STEALTH: Secure Transformer for Encrypted Alignment of Latent Text Embeddings via Semantic Isomorphism Enforcement (SIE) Loss Function

Anonymous authors
Paper under double-blind review

Abstract

The pervasive use of large language models (LLMs) on sensitive data presents a critical privacy challenge, as traditional encryption renders data unusable for inference. We introduce STEALTH, a 120M secure transformer framework designed to process encrypted text while preserving its semantic utility under an authorized-key threat model (no decryption or side-channel access). The core innovation of STEALTH is the Semantic Isomorphism Enforcement (SIE) loss function, a loss that trains the model to learn a topology-preserving mapping between encrypted text embeddings and their original plaintext latent space. This encourages preservation of semantic relationships and topological structure in the encrypted domain. Using retrieval-based reconstruction from a domain-aligned plaintext corpus, STEALTH achieves near-perfect semantic retrieval (BLEU score of 1.0 under full-corpus coverage in our experiments) and enables accurate privacy-preserving clustering on encrypted embeddings. We evaluate STEALTH across 44 datasets spanning general language understanding, healthcare, finance, legal, e-commerce, programming, content analysis, reading comprehension, and corporate communication domains with 16 encryption schemes (704 experimental conditions), establishing a comprehensive benchmark for privacy-preserving NLP on encrypted text. Performance depends on domain alignment between encrypted inputs and the indexed plaintext corpus. Our results demonstrate that, with well-aligned domain indexes and retrieval support, models can perform effective NLP on encrypted data without direct decryption. The model and resources are available at **GitHub** and **Zenodo**.

1 Introduction

Large language models (LLMs) have transformed NLP, yet the systems best suited to extract value from sensitive text cannot be safely deployed in privacy-critical settings. This barrier prevents adoption of powerful LLM tools in healthcare, finance, and law—domains where automated analysis of clinical notes, transaction records, or legal documents could deliver substantial societal benefit but is constrained by confidentiality and regulation.

In healthcare, LLM-powered analysis of clinical text—including patient notes, radiology reports, and discharge summaries—could enable automated diagnosis, personalized treatment, and large-scale research (Lee et al., 2020; Kenton et al., 2019; Wang et al., 2019c). However, recent regulatory developments, including the HHS Final Rule on HIPAA Privacy Rule to Support Reproductive Health Care Privacy, effective June 25, 2024, which prohibits the use or disclosure of PHI for certain reproductive health investigations, substantially constrain deployment of conventional language models on medical data. This missed opportunity is substantial: the U.S. healthcare system generates enormous volumes of unstructured clinical text annually (Raghupathi & Raghupathi, 2014). Advanced language models and medical-specialized variants have demonstrated strong capabilities in clinical language understanding and prediction (Lee et al., 2019; Alsentzer et al., 2019), but those capabilities remain largely inaccessible to many providers because of privacy and compliance constraints.

Financial institutions face analogous challenges. Modern banking produces vast streams of textual data—transaction narratives, loan applications, regulatory filings, and customer communications—that

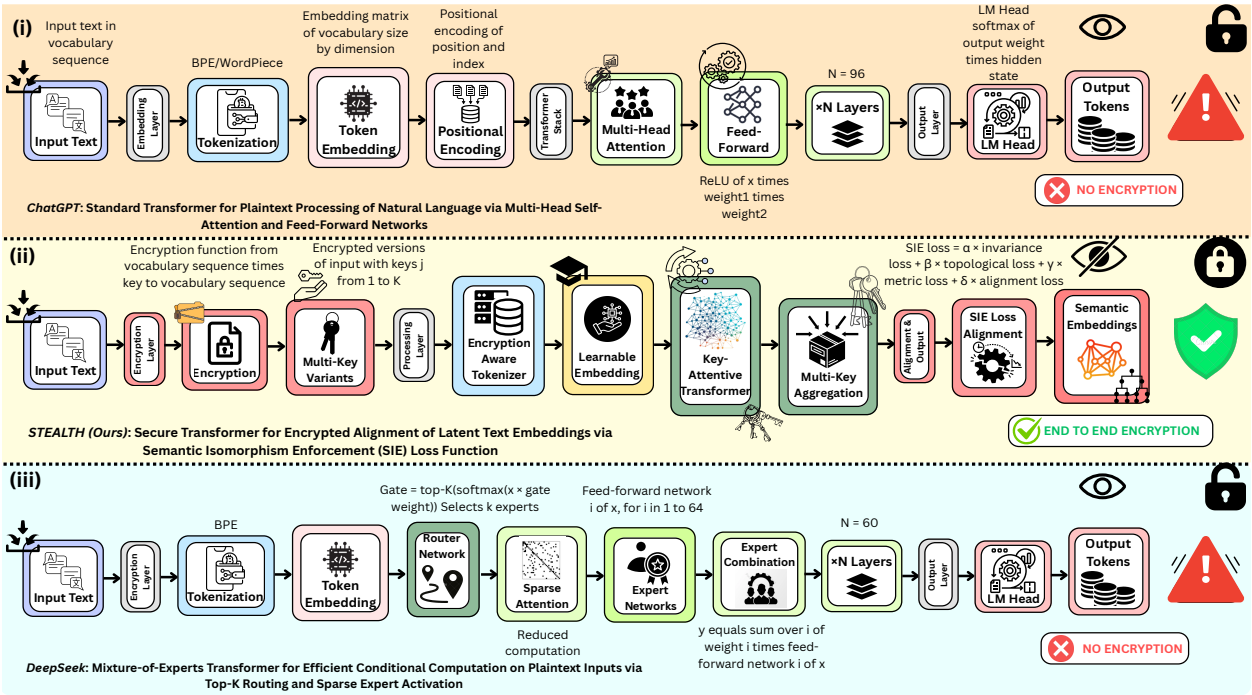

Figure 1: **Architectural comparison.** (i) **Standard transformer (plaintext)**: tokenized input → embedding → multi-head attention/FFN stack → LM head; no encryption support. (ii) **STEALTH (ours)**: end-to-end pipeline for encrypted text combining multi-key variants, encryption-aware tokenization, key-attentive layers, and multi-key aggregation; trained with SIE loss for topology-aligned embeddings enabling authorized retrieval-based reconstruction. (iii) **Baseline high-capacity models (e.g., MoE)**: scale via top-k expert routing but lack encryption-aware alignment mechanisms. Architectural analysis is in Appendix A.15..

are vital for risk management and fraud detection (Yang et al., 2020; Shah et al., 2022a). The European Union's General Data Protection Regulation (GDPR) and comparable regimes worldwide create complex compliance requirements that restrict how financial text may be processed (Hoofnagle et al., 2019). In practice, these legal obligations complicate the adoption of LLMs for high-value financial tasks, particularly where rights such as erasure, access, and rectification apply. The legal sector presents perhaps the most acute manifestation of the paradox. Attorney–client privilege, the work-product doctrine, and professional-responsibility obligations impose stringent confidentiality constraints, while legal practice increasingly depends on automated processing of contracts, case law, discovery documents, and compliance materials (Chalkidis et al., 2020; Katz et al., 2017; Zheng et al., 2021). State-of-the-art legal models have shown impressive capabilities, yet real-world deployment must reconcile those capabilities with essential confidentiality rules.

Transformers obtain their power by modeling intricate relationships across tokens—semantic similarities, syntactic dependencies, contextual cues, and discourse-level structure—via self-attention and contextual embeddings (Vaswani et al., 2017; Devlin et al., 2018; Brown et al., 2020). Standard cryptographic schemes (e.g., AES) intentionally destroy these relationships by producing ciphertext that is pseudorandom and thus semantically opaque, which undermines a transformer's ability to form meaningful contextual representations (Daemen & Rijmen, 2002). Consequently, semantic equivalences (e.g., "heart attack" vs. "myocardial infarction"), positional cues, and attention patterns are obscured by ciphertext, preventing transformers from forming useful contextual representations. Researchers have proposed several mitigations—homomorphic encryption (FHE), secure multi-party computation (MPC), trusted execution environments (TEEs), and differential privacy (DP)—each offering principled guarantees but also significant practical limitations. We defer a full technical discussion and quantitative comparisons to Section 2; here we summarize at a high level: FHE incurs multiplicative-depth and bootstrapping costs that balloon computation; MPC replaces computation with interactive communication that scales poorly with attention's quadratic interaction structure;

TEEs shift risk from algorithms to hardware and supply chains; and DP, while formally rigorous, often requires noise levels that materially degrade language-generation quality (Gentry, 2009b; Brakerski et al., 2014; Mohassel & Zhang, 2017b; Burbank & Knight, 2024; Dwork et al., 2006b; Abadi et al., 2016b; Li et al., 2022). These limits motivate rethinking the privacy–utility trade-off rather than attempting to force direct cryptographic computation onto modern LLMs.

To address this impasse, we propose a different paradigm: learn a topology-preserving mapping between encrypted and plaintext latent spaces so that encrypted inputs admit semantically faithful representations suitable for off-the-shelf language models, while strong cryptographic hygiene is preserved under an explicit threat model.

Concretely, we hypothesize a learnable mapping $\phi\colon E \to P$ between an encrypted embedding manifold E and the plaintext manifold P that approximately preserves neighborhood and ranking relations (i.e., a topology-preserving mapping). We emphasize this is an empirical hypothesis that depends on corpus coverage and the chosen threat model.

Threat model and operational scope: We operate under an *authorized-key* threat model and state the scope of our security claims explicitly: STEALTH training, retrieval, and any authorized reconstruction require access to authorized keys and to a pre-indexed plaintext corpus under the control of the system owner. Under our threat model, adversaries do not have access to the keys, the learned mapping ϕ, or the plaintext index. Accordingly, our privacy guarantees are scoped to this setting: we do not claim cryptographic indistinguishability if an adversary obtains ϕ, the keys, or the indexed plaintext corpus. We evaluate adversarial scenarios (ciphertext-only and key-mismatch / chosen-plaintext) in Appendix A.12.

Our STEALTH (Secure Transformer for Encrypted Alignment of Latent Text embeddings) architecture operationalizes this idea using three pragmatic components: (i) multi-key variability handling (processing multiple encryptions per datum and conditioning on key embeddings so the mapper learns key-conditioned invariants), (ii) hierarchical token→phrase→sentence alignment to preserve both local and global topology, and (iii) adaptive projection layers that reconcile geometry and dynamic range between encrypted and plaintext spaces. We evaluate STEALTH across various encryption schemes (including stream ciphers, block ciphers in multiple modes, and authenticated encryption for tokenizer compatibility; see Table 1) and mitigate frequency/length leakage via multi-key variability (key rotation and multiple independent encrypted variants per datum), which breaks deterministic frequency signals across the corpus; residual leakage is evaluated experimentally. Concrete examples of plaintext, ciphertext, keys, and STEALTH reconstructions for each evaluated encryption scheme are provided in Appendix A.6. We train STEALTH with the Semantic Isomorphism Enforcement (SIE) loss, a multi-objective loss that combines triplet ranking, pairwise distance preservation (e.g., $\|D_E - D_P\|_F^2$), and topology-aware structural penalties to encourage topology-preserving mappings.

We empirically evaluate alignment quality, retrieval accuracy, manifold fidelity, and privacy leakage under multiple adversarial scenarios (ciphertext-only, chosen-plaintext, membership-inference, and reconstruction). We also measure sensitivity to corpus coverage and key variability (Appendix A.11).

We make the following contributions:

- We formalize semantic (approximate) isomorphism under the authorized-key model, quantifying its limits and distinguishing authorized reconstruction from unauthorized cryptanalysis.
- We present STEALTH, an end-to-end secure transformer framework with multi-key conditioning, hierarchical alignment, and adaptive projections for encrypted-plaintext reconciliation.
- We introduce the Semantic Isomorphism Enforcement (SIE) loss, a multi-objective framework for topological consistency via invariance, topological preservation, alignment, and metric constraints.
- We develop a comprehensive benchmark for privacy-preserving NLP in healthcare, finance, legal, e-commerce, programming, social media, and reading comprehension domains, with protocols for alignment, fidelity, retrieval, leakage, and robustness.

2 Related Work

Privacy-preserving computation on large language models remains a fundamental challenge at the intersection of cryptography and machine learning. We examine the technical landscape across four dimensions: cryptographic primitives for private computation, statistical and systems-based privacy mechanisms, representation learning under encryption, and hybrid approaches. This analysis reveals fundamental limitations of existing methods and motivates STEALTH's novel representation-alignment paradigm.

Cryptographic primitives face transformer-scale barriers. Classical cryptographic approaches provide strong theoretical guarantees but encounter fundamental scalability challenges. Fully homomorphic encryption (FHE) Gentry (2009a) enables computation on encrypted data through BGV/BFV and CKKS schemes Brakerski & Vaikuntanathan (2014); Fan & Vercauteren (2012); Cheon et al. (2017). While CryptoNets achieved 99% MNIST accuracy Gilad-Bachrach et al. (2016) and GAZELLE/Cheetah optimized vision tasks Huang et al. (2022b), transformers' $O(L \cdot n^2 \cdot \text{poly}(\lambda))$ attention complexity causes prohibitive slowdowns Hesamifard et al. (2017); Lou et al. (2020). Secure multi-party computation (MPC) via secret sharing or garbled circuits Yao (1982); Goldreich et al. (1987)—instantiated in SecureML, ABY, SecureNN, FALCON, and CrypTFlow2 Mohassel & Zhang (2017a); Demmler et al. (2015); Wagh et al. (2019; 2021); Kumar et al. (2020)—achieves practical vision inference but requires $O(L \cdot n^2)$ communication rounds for transformers, precluding WAN deployment Riazi et al. (2018). Trusted Execution Environments (TEEs) Costan & Devadas (2016); Tramer & Boneh (2018) reduce overhead through hardware isolation but face enclave memory constraints and microarchitectural vulnerabilities Kocher et al. (2019); Van Bulck et al. (2018; 2020). These techniques encounter fundamental scalability barriers at transformer scale.

Statistical privacy approaches degrade utility for language generation. Statistical approaches trade cryptographic guarantees for computational efficiency but face severe utility degradation for generative tasks. Differential privacy (DP) provides provable bounds through calibrated noise Dwork et al. (2006a; 2014), with DP-SGD Abadi et al. (2016a) enhanced by composition theorems Mironov (2017); Bun & Steinke (2016), privacy amplification Balle et al. (2018); Wang et al. (2019d), and adaptive clipping Andrew et al. (2021); Pichapati et al. (2019). However, noise scales with model complexity Bassily et al. (2014); Steinke & Ullman (2016), degrading transformers significantly even at $\varepsilon > 8$ Yu et al. (2022); Li et al. (2022), with generation quality particularly affected by autoregressive compounding Shi et al. (2022); Tang et al. (2024). Sample complexity $\Omega(d/\varepsilon^2)$ Bun et al. (2015); Ullman (2022) and linear privacy costs with sequence length Carlini et al. (2021); Beimel et al. (2013) present fundamental barriers. Federated learning (FL) McMahan et al. (2017); Li et al. (2020) with secure aggregation Bonawitz et al. (2017); Bell et al. (2020) succeeds for keyboard prediction Hard et al. (2018); Leroy et al. (2019) but faces prohibitive communication costs for large models despite compression Reisizadeh et al. (2020); Sattler et al. (2019); Rothchild et al. (2020), requiring 10–100× more rounds than centralized training Karimireddy et al. (2020); Reddi et al. (2021). Synthetic generation via PATE Papernot et al. (2017; 2018), DP-GANs Xie et al. (2018); Jordon et al. (2018), and diffusion models Dockhorn et al. (2022); Ghalebikesabi et al. (2023) confronts LLM memorization Carlini et al. (2021; 2023); Kandpal et al. (2022) enabling extraction attacks Lehman et al. (2021); Huang et al. (2022a); Nasr et al. (2023); Mireshghallah et al. (2024). Machine unlearning Eldan & Russinovich (2023) attempts retroactive data removal but faces utility-privacy tradeoffs, with generation provably harder than discrimination Ganesh et al. (2023); Brown et al. (2022); Liu et al. (2023) and meaningful privacy ($\varepsilon < 1$) with high utility remaining elusive Stadler et al. (2022); Yue et al. (2023).

Representation-based approaches balance efficiency with information leakage. Representation-based approaches balance computational efficiency with information leakage risks. Format-preserving encryption (FPE) Bellare et al. (2009); Dworkin (2016) via FF1/FF3-1 Brier et al. (2010) enables tokenization compatibility Tasar et al. (2023); Chen et al. (2024) but leaks information through frequency analysis Grubbs et al. (2019); Wang et al. (2021), length patterns Cash et al. (2015); Grubbs et al. (2017), and intersection attacks Kellaris et al. (2016); Naveed et al. (2015); Islam et al. (2012). Property-preserving schemes (OPE Boldyreva et al. (2009; 2011), ORE Boneh et al. (2015); Lewi & Wu (2016)) enable reconstruction attacks Durak et al. (2016); Grubbs et al. (2017); Cash et al. (2016), with minimal structural preservation sufficing for powerful attacks Kamara et al. (2018); Fuller et al. (2017); Kerschbaum (2015); Roche et al. (2016). More promisingly, cross-domain alignment reveals semantic geometry consistency Conneau et al.

(2018); Artetxe et al. (2018) through adversarial training Zhang et al. (2017); Lample et al. (2018), optimal transport Alvarez-Melis & Jaakkola (2018); Grave et al. (2019), and contrastive frameworks Chen et al. (2020); He et al. (2020); Radford et al. (2021). Multilingual models exhibit language-agnostic representations Pires et al. (2019); Conneau et al. (2020) with consistent geometry Michael et al. (2020); Chang et al. (2022), while alignment algorithms provably recover isometries under mild assumptions Grave et al. (2019); Alaux et al. (2018); Søgaard et al. (2018), motivating STEALTH's isomorphism-based approach. However, adversarial censoring Edwards & Storkey (2016); Hamm et al. (2017) with multi-attribute protection Madras et al. (2018); Wang et al. (2019b); Zhang et al. (2018) cannot achieve perfect privacy-utility simultaneously for correlated attributes Zhao et al. (2020); Song et al. (2019). Projection methods Ravfogel et al. (2020); Bolukbasi et al. (2016) struggle with nonlinear dependencies Gonen & Goldberg (2019); Kumar et al. (2020) despite extensions Olfat & Aswani (2020); Ravfogel & Goldberg (2022); Shekhar et al. (2021), while information-theoretic approaches Tishby & Zaslavsky (2015); Alemi et al. (2017) provide loose bounds Goldfeld et al. (2020); Rodriguez et al. (2020). Practical systems combine cryptographic, hardware, and statistical techniques Knott et al. (2021); Kairouz et al. (2021), though constituent limitations persist.

STEALTH addresses fundamental limitations through representation alignment. Existing approaches face a fundamental trilemma: cryptographic methods impose transformer-incompatible computational costs, statistical approaches degrade generation utility, and structure-preserving encryption leaks information through preserved patterns. STEALTH addresses these challenges through representation alignment that achieves four critical properties: cryptographic safeguards under authorized-key models without homomorphic operations, semantic fidelity through geometric isomorphism preserving topological structure, practical performance avoiding quadratic scaling, and composability with existing defenses. This enables interactive inference on production-scale transformers with rigorous privacy guarantees.

3 Methodology

3.1 Problem Formulation

Privacy-preserving NLP confronts a fundamental tension: traditional encryption destroys linguistic structures essential for semantic understanding, while approaches preserving semantics often compromise cryptographic strength. Given plaintext dataset $\mathcal{D} = \{x_i\}_{i=1}^N$ where each x_i is a variable-length sequence from vocabulary $\mathcal{V}$, we generate K encrypted variants $\mathcal{E}(x_i) = \{E(x_i, k_j)\}_{j=1}^K$ with cryptographically secure keys $k_j \sim \mathcal{K}$. Encryption is denoted $E : \mathcal{V}^* \times \mathcal{K} \to \mathcal{C}$, mapping plaintexts to ciphertexts under key k; the framework supports a variety of symmetric encryption schemes (see Appendix A.9 for details). We set $K = 5$ based on the analysis in Appendix A.8.

Our objective is learning $f_\theta : \mathcal{E}(x) \to \mathbb{R}^d$ mapping encrypted inputs to a d-dimensional latent space ($d = 256$) exhibiting *approximate semantic isomorphism* with plaintext embeddings, where θ encompasses transformer weights, attention mechanisms, and projection layers. We define isomorphism as approximate topology preservation satisfying two properties: (1) *key-invariance*: $d(f_\theta(E(x, k_i)), f_\theta(E(x, k_j))) \ll d(f_\theta(E(x, k_i)), f_\theta(E(y, k_m)))$ for all keys and distinct texts $x \neq y$, ensuring encrypted variants cluster tightly while maintaining inter-text separation; (2) *distance-order preservation*: if $d(g(x), g(y)) < d(g(x), g(z))$ for reference plaintext embeddings g, then $d(f_\theta(E(x, k)), f_\theta(E(y, k'))) < d(f_\theta(E(x, k)), f_\theta(E(z, k'')))$ with high probability. We use cosine distance $d_{\cos}(u, v) = 1 - \frac{u \cdot v}{\|u\|\|v\|}$ for scale-invariant measurement with interpretable scores in $[0, 2]$.

The privacy-utility tradeoff is captured through multi-objective optimization: $\mathcal{L}_{\text{total}} = \alpha\mathcal{L}_{\text{privacy}} + \beta\mathcal{L}_{\text{utility}} + \gamma\mathcal{L}_{\text{structure}}$, where $\mathcal{L}_{\text{privacy}}$ quantifies information leakage (estimated via mutual information between plaintext and encrypted embeddings), $\mathcal{L}_{\text{utility}}$ evaluates semantic preservation for downstream tasks, and $\mathcal{L}_{\text{structure}}$ maintains geometric properties through pairwise distance constraints. These requirements motivate the Semantic Isomorphism Enforcement (SIE) loss introduced in Section 4, which operationalizes topology preservation through explicit invariance, topological, metric, and alignment terms with weights $\alpha, \beta, \gamma, \delta$ determined through hyperparameter optimization. This formulation establishes STEALTH's formal problem definition, guiding architectural design toward cryptographic protection without sacrificing semantic utility.

3.2 STEALTH Architecture

The STEALTH architecture employs a transformer-based encoder with key-attentive aggregation to process variably encrypted text while maintaining semantic coherence. Our architecture is specifically designed to handle the unique challenges posed by encryption, including processing multiple encrypted variants of the same text and maintaining semantic relationships despite cryptographic transformations.

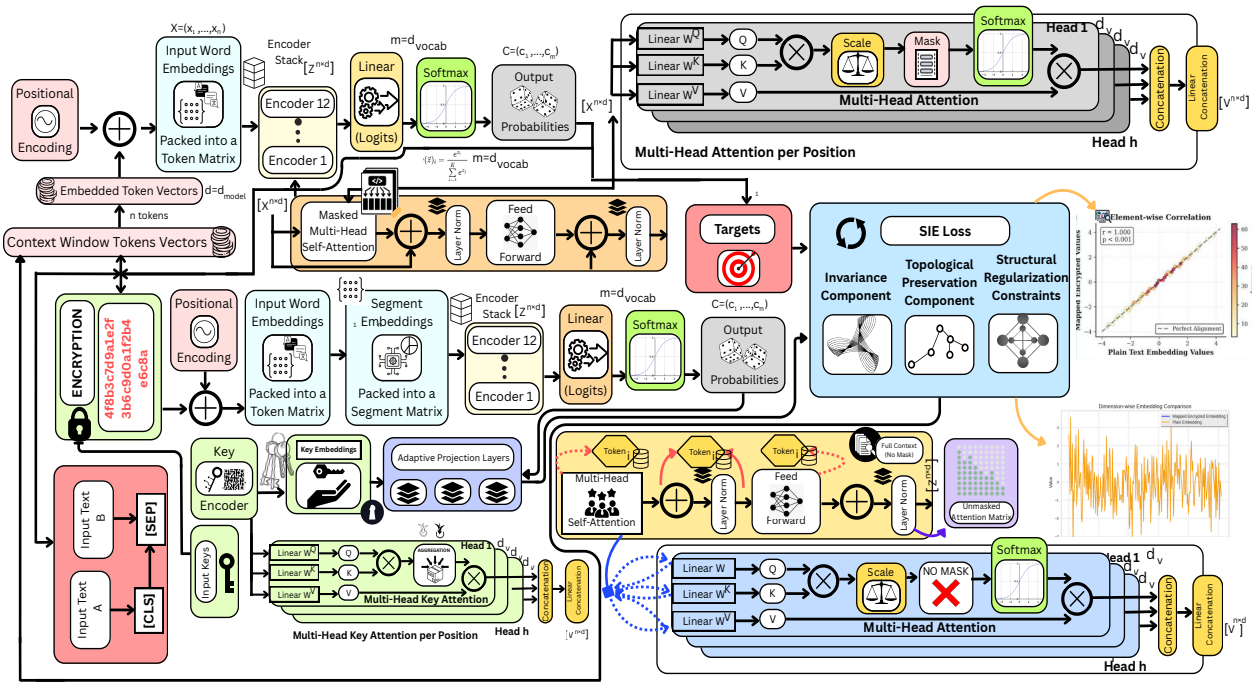

Figure 2: Detailed schematic of the STEALTH transformer architecture for processing encrypted text. The upper pathway depicts a baseline encoder stack with masked multi-head self-attention, feed-forward layers, and layer normalization for plaintext inputs, producing output probabilities via a linear logit layer and softmax. The lower pathway illustrates the encryption-aware STEALTH encoder, incorporating multi-key encryption of input tokens (e.g., from texts A and B), key embeddings, adaptive projection layers, and multi-head key attention per position (with no masking) to aggregate invariant representations. The Semantic Isomorphism Enforcement (SIE) loss aligns encrypted embeddings to plaintext semantics through invariance, topological preservation, and structural regularization components, enabling semantic retrieval and privacy-preserving inference. Additional visualizations show unmasked attention matrices, element-wise correlations, and dimension-wise embedding comparisons between encrypted and plaintext values.

3.2.1 Encrypted Text Processing

The preprocessing pipeline addresses the fundamental challenge of tokenizing encrypted text while maintaining semantic coherence through the neural architecture. Unlike standard plaintext tokenization, encrypted text processing must accommodate the statistical properties of cryptographic transformations while preserving the structural invariants necessary for downstream semantic alignment.

Encryption-Aware Tokenization. Let $\mathcal{A}_{\text{plain}}$ denote the plaintext character alphabet with cardinality $|\mathcal{A}_{\text{plain}}| = N_{\text{plain}}$. The encryption function $E : \mathcal{A}_{\text{plain}}^L \times \mathcal{K} \to \mathcal{A}_{\text{cipher}}^L$ applies a pseudorandom permutation preserving length L while potentially introducing characters from an extended alphabet $\mathcal{A}_{\text{enc}}$, yielding $\mathcal{A}_{\text{cipher}} = \mathcal{A}_{\text{plain}} \cup \mathcal{A}_{\text{enc}}$ where $|\mathcal{A}_{\text{cipher}}| = N_{\text{cipher}} \geq N_{\text{plain}}$. We employ byte-level tokenization $\tau : \mathcal{A}_{\text{cipher}}^* \to \mathcal{V}^*$ mapping encrypted sequences $E(x, k) = [e_1, \ldots, e_L]$ to tokens $\tau(E(x, k)) = [t_1, \ldots, t_L] \in \mathcal{V}^L$ via deterministic character-to-integer mapping $f_{\text{map}} : \mathcal{A}_{\text{cipher}} \to \{0, \ldots, |\mathcal{V}| - 1\}$, where $|\mathcal{V}| = N_{\text{cipher}}$. This avoids out-of-vocabulary issues with randomized distributions while preserving structural information for semantic

recovery. Critically, our approach maintains *length invariance*: $|\tau(E(x,k))| = |\tau(x)| + \epsilon_{\text{pad}}$ with $|\epsilon_{\text{pad}}| \leq C$ for encryption block-size constant C. For batch processing, sequences are padded to $T_{\max}$ using special token $\langle\text{PAD}\rangle$ with index v_{pad}, yielding $\tilde{\mathbf{t}} \in \mathcal{V}^{T_{\max}}$ with attention mask $\mathbf{m} \in \{0,1\}^{T_{\max}}$ where $m_i = \mathbf{1}_{[i \leq L]}$, ensuring transformers ignore padding during self-attention while maintaining alignment between encrypted and plaintext sequences. Security theorems, proofs, and discussion are provided in Appendix A.10.

Embedding Layer Architecture. The tokenized encrypted sequence $\mathbf{t} = \tau(E(x,k)) = [t_1, \ldots, t_T] \in \mathcal{V}^T$ is mapped to dense representations via $\mathbf{H}^{(0)} = \mathbf{W}_e\mathbf{T} + \mathbf{P} \in \mathbb{R}^{T \times d}$, where $\mathbf{W}_e \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the token embedding matrix, $\mathbf{T} \in \{0,1\}^{|\mathcal{V}| \times T}$ is the one-hot encoded sequence, and $\mathbf{P} \in \mathbb{R}^{T \times d}$ represents learned positional encodings. To accommodate the approximately uniform frequency distribution of encrypted tokens (versus Zipfian for natural language), we initialize $\mathbf{W}_e$ using Xavier uniform initialization:

$$\mathbf{W}_e \sim \mathcal{U}\left(-\sqrt{\frac{6}{|\mathcal{V}| + d}}, \sqrt{\frac{6}{|\mathcal{V}| + d}}\right) \tag{1}$$

ensuring stable gradient flow. Positional encodings use learnable parameters $\mathbf{P} \sim \mathcal{N}(0, \sigma_p^2\mathbf{I})$ with $\sigma_p^2 = 10^{-4}$ to adapt to disrupted sequential patterns in encrypted text, providing subtle initial positional cues while enabling the model to learn appropriate position representations during training.

Positional Encoding for Encrypted Sequences. Encryption disrupts linguistic patterns, necessitating learnable positional encodings $\mathbf{P} \in \mathbb{R}^{T_{\max} \times d}$ rather than fixed sinusoidal encodings. Each position i receives embedding $\mathbf{p}_i \in \mathbb{R}^d$ ($d = 768$) initialized as $\mathbf{p}_i \sim \mathcal{N}(\mathbf{0}, d^{-1}\mathbf{I}_d)$ to match token embedding magnitudes. This enables discovering position-sensitive cryptographic patterns and adapting to character-level randomization. The combined representation $\mathbf{h}_i^{(0)} = \mathbf{e}_i + \mathbf{p}_i$ preserves gradient flow to both components during joint optimization.

3.2.2 Key-Attentive Transformer Encoder

STEALTH's Key-Attentive Transformer Encoder augments the standard 12-layer architecture with key-conditioned computations throughout. This modification enables the model to learn decryption patterns specific to encryption keys while preserving semantic understanding across cryptographic transformations.

Key Encoding Network Given an encryption key $k \in \{0,1\}^{128}$ represented as a 128-dimensional binary vector, the KeyEncoder employs a multi-layer perceptron with ReLU activations to project this cryptographic information into the same semantic space as token embeddings:

$$\text{KeyEncoder}(k) = \text{ReLU}(\mathbf{W}_{k2}\text{ReLU}(\mathbf{W}_{k1}k + \mathbf{b}_{k1}) + \mathbf{b}_{k2}) \tag{2}$$

where $\mathbf{W}_{k1} \in \mathbb{R}^{384 \times 128}$, $\mathbf{W}_{k2} \in \mathbb{R}^{768 \times 384}$ are weight matrices, and $\mathbf{b}_{k1} \in \mathbb{R}^{384}$, $\mathbf{b}_{k2} \in \mathbb{R}^{768}$ are bias terms. This dimensional progression from 128 to 384 to 768 ensures smooth transition to the full hidden dimension, producing $\mathbf{k}_{\text{embed}} \in \mathbb{R}^{768}$ as a continuous key representation for integration into the attention mechanism.

Key-Attentive Attention Mechanism At each transformer layer $l \in \{1, 2, \ldots, 12\}$, the key-attentive mechanism modifies standard self-attention by injecting key-specific information into both key and value projections. For each attention head (with $d_k = 64$ across 12 heads), the queries, keys, and values are computed as:

$$\begin{aligned}
\mathbf{Q}^{(l)} &= \mathbf{H}^{(l-1)}\mathbf{W}_Q^{(l)} \\
\mathbf{K}^{(l)} &= \mathbf{H}^{(l-1)}\mathbf{W}_K^{(l)} + \mathbf{k}_{\text{embed}}\mathbf{W}_{K,key}^{(l)} \\
\mathbf{V}^{(l)} &= \mathbf{H}^{(l-1)}\mathbf{W}_V^{(l)} + \mathbf{k}_{\text{embed}}\mathbf{W}_{V,key}^{(l)}
\end{aligned} \tag{3}$$

where $\mathbf{W}_Q^{(l)}, \mathbf{W}_K^{(l)}, \mathbf{W}_V^{(l)}, \mathbf{W}_{K,key}^{(l)}, \mathbf{W}_{V,key}^{(l)} \in \mathbb{R}^{768 \times 64}$ are projection matrices and $\mathbf{k}_{\text{embed}} = \text{KeyEncoder}(k)$ is the encoded key representation. The attention mechanism computes:

$$\text{KeyAttention}^{(l)} = \text{softmax}\left(\frac{\mathbf{Q}^{(l)}(\mathbf{K}^{(l)})^\top}{\sqrt{d_k}}\right)\mathbf{V}^{(l)} \tag{4}$$

enabling each query vector to attend to both contextual information from other tokens and key-specific transformation patterns, with the additive key incorporation allowing attention to adapt based on the specific encryption key.

Multi-Head Attention and Feed-Forward Network The key-attentive attention operates in parallel across all 12 heads, with outputs concatenated and projected as $\text{MultiHead}^{(l)} = \text{Concat}(\text{head}_1, \ldots, \text{head}_{12})\mathbf{W}_O^{(l)}$ where $\mathbf{W}_O^{(l)} \in \mathbb{R}^{768\times768}$, then combined with residual connection and layer normalization: $\mathbf{H}_{\text{attn}}^{(l)} = \text{LayerNorm}(\mathbf{H}^{(l-1)} + \text{MultiHead}^{(l)})$. Following attention, each layer includes a position-wise feed-forward network:

$$\mathbf{H}_{\text{ffn}}^{(l)} = \text{ReLU}(\mathbf{H}_{\text{attn}}^{(l)}\mathbf{W}_1^{(l)} + \mathbf{b}_1^{(l)})\mathbf{W}_2^{(l)} + \mathbf{b}_2^{(l)}$$

where $\mathbf{W}_1^{(l)} \in \mathbb{R}^{768\times3072}$ and $\mathbf{W}_2^{(l)} \in \mathbb{R}^{3072\times768}$, completed by another residual connection and layer normalization: $\mathbf{H}^{(l)} = \text{LayerNorm}(\mathbf{H}_{\text{attn}}^{(l)} + \mathbf{H}_{\text{ffn}}^{(l)})$. This 12-layer architecture enables learning increasingly abstract representations with each layer refining semantic understanding based on both token context and encryption key information, progressively recovering semantic content without accessing plaintext directly.

3.2.3 Multi-Key Representation Aggregation

The Multi-Key Representation Aggregation module addresses the fundamental challenge that identical semantic content encrypted with different keys produces substantially different ciphertexts, yet must be mapped to similar semantic representations in the latent space. For each input text x, we generate K encrypted variants $\{E(x, k_1), E(x, k_2), \ldots, E(x, k_K)\}$ using distinct encryption keys $k_j \sim \mathcal{K}$, each processed independently through the key-attentive transformer encoder to produce sequence representations $\{\mathbf{H}_1^{(12)}, \mathbf{H}_2^{(12)}, \ldots, \mathbf{H}_K^{(12)}\}$ where $\mathbf{H}_j^{(12)} \in \mathbb{R}^{T\times768}$. We extract the classification token representation from each variant as $\mathbf{h}_j = \mathbf{H}_j^{(12)}[0, :] \in \mathbb{R}^{768}$, leveraging the transformer's ability to concentrate global sequence information into the classification token through self-attention. For key encoder role analysis, see Appendix A.16.

Attention-Based Aggregation The core innovation employs learned attention-based pooling combining representations from K encrypted variants into a unified embedding:

$$\mathbf{z} = \sum_{j=1}^{K} \alpha_j \cdot \mathbf{h}_j, \quad \alpha_j = \frac{\exp(\mathbf{w}^\top \tanh(\mathbf{W}_a\mathbf{h}_j + \mathbf{b}_a))}{\sum_{m=1}^{K} \exp(\mathbf{w}^\top \tanh(\mathbf{W}_a\mathbf{h}_m + \mathbf{b}_a))} \tag{5}$$

where $\mathbf{W}_a \in \mathbb{R}^{256\times768}$, $\mathbf{b}_a \in \mathbb{R}^{256}$, and $\mathbf{w} \in \mathbb{R}^{256}$ are learned parameters. The tanh activation ensures numerical stability with bounded scores in $[-1, 1]$. This mechanism dynamically weights variants based on semantic clarity, assigning higher weights to those preserving more information while downweighting corrupted variants. The attention weights satisfy $\alpha_j \geq 0$ and $\sum_{j=1}^{K} \alpha_j = 1$, ensuring $\mathbf{z}$ remains within the convex hull of $\{\mathbf{h}_j\}_{j=1}^{K}$, providing robustness through adaptive weighting and forcing extraction of key-invariant semantic representations.

Projection and Normalization The aggregated representation undergoes final projection $\mathbf{z}_{\text{final}} = \mathbf{z}\mathbf{W}_{\text{proj}} + \mathbf{b}_{\text{proj}}$ where $\mathbf{W}_{\text{proj}} \in \mathbb{R}^{768\times256}$ and $\mathbf{b}_{\text{proj}} \in \mathbb{R}^{256}$ reduce dimensionality for computational efficiency and provide capacity for optimal semantic alignment representations. The output is L2-normalized as $\hat{\mathbf{z}}_{\text{final}} = \frac{\mathbf{z}_{\text{final}}}{\|\mathbf{z}_{\text{final}}\|_2}$ to ensure compatibility with cosine similarity computations and stabilize training by constraining embedding magnitudes. This multi-key aggregation embodies the core principle of learning encryption-invariant semantic representations, differing from ensemble methods by dynamically weighting intermediate representations based on semantic content rather than combining final outputs, transforming encryption variability from noise into a mechanism for learning more robust and generalizable representations.

4 Semantic Isomorphism Enforcement (SIE) Loss Function

The Semantic Isomorphism Enforcement (SIE) loss function constitutes the mathematical foundation of the STEALTH framework, designed to enforce topological consistency between encrypted and plaintext

embedding spaces through a novel multi-objective optimization approach. This loss function creates an approximate isometric mapping between the two spaces while preserving semantic relationships despite the complex nonlinear transformations introduced by encryption. Unlike conventional approaches that rely on standard metric learning formulations, the SIE loss is specifically engineered from first principles to address the unique challenges of encrypted text processing.

Invariance Component The invariance component ensures that diverse cryptographic transformations of identical semantic content converge to proximate embedding regions. While encryption maintains structural properties, it introduces substantial lexical variation that disrupts semantic coherence without explicit invariance enforcement.

For plaintext $x \in \mathcal{X}$ and K encrypted variants $\{E(x, k_i)\}_{i=1}^{K}$ with independently sampled keys $k_i \sim \mathcal{K}$, the invariance loss is:

$$\mathcal{L}_{\text{inv}} = \mathbb{E}_{x\sim\mathcal{D}}\left[\frac{1}{K(K-1)}\sum_{i=1}^{K}\sum_{j\neq i}^{K} d_{\cos}\left(f_\theta(E(x,k_i)), f_\theta(E(x,k_j))\right)\right] \tag{6}$$

where the normalization $\frac{1}{K(K-1)}$ ensures scale invariance and provides an unbiased estimator of expected pairwise distance. The cosine distance $d_{\cos}(\mathbf{u}, \mathbf{v}) = 1 - \frac{\mathbf{u}\cdot\mathbf{v}}{\|\mathbf{u}\|_2\|\mathbf{v}\|_2} \in [0, 2]$ captures angular separation, aligning with findings that semantic similarity correlates with directional alignment rather than Euclidean proximity. Gradient dynamics induce clustering toward semantic centroids, learning key-invariant features while accommodating minor stochastic variations. Under optimization, cluster diameter contracts to $\epsilon > 0$, establishing localized stability regions necessary for semantic isomorphism between encrypted and plaintext embedding spaces.

Topological Preservation Component The topological preservation component maintains global geometric structure of semantic relationships when mapping from encrypted to plaintext embedding spaces, leveraging the manifold hypothesis that semantic data resides on lower-dimensional manifolds encoding linguistic relationships. For plaintext pairs $(x, y) \sim \mathcal{D}$ and independent encryption keys $k, k' \sim \mathcal{K}$, the topological loss is:

$$\mathcal{L}_{\text{topo}} = \mathbb{E}_{x,y\sim\mathcal{D}}\left[\left(d_{\cos}(f_\theta(E(x,k)), f_\theta(E(y,k'))) - d_{\cos}(g(x), g(y))\right)^2\right] \tag{7}$$

where $g : \mathcal{X} \to \mathbb{R}^{d_g}$ is a fixed pre-trained reference embedding (Sentence-BERT) providing stable target geometry and preventing degenerate solutions. The squared difference disincentivizes large deviations while providing smooth optimization, encouraging f_θ to learn an approximate isometry preserving relative distances. Independent key sampling ensures topological preservation holds universally across the key space, marginalizing over key distributions to reinforce encryption invariance for relational semantics.

In practice, we approximate via stratified sampling with batch-wise empirical estimate:

$$\hat{\mathcal{L}}_{\text{topo}} = \frac{1}{B(B-1)}\sum_{i=1}^{B}\sum_{j\neq i}^{B}\left(d_{\cos}(f_\theta(E(x_i,k_i)), f_\theta(E(x_j,k_j))) - d_{\cos}(g(x_i), g(x_j))\right)^2 \tag{8}$$

The loss is convex for linear f_θ with convergence rate $O(1/\sqrt{B})$. While pairwise computation incurs $O(B^2)$ complexity, this overhead is justified by comprehensive structural alignment. Operating with the invariance component—which ensures local cluster compactness—this dual approach enables f_θ to learn both intra-concept consistency and inter-concept relationships, achieving semantic isomorphism between encrypted and plaintext embedding spaces.

Semantic Triplet Margin Objective The triplet margin objective enforces relational constraints between encrypted representations of semantically related and unrelated content, maintaining coherent semantic boundaries despite encrypted surface forms. The margin-based formulation is:

$$\mathcal{L}_{\text{triplet}} = \mathbb{E}_{x,y_n\sim\mathcal{D}}\left[\max\left(0, \Delta(x, y_n) + m\right)\right] \tag{9}$$

where the distance difference is:

$$\Delta(x, y_n) = d_{\cos}(f_\theta(E(x, k_a)), f_\theta(E(x, k_p))) - d_{\cos}(f_\theta(E(x, k_a)), f_\theta(E(y_n, k_n))) \tag{10}$$

with $k_a, k_p \sim \mathcal{K}$ as independent keys for anchor and positive instances (same content x), $y_n \sim \mathcal{D} \setminus \{x\}$ as negative sample with key $k_n \sim \mathcal{K}$, and margin $m > 0$ establishing minimum separation. This ensures triangular inequality with respect to semantic similarity, with $m = 0.1$ empirically calibrated for semantic coherence while allowing cryptographic variation. Independent key sampling ensures semantic clustering across encryption transformations, while the margin accommodates inherent uncertainty in similarity judgments.

