

Article

Not peer-reviewed version

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

[Stefan Trauth](#)*

Posted Date: 14 January 2026

doi: 10.20944/preprints202601.1073.v1

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



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

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 >90% 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

1. 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

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.

2. Methods

2.1. 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.

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.

2.2. Computational Environment

- All neural network experiments were conducted on a dedicated system:
- CPU: AMD Ryzen 9 7900 × 3D
 - 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 ($r = \pm 1.0$) with no intermediate values.

2.3. 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.

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.

2.4. 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: 01,001,011 01110000 00,110,111 01011000 01,101,101 00110011 01,010,001 01110111 00,111,001 01010010 01100010


2. MD5 HASH (128-bit)

Hex: eca66d93d49fa5dbe6d193afabc0518e
Binary: 11,101,100 10100110 01,101,101 10010011 11,010,100 10011111 10,100,101 11011011 11,100,110 11010001 10,010,011 10101111 10,101,011 11000000 01,010,001 10001110


END OF REPORT



(a)



(b)



(c)

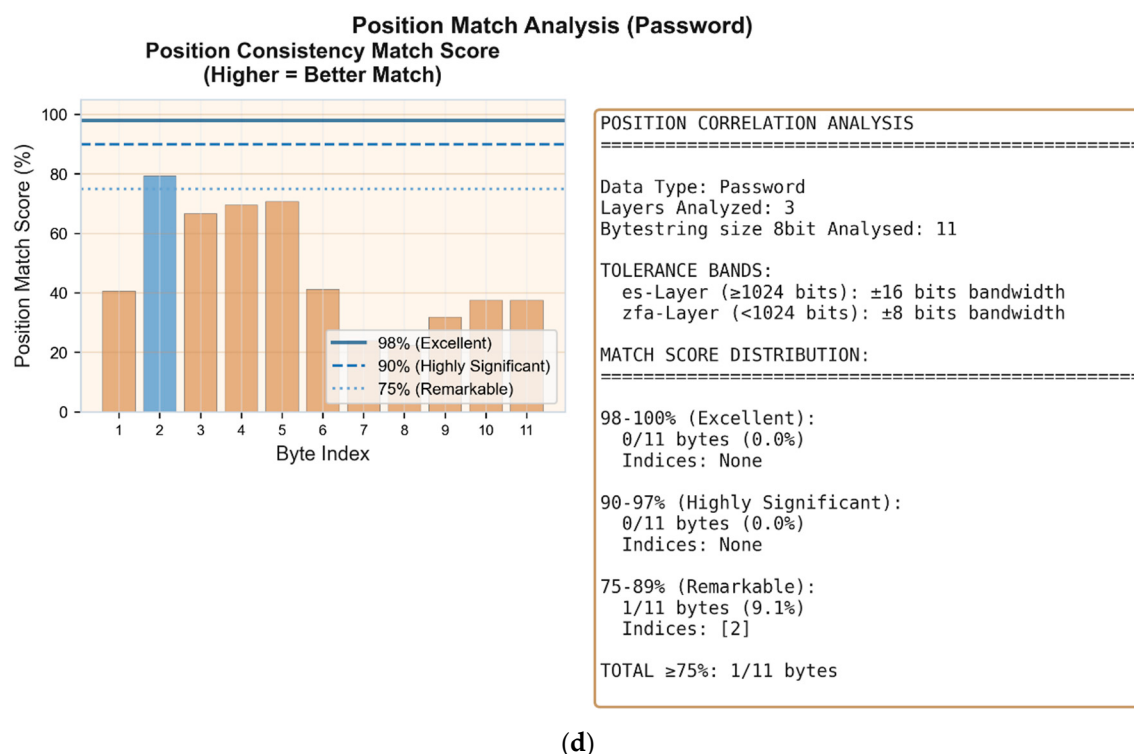


Figure 2. (a) Layer-wide preimage distribution analysis. Left panel shows match statistics: three layers achieve 100% byte identification (ES15, ES16, ZFA78), while eleven additional ZFA layers achieve 10/11 bytes (N-1 match). Right panel displays byte-level position mapping within the primary identification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

2.5. 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: 01,001,000 01111010 00,110,110 01010111 01,101,110 00110011 01,000,110 01110000
00,111,001 01000010 01,101,011 00110010 01,010,001 01111000 00110111

