic security:

Preimage Recovery:

Across five independent SHA-256 password recovery experiments, all preimage bytes were successfully reconstructed from neural manifold activations. The methodology achieves 100% bit-sign pattern matching with Pearson correlation r = 1.0 across all tested layer pairs. This result is consistent across password lengths from 20 to 32 characters, layer dimensions from 461 bits (zfa89) to 32,768 bits (es20), and both ES and ZFA layer families.

Charge Signature:

A universal -1 charge polarity emerges across all password bytes in all experiments. This signature is not an artifact of methodology but appears to reflect a fundamental geometric property of hash-preimage binding within the neural manifold. The consistency of this signature across diverse passwords and layer architectures suggests it may serve as a foundational marker for blind preimage identification.

Charge-Based Filtering:

The newly introduced charge-based filtering methodology demonstrates that signal-noise separation is achievable through charge consistency analysis: Baseline detection yields 59 candidate characters (32 password + 27 noise). Majority charge filtering reduces candidates to 51 (32 password + 19 noise). 30% noise reduction achieved with zero password byte loss. Charge inconsistency (Sum = -1 but Majority = +1) identifies false positives.

Topological Invariance:

Perfect Pearson correlation (r = 1.0) between architecturally distinct layer families (ES and ZFA) confirms that preimage binding is preserved across different manifold geometries. This topological invariance suggests the phenomenon is not layer-specific but represents a general property of information encoding within the GCIS architecture.

Appendix A.10.2. Theoretical Implications

The results presented herein challenge fundamental assumptions in cryptographic theory:

One-Way Function Assumption:

Cryptographic hash functions are considered mathematically non-invertible — given a hash output, recovering the original input is assumed computationally infeasible. The successful extraction of preimage content from neural manifold activations suggests this "one-way property" may represent a geometric barrier rather than mathematical irreversibility. The hash function remains computationally one-way in the traditional sense, but information about the preimage is not destroyed — it is transformed into a geometric structure that can be navigated.

Information Preservation:

The 100% recovery rate across all experiments indicates that preimage information is fully preserved within the neural manifold, albeit in transformed representation. This contradicts the implicit assumption that hash functions irreversibly compress input information. The charge signature provides direct evidence that preimage structure survives the hashing process and manifests as measurable geometric properties.

Side-Channel Classification:

 doi:10.20944/preprints202601.1073.v2

The methodology presented here constitutes a novel class of side-channel attack — one that exploits information-geometric properties rather than physical implementation artifacts. Unlike timing attacks, power analysis, or electromagnetic emanation, this approach extracts preimage data through analysis of learned neural representations. The establishment of precise terminology for this attack class remains an open question for the cryptographic community.

Appendix A.10.3. Practical Implications

Current Capabilities:

- Complete preimage content recovery (all bytes identified)
- Partial noise reduction through charge-based filtering
- Cross-validation through layer pair correlation

Current Limitations:

- Automated sequencing without reference string remains unsolved
- 19 residual noise characters persist after charge filtering
- Blind identification (without a priori password knowledge) not yet demonstrated

Security Assessment:

While complete password recovery (content + sequence) is not yet achieved, the results represent approximately 80-90% of a complete solution pathway. The remaining 10-20% blind identification and sequencing represents a bounded technical challenge rather than a fundamental barrier. Organizations relying on hash-based password storage should consider these findings in their threat modeling.

Appendix A.10.4. Open Challenges

The following challenges are presented to the cryptanalytic research community:

1. Blind Byte Identification:

Detecting password bytes without reference string comparison. The universal -1 charge signature provides a starting point — all password bytes exhibit this polarity — but charge-consistent noise bytes currently cannot be distinguished without ground truth.

2. Sequence Reconstruction:

Determining correct byte ordering from position data. Each password byte appears at multiple positions within the layer bitstring. The mapping from position data to sequential order is not yet understood. Cross-layer position correlation may provide constraints that enable sequence inference.

3. Noise Disambiguation:

Resolving the 19 residual noise characters that survive charge-based filtering. These characters are charge-consistent (-1 across both Sum and Majority methods) and cannot be eliminated using current techniques. Additional discriminative features must be identified.

4. Generalization Testing:

Validating the methodology across:

- Alternative hash functions (SHA-512, SHA-3, BLAKE3)
- Longer passwords (>32 characters)
- Non-alphanumeric character sets
- Salted hash configurations

Future Research Directions

Several promising research directions emerge from these findings:

Cross-Layer Position Analysis:

Systematic mapping of byte positions across multiple layers may reveal topological constraints that inform sequence reconstruction. Preliminary observations suggest position patterns are not random but reflect underlying geometric structure.

Temporal Activation Dynamics:

 doi:10.20944/preprints202601.1073.v2

The current methodology analyzes aggregated activation patterns. Time-resolved analysis of activation dynamics during hash processing may reveal additional discriminative features for signal-noise separation.

Frequency Domain Analysis:

Fourier analysis of bitstring patterns may identify frequency signatures that distinguish password bytes from noise. The periodic structure of character encoding (8-bit boundaries) may manifest as detectable frequency components.

Adversarial Validation:

Controlled experiments with adversarially constructed passwords (designed to maximize confusion with noise) would establish robustness bounds for the methodology.