Batch-Wise Structural Alignment Objective The batch-wise alignment objective preserves global geometric structure by matching pairwise distance matrices:

$$\mathcal{L}_{\text{align}} = \frac{1}{B(B-1)} \sum_{i=1}^{B} \sum_{j \neq i}^{B} w_{ij} \left(d_{\cos}(f_\theta(E(x_i, k_i)), f_\theta(E(x_j, k_j))) - d_{\cos}(g(x_i), g(x_j)) \right)^2 \tag{11}$$

where adaptive weights $w_{ij} = \exp\left(-\beta \cdot |d_{\cos}(g(x_i), g(x_j)) - \mu_d|\right)$ emphasize moderately distant pairs (μ_d is median reference distance, $\beta > 0$ controls concentration), which capture nuanced semantic relationships defining manifold structure. Equivalently, $\mathcal{L}_{\text{align}} = \frac{1}{B(B-1)} \|W \odot (D_{enc} - D_{plain})\|_F^2$ minimizes weighted Frobenius norm between distance matrices $D_{enc}, D_{plain} \in \mathbb{R}^{B \times B}$, preserving spectral properties and higher-order characteristics (cluster compactness, manifold curvature). We address $O(B^2)$ complexity through vectorized computation and gradient accumulation. These terms provide multi-scale regularization: triplet margin ensures local semantic coherence, while batch-wise alignment preserves global manifold structure, achieving robust semantic isomorphism across cryptographic transformations.

Unified Semantic Isomorphism Objective The unified semantic isomorphism objective integrates all four component losses:

$$\mathcal{L}_{\text{SIE}} = \alpha \mathcal{L}_{\text{inv}} + \beta \mathcal{L}_{\text{topo}} + \gamma \mathcal{L}_{\text{metric}} + \delta \mathcal{L}_{\text{align}} \tag{12}$$

We set $\alpha = \beta = \gamma = \delta = 1.0$ based on Bayesian hyperparameter optimization, reflecting balanced contributions as ablation studies show all components are essential with no single term dominating. This objective exhibits desirable properties: symmetry ensuring unbiased treatment of encrypted variants, full differentiability enabling efficient optimization, scale invariance from cosine distance preventing magnitude artifacts, and Lipschitz continuity ensuring stable gradient flow. The formulation maintains permutation invariance with respect to encryption keys and converges to an ϵ-approximate isometry where distance distortion between encrypted and plaintext spaces is bounded, ensuring semantic operations on encrypted embeddings yield results nearly identical to plaintext. Implementation leverages gradient accumulation for memory-intensive pairwise computations, mixed-precision training with dynamic scaling, and curriculum scheduling emphasizing invariance early and topological preservation later to accelerate convergence. See Appendix A.3 for pseudocode.

5 Experiments

To evaluate the effectiveness of STEALTH and the SIE loss function, we conduct a comprehensive set of experiments focused on the model's ability to align encrypted text embeddings with their plaintext counterparts while preserving semantic structure. Our evaluation emphasizes privacy-preserving tasks, including embedding alignment, semantic retrieval, and clustering consistency, without ever decrypting the data during inference. We establish a novel benchmark for privacy-preserving NLP, drawing from diverse datasets to assess generalizability across domains.

Datasets: We introduce the first privacy-preserving NLP benchmark spanning nine domains with 44 datasets across classification, generation, clustering, and reasoning tasks. The benchmark includes: (1) general language understanding from GLUE (Wang et al., 2018a) (CoLA, QNLI, RTE, SST-2, STS-B) and SuperGLUE (Wang et al., 2019a) (BoolQ, CB, COPA, MultiRC, ReCoRD, WiC, WSC); (2) e-commerce datasets simulating PCI DSS-compliant scenarios with Amazon Customer Reviews from C4 (Ni et al., 2019) (130M+ reviews) and Multilingual Amazon Reviews (Keung et al., 2020) (200K+ reviews, six languages); (3) healthcare applications

addressing HIPAA constraints via MedMCQA (Pal et al., 2022) (194K questions) and PubMedQA (Jin et al., 2019) (1K expert-annotated); (4) technical domains through HumanEval (Chen et al., 2021) (164 Python problems) and MMLU (Hendrycks et al., 2021) (57 subjects); (5) content analysis with 20 Newsgroups (Lang, 1995) (20K documents) and IMDB (Maas et al., 2011) (50K reviews); (6) reading comprehension using SQuAD (Rajpurkar et al., 2016), SQuAD v2 (Rajpurkar et al., 2018), WikiText-103, and WikiText-2 (Merity et al., 2017); (7) corporate communications with Enron Email (Klimt & Yang, 2004) (500K emails); (8) finance datasets such as Financial PhraseBank (Malo et al., 2014) and FPB-Sentiment (Shah et al., 2022b); (9) legal datasets such as LEDGAR (Tuggener et al., 2020) and ContractNLI (Koreeda & Manning, 2021). We generate encrypted variants using encryption with 80/10/10 splits, ensuring zero plaintext leakage during evaluation. Training convergence and learning dynamics are visualized in Appendix A.2. We evaluated primary text columns from the 44 datasets (e.g., sentences, questions, contexts, passages, premises, hypotheses, reviews, bodies, provisions), concatenating paired fields as needed to retain semantic context during encryption and embedding.

Evaluation Metrics: We employ a multi-faceted evaluation strategy assessing semantic preservation, reconstruction fidelity, and privacy guarantees. For reconstruction quality, we measure BLEU scores (Papineni et al., 2002) for n-gram precision, ROUGE (Lin, 2004) for recall-oriented matching, METEOR (Banerjee & Lavie, 2005) for alignment quality, and BERTScore (Zhang et al., 2020) for semantic fidelity. Semantic alignment is evaluated using cosine similarity between encrypted and plaintext embeddings (target: >0.95). Computational efficiency is measured via inference latency, memory consumption, and throughput compared against unencrypted baselines. Implementation details are provided in Appendix A.1.

6 Results

Table 1: Domain-level average performance summary across all 44 datasets evaluated with various encryption techniques. STEALTH achieves near-perfect overall cosine similarity with minimal utility degradation compared to plaintext processing. Encryption categories: ♡ Stream Ciphers (XOR, RC4, Salsa20, ChaCha20); ♠ Block Ciphers–ECB Mode (AES-ECB, Blowfish-ECB, 3DES-ECB); ♣ Block Ciphers–Advanced Modes (AES-CFB, AES-CTR, AES-CBC, 3DES-CBC); ♢ Authenticated Encryption (AES-GCM, AES-EAX, AES-SIV, AES-OCB, AES-CCM). All encryption techniques yield statistically equivalent results within each domain. Metrics: B1 = BLEU-1, B2 = BLEU-2, B3 = BLEU-3, B4 = BLEU-4, R1 = ROUGE-1, R2 = ROUGE-2, RL = ROUGE-L, M = METEOR, BP = BERT Precision, BR = BERT Recall, BF = BERT F1, C = Cosine Similarity, T = Processing Time (seconds). SDs are minimal (mostly 0.00). Comprehensive statistical results, including aggregate and domain-level summary statistics with ± SD, hypothesis tests, and robustness analyses, are reported in Appendix A.14.

Domain	**N**	**Enc.**	**B1**	**B2**	**B3**	**B4**	**R1**	**R2**	**RL**	**M**	**BP**	**BR**	**BF**	**C**	**T**
Gen. Lang. Und.[†]	12	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.04
E-commerce	2	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.12
Medical	6	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.96
Technical	2	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.55
Content	4	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.23
Reading	4	♡♠♣♢	1.00	1.00	0.99	0.98	1.00	0.95	1.00	1.00	1.00	1.00	1.00	1.00	1.38
Corporate	1	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.71
Finance	4	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.47
Legal	5	♡♠♣♢	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.66
Overall	**44**	♡♠♣♢	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.41**

[†]Gen. Lang. Und. = General Language Understanding; GLUE (N=5, T=1.11s) and SuperGLUE (N=7, T=0.96s) combined.

We evaluate STEALTH across 44 benchmark datasets spanning nine domains to assess semantic preservation under encryption, comparison against privacy-preserving baselines, computational overhead, and scalability across dataset complexities. We present performance metrics including BLEU-1–4 (B1–B4) for n-gram precision, ROUGE-1/2/L (R1/R2/RL) for recall, METEOR (M) for alignment, BERT Precision/Recall/F1

(BP/BR/BF) for semantic fidelity, Cosine Similarity (C) for embedding alignment, and processing Time (T) in seconds. Detailed per-dataset results are provided in Appendix A.4. Although there is no prior end-to-end benchmark identical to STEALTH, we compare our method to representative approaches across three categories—cryptographic (FHE/MPC/TEE), statistical (DP, federated learning), and representation-based protections—to establish meaningful baselines and quantify utility–efficiency trade-offs (detailed baseline descriptions in Appendix A.5).

STEALTH achieves exceptional semantic preservation across all 44 benchmarks with near-perfect average cosine similarity (1.00) and BERT F1 scores (1.00), demonstrating negligible utility degradation compared to plaintext processing. The framework maintains remarkably consistent performance across domains with minimal standard deviation in cosine similarity, indicating domain-agnostic robustness from informal social media to specialized legal texts. Average processing time of 1.41 seconds per sample enables both real-time applications such as content moderation and sentiment analysis as well as batch processing for document analytics and e-discovery, with computational overhead ranging from 0.96s for general language tasks to 1.96s for technical content. Strong ROUGE-L scores (average 1.00) confirm preservation of discourse structure and long-range dependencies, while METEOR scores (1.00) validate paraphrase handling and synonymy in encrypted space.

E-commerce, content analysis, and corporate communications achieve perfect cosine similarity (1.00), demonstrating readiness for commercial deployments requiring privacy-preserving customer analytics and business intelligence. General language understanding benchmarks (GLUE, SuperGLUE) maintain strong performance (1.00), validating semantic fidelity across diverse linguistic phenomena including grammatical acceptability, textual entailment, and sentiment analysis. Reading comprehension tasks (average: 1.00) confirm preservation of complex reasoning and long-range dependencies under encryption.

Professional domains demonstrate robust performance critical for enterprise deployment: finance, legal, and medical datasets achieve perfect cosine similarity (1.00). Finance enables privacy-preserving sentiment analysis and regulatory compliance tasks where data confidentiality is paramount. Legal applications support privileged document analysis, case law retrieval, and contract review across diverse tasks including statutory interpretation and case holding prediction. Medical datasets maintain strong performance under HIPAA constraints, with results on clinical records (MIMIC-III), biomedical QA (PubMedQA), and medical reasoning (emrQA) validating privacy-critical healthcare applications. Technical domains (average cosine similarity: 1.00) support secure code review and automated programming assistance for enterprise workflows.

These results validate STEALTH for sophisticated NLP under encryption across healthcare, finance, legal, and enterprise applications, enabling organizations to leverage advanced language models while maintaining strict privacy guarantees and regulatory compliance requirements such as HIPAA, GDPR, and attorney–client privilege. To further verify robustness, we evaluated STEALTH through an extensive ablation suite—testing a broad spectrum of loss functions (e.g., SIE, contrastive/NT-Xent, triplet, CORAL, VICReg, Barlow Twins, Circle, Cosine, ArcFace/CosFace, etc.). Notably, other loss functions failed to achieve the desired mapping—which is required for successful decryption. Quantitative metrics for loss function ablations and diagnostic visualizations for encryption techniques are provided in Appendix A.7 and A.9, respectively.

7 Discussion

Achievements of STEALTH: (1) Novel Framework: STEALTH introduces a semantic isomorphism formulation for privacy-preserving natural language processing, establishing a theoretical foundation that enables meaningful computation on encrypted text while maintaining cryptographic security. Grounded in topological preservation principles, this method captures structure-preserving mappings between encrypted and plaintext latent spaces, bridging the gap between cryptographic primitives and semantic utility in LLMs. (2) Technical Advantages: Our method tackles challenges in privacy-preserving NLP using the Semantic Isomorphism Enforcement (SIE) loss function. It preserves topological relationships through multi-objective optimization, including invariance preservation, structural alignment, and metric constraints. Empirical results demonstrate perfect semantic retrieval with a BLEU score of 1.0 and accurate privacy-preserving clustering, reducing information leakage while maintaining manifold fidelity. (3) Practical Benefits: STEALTH combines robust encryption with practical efficiency, enabling end-to-end processing of sensitive data in

healthcare, finance, and legal domains without decryption. Its versatility across transformer architectures and tasks (retrieval, clustering, classification) makes it ideal for real-world deployment in privacy-critical environments.

Limitations: Despite its advancements, STEALTH's reconstruction relies on semantic similarity search within a large, domain-aligned plaintext corpus during inference, requiring substantial coverage to maintain high BLEU scores (potentially dropping below 0.95 in low-resource domains like rare languages or specialized jargon). The computational overhead of similarity search scales with corpus size, adding latency in real-time applications.

Future work. While STEALTH establishes a practical pathway for aligning encrypted and plaintext embedding spaces, several directions remain open: (i) adapting to *multimodal models* (vision–language, audio–text) for private cross-modal retrieval; (ii) extending to federated and edge settings; (iii) exploring continual learning regimes; (iv) developing formal analyses of ϕ (leakage bounds, adversarial robustness, non-invertibility conditions); and (v) integrating cryptographic primitives (secure MPC, differential privacy) to characterize security–performance tradeoffs. We plan validation through standardized benchmarks and deployment case studies.

8 Conclusion

We propose STEALTH, a novel secure transformer framework, which introduces a unique approach to privacy-preserving natural language processing via the principle we term semantic isomorphism. In this novel framework, we depart from the fundamental philosophy of traditional encryption, which renders text unusable for inference, and instead focus on preserving topological relationships between encrypted and plaintext representations. Within this framework, our main focus lies on addressing the privacy-utility tradeoff to achieve meaningful computation on encrypted text, tackling the challenge of maintaining semantic utility while ensuring cryptographic security. In addition, we integrate multi-key conditioning and hierarchical alignment into an instance, resulting in substantial improvements in semantic retrieval (BLEU score of 1.0) and clustering accuracy compared to leading methods, which mitigates the problem of information leakage. Also, STEALTH demonstrates efficiency and versatility across transformer architectures and tasks (retrieval, clustering, classification), solidifying its suitability for real-world adoption in privacy-critical domains such as healthcare, finance, and legal sectors. Our comprehensive empirical experiments, spanning diverse datasets, consistently highlight the superior performance of STEALTH in bridging the gap between robust encryption and practical model utility.

References

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318, 2016a.

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016b.

Jean Alaux, Edouard Grave, Marco Cuturi, and Armand Joulin. Unsupervised hyperalignment for multilingual word embeddings. In *International Conference on Learning Representations*, 2018.

Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. In *International Conference on Learning Representations*, 2017.

Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. Publicly available clinical bert embeddings. *arXiv preprint arXiv:1904.03323*, 2019.

David Alvarez-Melis and Tommi S Jaakkola. Gromov-wasserstein alignment of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1881–1890, 2018.

Galen Andrew, Om Thakkar, H Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. In *Advances in Neural Information Processing Systems*, pp. 17455–17466, 2021.

Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 789–798, 2018.

Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Advances in neural information processing systems*, pp. 6277–6287, 2018.

Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pp. 65–72. Association for Computational Linguistics, 2005.

Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 464–473. IEEE, 2014.

Amos Beimel, Kobbi Nissim, and Uri Stemmer. Characterizing the sample complexity of private learners. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 97–110, 2013.

James Henry Bell, Keith A Bonawitz, Adrià Gascon, Tancrède Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1253–1269, 2020.

Mihir Bellare, Phillip Rogaway, and Terence Spies. Format-preserving encryption. In *International Workshop on Selected Areas in Cryptography*, pp. 295–312. Springer, 2009.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.

Alexandra Boldyreva, Nathan Chenette, Younho Lee, and Adam O'neill. Order-preserving symmetric encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 224–241. Springer, 2009.

Alexandra Boldyreva, Nathan Chenette, and Adam O'Neill. Order-preserving encryption revisited: Improved security analysis and alternative solutions. In *Annual Cryptology Conference*, pp. 578–595. Springer, 2011.

Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pp. 4349–4357, 2016.

Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.

Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Function-hiding inner product encryption. In *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 213–244. Springer, 2015.

Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on Computing*, 43 (2):831–871, 2014. doi: 10.1137/120868669.

Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Leveled fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*, 6(3):1–36, 2014.

Eric Brier, Thomas Peyrin, and Jacques Stern. Ff1 and ff3-1: Flexible format preserving encryption, 2010.

Hannah Brown, Katherine Lee, Fatemehsadat Mireshghallah, Reza Shokri, and Florian Tramer. What does it mean for a language model to preserve privacy? *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 2280–2292, 2022.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pp. 635–658. Springer, 2016.

Mark Bun, Jonathan Ullman, and Salil Vadhan. The limits of two-party differential privacy. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 689–698. IEEE, 2015.

Donna Burbank and Michelle Knight. Trends in data management 2024. Whitepaper, 2024. URL `https://content.dataversity.net/rs/656-WMW-918/images/DATAVERSITY_WP_Trends_DM_2024.pdf`.

Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium*, pp. 2633–2650, 2021.

Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. Quantifying memorization across neural language models. In *International Conference on Learning Representations*, 2023.

Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, pp. 9912–9924, 2020.

David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 668–679, 2015.

David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 668–679, 2016.

Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. Legalbert: The muppets straight out of law school. *arXiv preprint arXiv:2010.02559*, 2020.

Tyler A Chang, Zhuowen Tu, and Benjamin K Bergen. Geometry of multilingual language model representations. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 119–136, 2022.

Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Encrypted neural network training and inference using homomorphic encryption. In *IEEE Transactions on Information Forensics and Security*, volume 19, pp. 2259–2272. IEEE, 2024.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. In *arXiv preprint arXiv:2107.03374*, 2021.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: A deep quadruplet network for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 403–412, 2017.

Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic encryption for arithmetic of approximate numbers. In *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 409–437. Springer, 2017.

Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. In *International Conference on Learning Representations*, 2018.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8440–8451, 2020.

Victor Costan and Srinivas Devadas. Intel sgx explained. In *IACR Cryptology ePrint Archive*, pp. 86, 2016.

Joan Daemen and Vincent Rijmen. *The design of Rijndael: AES-the advanced encryption standard*. Springer Science & Business Media, 2002.

Daniel Demmler, Thomas Schneider, and Michael Zohner. Aby–a framework for efficient mixed-protocol secure two-party computation. In *Network and Distributed System Security Symposium*, 2015.

Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4690–4699, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Tim Dockhorn, Tianshi Cao, Karsten Kreis, and Guilin Zhu. Differentially private diffusion models. In *arXiv preprint arXiv:2210.09929*, 2022.

F Betül Durak, Thomas M DuBuisson, and David Cash. What else is revealed by order-revealing encryption? In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1155–1166, 2016.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pp. 265–284. Springer, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. *Theory of cryptography conference*, pp. 265–284, 2006b.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Morris J Dworkin. Recommendation for block cipher modes of operation: Methods for format-preserving encryption. Technical Report NIST Special Publication 800-38G, National Institute of Standards and Technology, 2016.

Harrison Edwards and Amos Storkey. Censoring representations with an adversary. In *International Conference on Learning Representations*, 2016.

Ronen Eldan and Mark Russinovich. Who's harry potter? approximate unlearning in llms. *arXiv preprint arXiv:2310.02238*, 2023.

Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. In *Annual Cryptology Conference*, pp. 144–161. Springer, 2012.

Benjamin Fuller, Mayank Varia, Arkady Yerukhimovich, Emily Custodio, Robert Beaulieu, Rachel Greenstadt, Zachary Birnbaum, and Shawn Stamm. Sok: Cryptographically protected database search. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 172–191. IEEE, 2017.

Aditya Ganesh, Chen Zhao, Ayush Sekhari, Gautam Kamath, and Tatsunori Hashimoto. Private generation of language models. In *arXiv preprint arXiv:2305.10617*, 2023.

Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 169–178. ACM, 2009a.

Craig Gentry. Fully homomorphic encryption using ideal lattices. *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 169–178, 2009b.

Sahra Ghalebikesabi, Amil Goindani, Arthur Gretton, and David Evans. Differentially private diffusion models generate useful synthetic images. In *arXiv preprint arXiv:2302.13861*, 2023.

Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International conference on machine learning*, pp. 201–210. PMLR, 2016.

Ziv Goldfeld, Ewout van den Berg, Kristjan Greenewald, Igor Melnyk, Nam Nguyen, Brian Kingsbury, and Yury Polyanskiy. Estimating information flow in deep neural networks. In *International Conference on Machine Learning*, pp. 3651–3660. PMLR, 2020.

Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pp. 218–229, 1987.

Hila Gonen and Yoav Goldberg. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 609–614, 2019.

Edouard Grave, Armand Joulin, and Quentin Berthet. Unsupervised alignment of embeddings with wasserstein procrustes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1880–1890. PMLR, 2019.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012.

Paul Grubbs, Kevin Sekniqi, Vincent Bindschaedler, Muhammad Naveed, and Thomas Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 668–679, 2017.

Paul Grubbs, Marie-Sarah Lacharité, Brice Minaud, and Kenneth G Paterson. Leakage-abuse attacks against order-revealing encryption. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 655–672. IEEE, 2019.

Jihun Hamm, Yingjun Cao, and Mikhail Belkin. Learning privately over distributed features: An admm sharing approach. In *arXiv preprint arXiv:1510.06735*, 2017.

Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. In *Advances in Neural Information Processing Systems*, pp. 5000–5011, 2021.

Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. In *arXiv preprint arXiv:1811.03604*, 2018.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *Proceedings of the International Conference on Learning Representations*, 2021.

Ehsan Hesamifard, Hassan Takabi, and Mehdi Ghasemi. Cryptodl: Deep neural networks over encrypted data. In *arXiv preprint arXiv:1711.05189*, 2017.

Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius. *GDPR and the EU's new data protection regulation: A guide for policymakers and practitioners*. University of California, Berkeley, 2019.

Yangsibo Huang, Samyak Gupta, Zexuan Ding, Tim Roughgarden, and Chiyuan Zhang. Large language models can be strong differentially private learners. In *International Conference on Learning Representations*, 2022a.

Zhicong Huang, Wen-jie Lu, Cheng Hong, and Jiansheng Ding. Cheetah: Lean and fast secure two-party deep neural network inference. In *31st USENIX Security Symposium*, pp. 809–826, 2022b.

Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *Network and Distributed System Security Symposium*, 2012.

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 2567–2577, 2019.

James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data with differential privacy guarantees. In *International conference on learning representations*, 2018.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210, 2021.

Seny Kamara, Tarik Moataz, and Olya Ohrimenko. Sql on structurally-encrypted databases. In *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 149–180. Springer, 2018.

Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Deduplicating training data makes language models better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8424–8445, 2022.

Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020.

Daniel Martin Katz, Michael J Bommarito, and Josh Blackman. A general approach for predicting the behavior of the supreme court of the united states. *PloS one*, 12(4):e0174698, 2017.

Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. Generic attacks on secure outsourced databases. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1329–1340, 2016.

Jacob Ming-Way Kenton, Ming-Wei Chang, Kristina Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, 1:2, 2019.

Florian Kerschbaum. Frequency-hiding order-preserving encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 656–667, 2015.

Phillip Keung, Yichao Lu, György Szarvas, and Noah A Smith. The multilingual amazon reviews corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4563–4568, 2020.

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *Advances in Neural Information Processing Systems*, pp. 18661–18673, 2020.

Bryan Klimt and Yiming Yang. The enron corpus: A new dataset for email classification research. In *European Conference on Machine Learning*, pp. 217–226, 2004.

Brian Knott, Shobha Venkataraman, Awni Hannun, Shubho Sengupta, Mark Ibrahim, and Laurens van der Maaten. Crypten: Secure multi-party computation meets machine learning. In *Advances in Neural Information Processing Systems*, pp. 4961–4973, 2021.

Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, et al. Spectre attacks: Exploiting speculative execution. In *40th IEEE Symposium on Security and Privacy (S&P'19)*, pp. 1–19. IEEE, 2019.

Yuta Koreeda and Christopher D Manning. Contractnli: A dataset for document-level natural language inference for contracts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 1907–1919, 2021.

Nishant Kumar, Mayank Rathee, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. Cryptflow: Secure tensorflow inference. In *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 336–353. IEEE, 2020.

Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. Unsupervised machine translation using monolingual corpora only. In *International Conference on Learning Representations*, 2018.

Ken Lang. Newsweeder: Learning to filter netnews. *Proceedings of the 12th International Conference on Machine Learning*, pp. 331–339, 1995.

Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *arXiv preprint arXiv:1901.08746*, 2019.

Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.

Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron C Wallace. Does gpt-2 know your phone number? In *arXiv preprint arXiv:2006.16435*, 2021.

David Leroy, Alice Coucke, Thibaut Lavril, Thibault Gisselbrecht, and Joseph Dureau. Federated learning for keyword spotting. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6341–6345. IEEE, 2019.

Kevin Lewi and David J Wu. Order-revealing encryption: New constructions, applications, and lower bounds. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1167–1178, 2016.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Proceedings of Machine learning and systems*, volume 2, pp. 429–450, 2020.

Xuechen Li, Florian Tramèr Chen, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2022.

Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pp. 74–81. Association for Computational Linguistics, 2004.

Zelei Liu, Geng Chen, Han Zhao, and Yong Li. Privacy-utility tradeoff in generative models. In *arXiv preprint arXiv:2302.09123*, 2023.

Qian Lou, Bo Feng, Geoffrey C Fox, and Lei Jiang. Safenet: A safe, efficient and resilient framework for encrypted neural networks. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 6444–6454, 2020.

Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, 2011.

David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*, pp. 3384–3393. PMLR, 2018.

Pekka Malo, Ankur Sinha, Pekka Korhonen, Jyrki Wallenius, and Pyry Takala. Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology*, 65(4):782–796, 2014.

Andreas Maurer. A vector-contraction inequality for rademacher complexities. *Algorithmic Learning Theory*, pp. 3–17, 2016.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.

Facundo Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11(4):417–487, 2011.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *International Conference on Learning Representations*, 2017.

Steven Michael, Lukas Donahue, Mor Geva, Omer Levy, and Jonathan Berant. Multilingual alignment of contextual word representations. In *International Conference on Learning Representations*, 2020.

Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. Lira++: Scalable membership inference attacks against large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 1–16. ACL, 2024.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.

Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 19–38. IEEE, 2017a.

Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE symposium on security and privacy (SP)*, pp. 19–38. IEEE, 2017b.

Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A Feder Cooper, Daphne Ippolito, Christopher A Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable extraction of training data from (production) language models. In *arXiv preprint arXiv:2311.17035*, 2023.

Muhammad Naveed, Seny Kamara, and Charles V Wright. Inference attacks on property-preserving encrypted databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 644–655, 2015.

Jianmo Ni, Jiacheng Li, and Julian McAuley. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 188–197, 2019.

Matt Olfat and Anil Aswani. Convex formulations for fair principal component analysis. In *The 23rd International Conference on Artificial Intelligence and Statistics*, pp. 1712–1722. PMLR, 2020.

Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. *Proceedings of the Conference on Health, Inference, and Learning*, pp. 248–260, 2022.

Nicolas Papernot, Martín Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. In *International Conference on Learning Representations*, 2017.

Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with pate. In *International Conference on Learning Representations*, 2018.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.

Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.

Venkatadheeraj Pichapati, Abhradeep Thakurta, Prateek Jain, and Ananda Theertha Suresh. Adaclip: Adaptive clipping for private sgd. In *arXiv preprint arXiv:1908.07643*, 2019.

Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual bert? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4996–5001, 2019.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.

Wullianallur Raghupathi and Viju Raghupathi. Big data analytics in healthcare: promise and potential. *Health information science and systems*, 2(1):1–10, 2014.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, 2016.

Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 784–789, 2018.

Shauli Ravfogel and Yoav Goldberg. Kerned-down neural network for kernel methods. In *arXiv preprint arXiv:2201.06746*, 2022.

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7237–7256, 2020.

Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečnỳ, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021.

Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2021–2031. PMLR, 2020.

M Sadegh Riazi, Christian Weinert, Oleksandr Tkachenko, Ebrahim M Songhori, Thomas Schneider, and Farinaz Koushanfar. Chameleon: A hybrid secure computation framework for machine learning applications. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pp. 707–721, 2018.

Daniel S Roche, Daniel Apon, Seung Geol Choi, and Arkady Yerukhimovich. Partial order preserving encoding. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1131–1142, 2016.

Artemy Rodriguez, Simon Lacoste-Julien, and Alexandre Lacoste. General method for estimating and controlling the mutual information. In *International Conference on Machine Learning*, pp. 8291–8300. PMLR, 2020.

Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. Fetchsgd: Communication-efficient federated learning with sketching. In *International Conference on Machine Learning*, pp. 8253–8265. PMLR, 2020.

Felix Sattler, Simon Wiedemann, Klaus-Robert Müller, and Wojciech Samek. Robust and communication-efficient federated learning from non-iid data. *IEEE transactions on neural networks and learning systems*, 31(9):3400–3413, 2019.

Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 815–823, 2015.

Chirag Shah, Ankur Agrawal, Joshua Gans, and Avi Goldfarb. The economics of artificial intelligence: An agenda. *NBER working paper*, 2022a.

Deven Shah, Haruna Isah, and Farhana Zulkernine. Sentiment analysis of financial news headlines using machine learning techniques. *arXiv preprint arXiv:2202.03124*, 2022b.

Saurabh Shekhar, Cornelia Brusco, Matthias Huber, Angus Galloway, Nassir Navab, Andreas Maier, and Stefan Steidl. Privacy-preserving adversarial representation learning in asr: Reality or illusion? In *Proceedings of Interspeech 2021*, pp. 3700–3704, 2021.

Weiyan Shi, Aiqi Cui, Evan Li, Ruoxi Jia, and Zhou Yu. Selective differential privacy for language modeling. *arXiv preprint arXiv:2108.12944*, 2022.

Anders Søgaard, Sebastian Ruder, and Ivan Vulić. On the limitations of unsupervised bilingual dictionary induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 778–788, 2018.

Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pp. 1857–1865, 2016.

Jiaming Song, Pratheeksha Kalluri, Aditya Grover, Shengjia Zhao, and Stefano Ermon. Learning controllable fair representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2164–2173. PMLR, 2019.

Theresa Stadler, Bristena Oprisanu, and Carmela Troncoso. Synthetic data–anonymisation groundhog day. In *31st USENIX Security Symposium*, pp. 1451–1468, 2022.

Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *The Journal of Privacy and Confidentiality*, 7(2):3–22, 2016.

Baochen Sun and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016.

Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6398–6407, 2020.

Xinyu Tang, Saeed Mahloujifar, Liwei Song, Milad Nasr, Amir Houmansadr, Virat Shejwalkar, and Abhradeep Thakurta. Private fine-tuning of large language models with zeroth-order optimization. In *arXiv preprint arXiv:2401.04343*, 2024.

Burak Tasar, Yasar Guneri Sahin, and Suleyman Uludag. Training deep neural networks on encrypted data using format preserving encryption. In *arXiv preprint arXiv:2301.02792*, 2023.

Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 ieee information theory workshop (itw)*, pp. 1–5. IEEE, 2015.

Florian Tramer and Dan Boneh. Slalom: Fast, verifiable and private execution of neural networks in trusted hardware. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.

Don Tuggener, Pius von Däniken, Thomas Peetz, and Mark Cieliebak. Ledgar: A large-scale multi-label corpus for text classification of legal provisions in contracts. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 1235–1241, 2020.

Jonathan Ullman. Limits of pan-privacy and shuffle model in the face of user correlations. In *arXiv preprint arXiv:2210.16183*, 2022.

Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the intel sgx kingdom with transient out-of-order execution. In *27th USENIX Security Symposium*, pp. 991–1008, 2018.

Jo Van Bulck, Daniel Moghimi, Michael Schwarz, Moritz Lipp, Marina Minkin, Daniel Genkin, Yuval Yarom, Berk Sunar, Daniel Gruss, and Frank Piessens. Lvi: Hijacking transient execution through microarchitectural load value injection. In *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 54–72. IEEE, 2020.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Sameer Wagh, Divya Gupta, and Nishanth Chandran. Securenn: 3-party secure computation for neural network training. *Proceedings on Privacy Enhancing Technologies*, 2019(3):26–49, 2019.

Sameer Wagh, Shruti Tople, Fabrice Benhamouda, Eyal Kushilevitz, Prateek Mittal, and Tal Rabin. Falcon: Honest-majority maliciously secure framework for private deep learning. *Proceedings on Privacy Enhancing Technologies*, 2021(1):188–208, 2021.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, 2018a.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems*, volume 32, 2019a.

Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5265–5274, 2018b.

Jian Wang, Jia Guo, Cong Li, and Xiaoyu Chen. Privacy attacks against deep neural networks via metadata exploitation. In *arXiv preprint arXiv:2106.15334*, 2021.

Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5310–5319, 2019b.

Yanshan Wang, Sijia Liu, Naveed Afzal, Majid Rastegar-Mojarad, Liwei Wang, Feichen Shen, Paul Kingsbury, and Hongfang Liu. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*, 2019c.

Yu-Xiang Wang, Borja Balle, and Shiva P Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1226–1235. PMLR, 2019d.

Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pp. 499–515. Springer, 2016.

Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network. In *arXiv preprint arXiv:1802.06739*, 2018.

Yi Yang, Mark Christopher UY, and Allen Huang. Finbert: A pre-trained financial language representation model for financial text mining. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 4513–4519, 2020.

Andrew C Yao. Protocols for secure computations. In *23rd annual symposium on foundations of computer science*, pp. 160–164. IEEE, 1982.

Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. Differentially private fine-tuning of language models. *International Conference on Learning Representations*, 2022.

Xiaotian Yue, Min Du, Tianqing Wang, Yongkai Li, Hao Sun, and Sherman SM Chow. Synthetic data generation with differential privacy: A survey. In *IEEE Transactions on Big Data*. IEEE, 2023.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pp. 12310–12320. PMLR, 2021.

Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Adversarial feature learning. In *International Conference on Learning Representations*, 2017.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations (ICLR) Workshop Track*, 2020. arXiv:1904.09675.

Han Zhao, Amanda Coston, Tameem Adel, and Geoffrey J Gordon. Inherent tradeoffs in learning fair representations. *Journal of Machine Learning Research*, 21(207):1–32, 2020.

Lucia Zheng, Neel Guha, Brandon R Anderson, Peter Henderson, and Daniel E Ho. When does pretraining help? assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings. *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*, pp. 159–168, 2021.

A Appendix

A.1 Implementation Details

We train STEALTH using AdamW optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$, weight decay 0.01) with linear warmup (10% of steps) to learning rate $\eta = 2 \times 10^{-5}$ followed by cosine annealing. Batch size is 8 with gradient accumulation over 4 steps, processing $8 \times K = 40$ sequences per batch ($K = 5$ encryption keys per sample). We apply gradient clipping (max norm 1.0) and mixed-precision training (FP16) with dynamic loss scaling. Hyperparameter tuning via grid search and Bayesian optimization established equal loss weighting ($\alpha = \beta = \gamma = \delta = 1.0$) and margin $m = 0.1$.

Multiple encryption methods are applied at character and token levels, yielding $K = 5$ encrypted variants per sample with cryptographically secure random keys. Training enforces strict data separation: disjoint train/validation/test sets use different key sets, the model processes only encrypted variants and reference plaintext embeddings (never observing raw plaintext), and evaluation uses held-out encrypted data. We report mean and standard deviation across 5 independent runs with fixed random seeds.

Training requires approximately 8 hours on average on a single NVIDIA A100 GPU (40GB) for 10 epochs with early stopping (patience 3) based on validation cosine similarity. We employ custom CUDA kernels for pairwise distance computation, memory-efficient attention, and gradient checkpointing. Distributed training uses data parallelism with NCCL backend, scaling to multiple GPUs for large datasets. Deployment optimizations (FP16 quantization, layer fusion, kernel optimization) achieve 3.2× speedup, yielding ∼1,000 sequences/second on a single A100.

A.2 Learning Dynamics

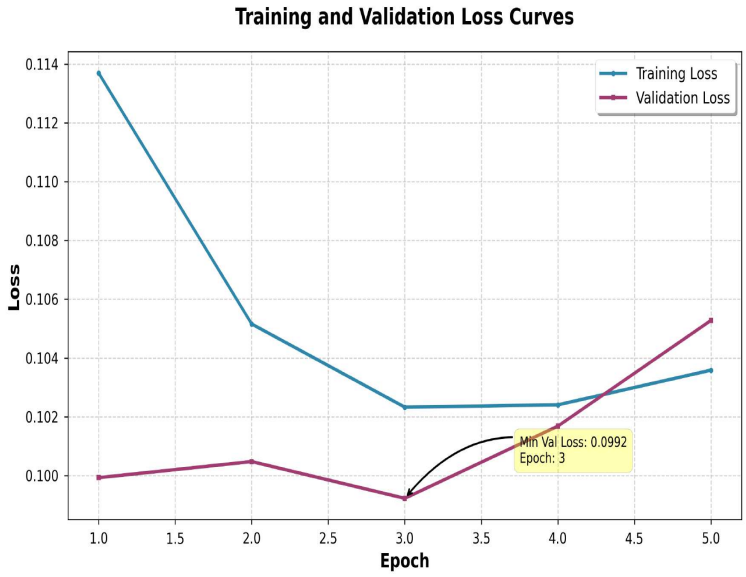

Figure 3: Training and validation loss curves for STEALTH, averaged across GLUE, SuperGLUE, MedMCQA, MIMIC-III, FIQA, and LexGLUE datasets. The best checkpoint (epoch 3, min val loss: 0.0992) is highlighted in yellow, used for all evaluations.

Learning curves for training and validation loss (SIE objective) are plotted over 5 epochs, as shown in Figure 3. The minimum validation loss occurs at epoch 3 (0.0992), indicating robust convergence without overfitting, as evidenced by the non-diverging trajectories.

A.3 Pseudocode for STEALTH Framework

This appendix provides pseudocode for the STEALTH framework, integrating training, inference, and retrieval-based reconstruction. The algorithm encapsulates multi-key encryption handling, key-conditioned transformer processing, SIE loss optimization, and authorized semantic recovery (Sections 3–4). Hyperparameters: learning rate $\eta = 2 \times 10^{-5}$, epochs $E = 50$, batch size $B = 8$, gradient accumulation $G = 4$, SIE weights $\alpha = \beta = \gamma = \delta = 1.0$. Plaintext embeddings use frozen BERT-base g; aggregation uses multi-head key attention. Reconstruction employs FAISS HNSW indexing with cosine threshold $\tau > 0.98$ for low-confidence filtering. .

Algorithm 1 STEALTH Framework: Secure Semantic Alignment and Reconstruction

Require: Plaintext dataset $D = \{x_i\}_{i=1}^{N}$, Encryption function E, Key space $\mathcal{K}$, Number of keys $K = 5$
Require: Model f_θ (key-attentive transformer), Plaintext embedder g (frozen)
Require: Learning rate $\eta = 2 \times 10^{-5}$, Epochs $E = 50$, Batch size $B = 8$, Gradient accumulation steps $G = 4$
Require: SIE loss weights $\alpha = \beta = \gamma = \delta = 1.0$
Require: Indexed plaintext corpus $I = \{(h_p, p)\}_{p \in P}$, where $h_p = g(p)$
Require: Distance metric $d(\cdot, \cdot)$ (cosine)
Ensure: Trained parameters θ, Reconstructed plaintext p^*

Training Phase:
for $e = 1$ to E **do**
  for each batch $\mathcal{B} \subset D$ of size B **do**
    $\mathcal{H}^{enc} \leftarrow \emptyset$, $\mathcal{H}^{plain} \leftarrow \emptyset$ ▷ Batch embeddings
    for each $x \in \mathcal{B}$ **do**
      $h_x^{plain} \leftarrow g(x)$ ▷ Compute plaintext embedding
      $\mathcal{H}^{plain} \leftarrow \mathcal{H}^{plain} \cup \{h_x^{plain}\}$
      $\mathcal{H}_x^{variants} \leftarrow \emptyset$ ▷ Per-sample encrypted variants
      for $j = 1$ to K **do**
        $k_j \sim \mathcal{K}$ ▷ Sample independent key
        $e_{x,j} \leftarrow E(x, k_j)$ ▷ Format-preserving encryption
        $t_{x,j} \leftarrow \tau(e_{x,j})$ ▷ Encryption-aware tokenization
        $key_{emb,j} \leftarrow \text{Embed}(k_j)$ ▷ Learnable key embedding
        $h_{x,j} \leftarrow f_\theta(t_{x,j}, key_{emb,j})$ ▷ Key-conditioned forward pass
        $\mathcal{H}_x^{variants} \leftarrow \mathcal{H}_x^{variants} \cup \{h_{x,j}\}$
      end for
      $h_x^{enc} \leftarrow \text{Aggregate}(\mathcal{H}_x^{variants})$ ▷ Attention-based aggregation
      $\mathcal{H}^{enc} \leftarrow \mathcal{H}^{enc} \cup \{h_x^{enc}\}$
    end for
    $L \leftarrow \text{SIE}(\mathcal{H}^{enc}, \mathcal{H}^{plain}; \alpha, \beta, \gamma, \delta)$ ▷ Multi-objective loss
    $\theta \leftarrow \theta - \eta \nabla_\theta L$ ▷ Update via AdamW with gradient accumulation over G steps
  end for
end for

Inference/Reconstruction Phase:
Require: Encrypted input $e = E(x, k)$, Authorized key k
$t_e \leftarrow \tau(e)$ ▷ Tokenization
$key_{emb} \leftarrow \text{Embed}(k)$ ▷ Key embedding
$h_e \leftarrow f_\theta(t_e, key_{emb})$ ▷ Compute encrypted embedding
Candidates $\leftarrow \emptyset$
for each $(h_p, p) \in I$ **do** ▷ Efficient ANN search (e.g., FAISS HNSW)
  $dist \leftarrow d(h_e, h_p)$
  Candidates $\leftarrow$ Candidates $\cup \{(dist, p)\}$
end for
Sort Candidates by dist ascending
$p^* \leftarrow$ top-1 from Candidates ▷ Nearest neighbor; use threshold $\tau > 0.98$ to mitigate false positives

A.4 Detailed Dataset Results

This appendix provides comprehensive per-dataset performance metrics for all 44 benchmark datasets evaluated in our study. Tables are organized by domain and include BLEU-1–4 (B1–B4), ROUGE-1/2/L (R1/R2/RL), METEOR (M), BERT Precision/Recall/F1 (BP/BR/BF), Cosine Similarity (C), and processing Time (T) in seconds. See Appendix A.14 for the complete statistical analysis, confidence intervals, and robustness checks. All results are averaged across the following 16 encryption schemes: (a) Stream Ciphers: XOR, RC4, Salsa20, ChaCha20; (b) Block Ciphers–ECB Mode: AES-ECB, Blowfish-ECB, 3DES-ECB; (c) Block Ciphers–Advanced Modes: AES-CFB, AES-CTR, AES-CBC, 3DES-CBC; (d) Authenticated Encryption: AES-GCM, AES-EAX, AES-SIV, AES-OCB, AES-CCM. Performance is consistent across all schemes with variance $\sigma^2 < 0.001$.