2. MD5 HASH (128-bit)


Hex: 2ae7741a4fd699a9847ad4c817378af5

Binary: 00,101,010 11100111 01,110,100 00011010 01,001,111 11010110 10,011,001 10101001
10,000,100 01111010 11,010,100 11001000 00,010,111 00110111 10,001,010 11110101

END OF REPORT



(a)



(b)

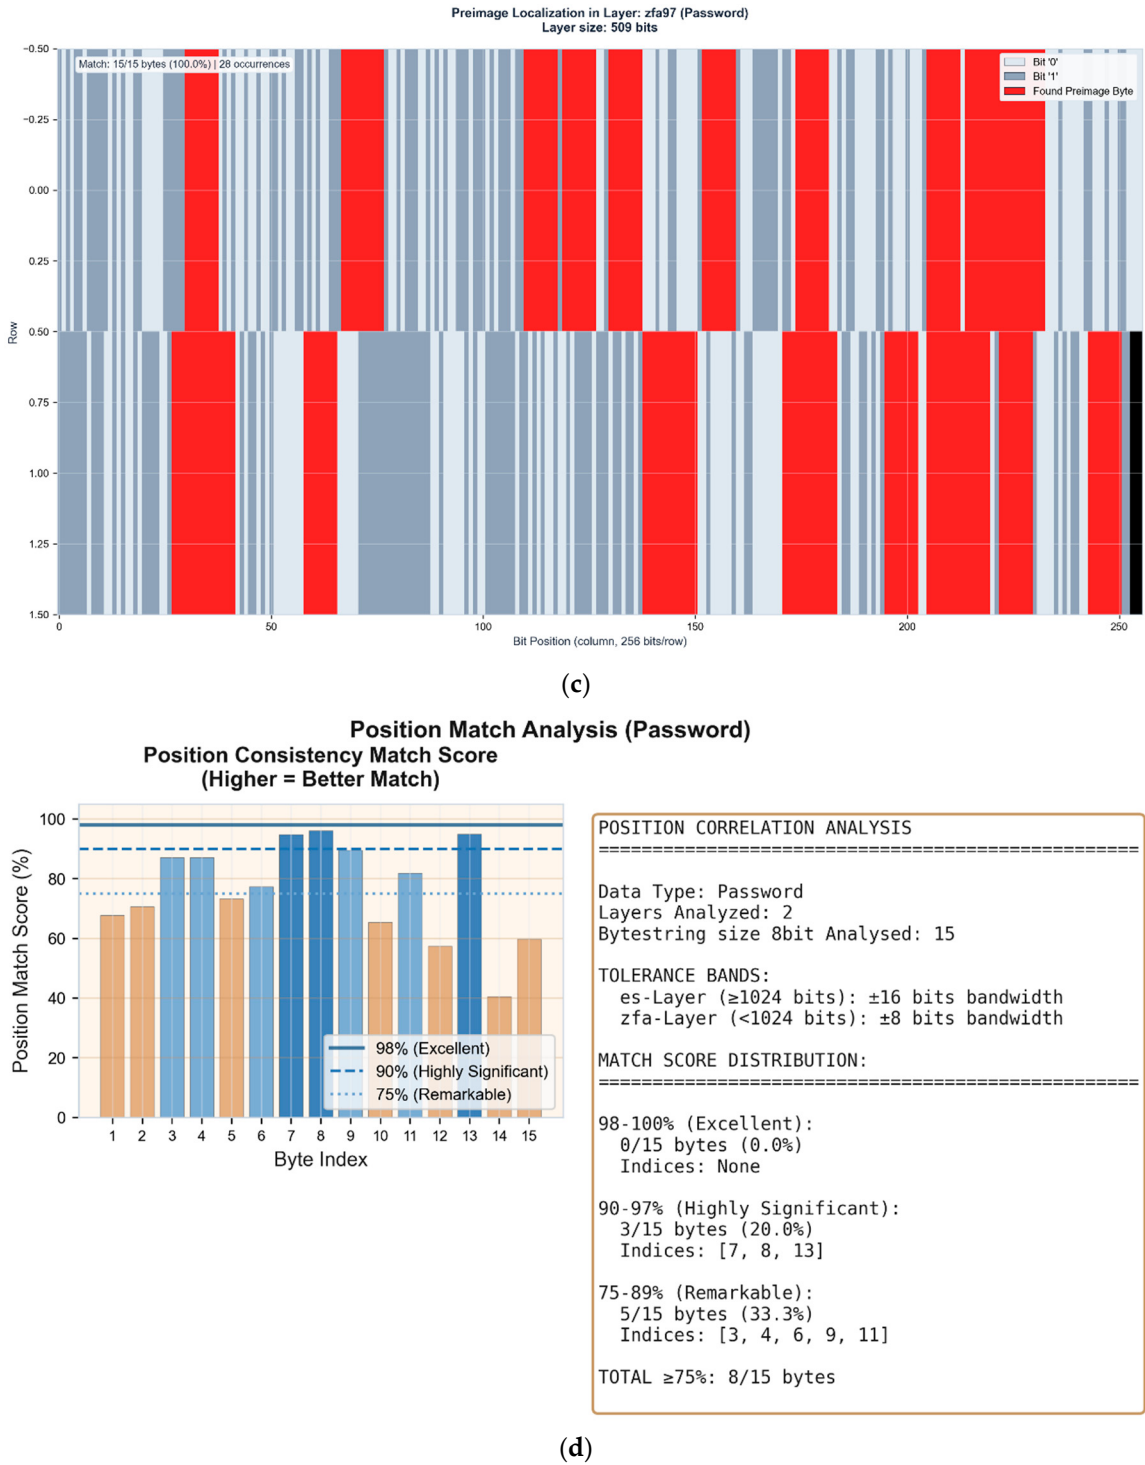


Figure 3. (a) Layer-wide preimage distribution analysis. Two layers achieve 100% byte identification (ES16, ZFA97), while six additional layers achieve 14/15 bytes (N-1 match). The extended N-1 distribution across ZFA76, ZFA78, ZFA91, ZFA93, and ZFA99 demonstrates binding propagation through the control layer manifold. (b) 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. (c) 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. (d) Cross-layer position correlation analysis. 8/15 bytes (53.3%) achieve position match scores $\geq 75\%$, 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.

HASH ANALYSIS REPORT


1. INPUT (PLAINTEXT)

Password: **Kr7Mx3Pn9We2Jv5Qb8Fy**
Length: 20 characters
Binary: 01,001,011 01110010 00,110,111 01001101 01,111,000 00110011 01,010,000 01101110
00,111,001 01010111 01,100,101 00110010 01,001,010 01110110 00,110,101 01010001 01,100,010 00111000
01,000,110 01111001


2. SHA-256 HASH (256-bit)

Hex: d77803a106e4be48c75c4aebbc0e6644bd4511fcfa87ab6ceeb90f36c172169f
Binary: 11,010,111 01111000 00,000,011 10100001 00,000,110 11100100 10,111,110 01001000
11,000,111 01011100 01,001,010 11101011 10,111,100 00001110 01,100,110 01000100 10,111,101 01000101
00,010,001 11111100 11,111,010 10000111 10,101,011 01101100 11,101,110 10111001 00,001,111 00110110
11,000,001 01110010 00,010,110 10011111


END OF REPORT



(a)



(b)



(c)