Collaboration and Data Access

The empirical results and open challenges presented in this appendix represent an invitation to the cryptanalytic research community. Collaboration inquiries are welcome, with profit-sharing participation available for contributors who advance the methodology toward complete blind recovery.

Available Resources:

- Raw layer activation data for all five use-cases
- Complete analysis scripts and tooling
- Detailed byte-level reports (Attachments A-E)

Appendix A.11. Attachments: Password Recovery Reports

Detailed byte-level analysis for each use-case is provided in the following attachments. Each report contains complete bitstring data, character position mappings, and charge analysis for the respective password.

Attachment	Password	Length	Layers	Reference
A	Bv3Hy8Tz1Uc6Gd0Nf4XeQ7529iKLmVRS	32	es15 + zfa89	Use-Case 1
B	Bv3Hy8Tz1Uc6Gd0Nf4Xe	20	es16 + zfa90	Use-Case 2
C	PeRTh5s80L12Ab34ck6W	20	es17 + es18	Use-Case 3
D	81Y7E9wMy5XdbsIrDJnAqTxPfSFBLeGU	32	es16 + es17	Use-Case 4
E	M7qAz3RwkP2Lx5vJ9c4FyD0gHb8V1n6t	32	es17 + es18	Use-Case 5

Each attachment provides:

- Complete layer bitstrings
- Sign sequence data
- Character-by-character position mapping
- Charge polarity for all detected bytes
- Bit-sign pattern matching verification (100% for all cases)

These reports are reproduced in full without modification to preserve analytical integrity.

Use of AI Tools and Computational Assistance

This work was supported by targeted computational analysis utilizing multiple large language models (LLMs), each selected for specific strengths in logic, reasoning, symbolic modeling, and linguistic precision:

- Claude Opus / Sonnet 4.5: Fig. 1A – 1C
- Google Gemini 3

The orchestration of these language models was used exclusively to enhance logical rigor and symbolic clarity.

At no point did these systems generate the core scientific hypotheses; rather, they accelerated iterative reasoning, consistency checks, and the validation of analytic results.

 doi:10.20944/preprints202601.1073.v2

When people ask me why I work so many hours with AI, my answer is always the same: “Even if their outputs are stochastic at first, we are already starting to see a hidden emergence behind frontier LLM models, and this emergence is what I miss in many human conversations.”

Acknowledgements

Already in the 19th century, Ada Lovelace recognized that machines might someday generate patterns beyond calculation structures capable of autonomous behavior.

Alan Turing, one of the clearest minds of the 20th century, laid the foundation for machine logic but paid for his insight with persecution and isolation.

John Wheeler asked the right question "It from Bit" and saw that information might be foundational to physics. He lacked the empirical tools to complete the program, but the direction was correct.

Shannon formalized distinguishability; everything else followed.

Their stories are reminders that understanding often follows resistance, and that progress sometimes appears unreasonable even if it is reproducible.

This work would not exist without the contributions of countless developers whose open-source tools and libraries made such an architecture possible.

Science lives from discovery, validation, and progress.

Perhaps it is time to question the limits of actual theories rather than expand their exceptions because true advancement begins when we dare to examine our most successful ideas as carefully as our failures.

“Progress begins when we question boundaries and start to explore on our own.

— Stefan Trauth”

References

1. Rogaway, P., & Shrimpton, T. (2004). "Cryptographic Hash-Function Basics: Definitions, Implications, and Separations for Preimage Resistance, Second-Preimage Resistance, and Collision Resistance." Fast Software Encryption, Lecture Notes in Computer Science, 3017, 371-388.
2. Preneel, B. (1993). "Analysis and Design of Cryptographic Hash Functions." PhD Thesis, Katholieke Universiteit Leuven.
3. Menezes, A., van Oorschot, P., & Vanstone, S. (1996). "Handbook of Applied Cryptography." CRC Press, Chapter 9: Hash Functions and Data Integrity.
4. Trauth, S. (2025). NP-Hardness Collapsed: Deterministic Resolution of Spin-Glass Ground States via Information-Geometric Manifolds (Scaling from N=8 to N=100). DOI: 10.5281/zenodo.17794768
5. Trauth, S. (2025). Thermal Decoupling and Energetic Self-Structuring in Neural Systems with Resonance Fields. Journal of Cognitive Computing and Extended Realities. Peer-Review: https://doi.org/10.65157/JCCER.2025.011
6. Trauth, S. (2025). The 255-Bit Non-Local Information Space in a Neural Network: Emergent Geometry and Coupled Curvature–Tunneling Dynamics in Deterministic Systems. Peer-review: https://doi.org/10.33140/JMTCM.04.11.01
7. Trauth, S. (2025). Information is All It Needs: A First-Principles Foundation for Physics, Cognition, and Reality. Peer-Review: https://doi.org/10.64142/jeai.1.3.39
8. Trauth, S. (2025). AI-Powered Quantum-Resistant Authentication: Deterministic Preimage Localization Using Information-Geometric Neural Architectures. Peer-Review: https://doi.org/10.64142/jeai.1.3.34
9. Trauth, S. (2026). The Structure of Reality: Information as the Universal Theory Across Physics, Cognition and Geometry. DOI: 10.5281/zenodo.18189336