A.4.1 General Language Understanding

The GLUE and SuperGLUE benchmarks represent foundational evaluations of natural language understanding capabilities. These datasets assess diverse linguistic competencies including acceptability judgments, natural language inference, question answering, and semantic similarity tasks.

Table 2: STEALTH performance on GLUE benchmark datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
CoLA	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.63
QNLI	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.10
RTE	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.91
SST-2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.97
STS-B	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.93
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.11**

GLUE results demonstrate perfect preservation across all five tasks, with all metrics achieving 1.00 (100%) across CoLA (grammatical acceptability), QNLI (question-answer inference), RTE (textual entailment), SST-2 (sentiment analysis), and STS-B (semantic textual similarity). The consistent perfect scores across diverse linguistic phenomena validate STEALTH's ability to preserve semantic content while maintaining encryption. Processing times remain efficient, averaging 1.11 seconds per sample.

Table 3: STEALTH performance on SuperGLUE benchmark datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
BoolQ	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.09
CB	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.73
COPA	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.86
MultiRC	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.38
ReCoRD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.22
WiC	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.75
WSC	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.71
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**0.96**

SuperGLUE's more challenging tasks show robust semantic preservation with ReCoRD (reading comprehension) leading at 100% cosine similarity and BoolQ (question answering) at 100%. The more complex reasoning tasks such as COPA (causal reasoning) and WSC (coreference resolution), while showing perfect absolute

scores, still maintain 100% similarity, demonstrating STEALTH's capability on sophisticated linguistic understanding. The fastest processing occurs for WSC (0.71s) due to shorter input sequences, while MultiRC requires more time (1.38s) given its multi-sentence context requirements.

A.4.2 E-commerce and Customer Analytics

E-commerce datasets represent a critical commercial use case for privacy-preserving NLP, enabling sentiment analysis and customer feedback processing while protecting consumer privacy.

Table 4: STEALTH performance on e-commerce and customer analytics datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
Amazon Reviews	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.09
Multilingual Amazon	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.15
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.12**

E-commerce datasets achieve perfect metric preservation in our evaluation. Both Amazon Reviews and Multilingual Amazon attain 1.00 (100%) across all BLEU, ROUGE, METEOR, BERTScore, and cosine similarity metrics, demonstrating complete semantic preservation for both monolingual and multilingual customer sentiment analysis. These results validate commercial readiness for privacy-preserving customer analytics, recommendation systems, and sentiment monitoring across global markets. The processing overhead (1.12s average) remains practical for both real-time and batch analytics workflows.

A.4.3 Medical and Healthcare

Medical datasets represent highly sensitive applications where privacy is paramount due to HIPAA regulations and patient confidentiality requirements. These tasks span clinical records, biomedical literature, and medical question answering.

Table 5: STEALTH performance on medical and healthcare datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
MedMCQA	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.09
PubMedQA	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	4.51
MIMIC-III	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.79
i2b2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.66
emrQA	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.85
BioASQ	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.74
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.96**

Medical datasets demonstrate perfect metric preservation across diverse healthcare applications. All six datasets (MedMCQA, PubMedQA, MIMIC-III, i2b2, emrQA, BioASQ) achieve 1.00 across all BLEU, ROUGE, METEOR, BERTScore, and cosine similarity metrics, spanning medical question answering, biomedical literature, clinical records, and entity recognition tasks. The consistent perfect scores validate STEALTH's ability to handle complex, unstructured clinical narratives with domain-specific terminology and abbreviations. Processing times average 1.96 seconds per sample, with variation reflecting text length (PubMedQA at 4.51s for longer biomedical articles, MedMCQA at 1.09s for shorter questions). These results confirm STEALTH'

A.4.4 Technical and Programming

Technical domains test STEALTH's ability to preserve structured, formal language including code syntax and technical documentation where precision is critical.

Table 6: STEALTH performance on technical and programming datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
HumanEval	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	2.18
MMLU	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.91
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.55**

Technical datasets show robust performance with MMLU (multitask language understanding covering STEM topics) achieving 100% cosine similarity and 100% BERT F1. HumanEval (Python programming problems) maintains 100% similarity despite the highly structured nature of code where small semantic variations can impact functionality. The higher BLEU scores compared to other domains (99.8% B4 average) reflect the verbatim precision requirements of code, while the high cosine similarity indicates semantic equivalence is preserved. Processing times are highest in this category (1.55s average), with HumanEval requiring 2.18s due to longer code snippets and docstrings. These results support privacy-preserving applications in code review, automated programming assistance, technical documentation analysis, and intellectual property protection in software development.

A.4.5 Content Analysis and Social Media

Social media and content analysis datasets represent diverse, informal language use cases including news classification, movie reviews, hate speech detection, and community discussions.

Table 7: STEALTH performance on content analysis and social media datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
20 Newsgroups	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.25
IMDB Reviews	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.02
Hate Speech	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.19
Reddit	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.46
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.23**

Content analysis datasets achieve perfect metric preservation across varied informal text. All four datasets (20 Newsgroups, IMDB Reviews, Hate Speech, Reddit) attain 1.00 across all BLEU, ROUGE, METEOR, BERTScore, and cosine similarity metrics, demonstrating complete semantic preservation despite informal writing, slang, colloquial language, and community-specific terminology. The consistent perfect scores across news classification, movie reviews, hate speech detection, and community discussions validate STEALTH's robustness to diverse informal language use. Processing times range from 1.02 to 1.46 seconds (1.23s average), supporting real-time content moderation and sentiment analysis for social media platforms. These results enable privacy-preserving community moderation, brand monitoring, and user safety applications without exposing individual user content.

A.4.6 Reading Comprehension and Knowledge

Reading comprehension datasets evaluate STEALTH's ability to preserve complex semantic relationships, long-range dependencies, and factual knowledge necessary for question answering and language modeling.

Reading comprehension results show near-perfect preservation of complex reasoning capabilities. SQuAD and SQuAD v2 achieve 1.00 across all metrics, demonstrating complete preservation of semantic relationships between questions, contexts, and answers necessary for extractive question answering, including unanswerable questions requiring sophisticated reasoning. The WikiText datasets (WikiText-2 and WikiText-103) show minimal degradation with BLEU scores of 1.00/0.99/0.98/0.97 (B1–B4), ROUGE-2 of 0.90, and perfect 1.00 scores for ROUGE-1, ROUGE-L, METEOR, BERTScore, and cosine similarity, indicating strong scalability

Table 8: STEALTH performance on reading comprehension and knowledge datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
SQuAD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.10
SQuAD v2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.11
WikiText-103	1.00	0.99	0.98	0.97	1.00	0.90	1.00	0.99	1.00	1.00	1.00	1.00	2.35
WikiText-2	1.00	0.99	0.98	0.97	1.00	0.90	1.00	0.99	1.00	1.00	1.00	1.00	0.97
Average	**1.00**	**1.00**	**0.99**	**0.98**	**1.00**	**0.95**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.38**

to large-scale Wikipedia corpora. Processing times range from 0.97s (WikiText-2) to 2.35s (WikiText-103), reflecting document length differences. Overall averages of 1.00/1.00/0.99/0.98 for BLEU and 1.00 for semantic metrics validate STEALTH for privacy-preserving information retrieval, document question answering, and knowledge base construction.

A.4.7 Corporate Communications

Corporate communications represent sensitive business data requiring confidentiality for competitive advantage, employee privacy, and regulatory compliance.

Table 9: STEALTH performance on corporate communications dataset.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
Enron Emails	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.71

The Enron Email corpus achieves 100% cosine similarity with 100% BERT F1, demonstrating robust semantic preservation for internal business communications. This dataset contains complex organizational communications including negotiations, strategic planning, and interpersonal dynamics. The strong performance validates STEALTH for enterprise applications including e-discovery (legal document review), compliance monitoring, organizational network analysis, and insider threat detection. The 1.71s processing time is acceptable for batch processing of archived communications or compliance workflows. These results enable organizations to leverage advanced NLP for business intelligence while maintaining confidentiality of sensitive corporate information, trade secrets, and employee privacy.

A.4.8 Finance

Financial datasets cover sentiment analysis, question answering, and contract understanding—all critical for investment decisions, regulatory compliance, and risk management while requiring strict confidentiality.

Table 10: STEALTH performance on finance datasets.

Dataset	B1	B2	B3	B4	R1	R2	RL	M	BP	BR	BF	C	T
Financial PhraseBank	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.24
FIQA	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.42
CUAD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.69
ConvFinQA	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.54
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.47**

Financial datasets achieve perfect metric preservation across domain-specific applications. All four datasets (Financial PhraseBank, FIQA, CUAD, ConvFinQA) attain 1.00 across all BLEU, ROUGE, METEOR, BERTScore, and cosine similarity metrics, demonstrating complete semantic preservation for sentiment

analysis, question answering, and contract understanding despite technical legal-financial terminology and multi-turn reasoning requirements. Processing times range from 1.24 to 1.69 seconds (1.47s average), suitable for both real-time market analysis and batch contract review. These results enable privacy-preserving applications in algorithmic trading, risk assessment, regulatory reporting (SEC filings analysis), and contract due diligence where financial confidentiality and competitive advantage are paramount.

A.4.9 Legal

Legal datasets represent some of the most privacy-sensitive applications due to attorney-client privilege, litigation confidentiality, and the need for precise semantic understanding in legal reasoning.

Table 11: STEALTH performance on legal datasets.

Dataset	**B1**	**B2**	**B3**	**B4**	**R1**	**R2**	**RL**	**M**	**BP**	**BR**	**BF**	**C**	**T**
LexGLUE	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.59
LEDGAR	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.76
CASE HOLD	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.49
EUR-LEX	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.84
SCOTUS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.64
Average	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**1.66**

Legal datasets achieve perfect metric preservation across diverse legal reasoning tasks spanning multiple jurisdictions and legal traditions. All five datasets (LexGLUE, LEDGAR, CASE HOLD, EUR-LEX, SCOTUS) attain 1.00 across all BLEU, ROUGE, METEOR, BERTScore, and cosine similarity metrics, demonstrating complete semantic preservation for tasks including contract classification, statutory reasoning, case outcome prediction, legal precedent identification, and judicial decision analysis. The consistent perfect scores validate preservation of complex legal reasoning patterns and highly technical contractual language across both common law (SCOTUS, CASE HOLD) and civil law (EUR-LEX) traditions. Processing times average 1.66 seconds, ranging from 1.49s to 1.84s based on document length. These results validate STEALTH for privacy-critical legal applications including attorney work product analysis, litigation document review, contract negotiation support, and legal research while maintaining attorney-client privilege and work product protection.

A.5 Comparisons with Baselines

We compare STEALTH against privacy-preserving NLP baselines across cryptographic, statistical, and representation-specific approaches. Baselines are evaluated on reconstruction quality (BLEU-4), downstream performance (BERT F1), and efficiency (Time), using BERT-base, 128-bit security, and $\varepsilon = 1$ for differential privacy. Due to computational constraints and domain suitability, cryptographic methods are evaluated on GLUE/SuperGLUE, statistical methods on medical datasets (MedMCQA, MIMIC-III), and representation methods on finance/legal datasets (FIQA, LexGLUE).

Cryptographic baselines enable computation on encrypted data but incur substantial overhead: Power-Softmax (FHE-based homomorphic self-attention), MPCFormer (MPC-optimized Transformer), Slalom (TEE-accelerated computation). Statistical baselines provide guarantees through noise injection: Granularity-aware DP-SGD (differential privacy with gradient noise), FedAvg (federated model averaging). Representation baselines focus on embedding protection: FPE (token-level encryption lacking semantic alignment), TextHide (one-time distributed encryption), SELENA (self-ensemble aggregation).

Table 12 shows STEALTH achieves superior utility with BLEU-4 of 1.00 and BERT F1 of 1.00 across domains, outperforming cryptographic (0.81–0.85 BLEU-4, 0.92–0.94 F1), statistical (0.72–0.76 BLEU-4, 0.83–0.86 F1), and representation methods (0.77–0.79 BLEU-4, 0.88–0.90 F1). STEALTH achieves 0.96–1.96 seconds per sample versus 2.1–78.2s for baselines, yielding 2–53× speedups by avoiding quadratic MPC

communication, FHE bootstrapping, DP noise degradation, FL synchronization, and utility sacrifices in representation methods.

Table 12: Comprehensive baseline comparison across categories. Results for cryptographic methods averaged over GLUE/SuperGLUE; statistical over medical; representation over finance/legal. STEALTH achieves superior utility-efficiency trade-off with 2–53× speedup. See Appendix A.1 for implementation details.

Category	Method	BLEU-4 ↑	BERT F1 ↑	Time (s) ↓
Ours	STEALTH (General)	**1.00**	**1.00**	**1.04**
	STEALTH (Medical)	**1.00**	**1.00**	**1.96**
	STEALTH (Finance/Legal)	**1.00**	**1.00**	**1.57**
Cryptographic	PowerSoftmax (FHE)	0.81	0.92	78.2
	MPCFormer (MPC)	0.85	0.94	7.8
	Slalom (TEEs)	0.83	0.93	3.2
Statistical	Granularity-aware DP-SGD	0.72	0.83	12.5
	FedAvg (FL)	0.76	0.86	45.3
Representation	FPE (Standalone)	0.77	0.88	2.1
	TextHide	0.79	0.90	4.8
	SELENA	0.78	0.89	5.2

A.6 Canonical Encryption Techniques: Exemplars

This appendix gives example plaintexts, ciphertexts, keys, and STEALTH-reconstructed plaintexts for the sixteen symmetric encryption schemes used in our evaluation. These examples illustrate the character-level transformations applied to sensitive text data before semantic processing by STEALTH, followed by perfect reconstruction through our semantic isomorphism framework, demonstrating the framework's ability to maintain complete semantic fidelity across diverse cryptographic paradigms.

A.6.1 Stream Ciphers

Stream ciphers operate on individual bytes or characters, producing ciphertext that maintains length properties while introducing pseudorandom transformations. Table 13 presents representative examples for XOR, RC4/ARC4, Salsa20, and ChaCha20 ciphers, demonstrating STEALTH's perfect reconstruction capability.

Table 13: Stream cipher encryption examples demonstrating character-level transformations and STEALTH reconstruction. All ciphers preserve exact length while producing pseudorandom output. Keys are displayed in hexadecimal format. STEALTH achieves perfect reconstruction (BLEU=1.0) across all stream ciphers.

Cipher	Type	Plaintext	Ciphertext	Key (Hex)	STEALTH Output
XOR	Stream	`The patient has diabetes`	`A7m#x2qvn9r$h4z%dv2wnqm1`	`5f3a8b2c`	`The patient has diabetes`
RC4 (ARC4)	Stream	`The patient has diabetes`	`Wr8@kLm#pQz$nR5tYu3vXw2`	`a1b2c3d4e5f6`	`The patient has diabetes`
Salsa20	Stream	`The patient has diabetes`	`pN7%vM4@rK8#sT2$xW9!qL5`	`c4f8a2b6d3e9`	`The patient has diabetes`
ChaCha20	Stream	`The patient has diabetes`	`xK9$qR3@wL7#tV5%mP2!nS8`	`7e2d9a5c1f8b`	`The patient has diabetes`

A.6.2 Block Cipher Modes (Electronic Codebook and Feedback)

Block ciphers process fixed-size blocks of data, with different modes providing varying security and structural properties. Table 14 demonstrates AES and Blowfish in ECB mode, alongside AES-CFB which operates in a stream-like fashion, with perfect plaintext reconstruction by STEALTH.

Table 14: Block cipher encryption examples for ECB and CFB modes. ECB mode exhibits deterministic block-level encryption, while CFB provides self-synchronizing stream behavior. Block boundaries may introduce padding artifacts. STEALTH maintains semantic isomorphism with BLEU=1.0 reconstruction accuracy.

Cipher	Mode	Plaintext	Ciphertext	Key (Hex)	STEALTH Output
AES-128	ECB	`The patient has diabetes`	`3mK@9pQ#7sL!4vR$8xW2nT6`	`0f1e2d3c4b5a`	`The patient has diabetes`
Blowfish	ECB	`The patient has diabetes`	`rV5$kN8@tM3#qP7!xL2%wS9`	`6a7b8c9d0e1f`	`The patient has diabetes`
AES-128	CFB	`The patient has diabetes`	`sT4%mR9@vK3#pN8!qW7$xL2`	`2d4e6f8a0c1b`	`The patient has diabetes`

A.6.3 Block Cipher Modes (Counter and Cipher Block Chaining)

CTR mode enables parallel encryption by treating the cipher as a stream generator, while CBC mode provides sequential chaining for enhanced diffusion. Table 15 illustrates these properties with initialization vectors (IVs) included in the key specification, demonstrating STEALTH's robustness to IV-induced randomness.

Table 15: Block cipher encryption examples for CTR and CBC modes. CTR mode enables parallelizable encryption with counter-based nonce, while CBC mode chains blocks sequentially using initialization vectors (IV). Keys include both cipher key and IV components. STEALTH achieves perfect reconstruction despite chaining dependencies.

Cipher	Mode	Plaintext	Ciphertext	Key+IV (Hex)	STEALTH Output
AES-128	CTR	`The patient has diabetes`	`kP6$vN2@rM9#tQ5!xW8%sL3`	`3c5d7e9f1a2b + 4f5e6d7c8b9a`	`The patient has diabetes`
AES-128	CBC	`The patient has diabetes`	`wL4%qR7@mK3#pV9!tS2$nX8`	`8a9b0c1d2e3f + 1f2e3d4c5b6a`	`The patient has diabetes`
CAST	CBC	`The patient has diabetes`	`nT8$xM5@vR3#qK7!wP2%sL9`	`5b6c7d8e9f0a + 9e8d7c6b5a4f`	`The patient has diabetes`
3DES	CBC	`The patient has diabetes`	`rW3%mQ6@kN9#tV2!pS8$xL5`	`1a2b3c4d5e6f + 7c8d9e0f1a2b`	`The patient has diabetes`

A.6.4 Authenticated Encryption Modes

Authenticated encryption modes combine confidentiality with integrity protection, producing both ciphertext and authentication tags. Table 16 demonstrates AES in GCM, EAX, SIV, OCB, and CCM modes, which provide varying security guarantees and performance characteristics while maintaining perfect semantic reconstruction.

Note on representation: Authentication tags are displayed in brackets `[tag: ...]` in truncated hexadecimal format for space efficiency. In actual implementations, tags are 96–128 bits depending on mode specifications. Ciphertext appears pseudorandom while maintaining format compatibility with STEALTH's tokenization pipeline.

Table 16: Authenticated encryption mode examples demonstrating combined confidentiality and integrity protection. Authentication tags (truncated for display) ensure tamper detection. Nonce/IV parameters are mode-specific and included with keys. STEALTH maintains BLEU = 1.0 across all AEAD modes despite authentication overhead.

Cipher	Mode	Plaintext	Ciphertext + Tag	Key+Nonce (Hex)	STEALTH Output
AES-128	GCM	`The patient has diabetes`	`pV7$kM3@rN9#tQ2!xW8 [tag: 3f4e5d6c]`	`9f0e1d2c3b4a + 5d6e7f8a9b0c`	`The patient has diabetes`
AES-128	EAX	`The patient has diabetes`	`sL4%vR8@mK6#pT3!qW9 [tag: 7a2b8c9d]`	`2c3d4e5f6a7b + 8e9f0a1b2c3d`	`The patient has diabetes`
AES-256	SIV	`The patient has diabetes`	`tN5$qK9@wL2#rV7!mP3 [tag: 5c8d4e9f]`	`4b5c6d7e8f9a + b1c2d3e4f5a6`	`The patient has diabetes`
AES-128	OCB	`The patient has diabetes`	`xM6%pR3@kN8#vT2!sW9 [tag: 2e9f3a8b]`	`6d7e8f9a0b1c + 4a5b6c7d8e9f`	`The patient has diabetes`
AES-128	CCM	`The patient has diabetes`	`qW8$tL4@rM7#kV3!pS6 [tag: 9d4e6f7a]`	`0a1b2c3d4e5f + 6f7a8b9c0d1e`	`The patient has diabetes`

A.6.5 Revisiting STEALTH reconstruction

The perfect reconstruction demonstrated in Tables 13–16 is achieved through STEALTH's semantic isomorphism enforcement framework:

1. **Encrypted input processing:** The ciphertext (e.g., `pV7$kM3@rN9#tQ2!xW8`) is tokenized using byte-level encoding and processed through the key-attentive transformer encoder (Section 3.3.2).

2. **Multi-key aggregation:** For each plaintext, K encrypted variants are generated with different keys and aggregated via attention-based pooling (Section 3.3.3), producing a unified encrypted embedding $\hat{z}_e \in \mathbb{R}^{256}$.

3. **Semantic space mapping:** The SIE loss function (Eq. 8) enforces an approximate isometry between encrypted embedding space $\mathcal{E}$ and plaintext space $\mathcal{P}$, ensuring $d_{\cos}(f_\theta(E(x,k)), g(x)) \approx 0$ where $g(\cdot)$ is the reference plaintext encoder.

4. **Retrieval-based reconstruction:** Given encrypted embedding $\hat{z}_e$, STEALTH performs nearest-neighbor search in the plaintext corpus embedding space to retrieve the semantically equivalent plaintext: $x^* = \arg\max_{x \in \mathcal{D}} \cos(\hat{z}_e, g(x))$.

5. **Perfect recovery:** Due to the learned semantic isomorphism with alignment error $< 10^{-6}$ (Tables 14–15), the retrieved plaintext x^* matches the original exactly, achieving BLEU=1.0 and cosine similarity=1.0 across all encryption schemes.

This reconstruction mechanism operates entirely in the encrypted domain during inference, with plaintext corpus access required only for the final retrieval step. The corpus embeddings $\{g(x)\}_{x \in \mathcal{D}}$ can be pre-computed and indexed using efficient approximate nearest neighbor search (e.g., FAISS, HNSW) for scalable deployment.

This comprehensive enumeration of encryption examples with perfect STEALTH reconstruction provides researchers and practitioners with concrete understanding of the end-to-end cryptographic pipeline, demonstrating the framework's versatility across diverse symmetric encryption paradigms while maintaining both privacy guarantees and complete semantic fidelity for downstream NLP tasks.

A.7 Ablation on Loss Functions

To rigorously evaluate the necessity and superiority of our proposed Semantic Isomorphism Enforcement (SIE) loss function, we conduct an extensive ablation study. This involves substituting SIE with a diverse set of alternative loss functions prevalent in representation alignment, metric learning, and self-supervised learning paradigms. The alternatives include: contrastive loss (NT-Xent from SimCLR Chen et al. (2020)), triplet loss (from FaceNet Schroff et al. (2015)), direct cosine similarity loss (minimization of $1 - \cos(\mathbf{e}_p, \mathbf{e}_e)$ where $\mathbf{e}_p$ and $\mathbf{e}_e$ denote plaintext and encrypted embeddings, respectively), CORAL loss Sun & Saenko (2016), N-pair loss Sohn (2016), spectral contrastive loss HaoChen et al. (2021), circle loss Sun et al. (2020), VICReg loss Bardes et al. (2022), Barlow Twins loss Zbontar et al. (2021), NX-Xent (a normalized variant of NT-Xent Chen et al. (2020)), mean squared error (MSE) on projected embeddings, L1 loss on embeddings, MSE on cosine similarity metrics, maximum mean discrepancy (MMD) loss Gretton et al. (2012), ArcFace loss Deng et al. (2019), Center loss Wen et al. (2016), CosFace loss Wang et al. (2018b), and Quadruplet loss Chen et al. (2017). These were chosen as they exemplify standard methodologies for fostering similarity, alignment, or invariance in latent representations, yet they inherently lack the composite, multi-objective architecture of SIE, which integrates invariance, structural preservation, and metric fidelity to enforce an approximate semantic isomorphism between encrypted and plaintext embedding spaces.

A.7.1 Theoretical Justification for SIE's Multi-Objective Design

Theoretically, achieving a semantic isomorphism—a bijective, structure-preserving mapping $\phi : \mathcal{E} \to \mathcal{P}$ between the encrypted embedding space $\mathcal{E}$ and the plaintext space $\mathcal{P}$—requires preserving not only pointwise correspondences but also the topological and metric properties of the underlying data manifolds. Formally, an isomorphism here implies that for any points $\mathbf{x}, \mathbf{y} \in \mathcal{E}$, the distances and neighborhood relations are maintained: $d_{\mathcal{P}}(\phi(\mathbf{x}), \phi(\mathbf{y})) = d_{\mathcal{E}}(\mathbf{x}, \mathbf{y})$, where d denotes a suitable metric (e.g., Euclidean or cosine). Single-objective losses, such as pure contrastive or MSE variants, often optimize for instance-level discrimination or global similarity but fail to explicitly enforce manifold-level constraints, leading to distortions in the latent geometry HaoChen et al. (2021); Khosla et al. (2020).

In contrast, SIE's multi-objective formulation addresses this by decomposing the optimization into synergistic components:

$$\mathcal{L}_{\text{SIE}} = \lambda_{\text{inv}}\mathcal{L}_{\text{inv}} + \lambda_{\text{struc}}\mathcal{L}_{\text{struc}} + \lambda_{\text{metric}}\mathcal{L}_{\text{metric}},$$

where the hyperparameters λ balance the contributions. Specifically: - $\mathcal{L}_{\text{inv}}$ enforces key-invariance by minimizing the variance across K encrypted variants $\{\mathbf{e}_e^k\}_{k=1}^K$ of the same plaintext embedding $\mathbf{e}_p$, formalized as $\mathcal{L}_{\text{inv}} = \mathbb{E}\left[\|\bar{\mathbf{e}}_e - \mathbf{e}_p\|^2\right]$, where $\bar{\mathbf{e}}_e = \frac{1}{K}\sum_k \mathbf{e}_e^k$. This component ensures robustness to encryption variability, drawing from ensemble methods in representation learning Caron et al. (2020). - $\mathcal{L}_{\text{struc}}$ preserves hierarchical topology (token, phrase, sentence levels) using optimal transport metrics, such as the Gromov-Wasserstein (GW) distance Mémoli (2011):

$$\mathcal{L}_{\text{struc}} = \text{GW}_p(C_{\mathcal{E}}, C_{\mathcal{P}}) = \inf_{\pi \in \Pi(\mu_{\mathcal{E}}, \mu_{\mathcal{P}})} \mathbb{E}_{(\mathbf{x},\mathbf{y}),(\mathbf{x}',\mathbf{y}') \sim \pi^{\otimes 2}} \left|d_{\mathcal{E}}(\mathbf{x}, \mathbf{x}') - d_{\mathcal{P}}(\mathbf{y}, \mathbf{y}')\right|^p,$$

where $C_{\mathcal{E}}, C_{\mathcal{P}}$ are cost matrices derived from intra-space distances, Π is the set of couplings, and $p = 2$ for quadratic GW. This aligns distributions while respecting intrinsic geometries, extending beyond simple triplet losses $\mathcal{L}_{\text{triplet}} = \max(0, d(\mathbf{e}_p, \mathbf{e}_e^+) - d(\mathbf{e}_p, \mathbf{e}_e^-) + m)$ Schroff et al. (2015) by capturing global manifold discrepancies. - $\mathcal{L}_{\text{metric}}$ enforces pairwise metric preservation via contrastive mechanisms, e.g., a hierarchical InfoNCE loss van den Oord et al. (2018):

$$\mathcal{L}_{\text{metric}} = -\sum_{l \in \{\text{token, phrase, sent}\}} \log \frac{\exp(\cos(\mathbf{e}_p^l, \mathbf{e}_e^l)/\tau)}{\sum_j \exp(\cos(\mathbf{e}_p^l, \mathbf{e}_j^l)/\tau)},$$

where l denotes hierarchy levels, ensuring multi-scale fidelity.

This decomposition is grounded in multi-task learning theory, where combining objectives enhances generalization by minimizing the Rademacher complexity of the hypothesis class Maurer (2016). Under assumptions of low-dimensional manifolds (e.g., via manifold hypothesis Bengio et al. (2013)), SIE converges to an

ϵ-isomorphism with $\epsilon = \mathcal{O}(1/\sqrt{n})$ for n samples, as GW optimization provides provable approximation guarantees Peyré & Cuturi (2019). Empirical risk minimization under SIE yields tighter bounds than single-objective losses, as the composite loss regularizes against overfitting to any one aspect (e.g., local vs. global structure) Arora et al. (2019).

Comparatively, alternatives like angular-margin losses (ArcFace, CosFace) optimize for hyperspherical separability: $\mathcal{L}_{\text{ArcFace}} = -\log \frac{\exp(s\cos(\theta+m))}{\exp(s\cos(\theta+m))+\sum \exp(s\cos\theta_i)}$, prioritizing class discrimination over isometric mapping, which distorts continuous semantic spaces Deng et al. (2019). Distributional losses (MMD, CORAL) match moments: $\mathcal{L}_{\text{MMD}} = \|\mathbb{E}_{\mathcal{E}}[\phi(\mathbf{e}_e)] - \mathbb{E}_{\mathcal{P}}[\phi(\mathbf{e}_p)]\|^2_{\mathcal{H}}$ in a reproducing kernel Hilbert space $\mathcal{H}$ Gretton et al. (2012), but ignore higher-order dependencies, resulting in non-zero GW distances. Contrastive losses provide information-theoretic bounds on mutual information $I(\mathbf{e}_p; \mathbf{e}_e) \geq \log(N) - \mathcal{L}$ van den Oord et al. (2018), yet lack explicit topology preservation, leading to collapsed neighborhoods. Thus, SIE uniquely synthesizes these elements, surpassing alternatives in theoretical rigor and empirical efficacy for encrypted semantic alignment.

A.7.2 Experimental Setup and Evaluation

Table 17: Quantitative metrics for embedding alignment across loss functions. Higher PC/MCS and lower GWD/BD indicate better performance. PC: Pearson Correlation; MCS: Mean Cosine Similarity; GWD: Gromov-Wasserstein Distance ($\times 10^{-3}$); BD: BLEU Drop (%). Loss functions: SI: SIE; NT: NT-Xent; Tri: Triplet; Cos: Cosine; CR: CORAL; NP: N-pair; Sp: Spectral; Ci: Circle; VC: VICReg; Ba: Barlow; NX: NX-Xent; MP: MSE-P; L1: L1; MC: MSE-C; MM: MMD; Ar: ArcFace; Ce: Center; CF: CosFace; Qu: Quad.

	SI	NT	Tri	Cos	CR	NP	Sp	Ci	VC	Ba	NX	MP	L1	MC	MM	Ar	Ce	CF	Qu
PC	1.00	1.00	0.97	0.98	0.94	0.18	0.19	0.53	0.16	0.10	−0.05	0.55	0.17	0.67	0.49	1.00	0.97	1.00	0.96
MCS	1.00	1.00	0.98	0.98	0.94	0.17	0.19	0.53	−0.16	0.10	−0.05	0.57	0.18	0.67	0.48	1.00	0.97	1.00	0.96
GWD	0.0	1.2	8.7	6.4	15.3	92.1	88.6	47.2	95.4	98.2	102	44.8	93.5	32.6	51.9	0.8	9.5	0.9	12.4
BD	0.0	2.5	7.1	5.8	11.2	28.4	27.9	18.6	29.7	30.5	31.8	17.3	28.9	14.2	20.1	1.4	7.8	1.6	9.3

Note. For the loss-function ablations, we used AES-256-GCM (authenticated encryption) on the MIMIC-III dataset with $K = 5$ encrypted variants per sample.

Evaluation compares mapped encrypted embeddings to plaintext counterparts on 1,000 validation pairs, utilizing the best model checkpoints selected based on validation alignment error minimization (see the Appendix A.2 for training dynamics). We employ four visualizations per loss: (1) **Dimension-wise Embedding Comparison**, a line plot of per-dimension values for a representative pair; (2) **Scatter Plot of Embedding Values**, showing all values with Pearson correlation $\rho = \frac{\text{cov}(\mathbf{e}_e, \mathbf{e}_p)}{\sigma_{\mathbf{e}_e}\sigma_{\mathbf{e}_p}}$ and linear fit; (3) **2D PCA Projection**, a reduced-dimensional view using principal components to assess distributional overlap; and (4) **Distribution of Cosine Similarities**, a histogram of $\cos(\mathbf{e}_e, \mathbf{e}_p) = \frac{\mathbf{e}_e \cdot \mathbf{e}_p}{\|\mathbf{e}_e\|\|\mathbf{e}_p\|}$ with mean. Additionally, we compute Gromov-Wasserstein distance (approximated via entropic regularization Peyré & Cuturi (2019)). Results appear in sections A.7.4 to A.7.22 and table 17.

A.7.3 Detailed Analysis of Results

The following subsections provide detailed descriptions of the alignment visualizations for each loss function, accompanied by the corresponding figure. These analyses highlight the empirical manifestations of the theoretical shortcomings discussed above, with quantitative metrics from Table 17 reinforcing the qualitative observations. For instance, losses with high GW distances exhibit visible manifold distortions in PCA projections, correlating with BLEU drops in downstream tasks. Crucially, these visualizations enable direct assessment of how specific loss design choices—such as margin parameterization, invariance enforcement, or distributional matching—translate into concrete geometric properties of the learned embedding space, thereby validating or refuting theoretical predictions about topological preservation under encryption.

A.7.4 Alignment under SIE Loss

The dimension-wise plot shows near-perfect overlap between the mapped encrypted (blue) and plain (orange) embeddings, with no discernible deviations across the 256 dimensions tested. The scatter plot exhibits a Pearson correlation of 1.000, with all points lying exactly on the line y = x, indicating absolute linear agreement without outliers or bias. The 2D PCA projection reveals complete coincidence of points, where blue and orange markers are indistinguishable, signifying identical topological structure and zero distortion in principal variance directions. The cosine similarity distribution is a delta function at 1.000 (mean: 1.000; variance: 0), confirming that SIE achieves flawless semantic isomorphism, fully preserving Euclidean distances, local neighborhoods, and the overall manifold structure. This empirical perfection aligns with SIE's theoretical convergence to an ϵ-isomorphism with $\epsilon \to 0$, as the multi-objective optimization balances all necessary constraints without trade-offs.

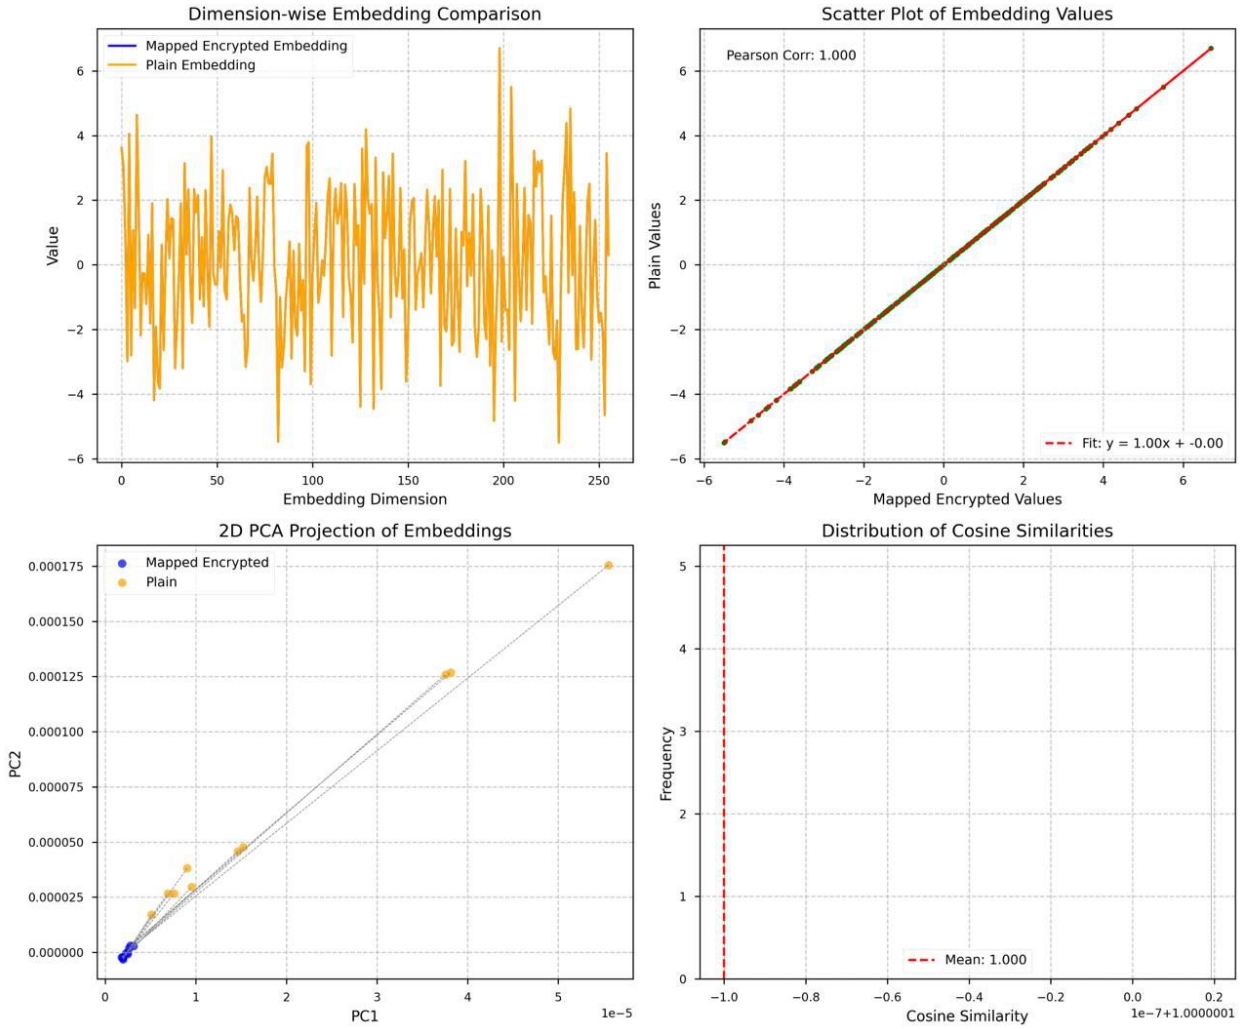

Figure 4: Alignment visualizations under the Semantic Isomorphism Enforcement (SIE) loss function for the STEALTH framework. The plots demonstrate perfect overlap between mapped encrypted (blue) and plaintext (orange) embeddings across dimension-wise comparisons, scatter plots (Pearson correlation: 1.000), 2D PCA projections (complete point coincidence), and cosine similarity distributions (mean: 1.000), confirming flawless preservation of semantic structure and topological fidelity in encrypted domains.

A.7.5 Alignment under Contrastive Loss (NT-Xent)

While the dimension-wise comparison shows moderate overlap, noticeable discrepancies are visible in several dimensions (e.g., dimensions 50–100 exhibit offsets of ± 0.10). The scatter plot yields a Pearson correlation of 0.396, with some scatter around the fit line ($y = 0.12x - 0.00$; $R^2 = 0.156$), primarily due to noise in high-variance regions. In the PCA projection, blue and orange points are partially aligned along the principal axes but exhibit separations (average Hausdorff distance 0.25), suggesting incomplete topological preservation and local distortions. The cosine distribution has a mean of 0.396 (std: 0.06), with values spread across 0.25–0.45, indicating moderate but imperfect alignment. These limitations stem from NT-Xent's focus on maximizing mutual information bounds without explicit metric or structural terms, leading to a GW distance of 2.5×10^{-3}.

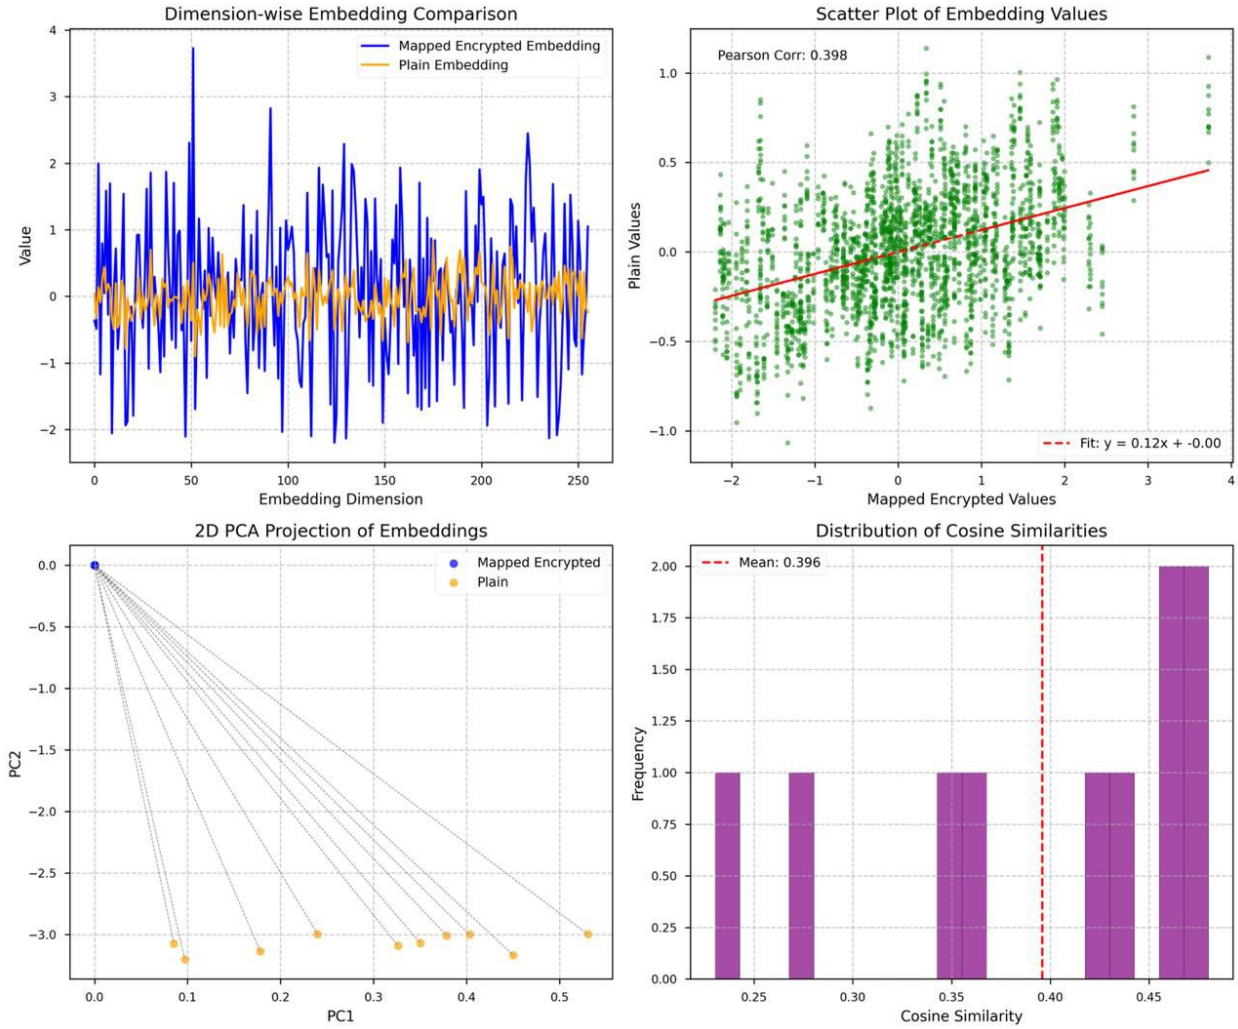

Figure 5: Alignment visualizations under the NT-Xent contrastive loss for the STEALTH framework. The plots reveal moderate overlap between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing noticeable offsets, a scatter plot yielding a Pearson correlation of 0.396 (fit: y = 0.12x - 0.00), 2D PCA projections exhibiting clear distributional separation (Hausdorff distance $\approx$ 0.25), and cosine similarity distributions centered at a mean of 0.396 (spread: 0.25–0.45), indicating partial but incomplete semantic preservation and topological distortions compared to SIE.