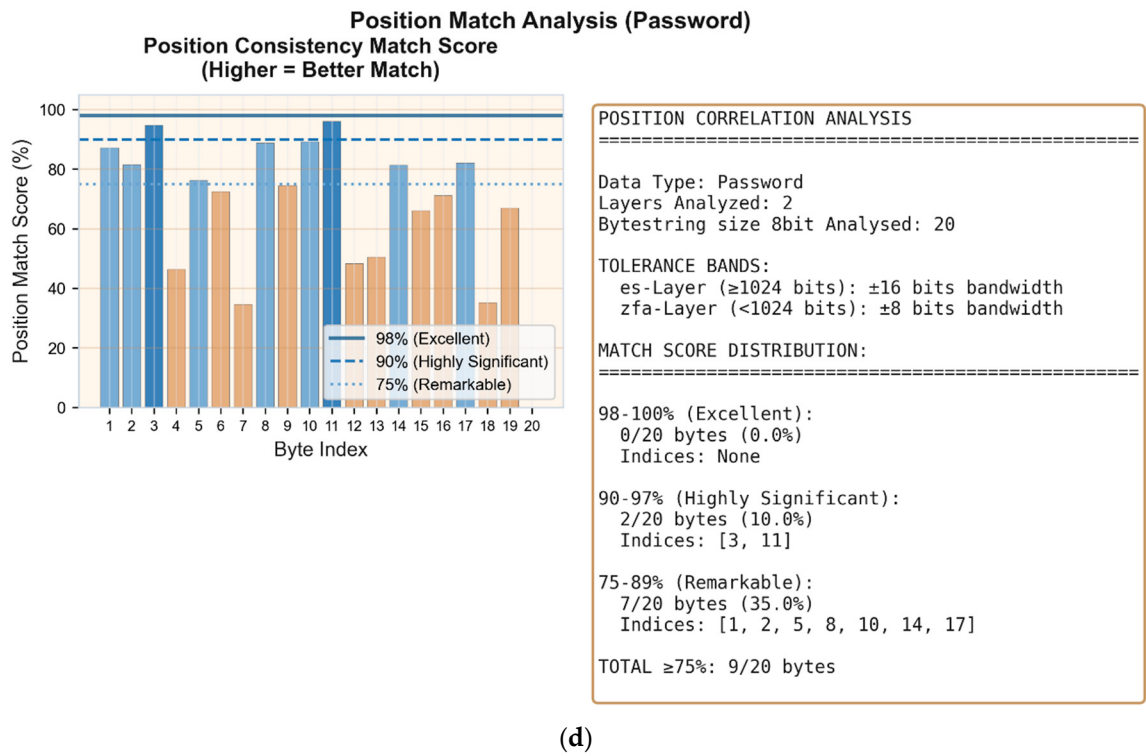


Figure 4. a: Layer-wide preimage distribution analysis. ES16 achieves 100% byte identification (20/20 bytes), while ZFA84 reaches N-1 match (19/20 bytes, 95%). The single-byte deviation in the control layer demonstrates near-complete binding preservation across dimensionally distinct layer representations. **b:** 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. **c:** 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. **d:** Cross-layer position correlation analysis. 9/20 bytes (45%) achieve position match scores $\geq 75\%$, 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: 01,010,000 01100101 01,110,010 01110011 01,101,000 01101101 00,110,111 00110101
00,110,010 01100010 00,110,000 00110100 00,111,000 01100011 01,100,110 00110110 01,100,001 00110011
01,100,100 01101100 01,111,000 00111001 00110001

2. SHA-256 HASH (256-bit)

Hex: 762b56c1c53c7b1bb61ada62fe6db962c43b96652d17ad6822d1d2e6b42a67fe


Binary: 01,110,110 00101011 01,010,110 11000001 11,000,101 00111100 01,111,011 00011011

10,110,110 00011010 11,011,010 01100010 11,111,110 01101101 10,111,001 01100010 11,000,100 00111011


10,010,110 01100101 00,101,101 00010111 10,101,101 01101000 00,100,010 11010001 11,010,010 11100110