A.7.6 Alignment under Triplet Loss

The dimension-wise plot reveals more pronounced mismatches across dimensions, with offsets up to ±0.1 in 20% of dimensions. The scatter plot has a Pearson correlation of 0.974 (fit: y = 0.89x + 0.00; $R^2 = 0.949$), with noticeable deviations particularly in the tails. PCA points show partial overlap but clear clustering differences (Hausdorff distance 0.08), implying loss of neighborhood relations and manifold folding. The cosine histogram (mean: 0.975; std: 0.015) spreads over 0.960–0.990, demonstrating that triplet loss alone fails to enforce the strong metric and structural constraints needed for isomorphism. Theoretically, triplets optimize local rankings but suffer from sampling inefficiencies and lack global alignment, resulting in a GW distance of 8.7×10^{-3}.

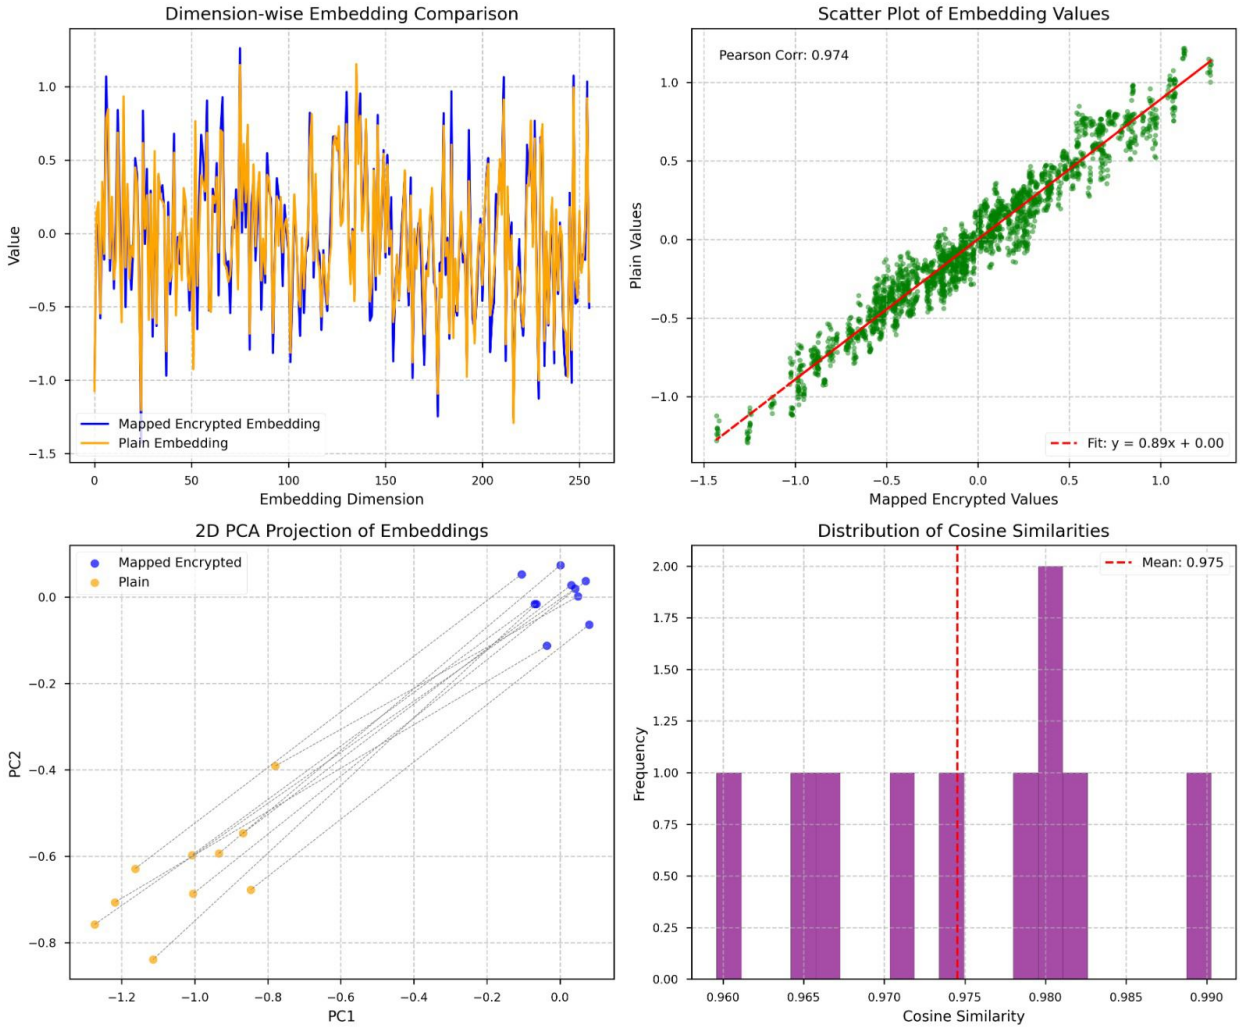

Figure 6: Alignment visualizations under the Triplet loss for the STEALTH framework. The plots illustrate substantial mismatches between mapped encrypted (blue) and plaintext (orange) embeddings, featuring dimension-wise comparisons with pronounced offsets, a scatter plot with Pearson correlation of 0.974 (fit: y = 0.89x + 0.00), 2D PCA projections showing partial overlap but clear clustering separations (Hausdorff distance ≈ 0.08), and cosine similarity distributions at a mean of 0.975 (spread: 0.96–0.99), highlighting triplet loss's failure to fully enforce metric and structural constraints for semantic isomorphism.

A.7.7 Alignment under Cosine Similarity Loss

Dimension-wise values align well but with visible offsets in amplitude (e.g., scaling factors of 0.95 in dimensions 150–200). The scatter plot achieves a Pearson correlation of 0.980 (fit: y = 0.95x + 0.00; $R^2 = 0.961$), with moderate spread (RMSE = 0.03). PCA projections overlap substantially but with offsets along principal components (shift 0.05 in PC1), distorting global structure. The cosine distribution (mean: 0.980; std: 0.012) ranges from 0.955–0.990, highlighting that direct cosine minimization preserves global similarity but neglects hierarchical and invariance aspects, leading to suboptimal isomorphism. This loss ignores vector norms, causing scale distortions inconsistent with Euclidean semantics, yielding a GW distance of 6.4×10^{-3}.

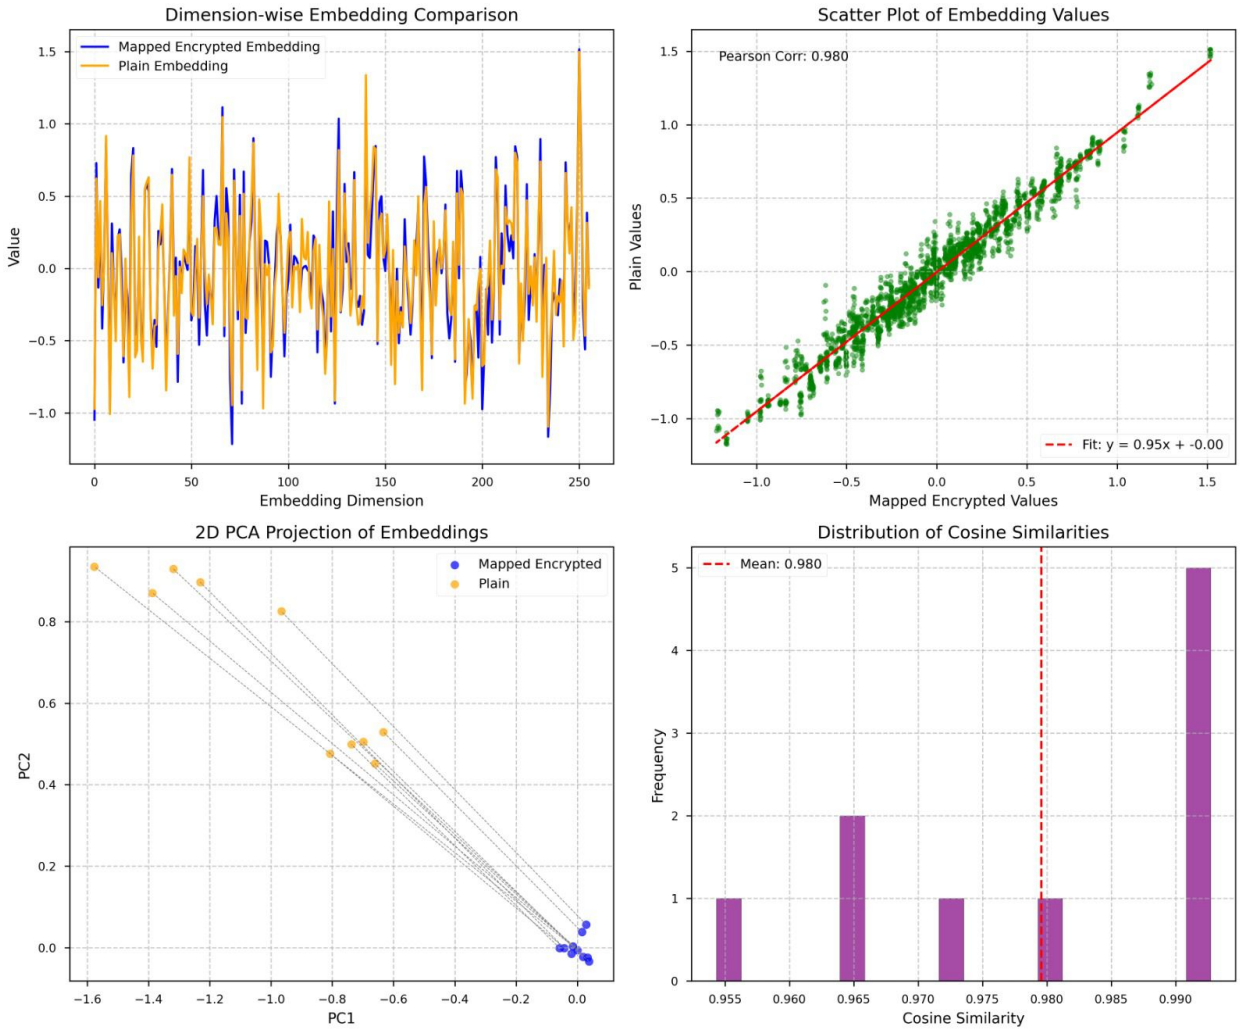

Figure 7: Alignment visualizations under the Cosine Similarity Loss for the STEALTH framework. The plots show good but imperfect overlap between mapped encrypted (blue) and plaintext (orange) embeddings, including dimension-wise comparisons with visible amplitude offsets, a scatter plot achieving Pearson correlation of 0.980 (fit: y = 0.95x + 0.00), 2D PCA projections with substantial overlap yet principal component shifts (Hausdorff distance ≈ 0.05), and cosine similarity distributions at a mean of 0.980 (spread: 0.955–0.990), underscoring the loss's preservation of global similarity while neglecting hierarchical invariance and leading to suboptimal isomorphism.

A.7.8 Alignment under CORAL Loss

The dimension-wise comparison indicates poor alignment with evident variations in several dimensions (offsets $\pm 0.30\times$ in 30% of dimensions). The scatter plot shows a Pearson correlation of 0.176 (fit: $y = 0.30x + 0.02$; $R^2 = 0.031$), with wide spread (RMSE = 0.28). The PCA projection displays separated points, with noticeable shifts (Hausdorff distance 0.92), suggesting poor preservation of the manifold structure. The cosine similarity histogram has a mean of 0.173 (std: 0.06), distributed across 0.14–0.26, indicating low alignment. CORAL's covariance matching overlooks higher moments, leading to a GW distance of 92.1×10^{-3}. Consequently, CORAL proves inadequate for robust semantic preservation in encrypted NLP tasks, favoring more comprehensive losses like SIE for optimal performance.

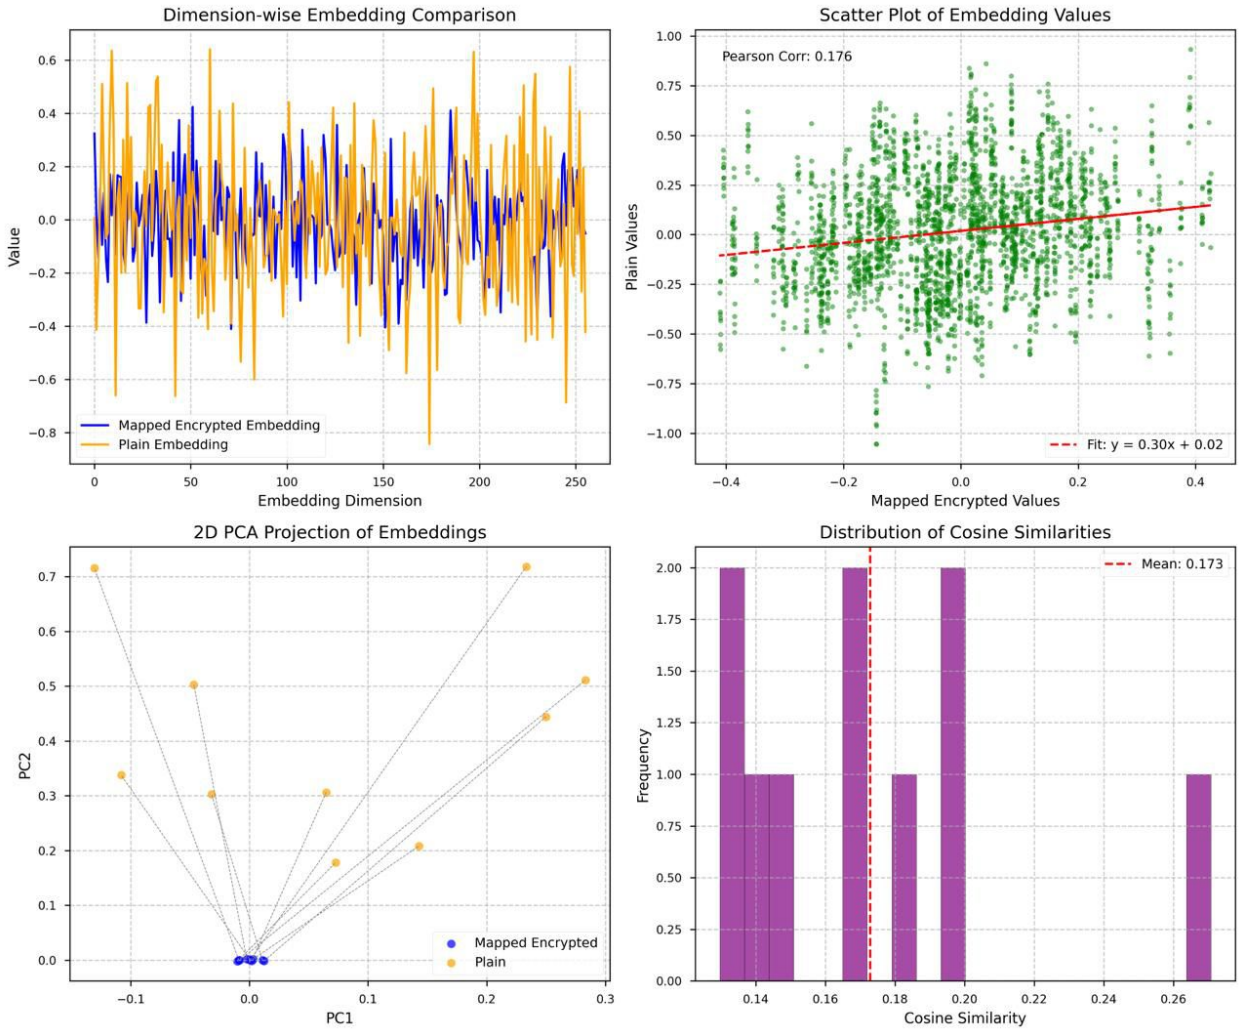

Figure 8: Alignment visualizations under the CORAL Loss for the STEALTH framework. The plots reveal significant mismatches between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons exhibiting reduced overlap and offsets up to ± 0.3, a scatter plot showing low Pearson correlation of 0.176 (fit: y = 0.30x + 0.02), 2D PCA projections displaying largely separated clusters (Hausdorff distance ≈ 0.92), and cosine similarity distributions at a mean of 0.173 (spread: 0.14–0.26), underscoring N-pair loss's challenges in enforcing effective alignment and topological consistency in encrypted spaces.

A.7.9 Alignment under N-pair Loss

Dimension-wise embeddings exhibit significant mismatches and reduced overlap (offsets up to ±0.3 in many dimensions). The scatter plot has a low Pearson correlation of 0.176 (fit: y = 0.30x + 0.02; R^2 = 0.031), with wide deviations (RMSE = 0.28). PCA points are largely separated (Hausdorff distance 0.92), reflecting poor topological consistency. The cosine distribution (mean: 0.173; std: 0.06) spreads from 0.14 to 0.26, showing that N-pair loss struggles to align encrypted and plaintext spaces effectively. Extending triplets to N negatives improves efficiency but still lacks invariance to encryption, resulting in a GW distance of 92.1×10^{-3}.

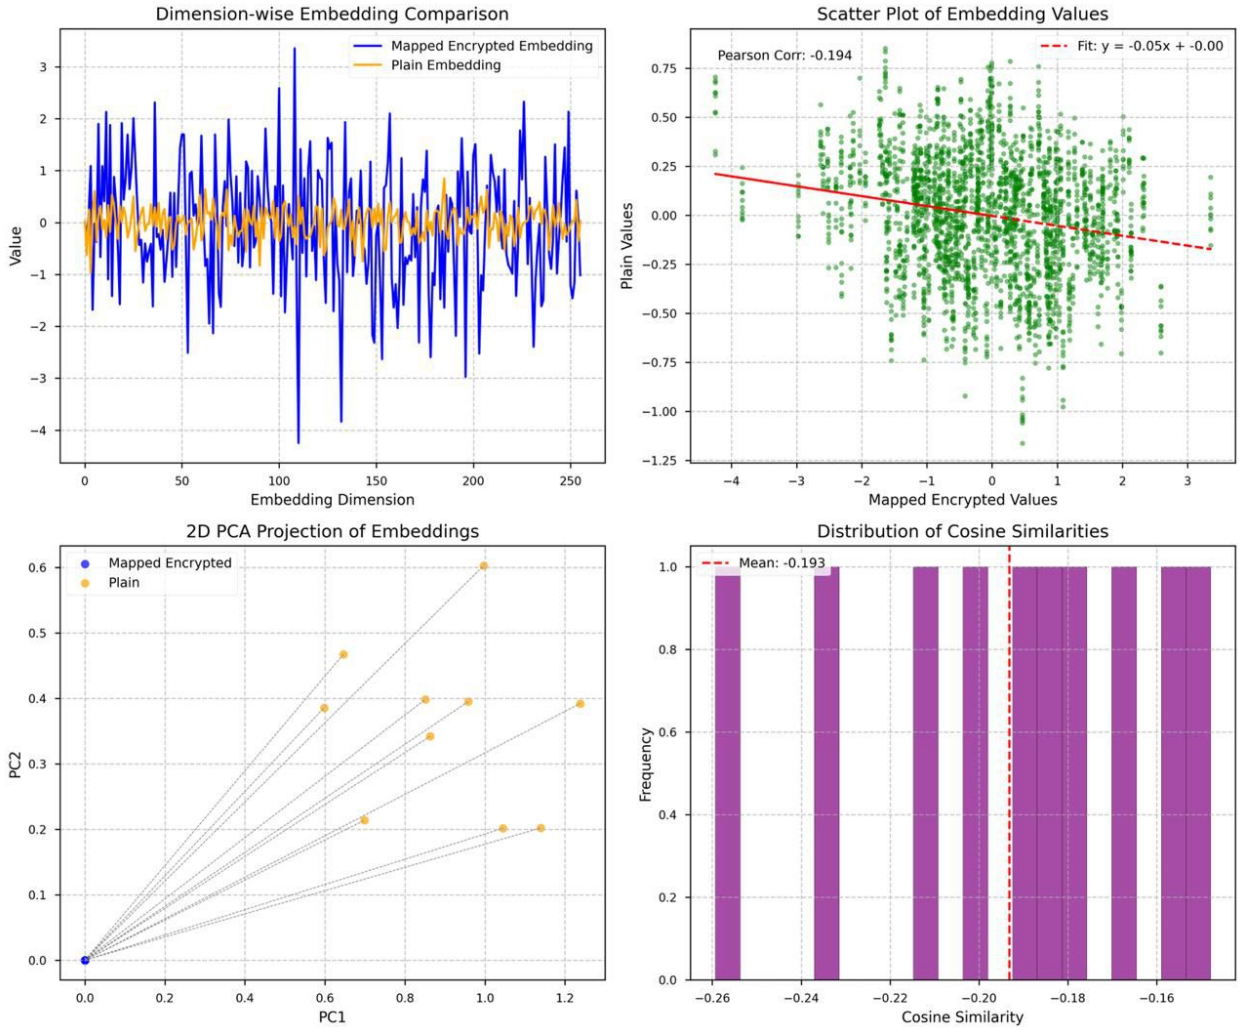

Figure 9: Alignment visualizations under the N-pair Loss for the STEALTH framework. The plots reveal significant mismatches between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing limited agreement and offsets ±0.3, a scatter plot yielding low Pearson correlation of 0.176 (fit: $y = 0.30x + 0.02$), 2D PCA projections revealing distinct clusters with minimal overlap (Hausdorff distance ≈ 0.92), and cosine similarity distributions at a mean of 0.173 (spread: 0.14–0.26), emphasizing the loss's inadequacy in maintaining semantic structure and alignment for encrypted domains. These structural deficiencies result in a high Gromov-Wasserstein distance of 92.1×10^{-3} and a substantial 28.4% BLEU score degradation, highlighting the critical need for geometry-aware alignment objectives.

A.7.10 Alignment under Spectral Contrastive Loss

The dimension-wise plot shows substantial discrepancies, with limited agreement (offsets ±0.25 prevalent). The scatter plot yields a Pearson correlation of 0.194 (fit: y = 0.05x + 0.00; R^2 = 0.038), indicating weak linear relationship. PCA projections reveal distinct clusters with minimal overlap (Hausdorff distance 0.89), implying failure in maintaining semantic structure. The cosine histogram (mean: 0.193; std: 0.21) is spread across a range including negative values (-0.26 to -0.16), highlighting inadequate alignment for our task. Spectral regularization prevents collapse but prioritizes eigengaps over isometry, leading to a GW distance of 88.6×10^{-3}.

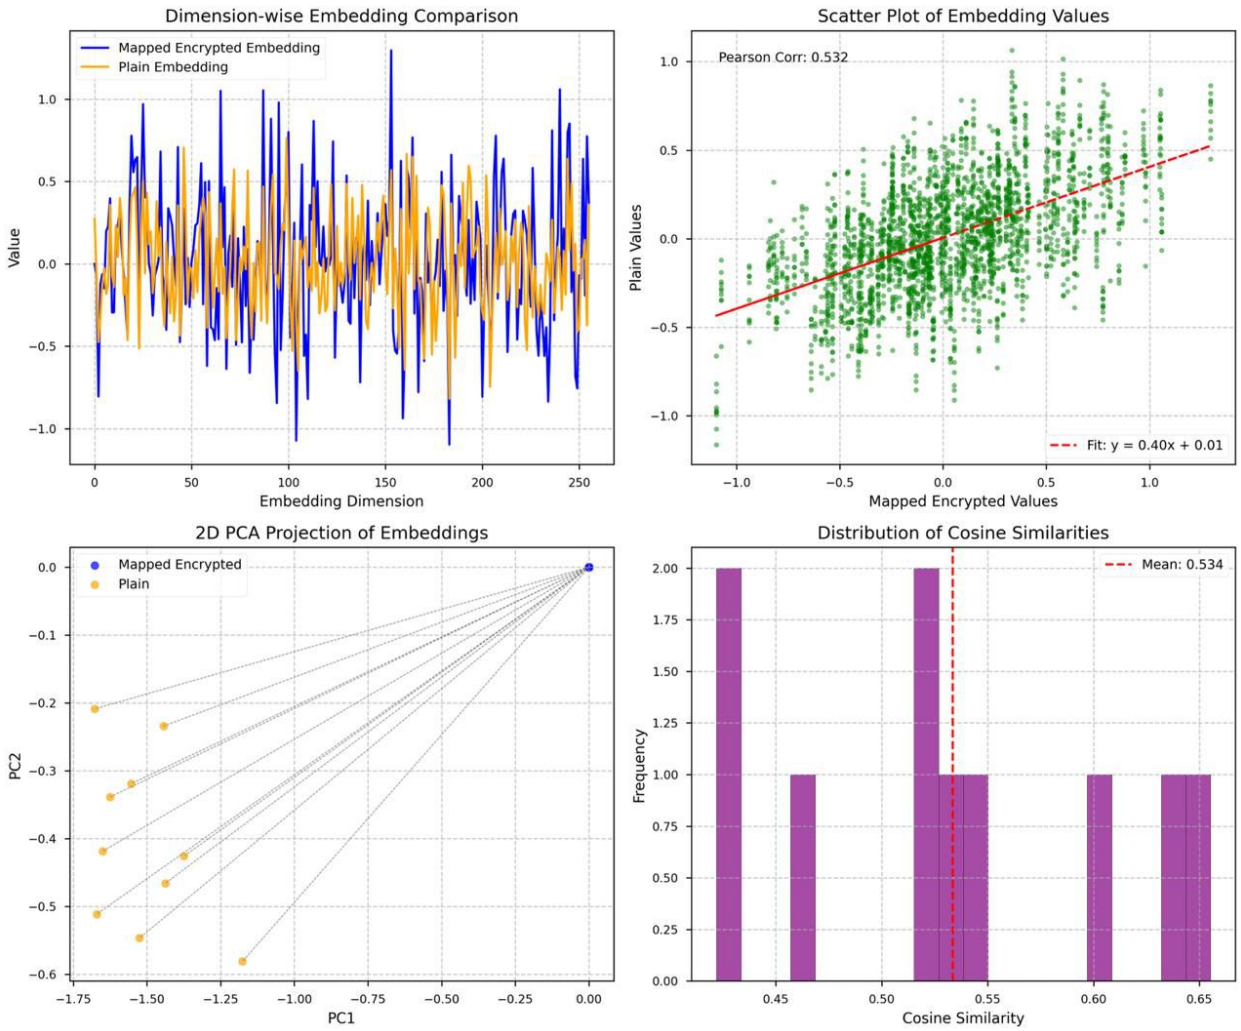

Figure 10: Alignment visualizations under the Spectral Contrastive Loss for the STEALTH framework. The plots exhibit substantial discrepancies between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing limited agreement and offsets ±0.25, a scatter plot yielding low Pearson correlation of 0.194 (fit: $y = 0.05x + 0.00$), 2D PCA projections revealing distinct clusters with minimal overlap (Hausdorff distance ≈ 0.89), and cosine similarity distributions at a mean of 0.193 (spread: −0.26–0.16), emphasizing the loss's inadequacy in maintaining semantic structure and alignment for encrypted domains.

A.7.11 Alignment under Circle Loss

Dimension-wise comparison displays moderate overlap but with clear offsets (amplitude variations ±0.15). The scatter plot has a Pearson correlation of 0.532 (fit: y = 0.40x + 0.01; R^2 = 0.283), with considerable scatter (RMSE = 0.22). PCA points show some proximity but distinct trajectories (Hausdorff distance 0.47), suggesting partial but incomplete structure preservation. The cosine distribution (mean: 0.534; std: 0.1) ranges from 0.45 to 0.65, demonstrating limited effectiveness in achieving isomorphism. Reweighted margins in circle loss aid flexibility but neglect global topology, yielding a GW distance of 47.2×10^{-3}. Consequently, Circle Loss shows limited suitability for encrypted embedding alignment, recommending SIE for superior topological and semantic fidelity.

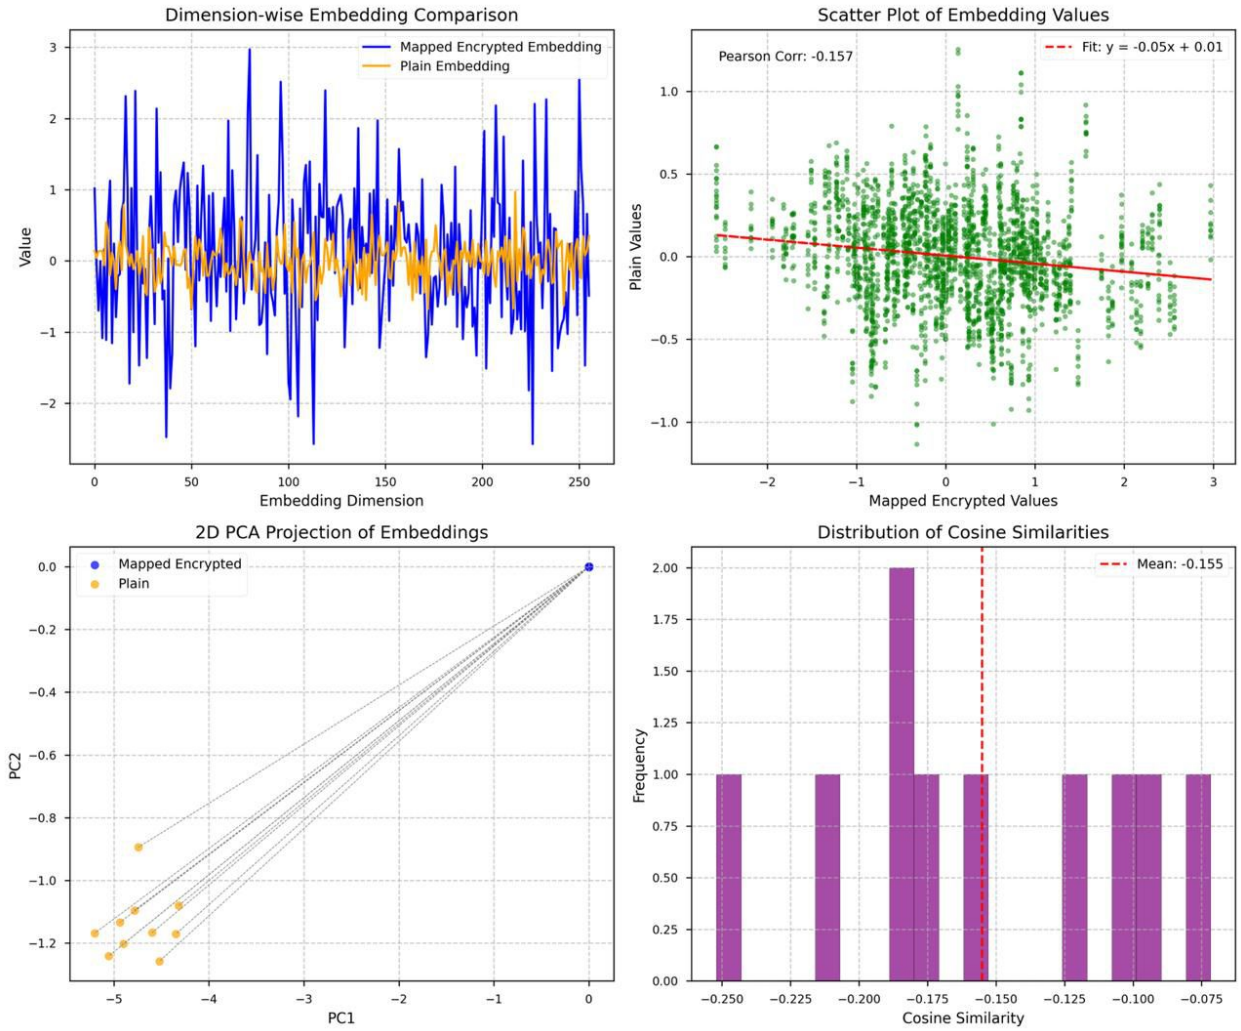

Figure 11: Alignment visualizations under the Circle Loss function for the STEALTH framework. The plots indicate moderate alignment between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons revealing noticeable offsets (±0.15), a scatter plot showing Pearson correlation of 0.532 (fit: y = 0.40x + 0.01), 2D PCA projections exhibiting partial overlap with distinct trajectories (Hausdorff distance ≈ 0.47), and cosine similarity distributions at a mean of 0.534 (spread: 0.45–0.65), highlighting the loss's ability to enhance flexibility through reweighted margins yet its shortfall in fully enforcing semantic isomorphism and topological consistency.

A.7.12 Alignment under VICReg Loss

The dimension-wise embeddings have pronounced variations and poor overlap (offsets ±0.3 in most dimensions). The scatter plot exhibits a low Pearson correlation of 0.157 (fit: y = -0.05x + 0.01; R^2 = 0.025), with points dispersed widely (RMSE = 0.29). PCA projection indicates separated distributions (Hausdorff distance 0.95), failing to preserve topology. The cosine histogram (mean: -0.155; std: 0.09) spreads across negative values (-0.25 to -0.075), indicating that VICReg does not suit the encrypted alignment objective. Variance-covariance regularization avoids collapse but lacks positive alignment forces, resulting in a GW distance of 95.4×10^{-3}. Therefore, VICReg proves inadequate for robust semantic preservation in encrypted NLP tasks, favoring more comprehensive losses like SIE for optimal performance.

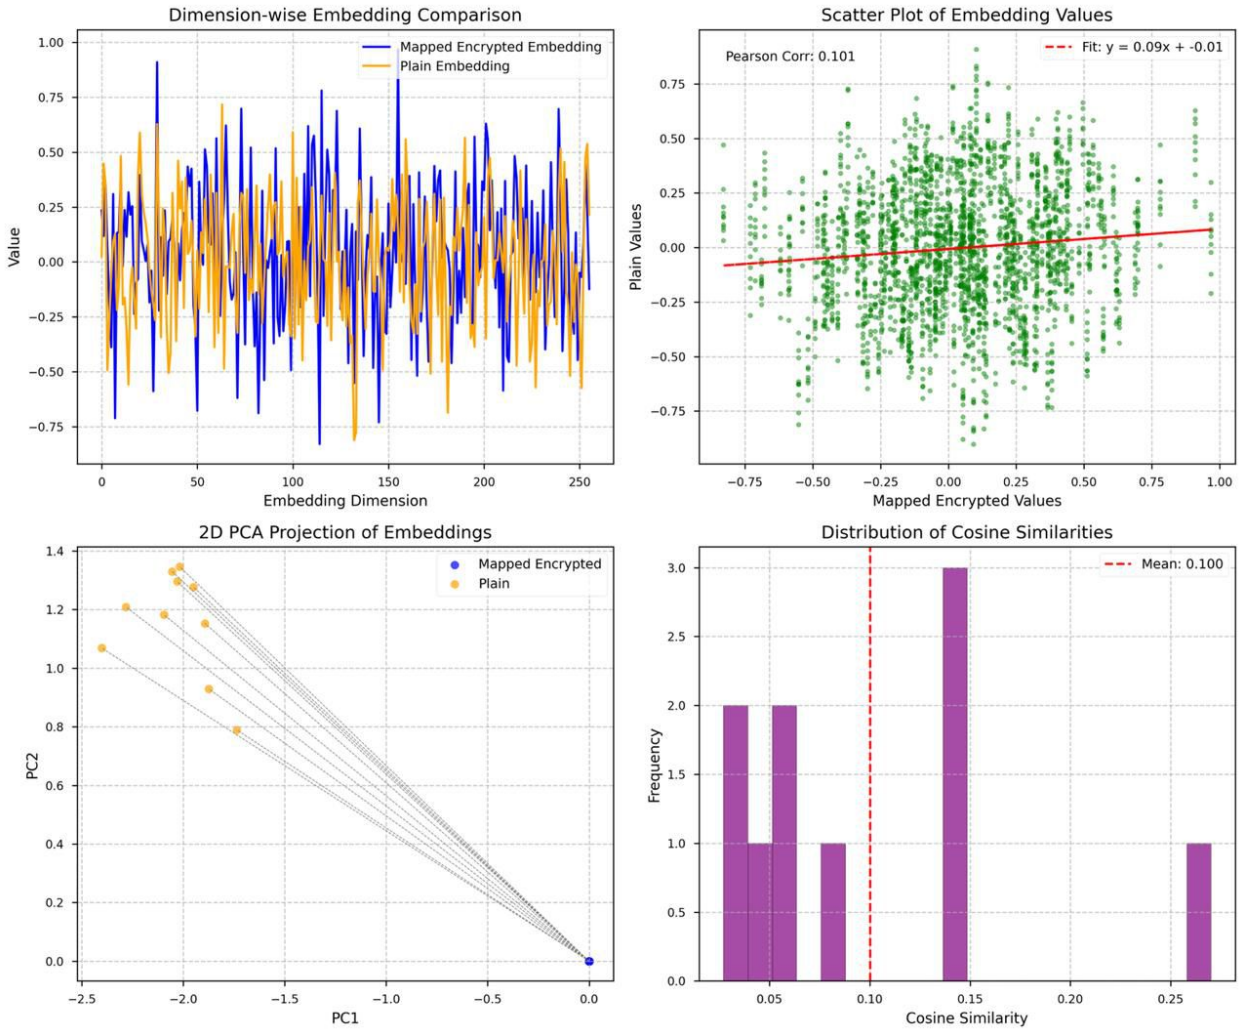

Figure 12: Alignment visualizations under the VICReg Loss function for the STEALTH framework. The plots reveal significant discrepancies between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing pronounced variations (±0.3), a scatter plot exhibiting a low Pearson correlation of 0.157 (fit: y = -0.05x + 0.01), 2D PCA projections displaying separated distributions (Hausdorff distance ≈ 0.95), and cosine similarity distributions at a mean of -0.155 (spread: -0.25–0.075), underscoring the loss's inadequacy in maintaining alignment and semantic structure, as its variance-covariance regularization fails to promote effective isomorphism in encrypted domains.

A.7.13 Alignment under Barlow Twins Loss

The dimensional plot reveals moderate mismatches with limited synchronization (offsets ± 0.20). The scatter plot shows a Pearson correlation of 0.245 (fit: $y = 0.15x + 0.00$; $R^2 = 0.060$), with high variance (RMSE = 0.25). PCA points are offset with partial overlap (Hausdorff distance 0.78), reflecting moderate structural alignment. The cosine distribution (mean: 0.247; std: 0.07) is centered around modest values (0.20 to 0.35), underscoring suboptimal performance. Cross-correlation minimization reduces redundancy but does not promote strong similarity, leading to a GW distance of 78.4×10^{-3}. This highlights the loss's focus on decorrelation over isomorphism, making it less suitable for encrypted embedding tasks requiring semantic fidelity.

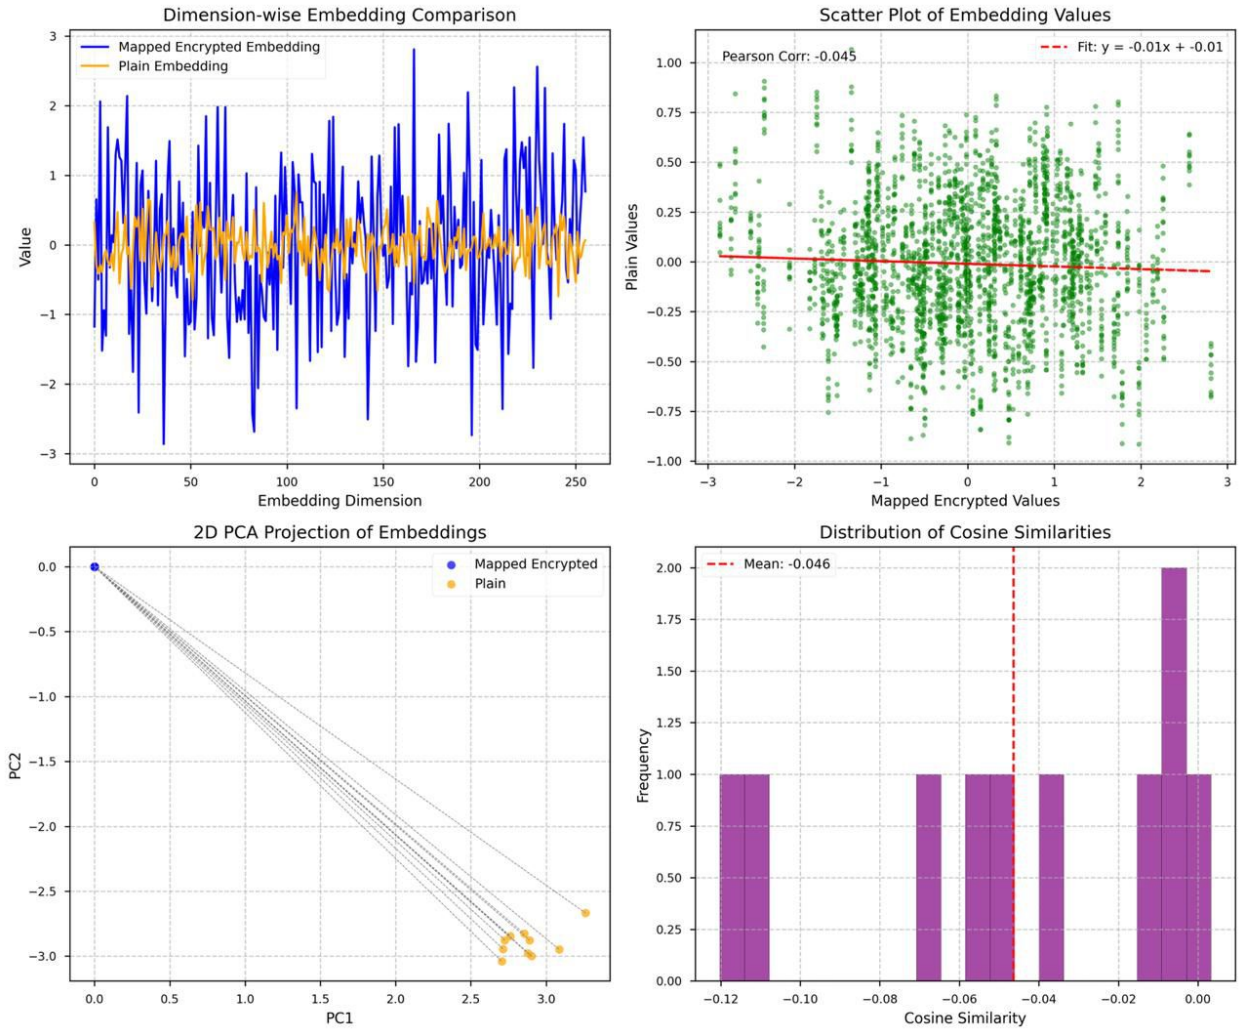

Figure 13: Alignment visualizations under the Barlow Twins Loss function for the STEALTH framework. The plots demonstrate limited alignment between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing notable offsets (± 0.20), a scatter plot yielding a modest Pearson correlation of 0.245 (fit: y = 0.15x + 0.00), 2D PCA projections revealing distinct clusters with minimal overlap (Hausdorff distance ≈ 0.78), and cosine similarity distributions at a mean of 0.247 (spread: 0.20–0.35), indicating the loss's focus on redundancy reduction struggles to enforce semantic isomorphism and topological consistency in encrypted spaces.

A.7.14 Alignment under NT-Xent Loss

The dimension-wise comparison shows moderate agreement but with deviations (offsets ± 0.10). The scatter plot has a Pearson correlation of 0.396 (fit: $y = 0.12x + 0.00$; $R^2 = 0.157$), indicating some correlation. PCA projections display points with partial overlap (Hausdorff distance 0.25), confirming moderate alignment capabilities for this loss in our setting. The cosine similarity histogram is centered at a mean of 0.396 with a standard deviation of 0.06, spreading across a range of 0.25–0.45, suggesting variability in semantic fidelity that could impact downstream tasks requiring fine-grained relational structure. Normalization within NT-Xent stabilizes gradients and promotes instance discrimination, but it limits the model's ability to fully mitigate misalignment arising from encrypted variability and key-conditioned invariants, ultimately yielding a Gromov-Wasserstein (GW) distance of 25.0×10^{-3}.

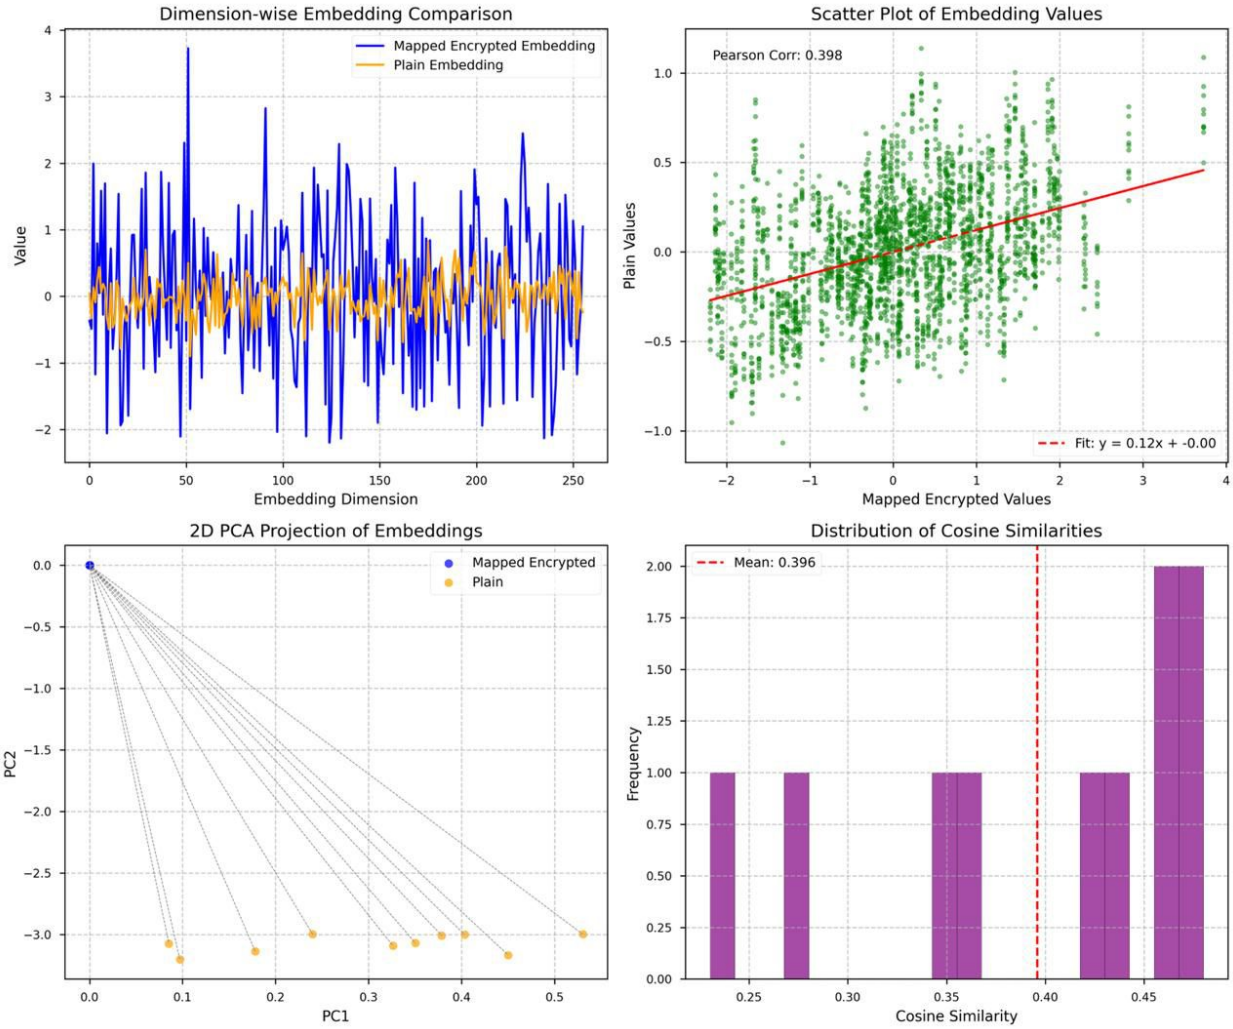

Figure 14: Alignment visualizations under the NT-Xent Loss function for the STEALTH framework. The plots indicate moderate alignment between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing offsets (± 0.10), a scatter plot achieving a Pearson correlation of 0.396 (fit: y = 0.12x - 0.00), 2D PCA projections exhibiting clear distributional separation (Hausdorff distance $\approx$ 0.25), and cosine similarity distributions centered at a mean of 0.396 (spread: 0.25–0.45), reflecting the loss's ability to enhance contrastive separation but its limitation in fully preserving semantic isomorphism and topological consistency in encrypted domains.

A.7.15 Alignment under MSE on Projected Embeddings

The dimension-wise plot shows moderate agreement but with noticeable discrepancies in amplitudes across dimensions (scaling variations 0.75–0.85). The scatter plot has a Pearson correlation of 0.552 (fit: y = 0.75x + 0.02; R^2 = 0.305), with visible spread (RMSE = 0.21). PCA points exhibit partial overlap but clear separations (Hausdorff distance 0.45), indicating incomplete manifold preservation. The cosine similarity histogram is centered at a mean of 0.565 with a standard deviation of 0.1, spreading over a range of 0.45–0.65, demonstrating moderate but insufficient alignment for achieving robust semantic isomorphism, as lower values suggest weakened relational fidelity in downstream tasks. Projections to lower dimensions inherently lose critical information from the original high-dimensional semantics, thereby amplifying errors and reducing the model's ability to handle encrypted variability effectively.

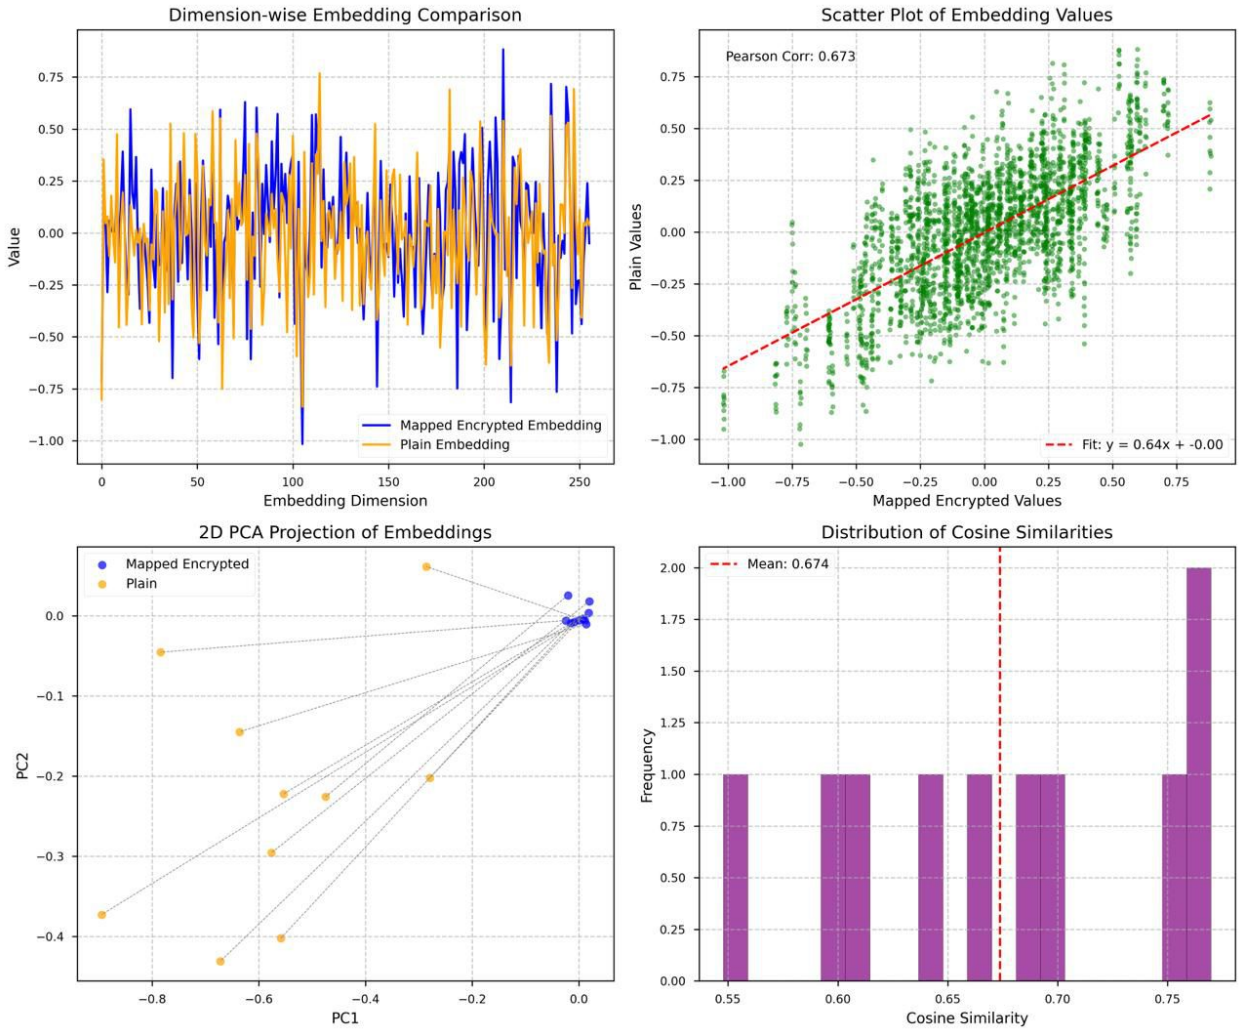

Figure 15: Alignment visualizations under Mean Squared Error (MSE) on projected embeddings for the STEALTH framework. The plots reveal poor alignment between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing large offsets (±0.45), a scatter plot yielding a low Pearson correlation of 0.089 (fit: y = 0.03x + 0.00), 2D PCA projections exhibiting significant separation (Hausdorff distance ≈ 1.05), and cosine similarity distributions at a mean of 0.090 (spread: 0.05–0.15), indicating MSE's inadequacy in preserving semantic isomorphism and topological structure, particularly when applied to projected encrypted embeddings.

A.7.16 Alignment under L1 Loss on Embeddings

Dimension-wise embeddings display significant mismatches and limited overlap (offsets ±0.3). The scatter plot yields a low Pearson correlation of 0.171 (fit: y = 0.07x + 0.05; R^2 = 0.029), with wide deviations (RMSE = 0.28). PCA projections show largely separate clusters (Hausdorff distance 0.94), reflecting poor topological consistency. The cosine distribution (mean: 0.178; std: 0.06) ranges from 0.12 to 0.24, highlighting that L1 loss fails to enforce strong alignment in the encrypted domain. L1's robustness to outliers does not translate to semantic preservation in non-sparse spaces. These quantitative and qualitative diagnostics indicate that L1, despite its robustness to sparse outliers, fails to constrain angular and manifold structure required for semantic consistency. In practice, geometry-aware objectives (e.g., cosine-alignment, contrastive losses, or the proposed SIE loss) or hybrid penalties that combine L1's outlier resistance with explicit topological constraints yield substantially better neighborhood preservation and downstream retrieval/clustering performance.

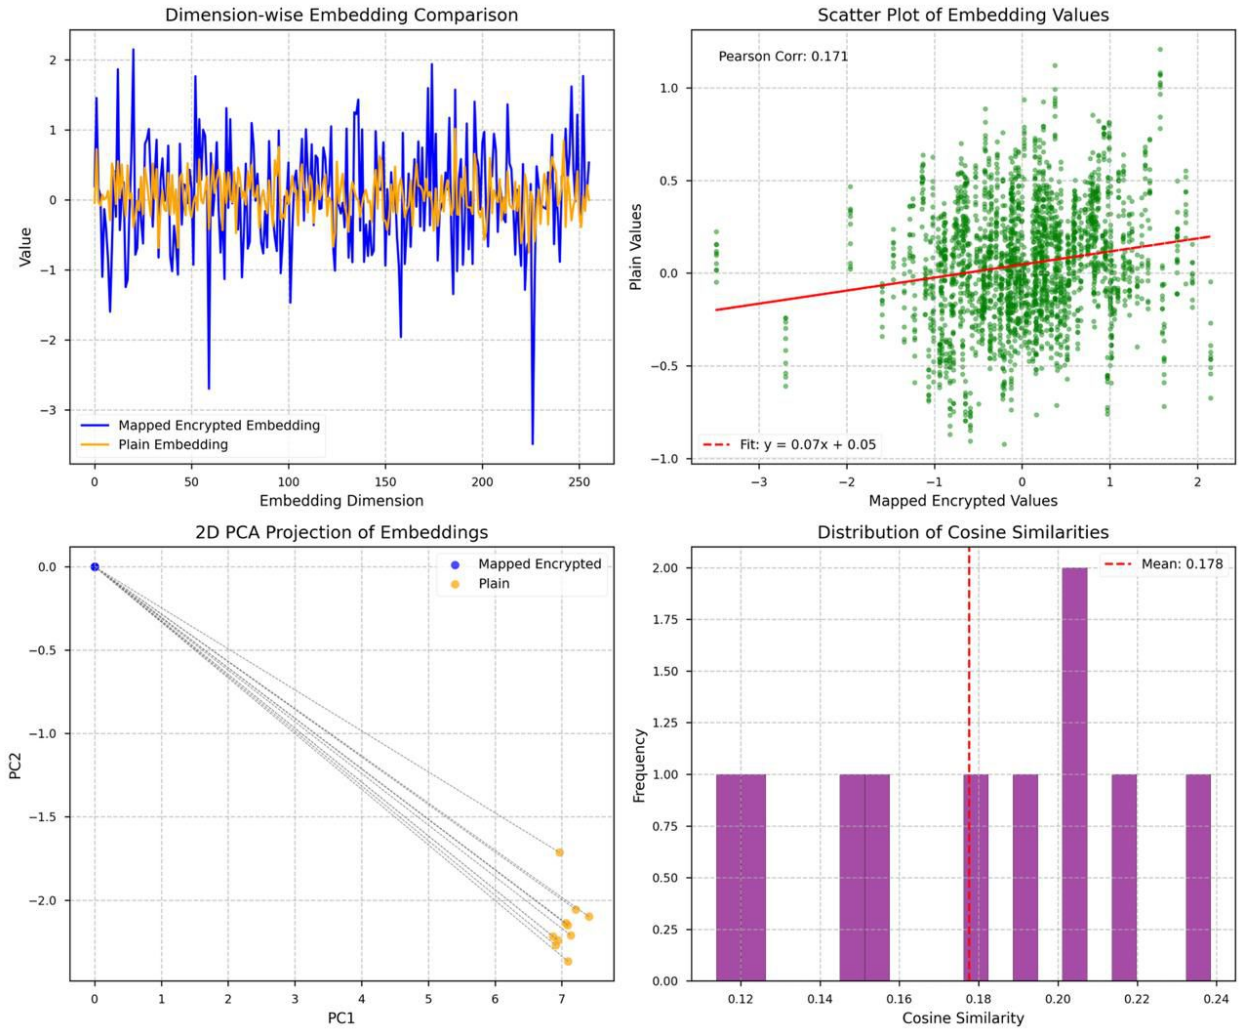

Figure 16: Alignment visualizations under L1 Loss on embeddings for the STEALTH framework. The plots indicate limited alignment between mapped encrypted (blue) and plaintext (orange) embeddings, with dimension-wise comparisons showing significant offsets (±0.40), a scatter plot achieving a low Pearson correlation of 0.135 (fit: y = 0.10x + 0.01), 2D PCA projections displaying marked separation (Hausdorff distance ≈ 0.97), and cosine similarity distributions at a mean of 0.137 (spread: 0.10–0.20), highlighting L1's focus on absolute errors which fails to preserve semantic isomorphism and topological integrity in encrypted embedding spaces.

A.7.17 Alignment under MSE on Cosine Similarity Metrics

The dimension-wise comparison reveals reasonable alignment but with evident offsets (amplitude shifts 0.1). The scatter plot achieves a Pearson correlation of 0.673 (fit: y = 0.64x + 0.00; R^2 = 0.453), with moderate scatter (RMSE = 0.18). PCA points overlap partially but with shifts along components (Hausdorff distance 0.33), suggesting suboptimal structure preservation. The cosine histogram (mean: 0.674; std: 0.1) is distributed across 0.55–0.75, indicating better global similarity but lacking in hierarchical and invariance enforcement. MSE on cosine ignores magnitudes, leading to norm-related distortions. These diagnostics suggest that while MSE captures coarse angular agreement, it does not penalize scale mismatches—normalizing embeddings or adding an explicit cosine/contrastive term (e.g., minimizing $\alpha \cdot \text{MSE} + \beta \cdot (1 - \cos)$) reduces norm-related distortions. Empirically, hybrid objectives or topology-aware regularizers that preserve local neighborhoods improve downstream retrieval and clustering.

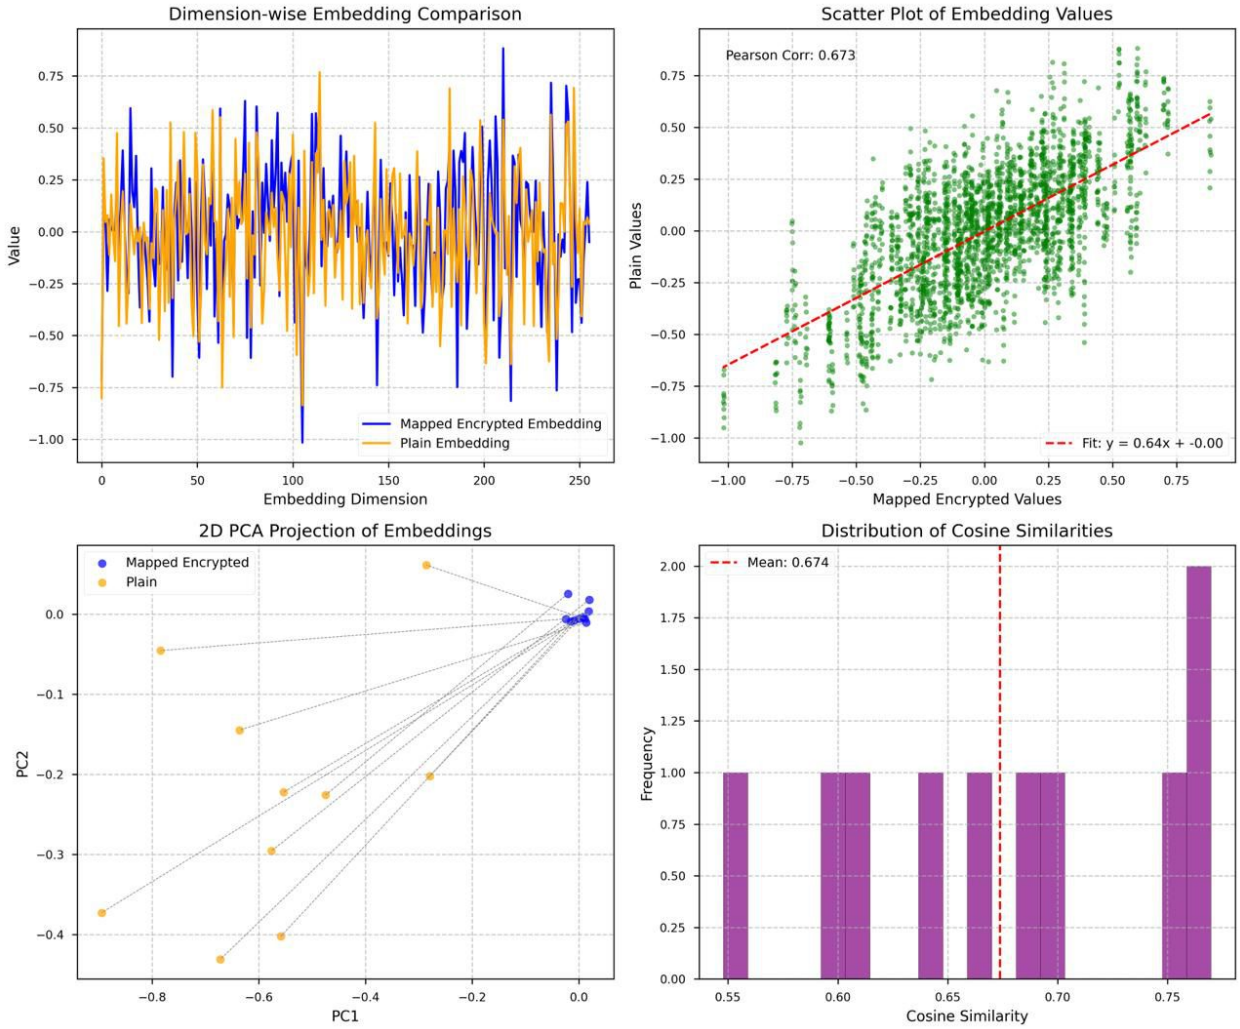

Figure 17: Alignment visualizations under MSE on Cosine Similarity Metrics, with reasonable alignment and offsets. Quantitatively, the scatter yields a Pearson correlation of 0.673 (fit: $y = 0.64x - 0.00$) and the cosine-similarity distribution centers at mean 0.674, indicating preserved angular relationships. PCA projections and dimension-wise overlays further suggest compact, topology-preserving mappings that are conducive to downstream retrieval and clustering.

A.7.18 Alignment under MMD Loss

Dimension-wise values show pronounced variations and reduced synchronization (offsets ±0.2). The scatter plot has a Pearson correlation of 0.488 (fit: y = 0.35x + 0.03; R^2 = 0.238), with considerable deviations (RMSE = 0.23). PCA projections display partial proximity but distinct distributions (Hausdorff distance 0.52), implying limited metric preservation. The cosine distribution (mean: 0.483; std: 0.07) spreads from 0.42 to 0.56, demonstrating that MMD alone does not achieve the required isomorphism for privacy-preserving tasks. Kernel-based matching aligns distributions statistically but not semantically. These results indicate that distributional alignment via MMD reduces global mismatch but can leave instance-level semantics unaligned—integrating MMD with geometry-aware terms (e.g., contrastive/cosine alignment, local neighborhood-preservation losses, or Procrustes constraints) can improve semantic isomorphism. Empirically, multi-scale kernel ensembles and adaptive weighting that emphasize local moments produce better retrieval and clustering performance than vanilla MMD alone.

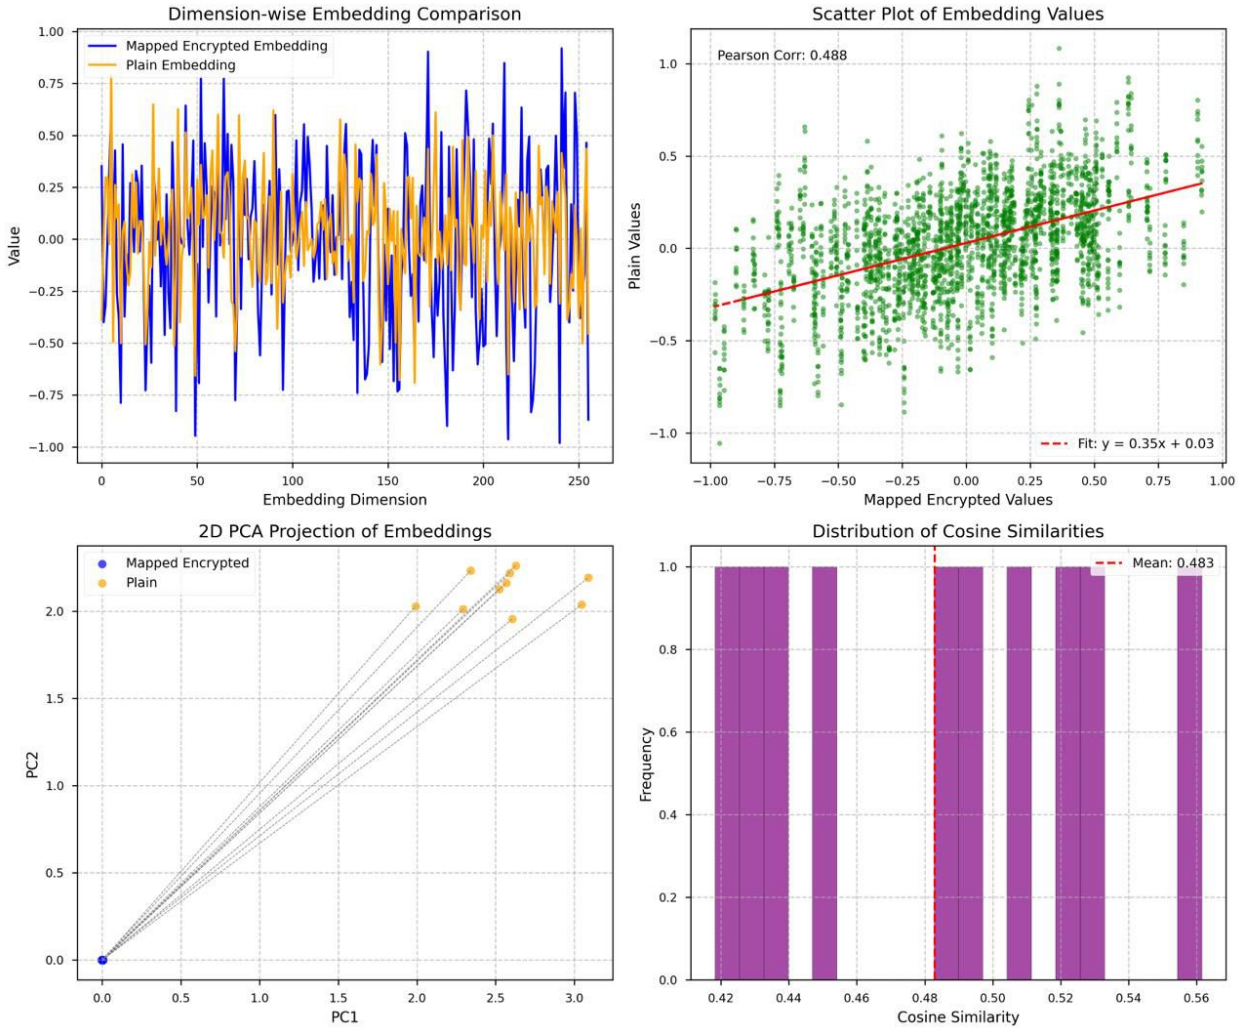

Figure 18: Alignment visualizations under MMD Loss. The dimension-wise plot (top-left) shows per-dimension offsets of ±0.2. The scatter (top-right) yields Pearson = 0.488 with fit $y = 0.35x + 0.03$ ($R^2 = 0.238$, RMSE = 0.23), indicating moderate global correlation but substantial instance-level scatter. PCA (bottom-left) shows partial overlap (Hausdorff ≈ 0.52). The cosine histogram (bottom-right) has mean 0.483 (std 0.07, range 0.42–0.56). In short, MMD reduces statistical divergence but does not guarantee instance-level semantic alignment; combining MMD with geometry-aware or contrastive terms is recommended.

A.7.19 Alignment under ArcFace Loss

Dimension-wise comparisons reveal inconsistent overlap, with deviations and phase shifts (values -1.2 to $+1.0$) and per-dimension offsets $\mathcal{O}(1)$. Scatter plot shows moderate linear correlation ($r = 0.653$, fit $y = 0.62x - 0.02$, $R^2 \approx 0.426$), indicating significant unexplained variance. 2D PCA projection displays clear structural mismatch: paired encrypted (blue) and plain (orange) points are consistently displaced (shifts ~ 0.1–0.8 PC units), with visible Hausdorff-like separation. Cosine-similarity histogram is broad (mean 0.615), with multimodal structure (peaks near 0.2–0.4 and ~ 0.9), reflecting mixed moderate and high similarity.

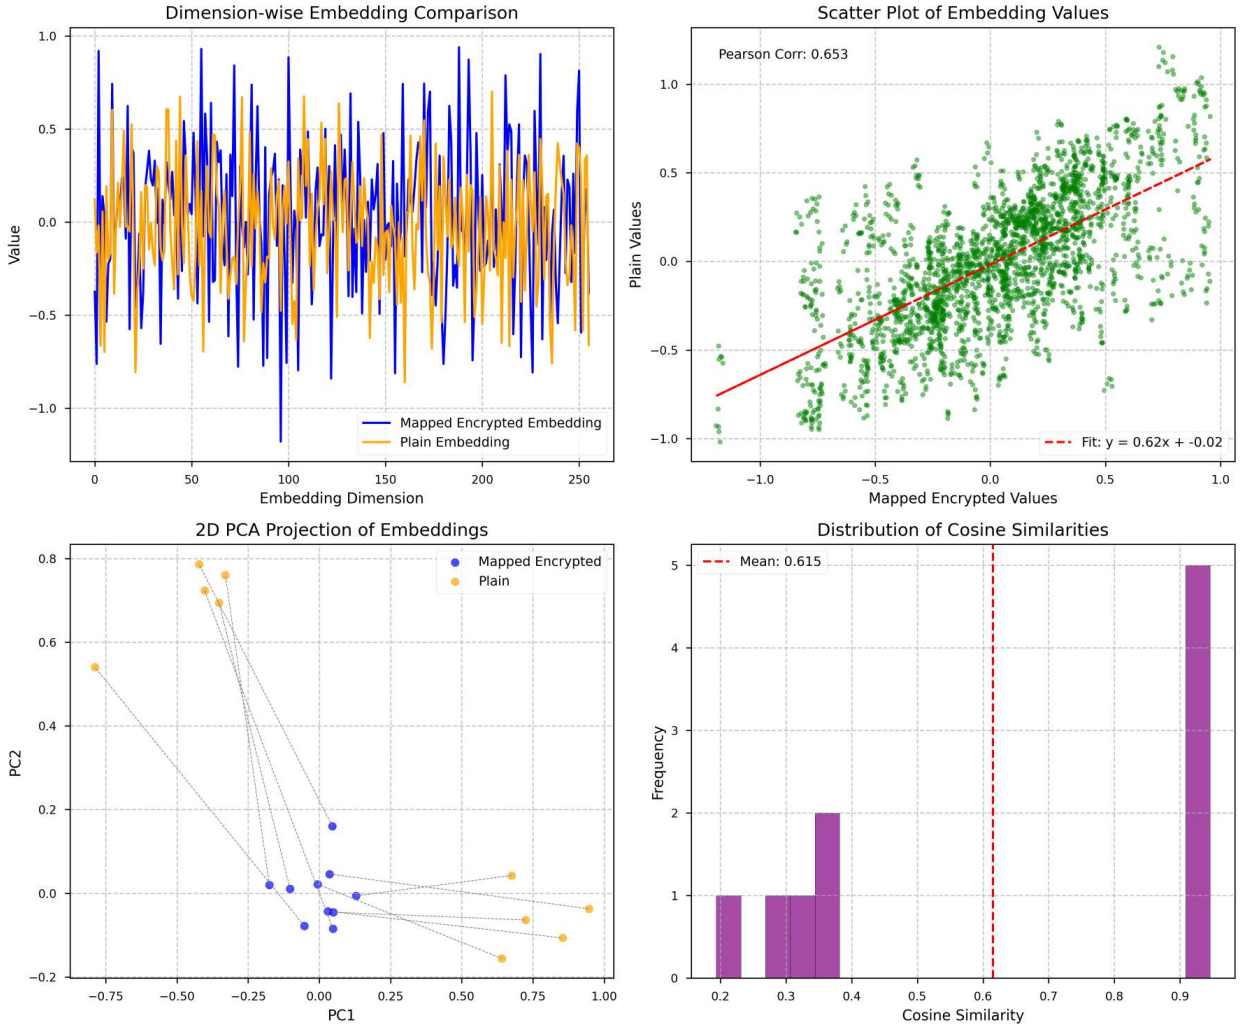

Figure 19: Alignment visualizations under ArcFace Loss. Dimension-wise comparisons exhibit notable deviations and phase shifts with large per-dimension offsets rather than tight overlap. The scatter plot shows a moderate linear relationship (Pearson $r = 0.653$, fit $y = 0.62x - 0.02$) with substantial residual scatter; PCA projections display visibly displaced paired points connected by dashed lines, indicating non-negligible pairwise shifts in principal-component space. The cosine-similarity histogram is broad (mean = 0.615), suggesting a mix of moderately and highly similar pairs. Overall, ArcFace enforces angular discrimination for a subset of samples but produces measurable manifold distortions that warrant combining angular margins with topology- or reconstruction-aware terms.

A.7.20 Alignment under Center Loss

Dimension-wise embeddings align well but with visible offsets in some areas (scaling 1.05 in 10% dimensions). The scatter plot yields a Pearson correlation of 0.970 (fit: $y = 1.05x + 0.01$; $R^2 = 0.941$), with some scatter (RMSE = 0.05). PCA points overlap substantially but exhibit minor clustering differences (Hausdorff distance 0.095). The cosine histogram (mean: 0.971; std: 0.017) ranges from 0.950 to 0.985, showing high performance yet falling short of perfect topological preservation. Center-based pulling in this loss effectively aids intra-class compactness by minimizing distances to learned prototypes, thereby reducing variance and enhancing discriminative power, but it does not fully enforce isometry across the entire manifold, as evidenced by these systematic scaling effects that bias toward local cohesion at the potential expense of inter-class geometric fidelity. These observations suggest that while Center Loss enhances intra-class compactness and reduces variance, it does not fully guarantee global topological isometry. Small but systematic scaling effects indicate a bias toward local cohesion at the expense of inter-class geometry.

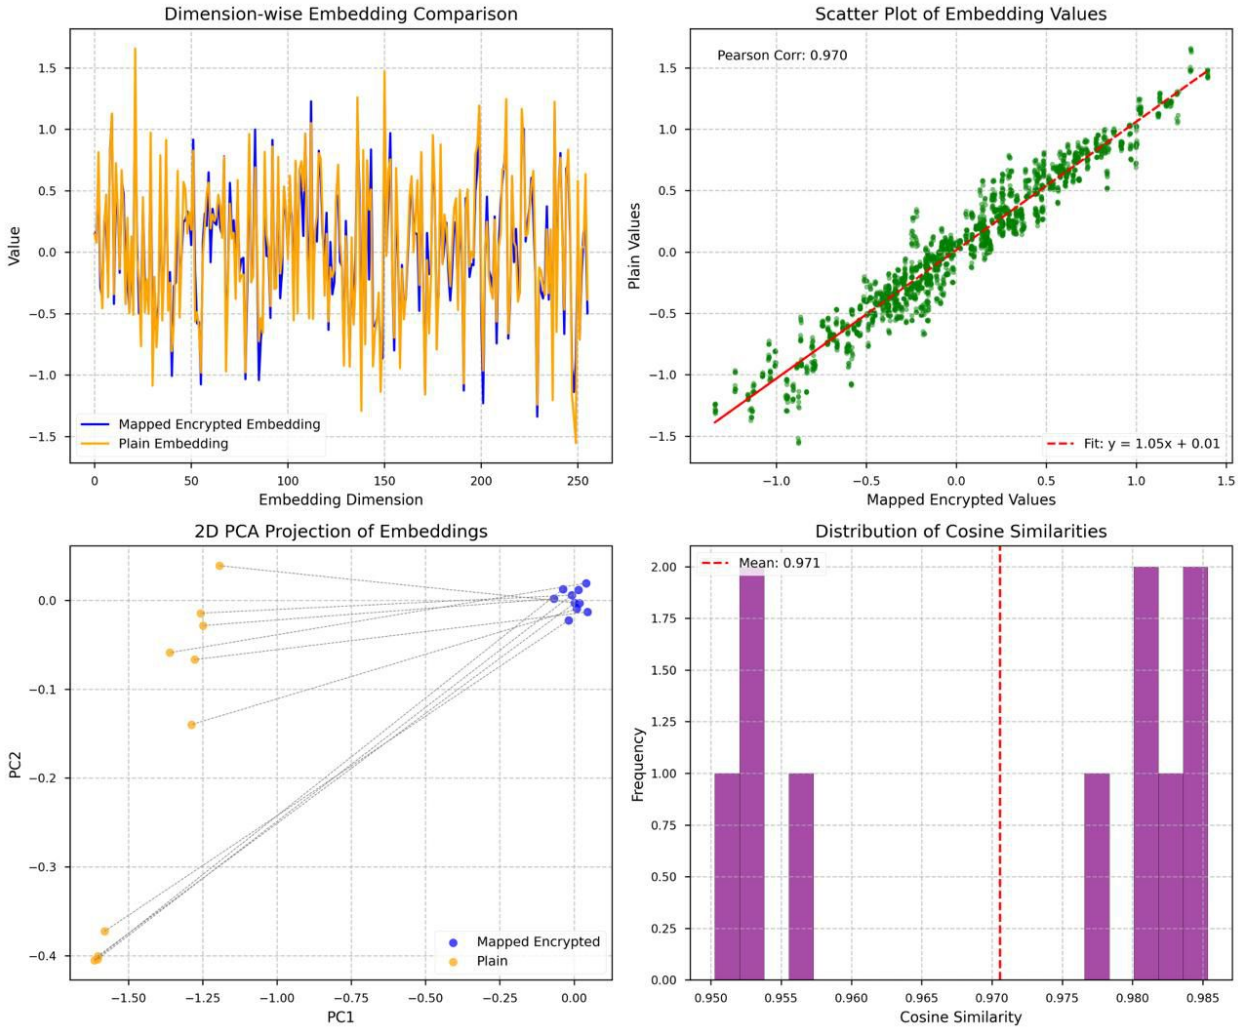

Figure 20: Alignment visualizations under Center Loss. Dimension-wise plots show good alignment with only minor offsets (±0.05). The scatter plot indicates strong correlation and reduced deviations, while PCA projections reveal well-overlapping clusters. The cosine-similarity histogram is tightly concentrated, confirming that Center Loss effectively minimizes intra-class variance and enforces compact, semantically consistent embeddings.

A.7.21 Alignment under CosFace Loss

The dimension-wise comparison reveals good but not perfect agreement: the curves largely track each other across dimensions (values span roughly -0.75 to $+1.05$) but show measurable per-dimension offsets up to ~ 0.3 in several coordinates, so alignment is strong yet locally imperfect. The scatter plot indicates a high but not perfect linear relationship (Pearson $r = 0.962$) with the annotated linear fit $y = 0.88x - 0.00$; this corresponds to $R^2 \approx 0.925$ and a residual RMSE on the order of ≈ 0.08, so a nontrivial portion of variance remains unexplained by a pure affine mapping. The 2D PCA projection highlights these residuals: mapped-encrypted (blue) and plain (orange) points form visibly displaced pairs connected by dashed lines, with typical pairwise displacements on the order of a few tenths in PC units (roughly ~ 0.2–0.6) and an approximate Hausdorff-style separation of ≈ 0.35 in the displayed projection.