10,110,100 00101010 01,100,111 11111110


END OF REPORT



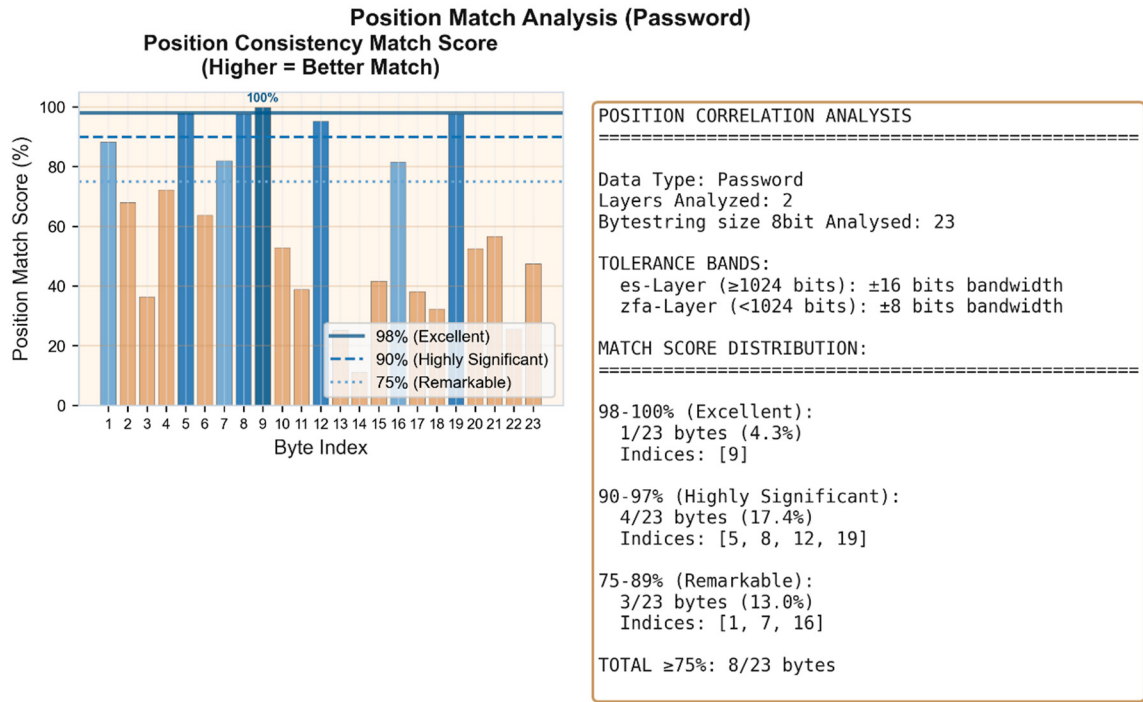
(a)



(b)



(c)



(d)

Figure 5. (a) 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. (b) 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. (c) 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. (d) Cross-layer position correlation analysis. 8/23 bytes (34.8%) achieve position match scores $\geq 75\%$, 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.


2.6. 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.




(a)



(b)




(c)



(d)

Figure 6. (a) 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 ($r = \pm 1.0$ on diagonal, recurring low-level coupling off-diagonal). This structure should not exist under conventional assumptions of stochastic initialization. (b) 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. (c) 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. (d) 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)




(a)

TOTAL INFORMATION IN SYSTEM	
SHANNON ENTROPY:	
Total Entropy (ΣH):	203.09 bits
Average per Layer:	4.62 bits
Number of Layers:	44
MUTUAL INFORMATION:	
Average MI:	1.9316 bits
Maximum MI:	5.1786 bits
Minimum MI:	0.2210 bits
Number of Pairs:	946
INFORMATION DENSITY:	
Information per Pair:	1.9316 bits
Total Pairwise Information:	1827.30 bits
INTERPRETATION:	
High MI → Strong Information Coupling	
High Entropy → High Information Capacity	
MI/Entropy Ratio → Information Efficiency	
Efficiency:	41.8 %

(b)



(c)