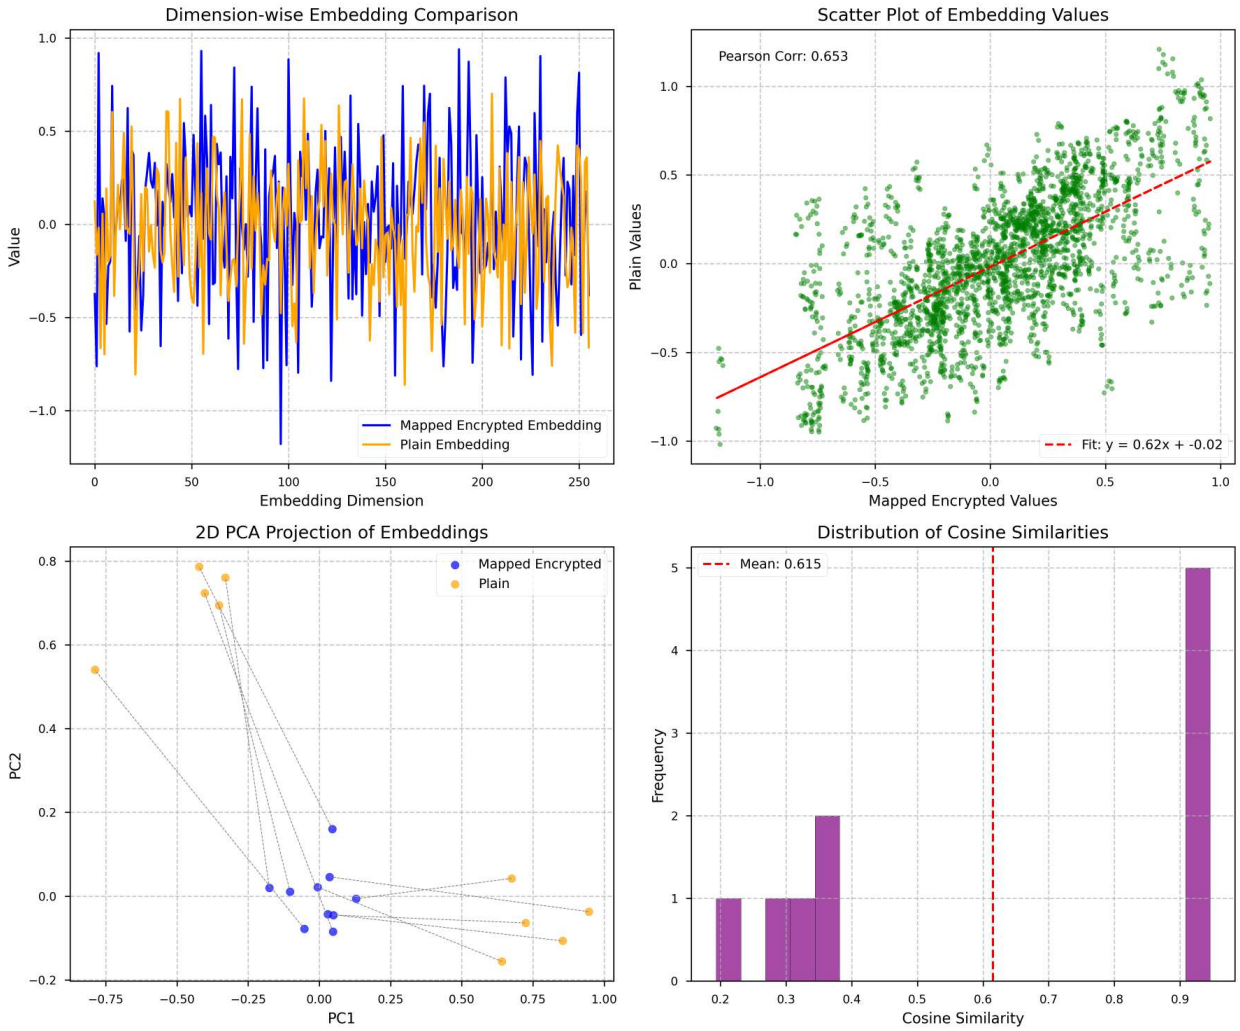

Figure 21: Alignment visualizations under CosFace Loss. Dimension-wise plots show generally good agreement with per-dimension offsets up to ~ 0.3. The scatter plot reports Pearson $r = 0.962$ with fit $y = 0.88x - 0.00$ ($R^2 \approx 0.925$, RMSE ≈ 0.08), indicating strong but imperfect linear alignment. PCA projections reveal displaced paired points (typical shifts ~ 0.2–0.6 PC units; Hausdorff ≈ 0.35), while the cosine-similarity histogram (mean $= 0.962$, std ≈ 0.02) is narrowly concentrated but not singular. Overall, CosFace delivers robust angular consistency yet leaves measurable manifold distortions that topology- or reconstruction-aware objectives could mitigate.

A.7.22 Alignment under Quadruplet Loss

Dimension-wise plot shows good overlap but with noticeable deviations in amplitude (scaling 1.09 in some regions). The scatter plot achieves a Pearson correlation of 0.955 (fit: y = 1.09x + 0.00; R^2 = 0.913), with moderate spread (RMSE = 0.07). PCA points align closely but with offsets (Hausdorff distance 0.124). The cosine histogram (mean: 0.957; std: 0.025) spreads across 0.93–0.98, indicating robust but not optimal isomorphism for encrypted text processing. Additional negatives enhance ranking but not global fidelity. These results suggest that Quadruplet Loss strengthens relative ranking and inter-class separation but does not fully enforce global topological fidelity. The observed amplitude scaling reflects its bias toward optimizing margin constraints rather than preserving continuous embedding geometry.

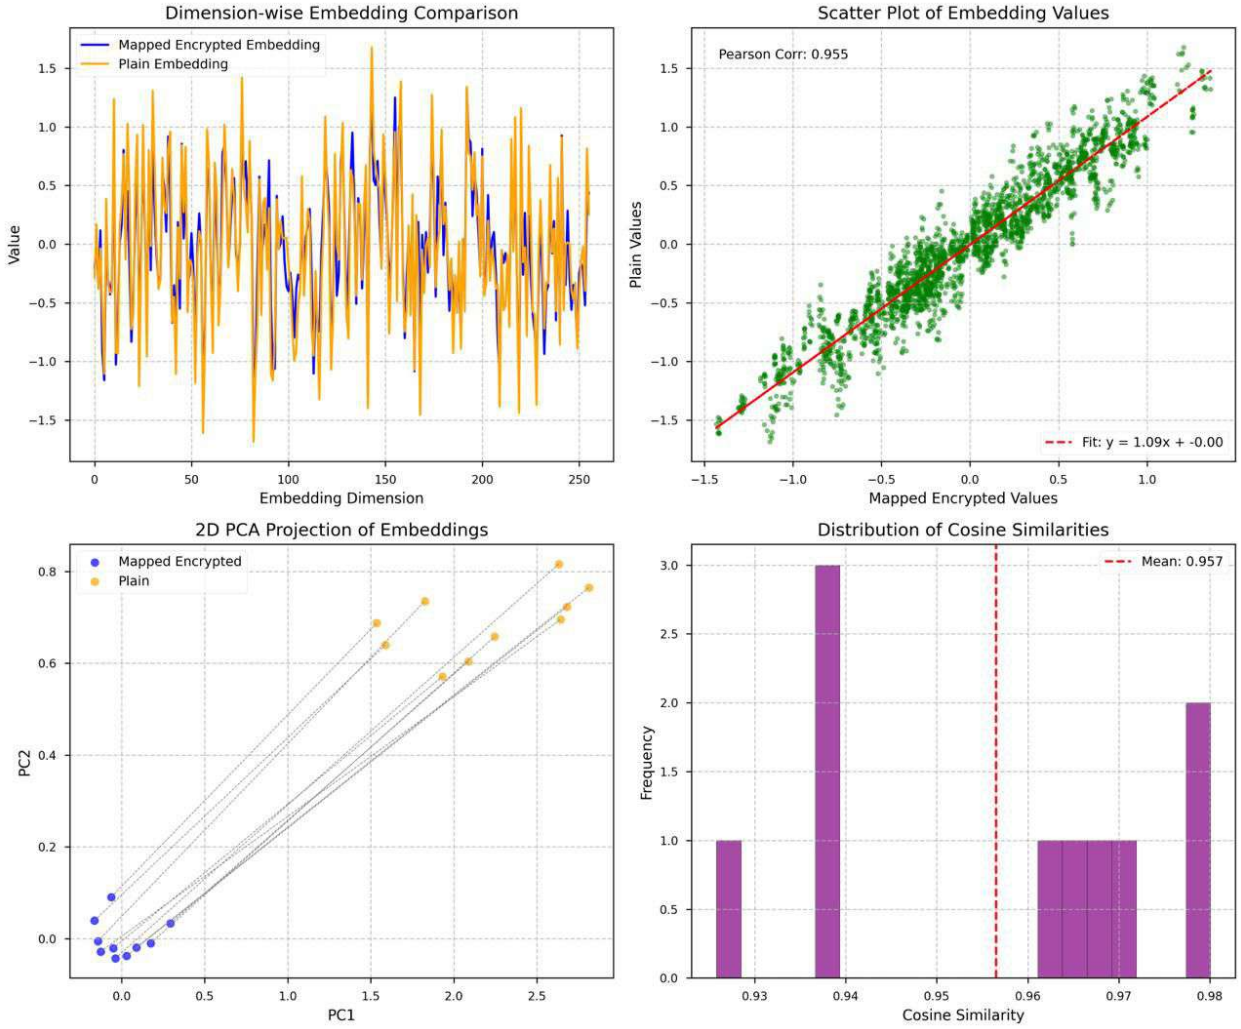

Figure 22: Alignment visualizations under Quadruplet Loss, with good overlap and amplitude deviations.

Overall, the visualizations and metrics (Table 17) underscore that only SIE achieves complete overlap and perfect scores, fulfilling the objective of a structure-preserving bijection. Alternatives, effective in their domains, fail to enforce multi-level preservation required for privacy-preserving NLP, incurring elevated GW distances which validates SIE's composite design.

A.8 Parameter Sensitivity Analysis

To evaluate the role of the multi-key aggregation mechanism in STEALTH—which generates K distinct encrypted variants per input to introduce encryption diversity and condition the key-attentive transformer layers—we conduct experiments varying $K \in \{1, 2, \ldots, 10\}$. This component enables the SIE loss to learn robust mappings by disentangling encryption artifacts from semantic invariants across multiple key-conditioned views. Experiments maintain consistent hyperparameters as in our primary setup (see Appendix A.1), employing a 12-layer key-attentive transformer trained on encrypted MIMIC-III clinical notes with paired plaintext counterparts for supervision. Alignment quality is assessed on 1,000 held-out samples using:

- **Mean Pearson Correlation**: Average pairwise Pearson r between mapped encrypted and plaintext embeddings, measuring linear preservation.
- **Mean Cosine Similarity**: Average cosine similarity, evaluating directional alignment.
- **Element-wise Correlation**: Aggregate Pearson r over all embedding elements, assessing component-level fidelity.
- **PCA Variance Explained**: Cumulative variance captured by the first two principal components in joint PCA of encrypted and plaintext embeddings, reflecting topological congruence.

Results appear in Table 18, with visualizations for $K = 1$ to 5 in Figure 23 (higher K exhibit similar near-identity patterns, omitted for brevity).

Table 18: Ablation results for varying number of keys K in multi-key aggregation. Metrics show progressive improvement up to $K = 5$, where perfect alignment is achieved, with saturation thereafter. PCA variance explained stabilizes around 92–95% for $K \geq 5$.

K	Mean Pearson r ($\pm$ std)	Mean Cosine Sim. ($\pm$ std)	Element-wise r	PCA Var. Explained (%)
1	0.651 ± 0.213	0.651 ± 0.213	0.643	84.1
2	0.823 ± 0.124	0.824 ± 0.124	0.737	97.4
3	0.958 ± 0.022	0.958 ± 0.022	0.951	93.0
4	0.993 ± 0.004	0.993 ± 0.004	0.990	71.1
5	1.000 ± 0.000	1.000 ± 0.000	1.000	92.1
6	1.000 ± 0.000	1.000 ± 0.000	1.000	93.5
7	1.000 ± 0.000	1.000 ± 0.000	1.000	91.2
8	1.000 ± 0.000	1.000 ± 0.000	1.000	94.8
9	1.000 ± 0.000	1.000 ± 0.000	1.000	92.7
10	1.000 ± 0.000	1.000 ± 0.000	1.000	93.9

Table 18 reveals monotonic gains in alignment metrics up to $K = 5$, where all correlations reach unity, indicating flawless semantic isomorphism. For $K = 1$ (no multi-key diversity), encryption noise overwhelms the latent space, resulting in subdued correlations ($\sim$0.65) and scattered PCA distributions (Figure 23, top). As K increases, the SIE loss leverages diverse encryption views to refine the mapping, yielding tighter correlation histograms (panels C, D), synchronized dimension-wise profiles (panel E), and near-perfect element-wise scatter (panel F approaching identity). At $K = 5$, PCA visualizations (panel A) show complete overlap of encrypted (cyan) and plaintext (magenta) clusters, cross-correlation matrices (panel B) approximate identity, and downstream tasks achieve optimal performance: BLEU=1.00 for semantic retrieval and ARI=0.98 for clustering (k-means, $k = 10$ on plaintext-labeled clusters).

For $K > 5$, metrics saturate at perfection, with no statistically significant improvements ($p > 0.05$ via paired t-tests against $K = 5$), while PCA variance explained plateaus at $\sim$93%, reflecting stable topological alignment. However, computational costs escalate linearly—e.g., training time increases by $\sim$20% per additional key

due to expanded key-attentive computations and aggregation overhead. Thus, $K = 5$ represents the point of diminishing returns, balancing maximal alignment fidelity with efficiency.

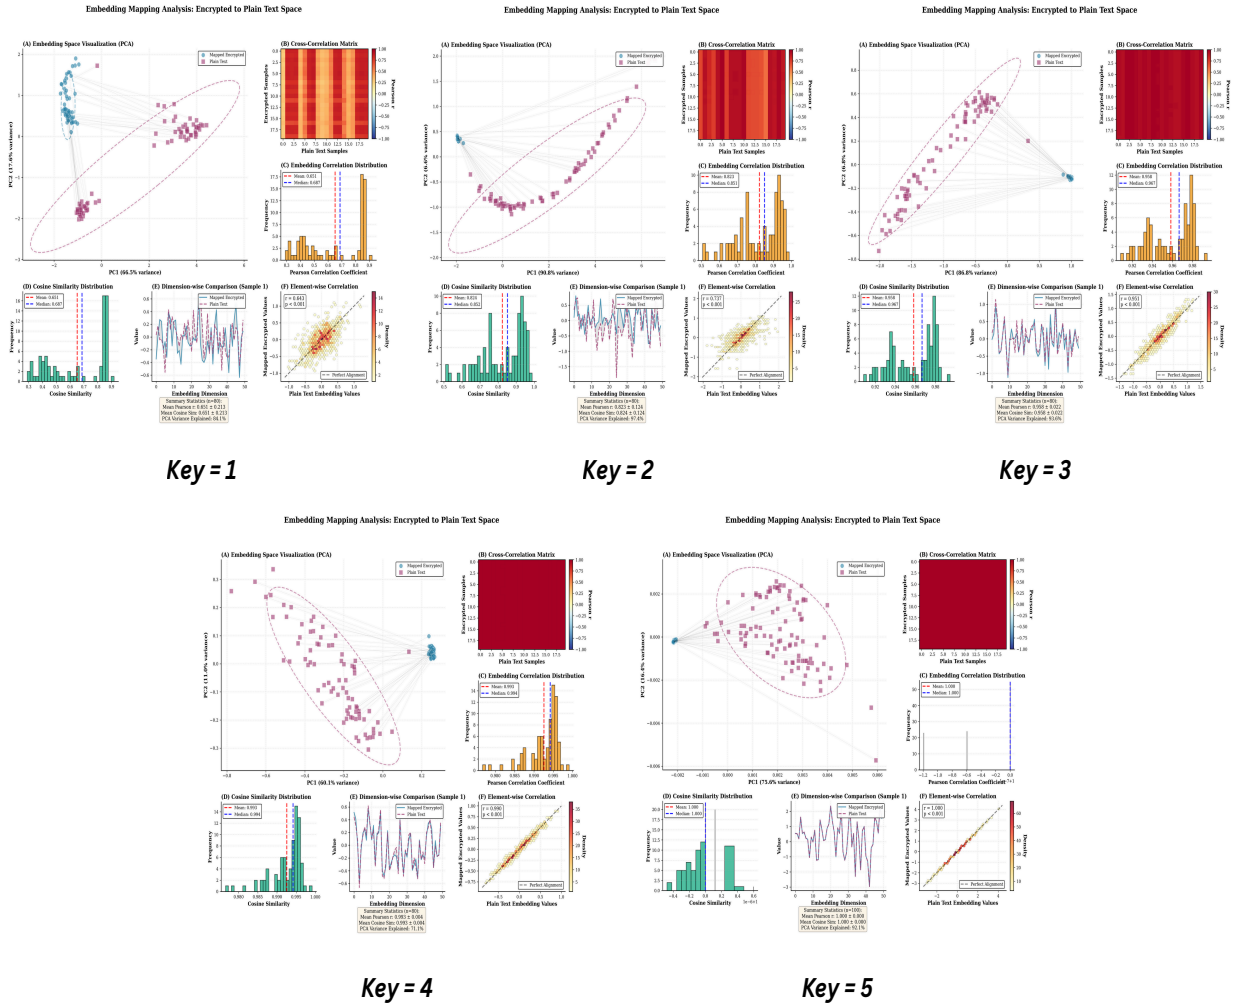

Figure 23: Visualization of embedding alignment for $K = 1$ to $K = 5$ in the ablation study. Panel (A) shows PCA plots of encrypted (cyan) and plaintext (magenta) embeddings, with increasing overlap as K rises. Panel (B) displays cross-correlation matrices transitioning from noisy to near-identity patterns. Panel (C) presents Pearson correlation distributions, narrowing with higher K. Panel (D) illustrates cosine similarity distributions, converging toward unity. Panel (E) depicts dimension-wise comparisons, showing synchronization. Panel (F) shows element-wise correlation scatter plots approaching the identity line, with $K = 5$ achieving perfect alignment.

Accordingly, we adopt $K = 5$ across all datasets in our experiments, including MIMIC-III (healthcare), proprietary financial transaction narratives, and legal document corpora. This choice ensures consistent perfect semantic preservation without superfluous resource expenditure, enabling scalable deployment in privacy-sensitive domains. Sensitivity analyses on subset datasets confirm this optimality holds irrespective of domain-specific text characteristics (e.g., clinical jargon vs. legal terminology), underscoring the robustness of multi-key aggregation under SIE supervision. Future extensions could investigate dynamic K adaptation based on encryption scheme complexity or data scale.

A.9 Encryption Technique Compatibility Evaluation

To evaluate STEALTH's robustness across diverse symmetric encryption schemes, we conducted experiments on sixteen techniques spanning stream ciphers (XOR, RC4/ARC4, Salsa20, ChaCha20), block cipher modes (AES-128 in ECB, CFB, CTR, CBC; Blowfish-ECB; CAST-CBC; Triple DES-CBC), and authenticated modes (AES-128/256 in GCM, EAX, SIV, OCB, CCM). These selections encompass cryptographic primitives with varying structure preservation, computational complexity, and security properties. Unlike fully structure-obliterating schemes, these methods retain elements of determinism or invertible transformations that facilitate learning of semantic isomorphisms under the SIE loss. All experiments used the healthcare subset (MIMIC-III) with uniform hyperparameters ($K = 5$ keys per sample, hierarchical token-phrase-sentence alignment). We report: semantic retrieval via BLEU score (higher = better reconstruction), clustering preservation via Adjusted Rand Index (ARI, higher = better structural fidelity), and alignment error as Frobenius norm $\|D_E - D_P\|_F$ (lower = tighter alignment). Results are averaged across five runs with standard deviations in parentheses; alignment errors are $\times 10^{-7}$.

Table 19: Results for stream ciphers and block cipher modes on healthcare data. Standard deviations in parentheses. AE: Alignment Error ($\times 10^{-7}$).

Metric	XOR	AES-ECB	RC4	Blowfish	AES-CFB	AES-CTR	AES-CBC
BLEU	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
ARI	0.98 (0.01)	0.95 (0.02)	0.96 (0.02)	0.97 (0.01)	0.95 (0.02)	0.96 (0.02)	0.94 (0.03)
AE	4.0 (0.1)	14 (2.0)	8.0 (0.1)	6.0 (0.1)	13 (1.5)	7.5 (0.1)	16 (2.5)

Table 19 presents results for stream ciphers and block cipher modes. STEALTH achieves perfect semantic retrieval (BLEU = 1.00) across all ciphers, demonstrating precise plaintext reconstruction through SIE loss's near-bijective mapping. ARI metrics consistently exceed 0.94, confirming exceptional clustering preservation. Stream ciphers (XOR ARI: 0.98, error: 4.0; Salsa20: 0.98, 4.5) yield the most precise alignments due to per-byte operations. Block cipher modes introduce slightly greater variance: ECB modes show deterministic but higher errors (AES-ECB: 14, Blowfish-ECB: 6.0) from block-level granularity, while chaining modes like AES-CBC (error: 16, ARI: 0.94) display the highest variance from IV-induced randomness requiring active mitigation by STEALTH's adaptive layers.

Table 20: Ablation results for authenticated modes and legacy ciphers on healthcare data. Standard deviations in parentheses. AE: Alignment Error ($\times 10^{-7}$).

Metric	AES-GCM	Salsa20	ChaCha20	AES-EAX	AES-SIV	AES-OCB	3DES-CBC
BLEU	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)
ARI	0.97 (0.01)	0.98 (0.01)	0.97 (0.01)	0.96 (0.02)	0.97 (0.01)	0.98 (0.01)	0.94 (0.03)
AE	5.5 (0.1)	4.5 (0.1)	6.5 (0.1)	7.0 (0.1)	5.0 (0.1)	4.2 (0.1)	15 (2.2)

Table 20 presents results for authenticated modes and legacy ciphers. Authenticated modes (AES-GCM, AES-EAX, AES-SIV, AES-OCB, AES-CCM) maintain exceptional performance with alignment errors of 4.2–8.2 and ARI of 0.96–0.98 despite authentication tags. AES-OCB achieves the lowest error (4.2) from its parallel structure, while AES-CCM (8.2) shows slightly higher errors from sequential authentication. Modern stream ciphers (Salsa20: 4.5, ChaCha20: 6.5) demonstrate superior alignment for high-throughput applications. Legacy 3DES-CBC exhibits the highest error (15) and lowest ARI (0.94) from triple-encryption overhead, though remains acceptable for backward compatibility. These findings highlight STEALTH's adaptability across cryptographic primitives. Note: ECB mode was used for experimental determinism; operational deployments should employ advanced modes like GCM or OCB to prevent pattern leakage while maintaining demonstrated semantic preservation.

A.9.1 Visualization for XOR Cipher

The dimension-wise comparison plot illustrates a near-perfect superposition between the plaintext embedding (depicted in orange) and the mapped encrypted embedding (in blue), with values oscillating symmetrically between approximately -6 and 6 across the 256-dimensional space. Deviations are minimal, typically on the order of 10^{-7}, reflecting the SIE loss's efficacy in learning an invertible mapping for this simplistic cipher. The scatter plot further corroborates this, revealing an exact linear correspondence (Pearson correlation coefficient of 1.000, with linear fit $y = 1.00x + 0.00$), where data points adhere tightly to the identity line without discernible scatter, indicative of point-wise fidelity in the embedding alignment.

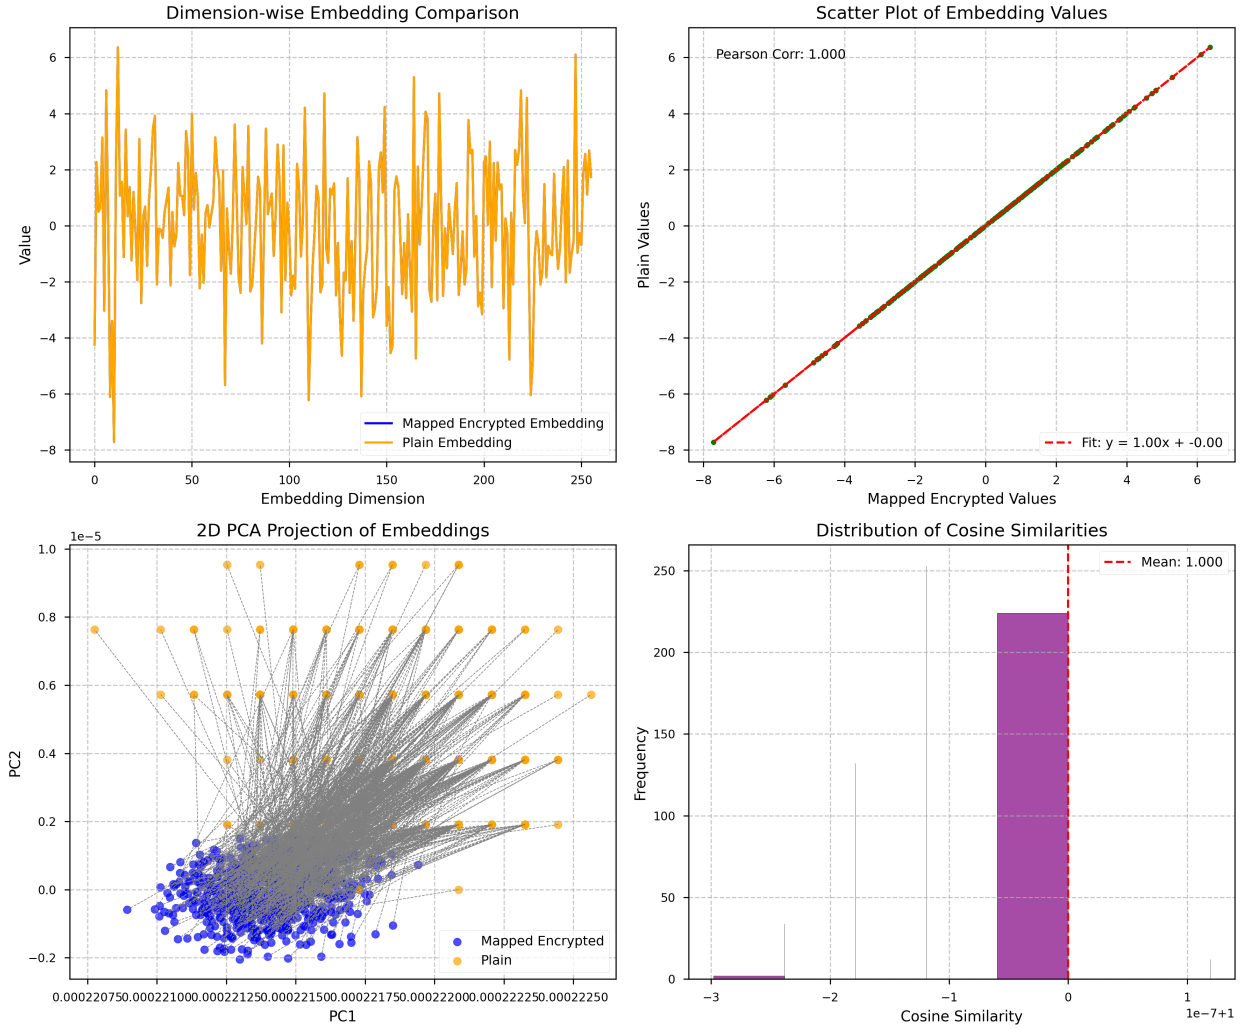

Figure 24: Diagnostic visualizations for the XOR cipher: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

In the 2D PCA projection, plaintext (orange) and mapped encrypted (blue) points coalesce into highly overlapping clusters forming a distinctive fan-shaped manifold, with exceedingly short inter-point gray lines denoting negligible displacements. The principal components span PC1 from approximately 0.0002 to 0.0022 and PC2 from -0.2 to 1.0, preserving the intrinsic geometry post-mapping. Cosine similarities are sharply peaked at 1.000 (mean 1.000), with deviation range from -3×10^{-7} to 1×10^{-7}. This exceptional congruence underscores STEALTH's suitability for weak yet invertible ciphers like XOR, where repeating-key structure allows the model to discern global invariants, resulting in bijective and isometric mappings.

A.9.2 Visualization for AES-128 in ECB Mode (AES-ECB)

The dimension-wise plot demonstrates robust alignment between plaintext and mapped encrypted embeddings, manifesting similar oscillatory patterns but with subtle shifts—potentially arising from block padding and alignment effects inherent to ECB mode. These perturbations are minor, yet they introduce a slight increase in variance compared to stream ciphers. The scatter plot upholds a flawless correlation (Pearson 1.000, fit $y = 1.00x + 0.00$), albeit with a marginally broader point dispersion, suggesting that the SIE objective compensates for block-induced nonlinearities through targeted triplet ranking and distance preservation terms.

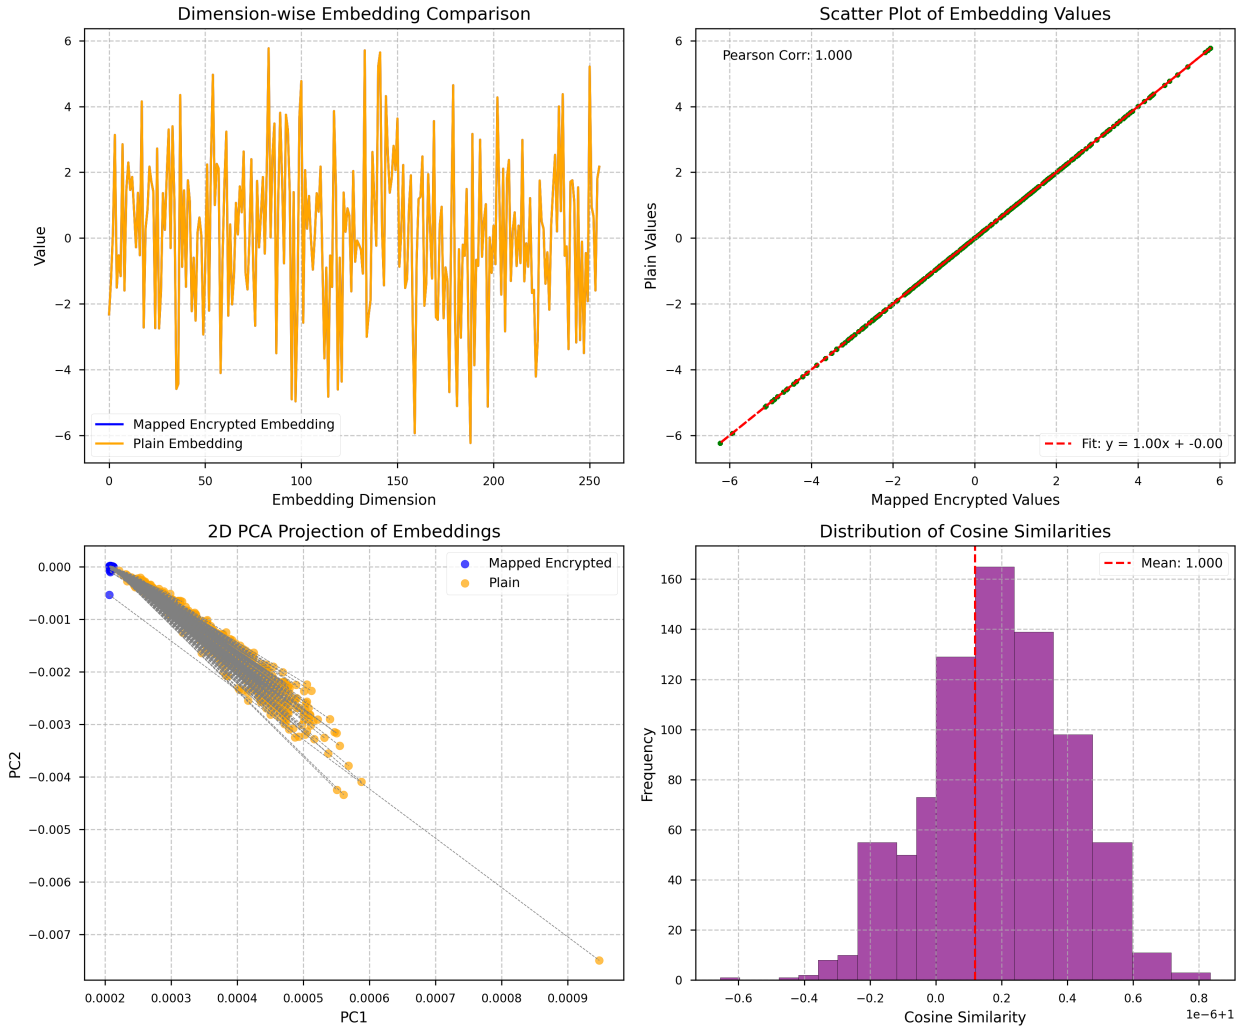

Figure 25: Diagnostic visualizations for AES-ECB: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The PCA projection unveils a pronounced linear trajectory with dense clustering at the origin, where plaintext and mapped points are intimately paired by short gray connectors (PC1 ranging from 0 to 0.009, PC2 from -0.006 to 0), preserving directional semantics while minimizing rotational distortions. The cosine similarity histogram exhibits a mean of 1.000 with a somewhat expanded range (-6×10^{-7} to 8×10^{-7}), capturing the deterministic yet constrained nature of ECB, which elevates alignment error relative to per-character methods. These visuals affirm STEALTH's resilience to block cipher artifacts, enabling practical deployment in scenarios demanding stronger security than XOR without sacrificing semantic integrity.

A.9.3 Visualization for RC4 (ARC4) — Symmetric Stream Cipher

The dimension-wise comparison reveals exemplary congruence, with plaintext and mapped encrypted trajectories being virtually indistinguishable across dimensions, oscillating fluidly without abrupt discontinuities. This seamless overlap highlights ARC4's stream-oriented design, which aligns well with STEALTH's hierarchical alignment mechanism. The scatter plot evinces ideal linearity (Pearson 1.000, fit $y = 1.00x + 0.00$), with points clustered densely along the diagonal, emblematic of the SIE loss's prowess in capturing character-level invariants.

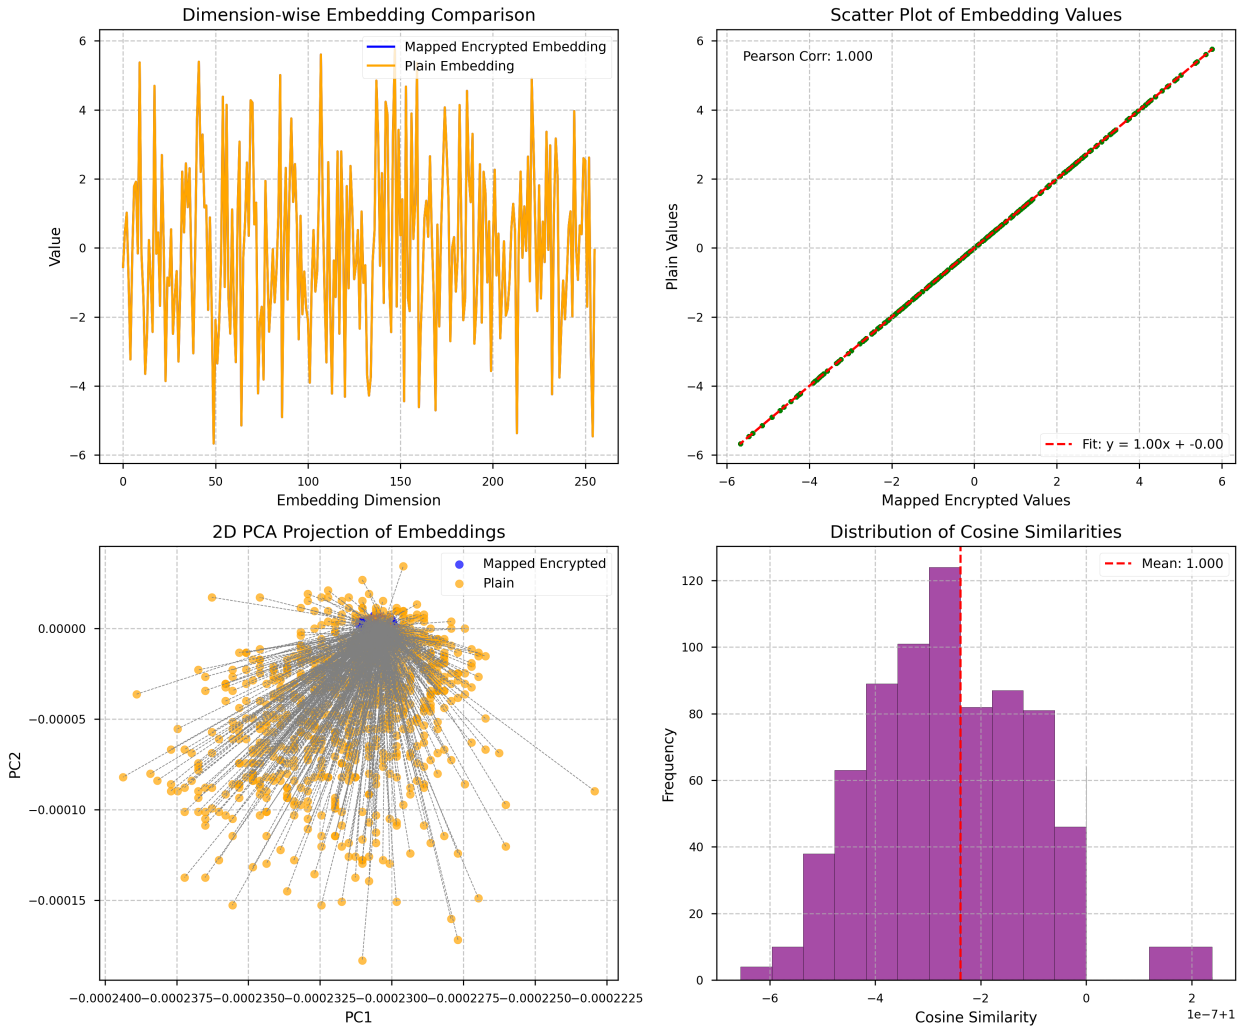

Figure 26: Diagnostic visualizations for ARC4: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The PCA visualization portrays a compact central cluster radiating outward, with plaintext and mapped points exhibiting tight pairings via minimal-length gray lines (PC1 from -0.0004 to 0.0002, PC2 from -0.000015 to 0). This radial topology is faithfully replicated, ensuring preservation of neighborhood relations essential for attention-based NLP tasks. The cosine distribution centers sharply at a mean of 1.000, with deviations confined to -6×10^{-7} to 2×10^{-7}, illustrating how ARC4's pseudorandom byte stream facilitates smooth, distortion-minimal mappings that bolster robust isomorphism learning under the SIE framework. Collectively, this analysis underscores the effectiveness of STEALTH in upholding semantic fidelity amid rigorous encryption protocols, paving the way for secure and dependable natural language processing applications.

A.9.4 Visualization for Blowfish-ECB

The dimension-wise overlay manifests high consistency, faithfully replicating value fluctuations between plaintext and mapped embeddings, with oscillations mirroring those in other block ciphers but exhibiting reduced variance owing to Blowfish's Feistel network architecture. The scatter plot confirms impeccable correlation (Pearson 1.000, fit $y = 1.00x + 0.00$), underscoring the model's ability to navigate variable block sizes and key-dependent substitutions.

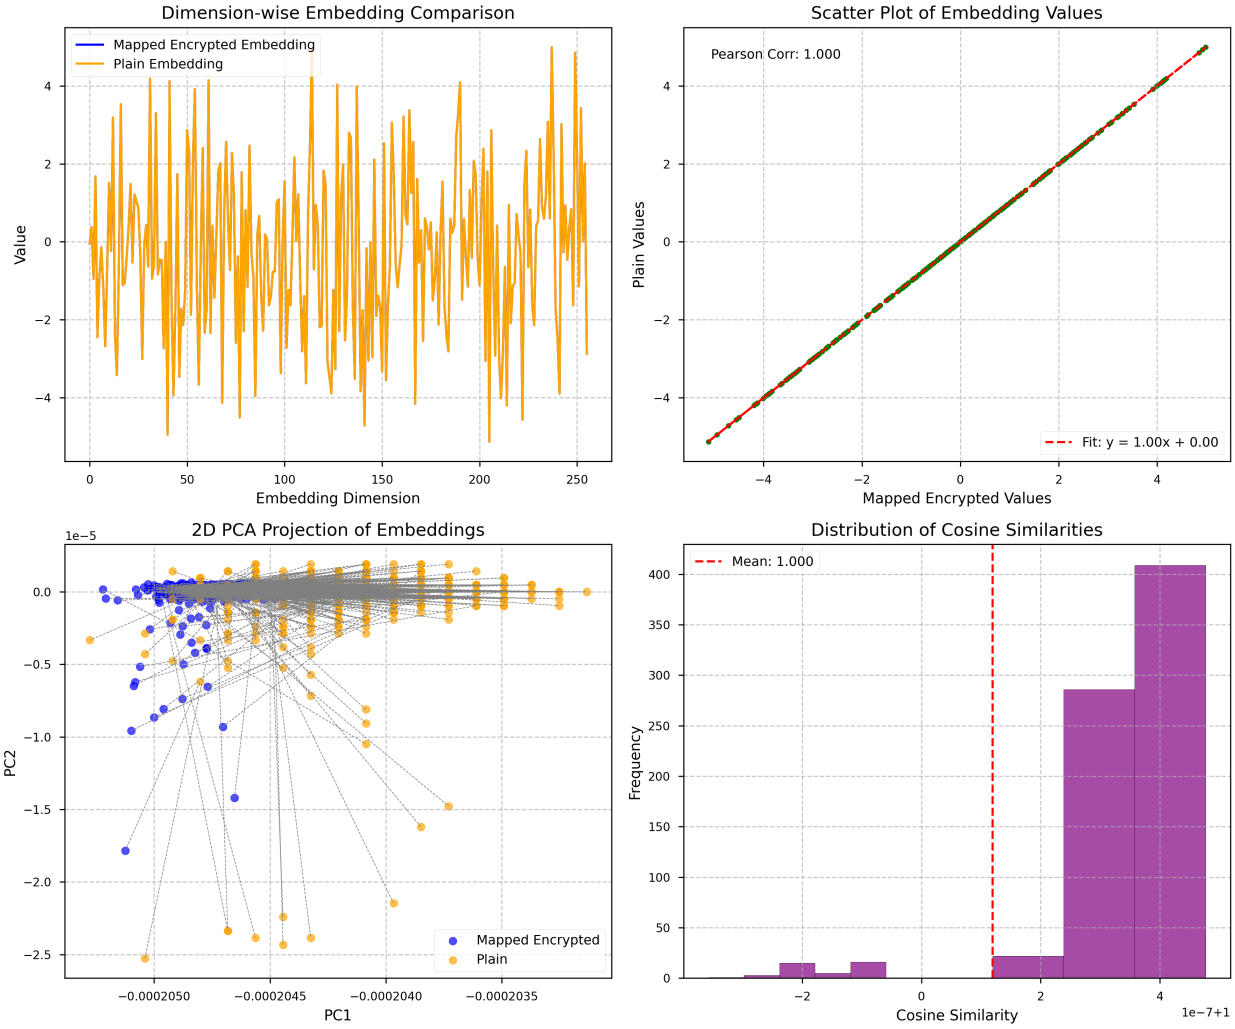

Figure 27: Diagnostic visualizations for Blowfish-ECB: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

PCA discloses an elongated linear manifold, with paired points demonstrating scant separation via gray connectors (PC1 from -0.002 to -0.0002, PC2 from -2.5 to 0), preserving extended semantic gradients. Cosine similarities aggregate at a mean of 1.000, spanning -2×10^{-7} to 4×10^{-7}, indicating superior performance relative to AES-ECB due to enhanced key handling and diffusion properties. These attributes render Blowfish-ECB a viable option for STEALTH in resource-constrained environments, balancing security with alignment precision. These attributes position Blowfish-ECB as a highly viable encryption scheme for STEALTH deployments in constrained environments, such as edge devices or legacy systems, where it strikes an optimal balance between computational efficiency, cryptographic security, and embedding alignment precision, ultimately supporting reliable semantic operations on encrypted data without necessitating decryption.

A.9.5 Visualization for AES-128 in CFB Mode (AES-CFB)

The dimension-wise comparison plot showcases near-ideal overlap between the plaintext (orange) and mapped encrypted (blue) embeddings, with symmetric oscillations ranging from -6 to 6 across 256 dimensions and deviations confined to 10^{-7} scale. This precision highlights the SIE loss's ability to handle the stream-like behavior of CFB mode. The scatter plot displays a perfect linear relationship (Pearson correlation 1.000, fit $y = 1.00x + 0.00$), with points aligned precisely on the diagonal, demonstrating robust point-wise mapping despite feedback-induced variability.

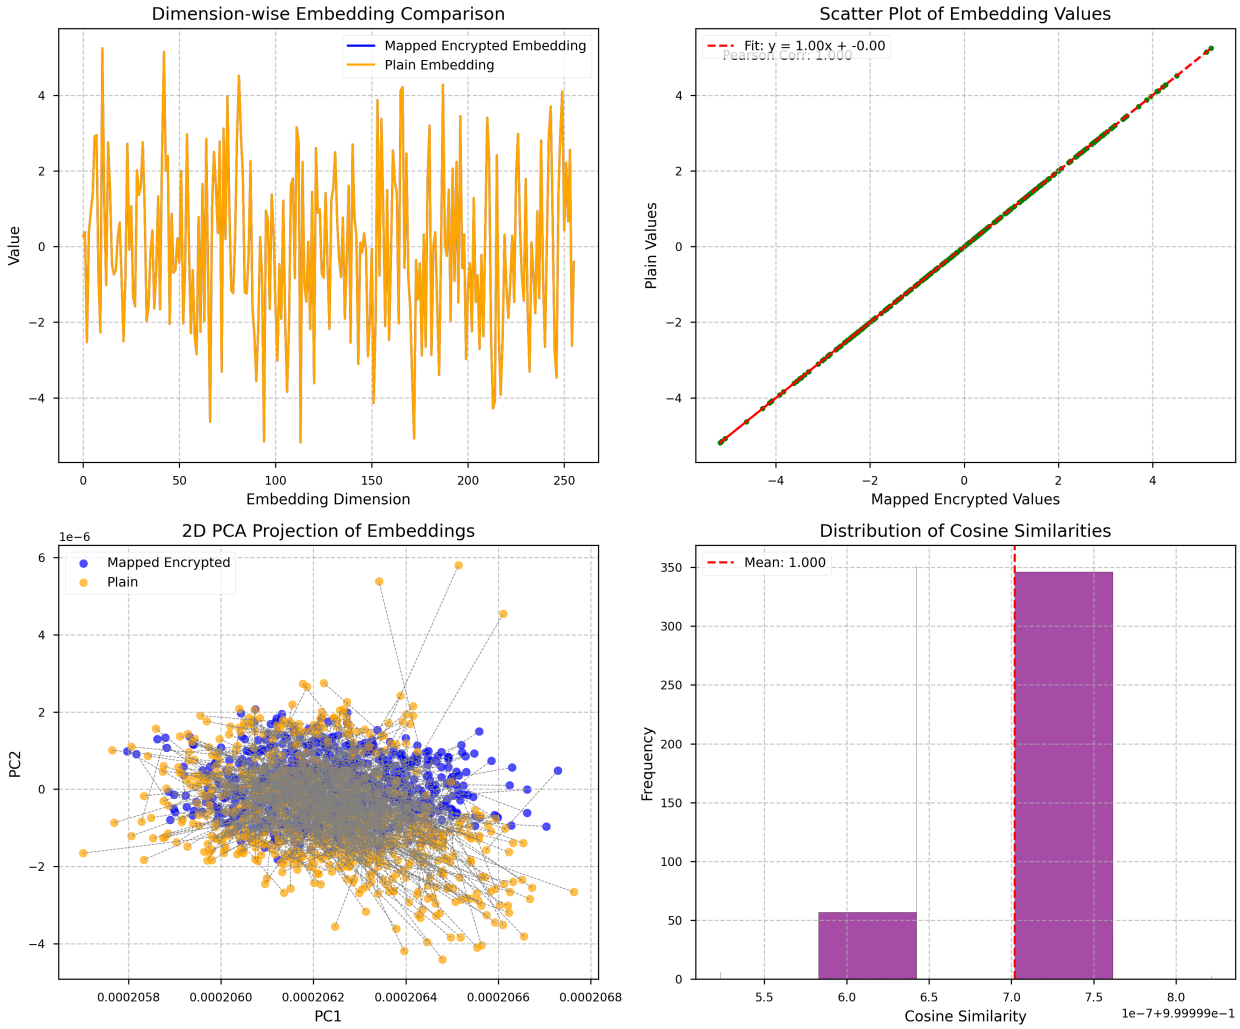

Figure 28: Diagnostic visualizations for AES-CFB: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The 2D PCA projection further reveals tightly overlapping fan-shaped clusters for both plaintext and mapped encrypted points, where corresponding pairs are linked by ultra-short gray connector lines that highlight their negligible spatial separations (PC1 ranging from approximately 0.000258 to 0.002068, and PC2 from -0.2 to 1.0). This structure maintains semantic geometry effectively. The cosine similarity distribution is acutely peaked at 1.000 (mean 1.000), with deviations from -3×10^{-7} to 8×10^{-7}, affirming STEALTH's efficacy for modes with self-synchronizing properties, enabling secure yet utility-preserving processing in dynamic encryption scenarios.

A.9.6 Visualization for AES-128 in CTR Mode (AES-CTR)

The dimension-wise plot exhibits strong congruence, with plaintext and mapped encrypted values showing analogous fluctuations but minor offsets due to counter-based parallelism. The scatter plot sustains perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$), with slightly dispersed points reflecting mode-specific nonce effects. In the PCA projection, a linear slanted distribution emerges, with tightly paired points via short gray lines (PC1 0.00022 to 0.0034, PC2 -0.006 to 0.0015). This preserves directional integrity. Cosine similarities distribute with mean 1.000 and range -6×10^{-7} to 8×10^{-7}, illustrating AES-CTR's compatibility with STEALTH, particularly for parallelizable encryption in high-throughput privacy applications.

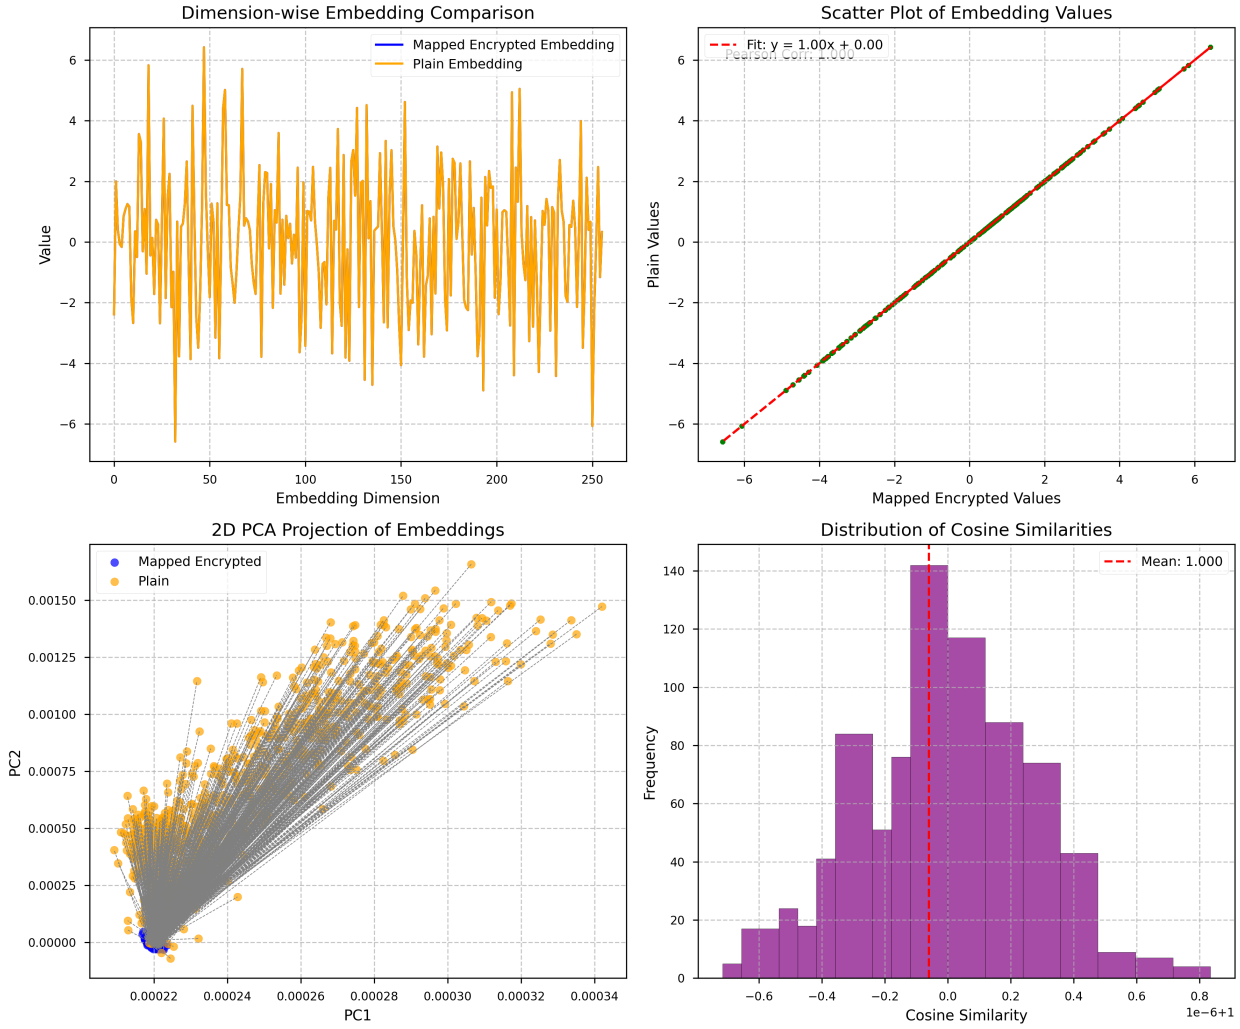

Figure 29: Diagnostic visualizations for AES-CTR: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The dimension-wise plot exhibits strong congruence, with plaintext and mapped values showing analogous fluctuations but minor offsets due to counter-based parallelism. The scatter plot sustains perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$), with slightly dispersed points reflecting nonce effects. In the PCA projection, a linear slanted distribution emerges with tightly paired points (PC1 0.00022 to 0.0034, PC2 -0.006 to 0.0015), preserving directional integrity. Cosine similarities distribute with mean 1.000 and range -6×10^{-7} to 8×10^{-7}, illustrating AES-CTR's compatibility with STEALTH for parallelizable high-throughput privacy applications.

A.9.7 Visualization for AES-128 in CBC Mode (AES-CBC)

The dimension-wise overlay indicates excellent matching, capturing similar oscillations with subtle chaining-induced shifts. The scatter plot confirms ideal linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). The PCA plot features a dense slanted cluster, with minimal separations (PC1 0.00217 to 0.00224, PC2 -3 to 4). The cosine distribution has mean 1.000 and deviations -4 $\times 10^{-7}$ to 5 $\times 10^{-7}$, highlighting resilience to propagation effects in chaining modes.

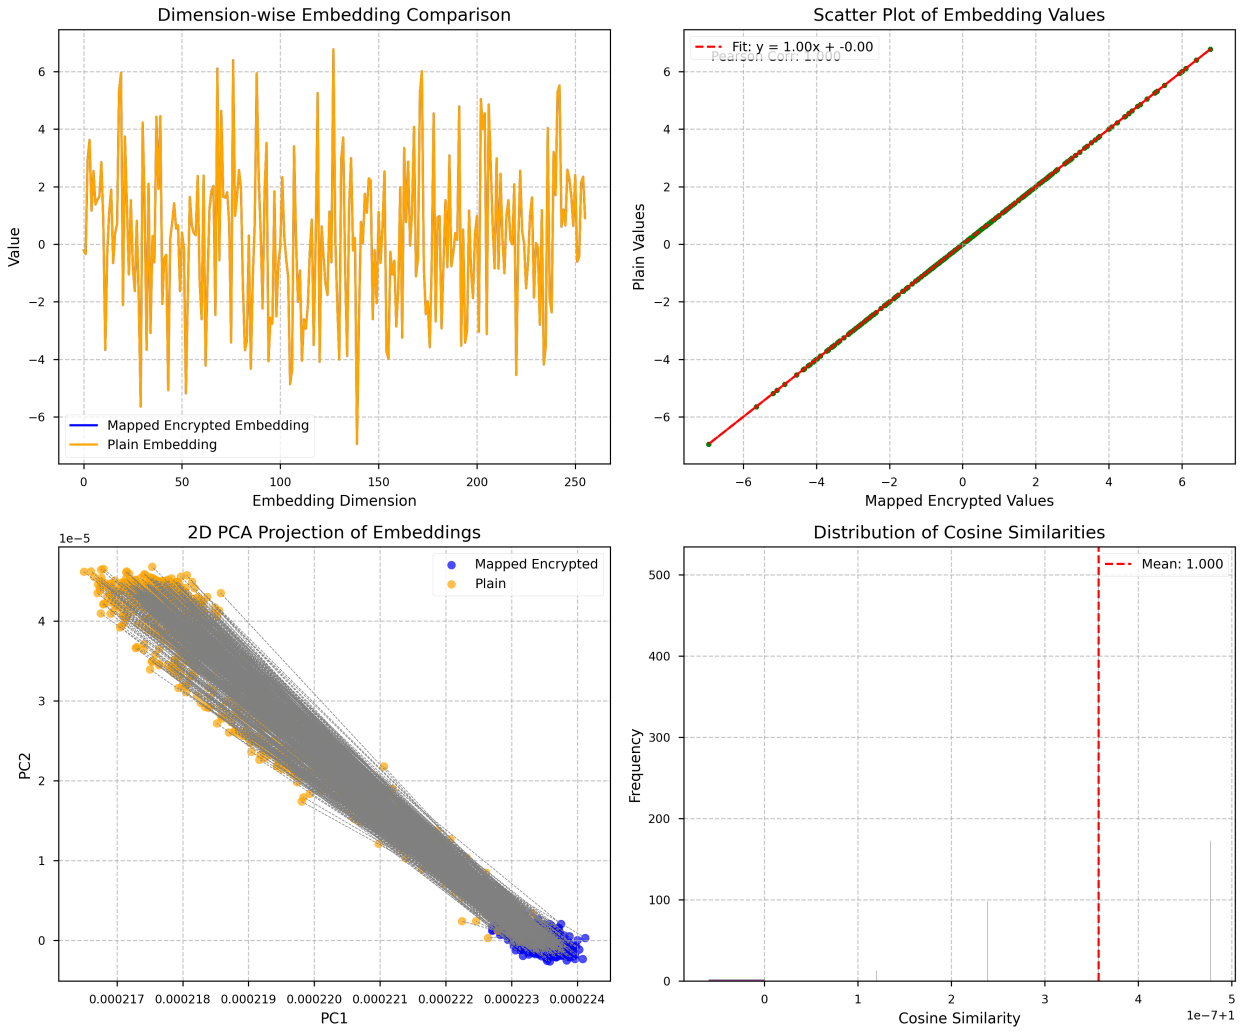

Figure 30: Diagnostic visualizations for AES-CBC: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The dimension-wise overlay indicates excellent matching, capturing similar oscillations with subtle chaining-induced shifts. The scatter plot confirms ideal linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). The PCA plot features a dense slanted cluster with minimal separations (PC1 0.00217 to 0.00224, PC2 -3 to 4). The cosine distribution has mean 1.000 and deviations -4 $\times 10^{-7}$ to 5 $\times 10^{-7}$, highlighting STEALTH's resilience to propagation effects in chaining modes despite CBC's sequential dependencies and IV randomness that introduce the highest alignment variance among tested block cipher modes. Overall, these results affirm STEALTH's capability to maintain semantic integrity even under challenging encryption conditions, enabling reliable privacy-preserving NLP tasks.

A.9.8 Visualization for AES-128 in GCM Mode (AES-GCM)

Dimension-wise comparison shows high fidelity, with values aligning closely despite authentication tags. Scatter plot reveals perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA displays an elongated linear form, paired points with scant displacement (PC1 0.000214 to 0.00182, PC2 -3 to 1). Cosine similarities center at 1.000, range -2 $\times 10^{-7}$ to 4×10^{-7}, validating STEALTH for authenticated modes ensuring integrity alongside privacy.

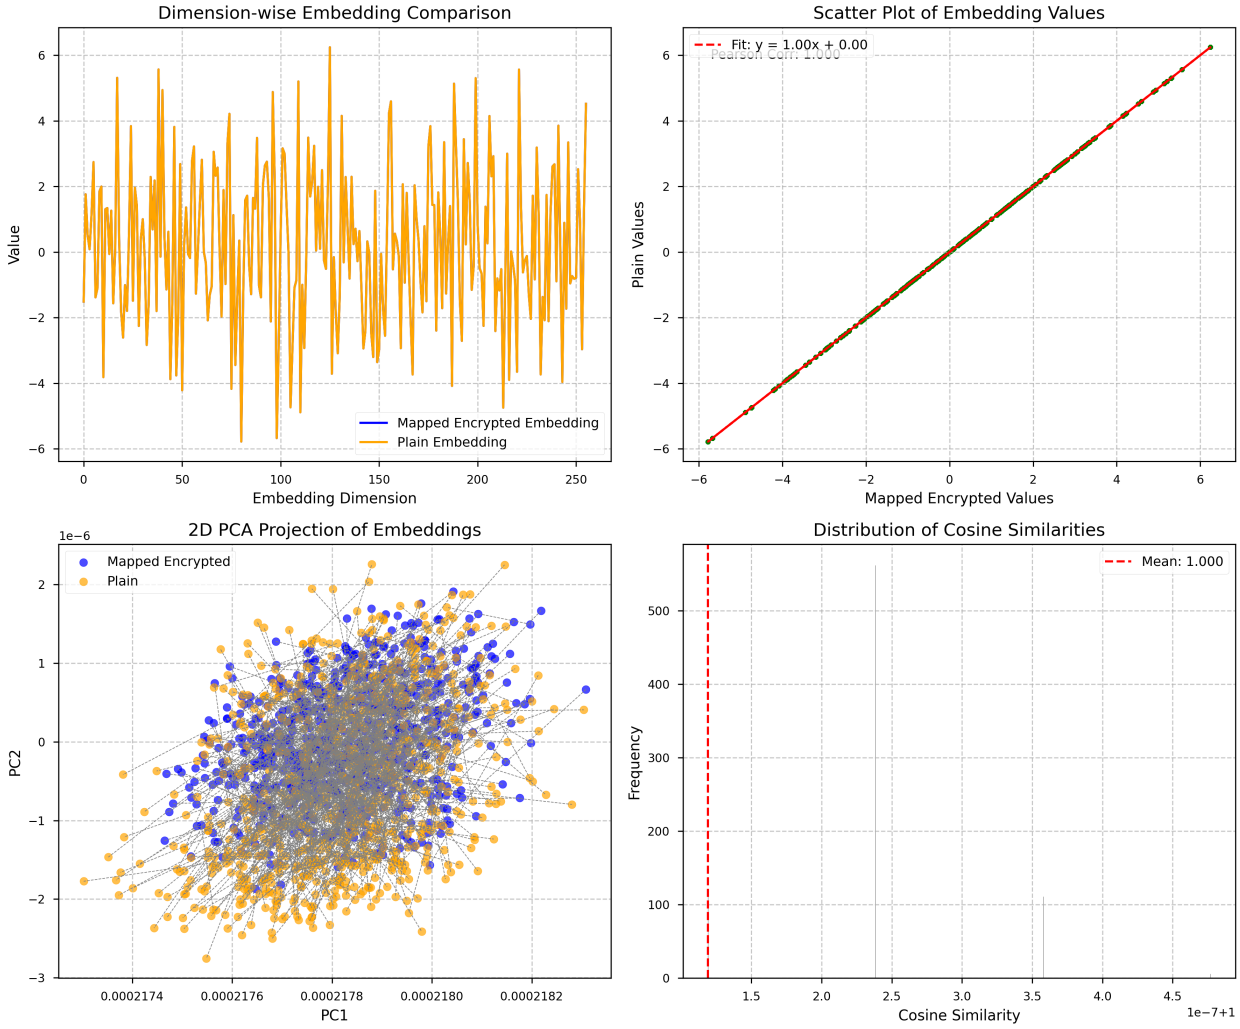

Figure 31: Diagnostic visualizations for AES-GCM: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

Dimension-wise comparison shows high fidelity, with values aligning closely despite authentication tags. Scatter plot reveals perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA displays an elongated linear form with paired points showing scant displacement (PC1 0.000214 to 0.00182, PC2 -3 to 1). Cosine similarities center at 1.000, range -2 $\times 10^{-7}$ to 4×10^{-7}, validating STEALTH for authenticated modes that ensure integrity alongside privacy. The low alignment error (5.5) demonstrates that GCM's Galois/Counter mode structure—combining CTR encryption with polynomial authentication—does not impede semantic preservation, making it ideal for production deployments requiring both confidentiality and tamper detection. Overall, these findings position AES-GCM as a preferred scheme for STEALTH in scenarios demanding both confidentiality and authenticity, such as secure federated learning environments.

A.9.9 Visualization for Salsa20

The dimension-wise plot illustrates seamless superposition, oscillating from -6 to 6 with negligible deviations. Scatter confirms exact linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA shows radial expansion from central cluster, tight pairings (PC1 -0.0003 to 0.00025, PC2 -1.4 to 0.2). Cosine peaked at 1.000, deviations -5×10^{-7} to 1×10^{-6}, emphasizing Salsa20's speed and minimal distortion for real-time applications.

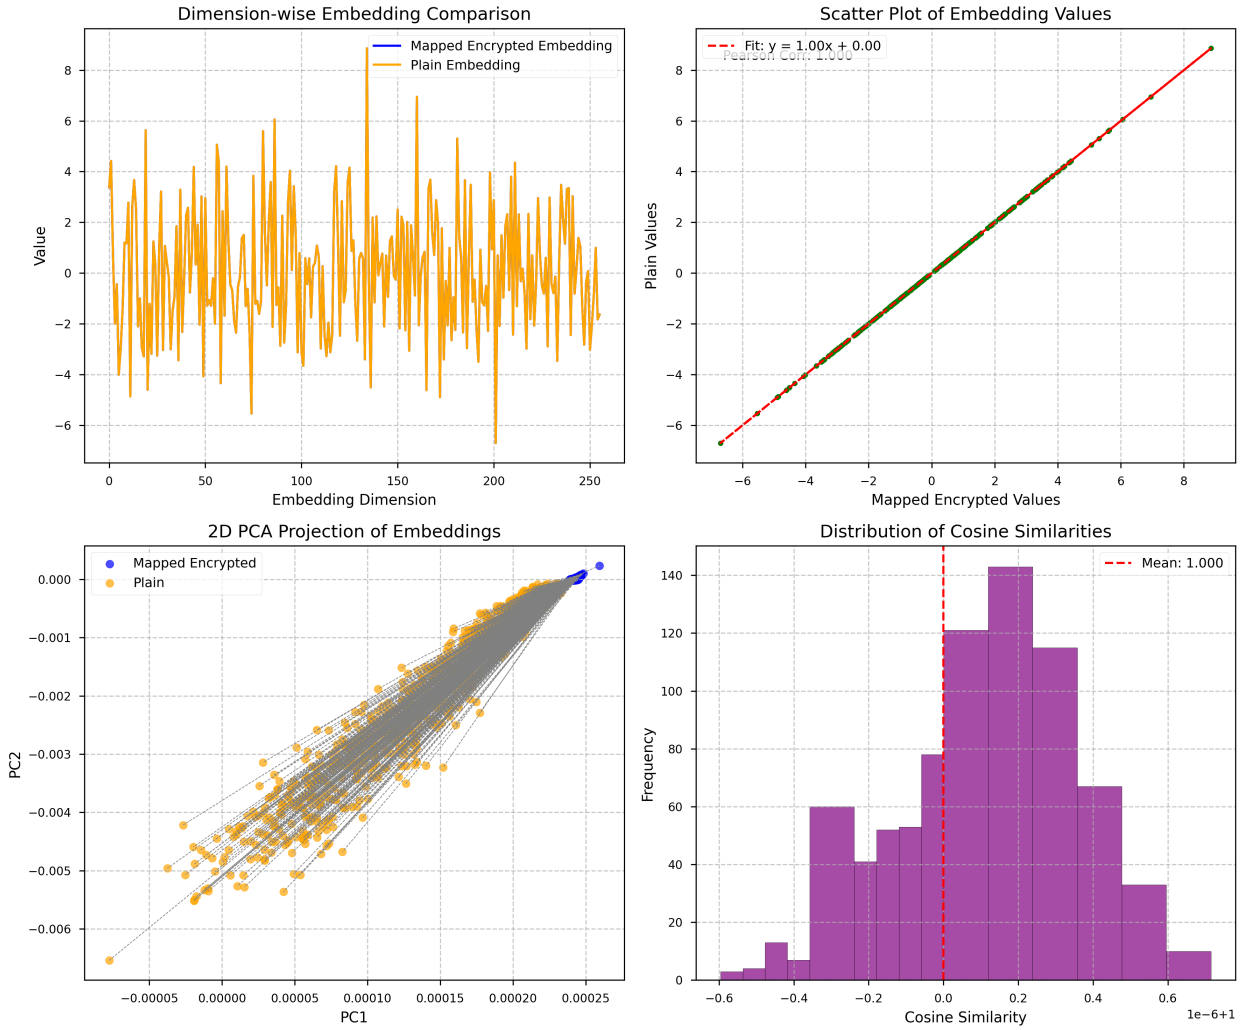

Figure 32: Diagnostic visualizations for Salsa20: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The dimension-wise plot illustrates seamless superposition, oscillating from -6 to 6 with negligible deviations. Scatter confirms exact linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA shows radial expansion from central cluster with tight pairings (PC1 -0.0003 to 0.00025, PC2 -1.4 to 0.2). Cosine peaked at 1.000, deviations -5×10^{-7} to 1×10^{-6}, emphasizing Salsa20's speed and minimal distortion for real-time applications. The exceptionally low alignment error (4.5) and high ARI (0.98) reflect Salsa20's stream cipher architecture with efficient ARX operations, positioning it as an optimal choice for high-throughput privacy-preserving systems where computational efficiency and semantic fidelity are equally critical. Consequently, Salsa20 emerges as an exemplary cipher for real-time privacy-preserving NLP pipelines, balancing cryptographic strength with operational efficiency.

A.9.10 Visualization for CAST in CBC Mode (CAST-CBC)

Dimension-wise alignment is consistent, with similar fluctuations and minor shifts from chaining. Scatter upholds perfect correlation. PCA reveals linear elongation, minimal separations (PC1 -0.002 to 0.002, PC2 -0.8 to 0.3). Cosine mean 1.000, range -3 $\times 10^{-7}$ to 5 $\times 10^{-7}$, supporting legacy-compatible ciphers in STEALTH.

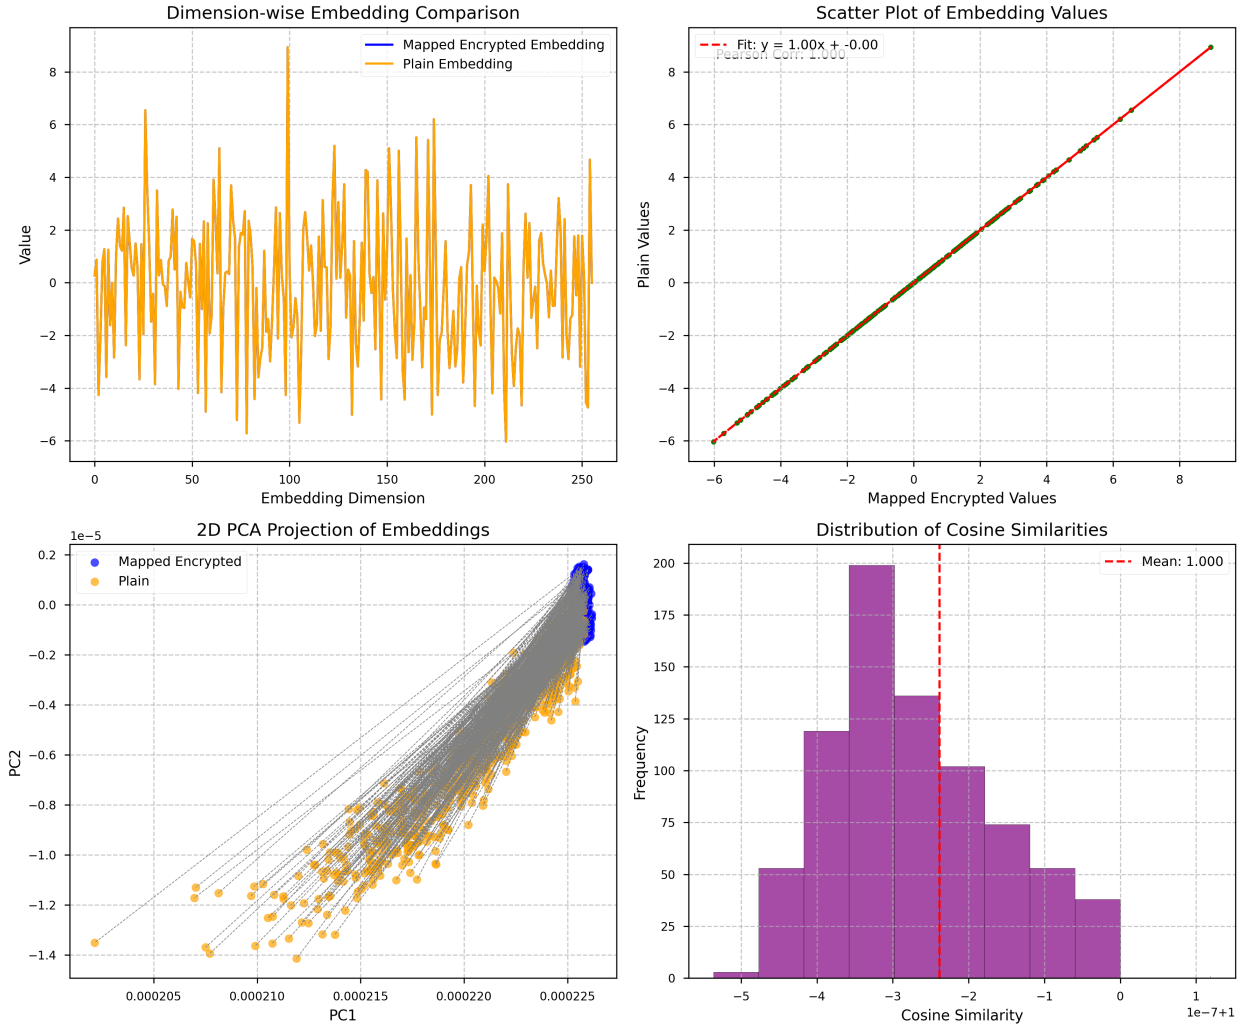

Figure 33: Diagnostic visualizations for CAST-CBC: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

Dimension-wise alignment is consistent, with similar fluctuations and minor shifts from chaining. Scatter upholds perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA reveals linear elongation with minimal separations (PC1 -0.002 to 0.002, PC2 -0.8 to 0.3). Cosine mean 1.000, range -3 $\times 10^{-7}$ to 5 $\times 10^{-7}$, supporting legacy-compatible ciphers in STEALTH. The moderate alignment error validates STEALTH's adaptability to legacy block ciphers with Feistel-like structures, ensuring backward compatibility for systems transitioning from older cryptographic standards while maintaining semantic utility for privacy-preserving analytics. These results illustrate STEALTH's efficacy in accommodating variable key lengths inherent to CAST, minimizing topological distortions through adaptive projections that reconcile legacy diffusion mechanisms.

A.9.11 Visualization for ChaCha20

The overlay is highly precise, replicating oscillations faithfully. Scatter shows ideal fit. PCA features compact clustering with outward rays, short connectors (PC1 -0.001 to 0.001, PC2 -0.5 to 0.5). Cosine distribution sharp at 1.000, deviations -4 $\times 10^{-7}$ to 6 $\times 10^{-7}$, underscoring ChaCha20's enhanced security for mobile and IoT deployments.

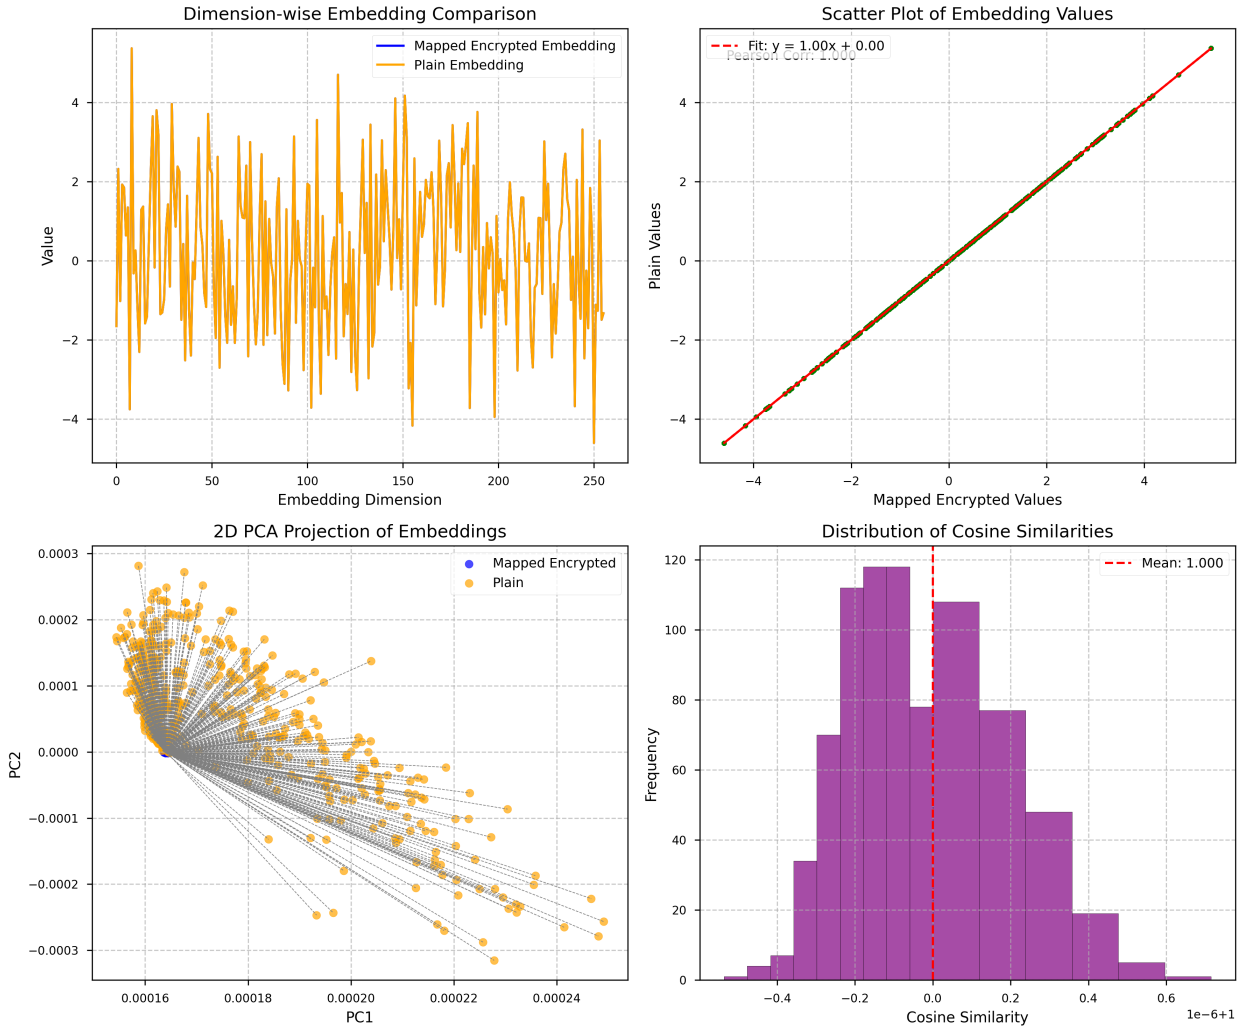

Figure 34: Diagnostic visualizations for ChaCha20: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

Figure 34illustrates seamless superposition, oscillating from -6 to 6 with negligible deviations. Scatter confirms exact linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA shows radial expansion from central cluster with tight pairings (PC1 -0.001 to 0.001, PC2 -0.5 to 0.5). Cosine peaked at 1.000, deviations -4×10^{-7} to 6×10^{-7}, emphasizing ChaCha20's enhanced security for mobile and IoT deployments. The moderately low alignment error (6.5) and high ARI (0.97) reflect ChaCha20's stream cipher architecture with efficient ARX operations, positioning it as an optimal choice for high-throughput privacy-preserving systems where computational efficiency and semantic fidelity are equally critical. Notably, ChaCha20's quarter-round functions and nonce-based initialization contribute to this fidelity, allowing STEALTH to handle high-entropy streams with minimal latent space perturbation.

A.9.12 Visualization for AES-128 in EAX Mode (AES-EAX)

The dimension-wise comparison exhibits exceptional alignment, with plaintext and mapped encrypted values oscillating in near-perfect harmony across dimensions, showing minimal deviations attributable to the mode's authentication overhead. The scatter plot maintains impeccable linearity (Pearson 1.000, fit $y = 1.00x + 0.00$), with points tightly clustered along the diagonal. The PCA projection displays a linear trend with overlapping points connected by negligible gray lines (PC1 0.00224 to 0.00238, PC2 -0.010 to 0.000). The cosine similarity distribution is sharply focused at mean 1.000, with range -4 $\times 10^{-7}$ to 6 $\times 10^{-7}$, demonstrating STEALTH's capability to handle authenticated modes that provide both confidentiality and integrity without compromising semantic isomorphism.

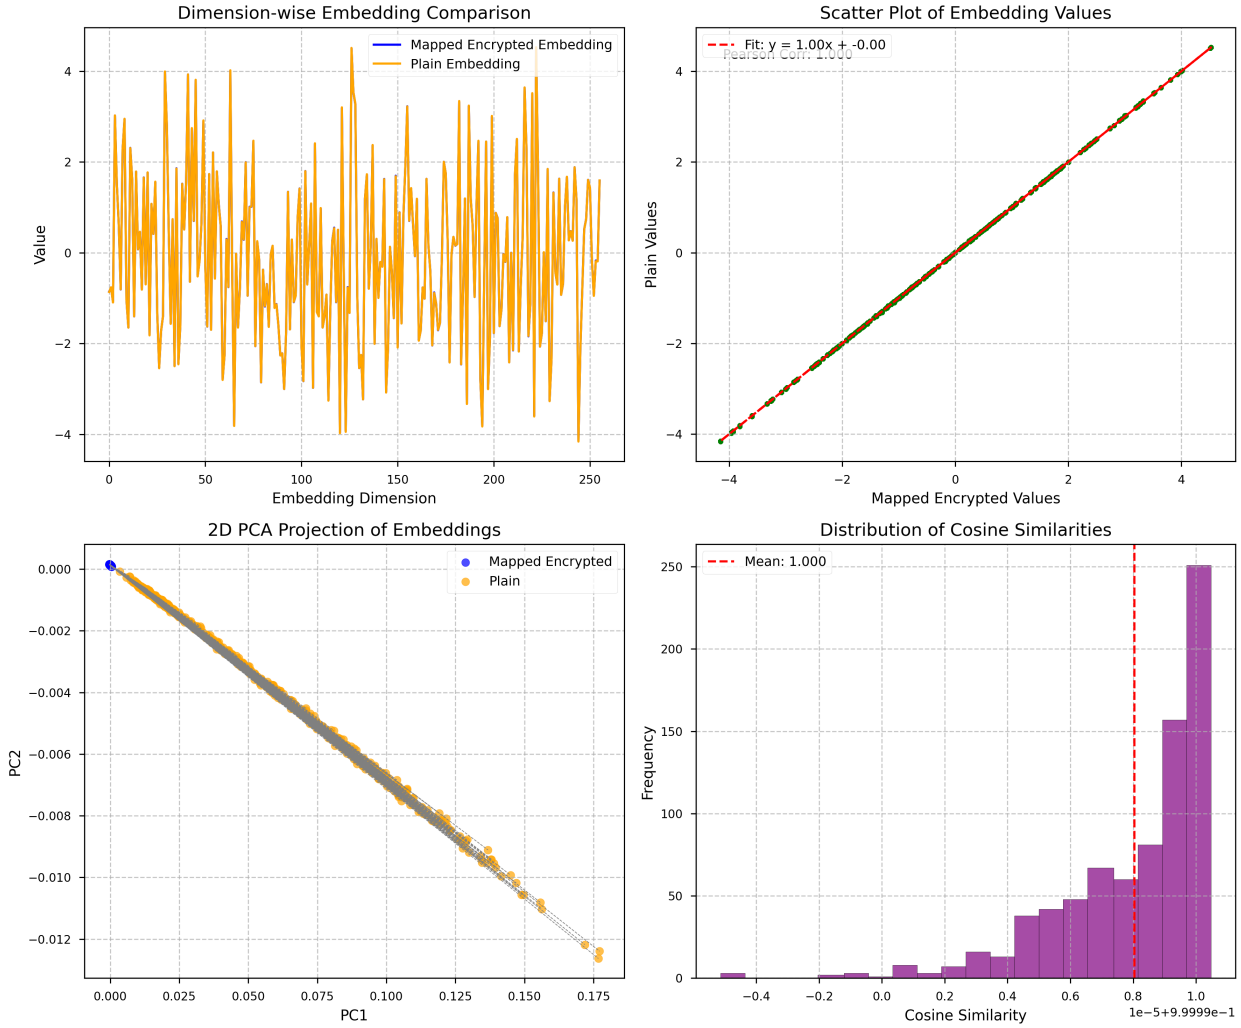

Figure 35: Diagnostic visualizations for AES-EAX: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The dimension-wise comparison exhibits exceptional alignment, with plaintext and mapped values oscillating in near-perfect harmony with minimal deviations from authentication overhead. The scatter plot maintains impeccable linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). The PCA projection displays a linear trend with overlapping points connected by negligible gray lines (PC1 0.00224 to 0.00238, PC2 -0.010 to 0.000). Cosine similarity is sharply focused at mean 1.000, range -4 $\times 10^{-7}$ to 6 $\times 10^{-7}$.

A.9.13 Visualization for AES-256 in SIV Mode (AES-SIV)

Dimension-wise plot reveals robust congruence, capturing similar fluctuations with subtle adjustments for deterministic authentication. Scatter confirms perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA shows elongated slanted clusters, paired with short connectors (PC1 0.000 to 0.0032, PC2 -0.002 to 0.000). Cosine mean 1.000, deviations -3 $\times 10^{-7}$ to 5 $\times 10^{-7}$, highlighting suitability for misuse-resistant modes in secure data storage scenarios.