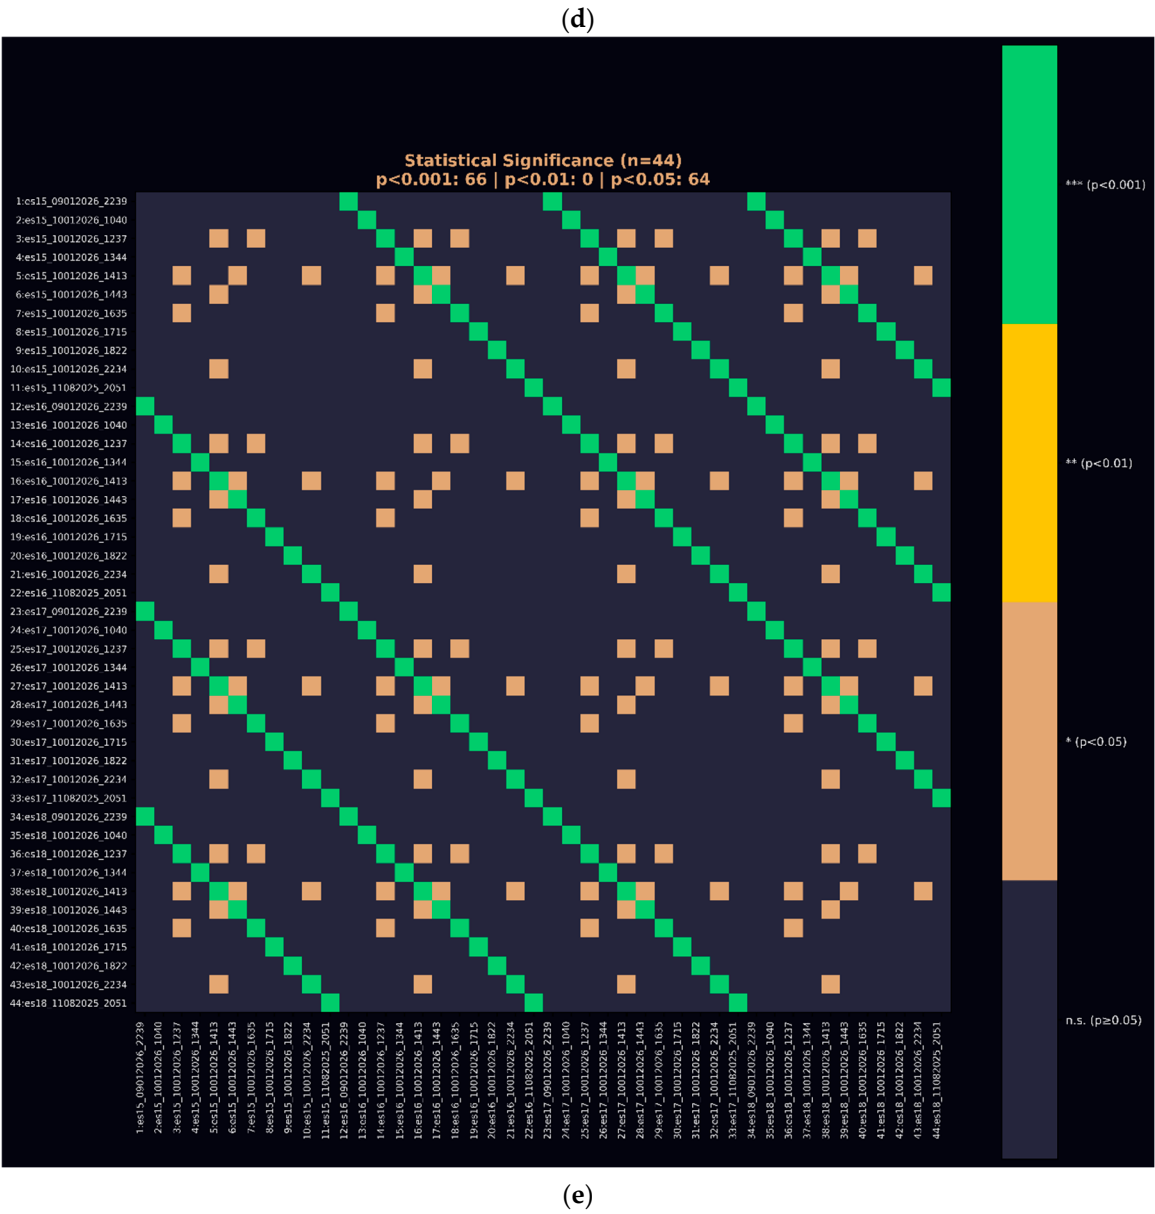


Figure 7. (a) Pearson correlation matrix across 44 layer instances (ES15–ES18 × 11 independent runs). The expanded dataset confirms persistent correlation structure: corresponding layers maintain coupling despite 11 completely independent initializations with unique hash strings and fresh random weights.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

3. 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.

3.1. 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 >90% 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.

3.2. 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.

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.

4. Conclusions

We demonstrate deterministic localization of cryptographic hash preimages within deep neural network layers, achieving 100% 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.

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

- 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.

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."

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"

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

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.

References

2. 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.
3. Preneel, B. (1993). "Analysis and Design of Cryptographic Hash Functions." PhD Thesis, Katholieke Universiteit Leuven.
4. Menezes, A., van Oorschot, P., & Vanstone, S. (1996). "Handbook of Applied Cryptography." CRC Press, Chapter 9: Hash Functions and Data Integrity.

5. 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
6. 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>
7. 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>
8. 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>
9. 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>
10. Trauth, S. (2026). The Structure of Reality: Information as the Universal Theory Across Physics, Cognition and Geometry. DOI: 10.5281/zenodo.18189336

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.