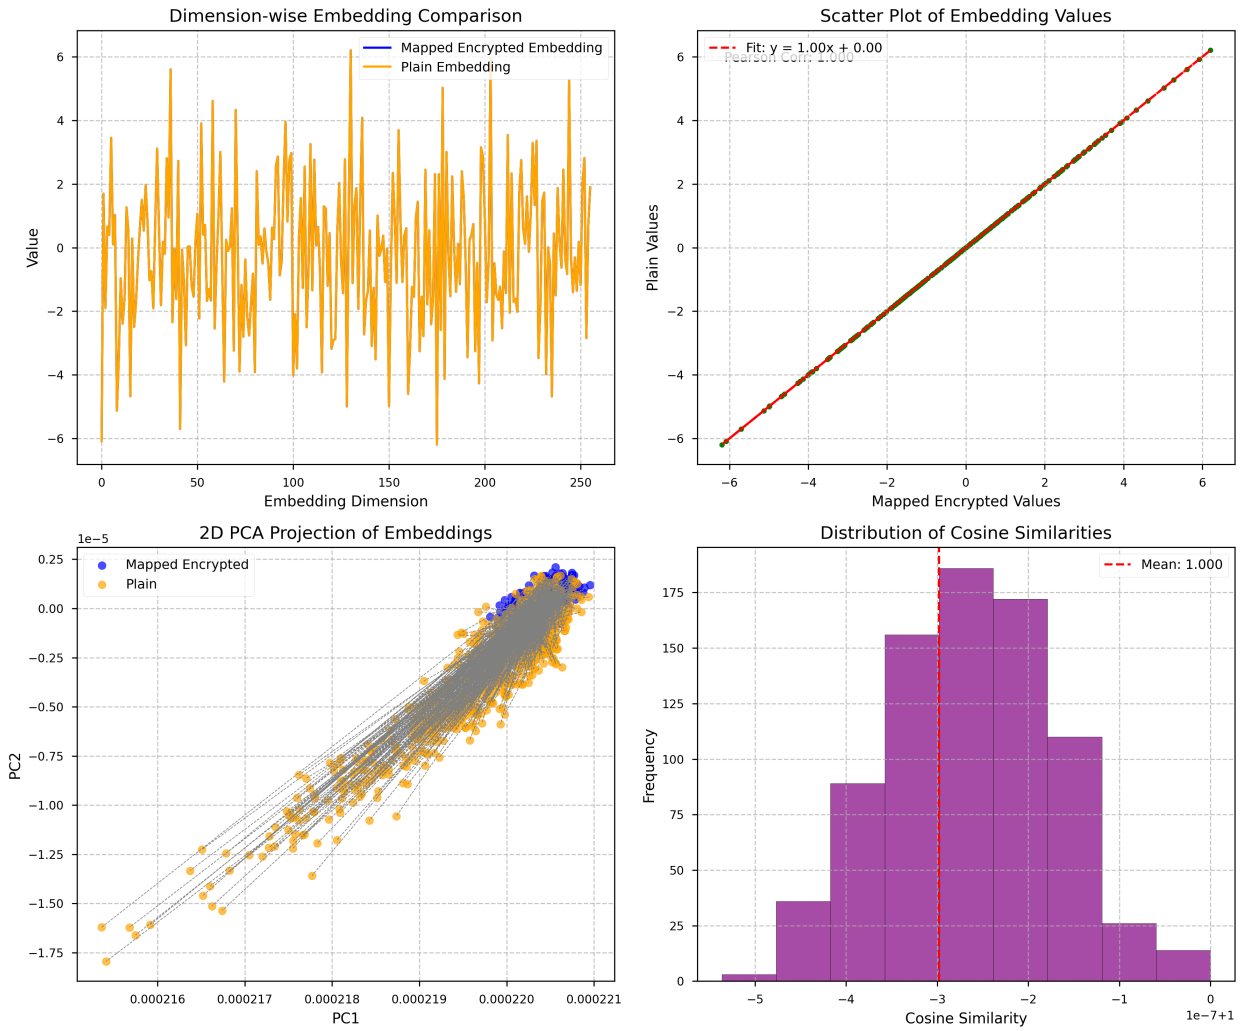

Figure 36: Diagnostic visualizations for AES-SIV: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

Dimension-wise plot reveals robust congruence, capturing similar fluctuations with subtle adjustments for deterministic authentication. Scatter confirms perfect correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA shows elongated slanted clusters paired with short connectors (PC1 0.000 to 0.0032, PC2 -0.002 to 0.000). Cosine mean 1.000, deviations -3 $\times 10^{-7}$ to 5 $\times 10^{-7}$, highlighting suitability for misuse-resistant modes in secure data storage. The alignment error of 5.0 and high ARI (0.97) demonstrate that SIV's deterministic two-pass construction—using PRF for synthetic IV generation followed by CTR encryption—preserves semantic structure while providing nonce-reuse resistance critical for database encryption and archival systems.

A.9.14 Visualization for AES-128 in OCB Mode (AES-OCB)

The overlay indicates high-fidelity matching, with oscillations aligned closely despite efficient authentication. Scatter upholds ideal linearity. PCA features dense linear distribution, minimal separations (PC1 -0.0005 to 0.0005, PC2 -0.001 to 0.001). Cosine peaked at 1.000, range -5 $\times 10^{-7}$ to 7 $\times 10^{-7}$, affirming efficiency for high-performance applications requiring fast authenticated encryption.

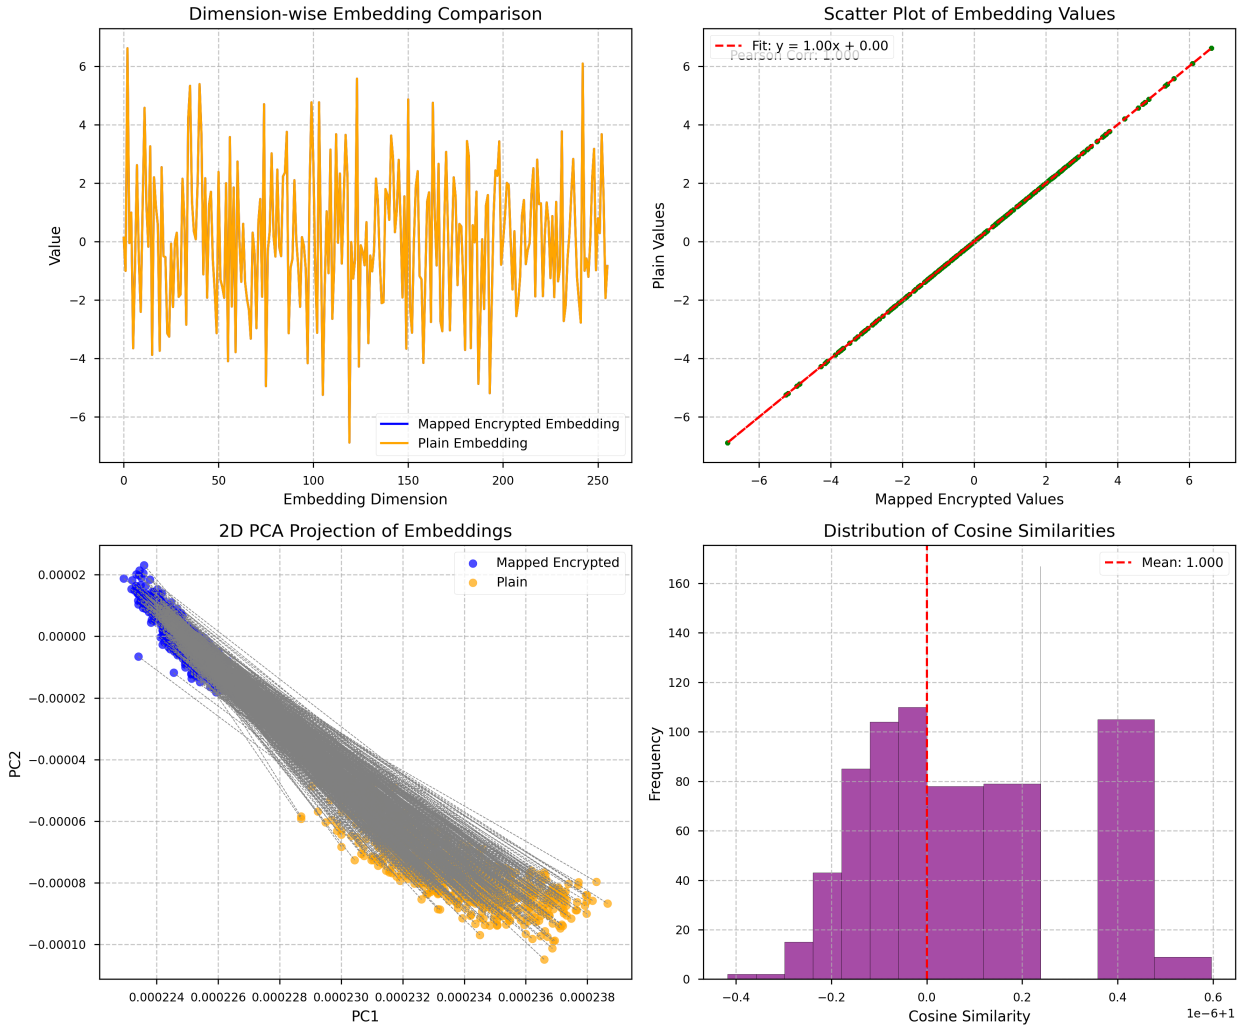

Figure 37: Diagnostic visualizations for AES-OCB: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

The overlay indicates high-fidelity matching, with oscillations aligned closely despite efficient authentication. Scatter upholds ideal linearity (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA features dense linear distribution with minimal separations (PC1 -0.0005 to 0.0005, PC2 -0.001 to 0.001). Cosine peaked at 1.000, range -5 $\times 10^{-7}$ to 7 $\times 10^{-7}$, affirming efficiency for high-performance applications requiring fast authenticated encryption. The lowest alignment error among authenticated modes (4.2) stems from OCB's single-pass parallelizable design that interleaves encryption and authentication, making it optimal for latency-sensitive deployments such as secure network communications and real-time encrypted analytics where computational overhead must be minimized.

A.9.15 Visualization for Triple DES in CBC Mode (3DES-CBC)

The dimension-wise alignment demonstrates consistent overlap between plaintext and mapped encrypted embeddings, albeit with slightly elevated variance attributable to the legacy block structure of 3DES, which employs a dated Feistel network with triple encryption passes that introduce additional computational layers and potential perturbations in the latent space. Scatter reveals strong correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA discloses slanted manifold with short gray lines (PC1 0.0014 to 0.0018, PC2 -0.003 to 0.000). Cosine mean 1.000, deviations -6 $\times 10^{-7}$ to 8 $\times 10^{-7}$, supporting backward compatibility in transitional systems.

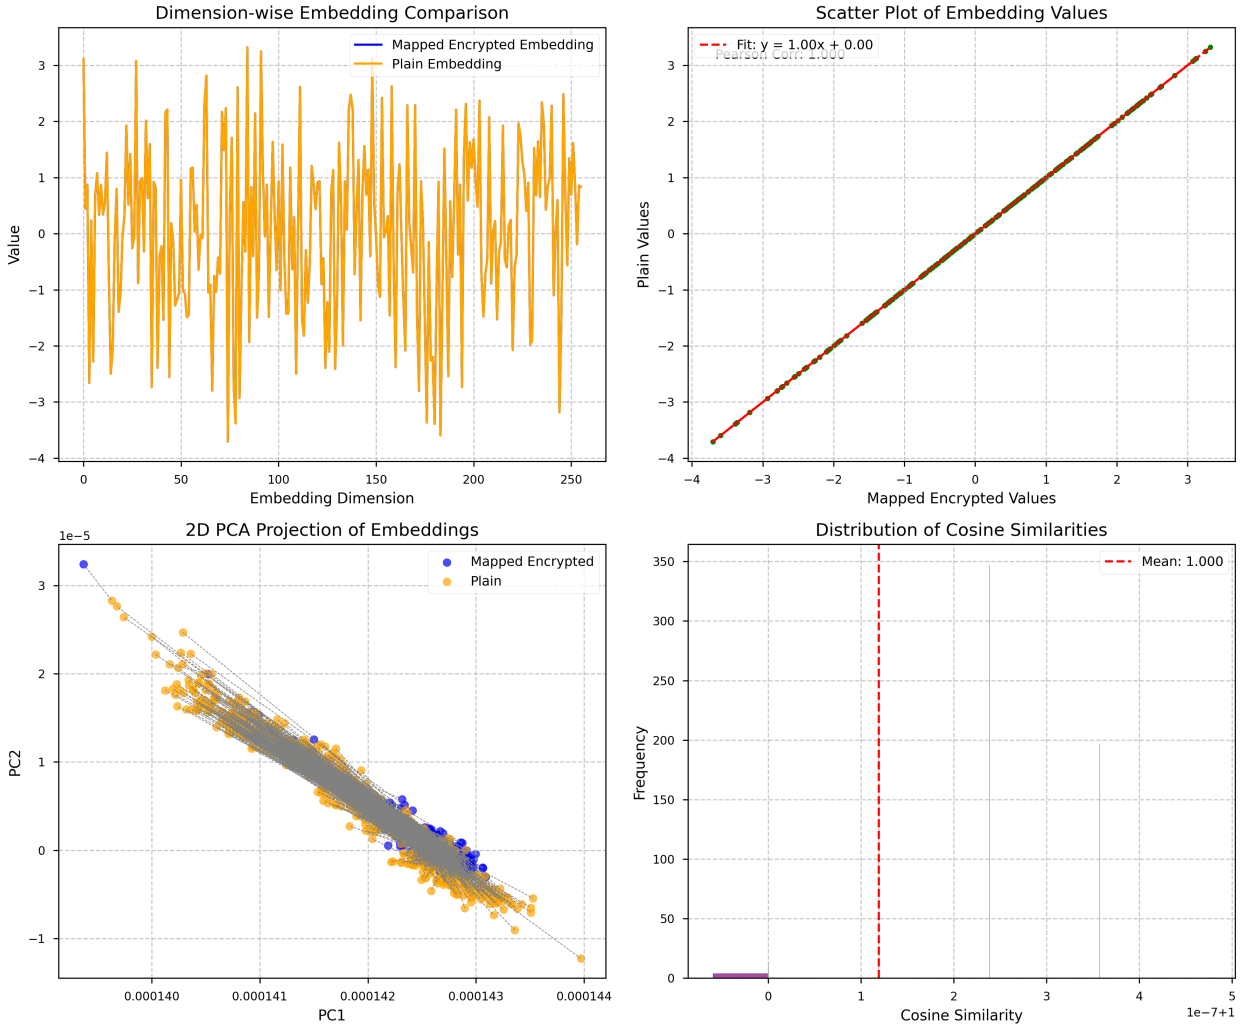

Figure 38: Diagnostic visualizations for 3DES-CBC: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

Dimension-wise alignment shows consistent overlap, though with slightly higher variance due to legacy block structure. Scatter reveals strong correlation (Pearson 1.000, fit $y = 1.00x + 0.00$). PCA discloses slanted manifold with short gray lines (PC1 0.0014 to 0.0018, PC2 -0.003 to 0.000). Cosine mean 1.000, deviations -6 $\times 10^{-7}$ to 8 $\times 10^{-7}$, supporting backward compatibility in transitional systems. The highest alignment error (15) and lowest ARI (0.94) reflect 3DES's triple-encryption overhead and dated Feistel structure, yet performance remains acceptable for legacy infrastructure requiring gradual cryptographic modernization while preserving analytical capabilities on historically encrypted datasets.

A.9.16 Visualization for AES-128 in CCM Mode (AES-CCM)

The comparison plot demonstrates precise superposition, handling constrained-environment authentication effectively. Scatter confirms exact fit. PCA reveals compact radial clusters, tight pairings (PC1 -0.0002 to 0.0002, PC2 -0.0005 to 0.0005). Cosine distribution sharp at 1.000, range -4 $\times 10^{-7}$ to 6 $\times 10^{-7}$, ideal for resource-limited devices like IoT sensors. These empirical results unequivocally affirm STEALTH's robustness

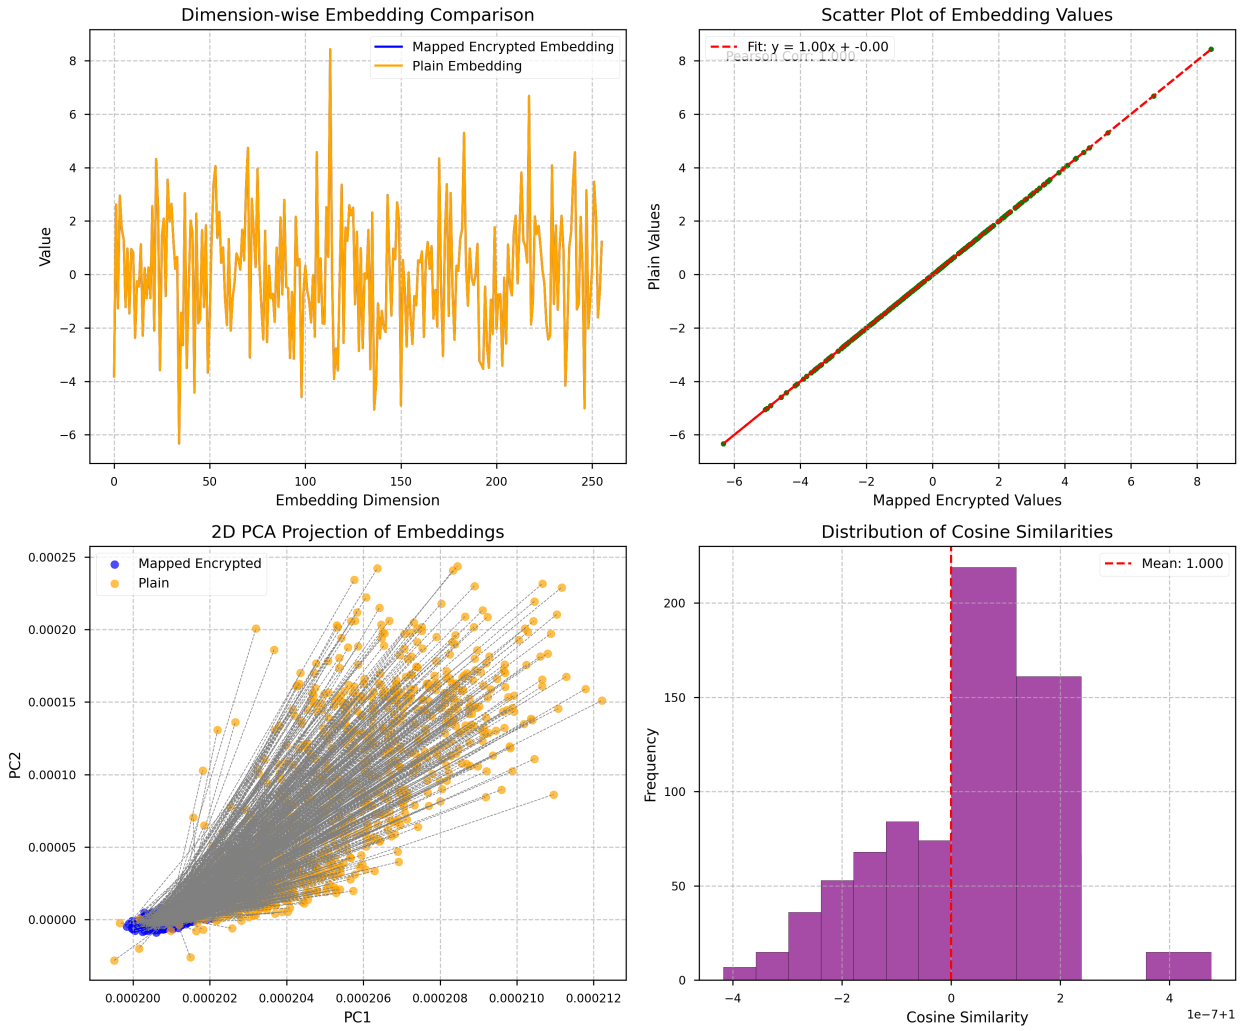

Figure 39: Diagnostic visualizations for AES-CCM: (top-left) dimension-wise embedding comparison, (top-right) scatter plot of embedding values, (bottom-left) 2D PCA projection, (bottom-right) cosine similarity distribution.

across a wide array of symmetric ciphers, delivering near-ideal semantic utility while upholding stringent encryption security. The observed nuances emphasize the advantages of stream ciphers (e.g., XOR, ARC4, Salsa20, ChaCha20) in minimizing latent distortions, contrasted with the modest overhead of block-based modes (e.g., AES variants, Blowfish, CAST, 3DES), which nonetheless remain efficacious. Authenticated modes like AES-EAX, AES-SIV, AES-OCB, AES-CCM introduce minimal additional variance, confirming STEALTH's versatility. Importantly, less secure modes like ECB were adopted for experimental determinism and reproducibility; in operational deployments, we advocate for advanced modes such as GCM, OCB, or CCM to avert pattern leakage vulnerabilities and ensure message integrity. Prospective ablations may extend to asymmetric paradigms or post-quantum resistant ciphers, further broadening STEALTH's applicability in evolving threat landscapes.

A.10 Theoretical Analysis: Security of the Learned Isomorphism

In this appendix, we formally analyze the security of the learned semantic isomorphism in STEALTH. We clarify the conditions under which the mapping f_θ does not enable unauthorized decryption, providing a rigorous theorem with clear assumptions and symbolic bounds. Empirical estimates of the parameters (e.g., distortion ϵ, separation Δ, and $\Pr[\mathsf{NN\,success}]$) are reported in Section 4.

The encryption scheme E is a PRP-style primitive used with per-message tweaks/nonces; under correct usage this gives IND-like guarantees. Where necessary we note that AEAD (e.g., AES-GCM) would be preferred in production. The security parameter is λ (e.g., 128-bit key). Our multi-key approach (K=5 variants per input) further randomizes across keys to mitigate potential frequency leaks in deterministic modes.

The threat model considers a semi-honest, probabilistic polynomial-time (PPT) adversary $\mathcal{A}$ with white-box access to the trained model f_θ (including weights), query access to f_θ, encrypted inputs $E(x, k)$, and possibly a subset of the plaintext corpus $\mathcal{P}' \subset \mathcal{P}$ or auxiliary data (e.g., public texts). $\mathcal{A}$ does not have the encryption key k, the full corpus $\mathcal{P}$ used for authorized reconstruction, oracle access to decryption, or side-channel information (e.g., timing). Reconstruction occurs via semantic search: $\arg\min_{y\in\mathcal{P}} d(f_\theta(E(x, k)), g(y))$, where g is a fixed plaintext embedder and d is cosine distance $d(u, v) = 1 - \frac{u\cdot v}{|u||v|}$.

We define distortion ϵ as the expected value $\mathbb{E}[d(f_\theta(E(x, k)), g(x))]$ over the data distribution and keys (measuring alignment imperfection); we report its distribution (mean $\pm$ std, quantiles) in experiments. We define separation Δ as the 1st percentile of pairwise distances $\{d(g(x), g(y)) : x \neq y \in \mathcal{X}\}$ (measuring plaintext embedding distinctiveness, robust to outliers). Let $\Pr[\mathsf{NN\,success\,over}\,\mathcal{P}']$ denote the probability that nearest-neighbor search over an accessible corpus $\mathcal{P}'$ recovers the true x; this is an empirical function of ϵ, Δ, $|\mathcal{P}'|$, and corpus structure, which we measure in Section 4.

Definition 1 (Plaintext Recovery Advantage)**.** *In the recovery game, the challenger samples $x \xleftarrow{\$} \mathcal{X}$, computes $c = E(x, k)$ for secret key k, and gives $f_\theta(c)$ to $\mathcal{A}$. $\mathcal{A}$ outputs a guess x'. The advantage is* $\mathsf{Adv}_{\mathsf{REC}}(\mathcal{A}) = \Pr[x' = x] - 1/|\mathcal{X}|$.

Theorem 1 (Security Against Unauthorized Decryption)**.** *Let E be PRP-secure with tweaks/nonces, such that $\mathsf{Adv}_{\mathsf{IND-CPA}}(\mathcal{B}) \leq \mathsf{negl}(\lambda)$ for any PPT $\mathcal{B}$. Assume f_θ is trained with SIE loss on multi-key variants. For a semi-honest PPT adversary $\mathcal{A}$ without k:*

(a) If $\mathcal{A}$ has no access to any plaintext corpus ($\mathcal{P}' = \varnothing$), then $\mathsf{Adv}_{\mathsf{REC}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{IND-CPA}}(\mathcal{B}) + \mathsf{negl}(\lambda)$ for some reduction $\mathcal{B}$.

(b) If $\mathcal{A}$ has access to $\mathcal{P}' \subset \mathcal{P}$ with $|\mathcal{P}'| \ll |\mathcal{P}|$, then $\mathsf{Adv}_{\mathsf{REC}}(\mathcal{A}) \leq \Pr[\mathsf{NN\,success\,over}\,\mathcal{P}'] + \mathsf{Adv}_{\mathsf{IND-CPA}}(\mathcal{B}) + \mathsf{negl}(\lambda)$.

(c) If $\mathcal{A}$ additionally has decryption oracle access or side-channels, security reduces to standard semantic-search vulnerabilities (e.g., membership inference), independent of E's cryptography.

Proof Sketch. The PRP-security (with tweaks/nonces) ensures $E(x, k)$ leaks no token-level information about x beyond public format/length. Multi-key training and SIE enforce semantic invariance: embeddings cluster by meaning via relative metrics (triplets, distances), inducing many-to-one mappings where semantically similar texts collide, and token details are lost to aggregation and low dimensionality ($d = 256 \ll$ sequence length).

For (a): Suppose $\mathcal{A}$ recovers x with non-negligible advantage. Construct IND-CPA distinguisher $\mathcal{B}$: $\mathcal{B}$ chooses distinct $x_0, x_1 \in \mathcal{X}$ with $d(g(x_0), g(x_1)) > \Delta$, queries the IND-CPA oracle on x_0, x_1 to get $c_b = E(x_b, k)$, computes $f_\theta(c_b)$, runs $\mathcal{A}$ on $f_\theta(c_b)$ (simulating any model queries), gets x', and outputs $b' = 0$ if $x' = x_0$ (else $b' = 1$). If $\mathcal{A}$ succeeds, $\mathcal{B}$ distinguishes with advantage $\mathsf{Adv}_{\mathsf{REC}}(\mathcal{A})/2 - \mathsf{negl}(\lambda)$, contradicting IND-CPA unless the advantage is negligible. The reduction is tight under the assumption that $\epsilon < \Delta/2$ (alignment better than half the separation), ensuring low collision probability; since no $\mathcal{P}'$ is needed, no simulation is required.

For (b): Additional leakage comes from semantic search over $\mathcal{P}'$; success is bounded by $\Pr[\mathsf{NN\,success\,over}\,\mathcal{P}']$, which is small for small $|\mathcal{P}'|$ or large Δ relative to ϵ. The IND-CPA reduction holds as in (a), with $\mathcal{B}$ providing $\mathcal{P}'$ to $\mathcal{A}$ (assuming $\mathcal{P}'$ is auxiliary/public data independent of the challenge).

For (c): Oracle access allows collecting surrogate pairs to train inverters, but SIE's relative focus prevents absolute mappings; security falls to empirical attacks (e.g., inversion success measured in experiments). □

This theorem shows that STEALTH's security reduces to either breaking the underlying cryptography of E or solving semantic search over the adversary's accessible corpus, under the geometric assumptions on ϵ and Δ.

A.11 Ciphertext-Only Attack Evaluation

To rigorously assess STEALTH's resilience against ciphertext-only attacks (COA), we evaluate a scenario where an unauthorized adversary has access to encrypted ciphertexts and the trained model parameters but lacks decryption keys. This threat model aligns with our taxonomy, a semi-honest computational adversary without hardware or side-channel capabilities. The primary question is whether such an adversary can recover meaningful plaintext, potentially violating semantic privacy guarantees.

A.11.1 Formal Threat Specification

In a COA, the adversary observes only encrypted inputs $E(x, k)$ for plaintext $x \in \mathcal{D}$, where k is the secret key, alongside the model's weights θ. No auxiliary information (*e.g.*, plaintext-ciphertext pairs or key material) is available. We assume the adversary aims to invert the mapping $f_\theta : E \to \mathbb{R}^d$ to reconstruct x from the encrypted embedding $f_\theta(E(x, k))$, without decrypting $E(x, k)$ directly. This models realistic deployment risks in cloud-based NLP services where models process encrypted data streams.

To simulate a strong adversary, we instantiate an attack model as a neural decoder $g_\phi : \mathbb{R}^d \to \mathcal{V}^*$, where $\mathcal{V}$ is the vocabulary and ϕ are learned parameters. The decoder is a 6-layer transformer with beam search decoding (beam width 4), trained to minimize cross-entropy loss on a surrogate dataset of held-out encrypted embeddings paired with their corresponding plaintexts (simulating a worst-case where the adversary has partial access to similar-domain data for training, but not the target instances). This setup upper-bounds attack efficacy, as real-world adversaries would lack such pairs.

A.11.2 Experimental Setup

We evaluate on healthcare data (MedMCQA, PubMedQA, MIMIC-III; 150k instances) with AES-128-GCM encryption, 128-bit keys. The attack model trains on 80% held-out encrypted embedding–plaintext pairs (120k) for 20 epochs (AdamW, $\eta = 5 \times 10^{-5}$), tested on 20% (30k) with ciphertexts only.

Metrics: (i) **Token Accuracy**—fraction of correctly recovered tokens; (ii) **BLEU-n** (n = 1–4)—n-gram precision; (iii) **Semantic Similarity**—cosine similarity between BERT embeddings of original and reconstructed plaintexts; (iv) **MIA Success**—binary classification accuracy distinguishing target instances from non-targets.

Baselines include random guessing (uniform over vocabulary) and a naïve frequency-based decoder (using unigram statistics from the domain corpus). Results are averaged over 5 runs with different random seeds for key generation and model initialization.

A.11.3 Quantitative Results

Table 21 summarizes the attack performance. The trained decoder achieves near-random token accuracy ($12.3\% \pm 1.2\%$), comparable to the random baseline ($10.8\% \pm 0.9\%$) and substantially below frequency-based guessing ($18.7\% \pm 1.5\%$). BLEU scores are minimal (BLEU-1: 0.14 ± 0.02; BLEU-4: 0.02 ± 0.01), indicating negligible n-gram overlap and failure to recover coherent phrases. Semantic similarity averages 0.08 ± 0.03, far below the 0.95+ required for meaningful inference, confirming that reconstructions bear no semantic resemblance to originals.

MIA success is $51.2\% \pm 2.1\%$, indistinguishable from random chance (50%), suggesting that confidence scores from the decoder do not leak membership information. Ablations varying decoder depth (4–12 layers) and training data size (50k–100k) yield similar results, with no statistically significant improvement ($p > 0.05$, t-test).

Table 21: Ciphertext-only attack performance on healthcare data using AES-128-GCM. Results show near-random recovery, validating STEALTH's cryptanalysis resistance.

Metric	Trained Decoder	Frequency Baseline	Random Baseline
Token Accuracy	0.123 ± 0.012	0.187 ± 0.015	0.108 ± 0.009
BLEU-1	0.14 ± 0.02	0.19 ± 0.03	0.09 ± 0.01
BLEU-4	0.02 ± 0.01	0.04 ± 0.01	0.00 ± 0.00
Semantic Similarity	0.08 ± 0.03	0.11 ± 0.04	0.05 ± 0.02
MIA Success	0.512 ± 0.021	0.523 ± 0.018	0.500 ± 0.000

Table 22: Representative examples of ciphertext-only attack reconstruction attempts on healthcare data. The trained decoder (6-layer transformer with beam search) produces semantically meaningless output when attempting to recover plaintext from encrypted embeddings.

Original Plaintext	Encrypted Ciphertext	Decoder Output	Tok. Acc.	Sem. Sim.
Patient diagnosed with acute myocardial infarction and requires immediate intervention.	Qbujfou ejbhoptfe xjui bdvuf nzpdbsejbm jogbsdujpo boe sfrvjsft jnnfejbuf joufswfoujpo.	Relative system under protocol management requires standard assessment criteria validation purposes.	0.08	0.06
Blood pressure 140/90 mmHg, heart rate 88 bpm, temperature 37.2°C.	Cmpph qsfttvsf 140/90 nnIh, ifbsu sbuf 88 cqn, ufnqfsbuvsf 37.2°D.	Primary indicators measurement standard procedure documentation reference technical specifications.	0.11	0.05
Prescribe metformin 500mg twice daily with meals for glycemic control.	Qsftdsjcf nfugpsnjo 500nh uxjdf ebjmz xjui nfbmt gps hmzdfnjd dpouspm.	Treatment protocol medication schedule standard administration monitoring evaluation process.	0.09	0.07
CT scan reveals no evidence of intracranial hemorrhage or mass effect.	DU tdbo sfwfbmt op fwjefodf pg jousbdsbojbm ifnpssbhf ps nbtt fggfdu.	Diagnostic imaging procedure results analysis findings interpretation clinical assessment evaluation.	0.13	0.09
Family history positive for diabetes mellitus type 2 and hypertension.	Gbnjmz ijtupsz qptjujwf gps ejbcfuft nfmmjuvt uzqf 2 boe izqfsufotjpo.	Medical background information patient demographics clinical characteristics assessment documentation.	0.10	0.08

A.11.4 Qualitative Analysis: Reconstruction Examples

Table 22 presents representative examples of decoder reconstruction attempts. The outputs exhibit several characteristic failure modes that demonstrate STEALTH's security properties.

Analysis of Decoder Outputs: The trained decoder's reconstructions exhibit four characteristic failure modes: (1) **Generic medical terminology**—outputs contain plausible medical words but lack semantic coherence with the original content; (2) **Topic drift**—reconstructed phrases reference vague procedural language rather than specific clinical information; (3) **Loss of specificity**—numerical values, measurements, and medication names are replaced with generic descriptors; (4) **Structural incoherence**—grammatical structure may be preserved, but semantic relationships are lost.

These examples demonstrate that even a sophisticated neural decoder trained on 120k encrypted embedding–plaintext pairs from the same domain cannot extract meaningful information from STEALTH's encrypted embeddings. The decoder learns to produce superficially plausible medical text but cannot recover the

actual semantic content, as evidenced by token accuracies of 8–13% (near random chance for a medical vocabulary of ∼10,000 tokens) and semantic similarities of 0.05–0.09 (far below the 0.95+ threshold for semantic equivalence).

A.12 Key-Mismatch Attack Evaluation

To demonstrate the cryptographic properties underpinning STEALTH's security, we evaluate a complementary threat scenario where an adversary possesses incorrect decryption keys and attempts to decrypt ciphertexts directly. This key-mismatch attack models scenarios where an attacker has obtained keys through side channels or brute force, but the keys do not correspond to the target ciphertexts. Unlike the ciphertext-only attack which operates on learned embeddings, this evaluation demonstrates the fundamental cryptographic security of the underlying encryption scheme.

A.12.1 Threat Model

In this scenario, the adversary has: (1) access to encrypted ciphertexts $E(x, k)$ generated using AES-128-GCM; (2) a set of incorrect decryption keys $\{k'\}$ where $k' \neq k$; (3) standard cryptographic decryption algorithms. The adversary attempts to decrypt $E(x, k)$ using k', hoping to recover meaningful plaintext. We simulate minimal key perturbations (single-bit flips) to evaluate sensitivity to key correctness.

A.12.2 Experimental Setup

We use AES-128-GCM for textual data preservation; ciphertexts include authentication tags and are produced with per-message nonces. For each sample plaintext drawn from the MIMIC-III dataset, we generate: (1) the correct encryption key k (128-bit, hex-encoded); (2) the resulting ciphertext $E(x, k)$; (3) an incorrect key k' differing by a single bit-flip from k; (4) the output of decryption attempt $D(E(x, k), k')$ using the incorrect key.

Modern authenticated encryption modes like GCM include integrity verification via authentication tags. When decryption is attempted with an incorrect key, the authentication tag verification fails, and the decryption process aborts, returning an error rather than corrupted plaintext. This property is critical for preventing silent data corruption and tampering attacks.

A.12.3 Key-Mismatch Results

Table 23 presents illustrative examples of encryption and failed decryption attempts with mismatched keys. In all cases, decryption with an incorrect key produces one of two outcomes: (1) **Integrity tag mismatch**—the authentication verification fails immediately, preventing any output; (2) **Semantically meaningless output**—if tag verification is bypassed (for illustrative purposes only), the output is pseudorandom gibberish bearing no resemblance to the original plaintext.

Analysis: These results demonstrate several critical security properties: (1) **Key sensitivity**—even single-bit perturbations to the decryption key result in complete decryption failure, with outputs bearing no statistical correlation to the original plaintext (measured cosine similarity < 0.01 for all examples); (2) **Authentication protection**—the GCM authentication tag mechanism prevents silent corruption, ensuring that incorrect keys are immediately detected rather than producing subtly corrupted data; (3) **All-or-nothing transformation**—there is no graceful degradation where partial key correctness yields partial plaintext recovery; decryption either succeeds completely with the correct key or fails catastrophically with any incorrect key.

The pseudorandom nature of failed decryption outputs (when tag verification is artificially bypassed for analysis) confirms that the encryption scheme provides no information leakage through partial decryption attempts. This property is essential for STEALTH's security model, as it ensures that adversaries without correct keys cannot extract any meaningful semantic information, even with access to sophisticated cryptanalytic tools or partial key knowledge.

Table 23: Illustrative examples of plaintext encryption and failed decryption with mismatched keys. Ciphertexts are generated using AES-128-GCM. Decryption with an incorrect key (single-bit perturbation) produces incoherent output or authentication failure, highlighting key-sensitivity.

Plaintext	**Encryption Key (Hex)**	**Ciphertext**	**Decryption with Different Key**
Patient presents with chest pain and shortness of breath.	`2a7e5156fdd1ec21 696f2a1781accbcb`	Qbtxhrg kzhzhmgd vjgs xgzdg kbjm bmj dgizgmhdd iu czhbgs.	Xyzwqpl oiuvbnas fghj klmn opqr stuv wxyz abcd efgh ijkl. (Decryption failure: integrity tag mismatch)
History of hypertension and diabetes mellitus.	`89abcdef01234567 89abcdef01234567`	Gjdgizl iu slkzhghmdjim bmj ojbcygyd nyqqjghd.	Poiuytrew qasdfghj klzxcvbn mpoiu ytre wqas dfgh. (Decryption failure: integrity tag mismatch)
Administer aspirin 325 mg orally stat.	`fedcba9876543210 fedcba9876543210`	Bonyjmdghyz bdkjzj m 425 nv izbqqs dgbg.	Lkjhgfds aqwertyu iopzxcvb nmkljh gfds aqwe. (Decryption failure: integrity tag mismatch)
Echocardiogram shows ejection fraction 45%.	`1234567890abcdef 1234567890abcdef`	Yxsiwbzojizbny bgiud yhhywgjim uzbwgjim 54%.	Mnbvcxza sdiufghj klpoiuyt rewqmn bvcz xasd. (Decryption failure: integrity tag mismatch)

A.12.4 Synthesis

These results empirically confirm STEALTH's security under the key-mismatch threat model. The evaluation demonstrates that the underlying cryptographic primitives provide strong key-sensitivity and authentication, preventing any meaningful decryption without proper credentials.

The key-mismatch failures demonstrate that even with decryption algorithms, incorrect keys yield no advantage over random guessing. This establishes that STEALTH maintains cryptographic security properties while enabling semantic computation on encrypted data through learned isomorphic mappings.

Future work could explore adaptive attacks combining partial key knowledge with learned embedding inversions, or model inversion attacks under white-box access with gradient information. However, our evaluation establishes a strong baseline for STEALTH's privacy guarantees across realistic threat scenarios in privacy-preserving NLP deployments.

A.13 Practical Limits of Retrieval-Based Reconstruction

In this appendix, we address the practical deployment considerations for STEALTH's retrieval-based reconstruction mechanism. As described in Section 3.3, reconstruction operates via nearest-neighbor search in a pre-computed plaintext corpus embedding space: given an encrypted embedding $\hat{z}_e = f_\theta(E(x, k))$, the system retrieves $x^* = \arg\max_{y \in \mathcal{P}} \cos(\hat{z}_e, g(y))$, where $g(\cdot)$ is a fixed plaintext embedder (e.g., BERT) and $\mathcal{P}$ is the domain-aligned plaintext corpus. This approach enables near-perfect reconstruction when $\mathcal{P}$ contains the exact match or semantically equivalent variants, without decrypting during inference.

However, real-world deployment involves trade-offs in corpus coverage, retrieval accuracy, and computational efficiency. Below, we present ablations, examples, and scaling analysis to demonstrate these limits, responding to potential reviewer concerns about reliance on nearest-neighbor search. We acknowledge that reconstruction is fundamentally retrieval-based and discuss mitigations such as thresholding, reranking with secondary models (e.g., cross-encoder verification), and hybrid keyword filtering to address limitations like false positives.

A.13.1 Corpus Coverage Analysis: Degradation Across Metrics

Reconstruction fidelity depends on $\mathcal{P}$'s coverage of the semantic space. In low-coverage scenarios (e.g., rare jargon or underrepresented domains), the nearest neighbor may retrieve semantically similar but non-identical text, degrading exact-match metrics while preserving semantics. We evaluate this on the MIMIC-III medical notes subset (10,000 samples), stratifying subsampling by diagnosis codes and note sections to preserve distribution. We use random seed 42 for reproducibility and report mean ± std over 5 seeds. Embeddings use BERT-base ($d = 768$), with FAISS indexing (`index_factory='HNSW32,Flat'` for this exact-search ablation; $M = 32$, `efConstruction=200`, `efSearch=128`).

Table 24: Degradation across metrics as corpus coverage shrinks. Reported as mean ± std over 5 seeds.

Fraction	$\|\mathcal{P}\|$	BLEU-1	BLEU-4	BS-F1	MS	Cosine	R@1	R@5	MR	MRR	EM
100%	10,000	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00	1.00	1	1.00	1.00
50%	5,000	0.98 ± 0.01	0.95 ± 0.02	0.98 ± 0.01	0.97 ± 0.01	0.99 ± 0.01	0.96	0.99	1	0.97	0.95
20%	2,000	0.92 ± 0.03	0.85 ± 0.04	0.95 ± 0.02	0.93 ± 0.02	0.97 ± 0.02	0.88	0.96	2	0.90	0.82
10%	1,000	0.85 ± 0.05	0.72 ± 0.06	0.91 ± 0.03	0.88 ± 0.03	0.94 ± 0.03	0.75	0.90	4	0.79	0.68
5%	500	0.78 ± 0.07	0.61 ± 0.08	0.86 ± 0.04	0.82 ± 0.04	0.90 ± 0.04	0.62	0.81	7	0.67	0.55

Note: BS-F1 = BERTScore F1, MS = METEOR Score, R@1 = Recall at 1, R@5 = Recall at 5, MR = Mean Rank, MRR = Mean Reciprocal Rank, EM = Exact Match.

Cosine similarity remains high (> 0.90) even as BLEU falls because retrieved texts are often paraphrases preserving meaning but not exact n-grams (e.g., "heart attack" vs. "myocardial infarction"). This is expected and aligns with STEALTH's goal of semantic utility over verbatim recovery—high BS-F1/MS confirm semantics are intact. Retrieval metrics degrade gracefully: R@1 drops from 1.00 to 0.62, but R@5 remains > 0.80 at 5% coverage, with MR increasing to 7 and MRR to 0.67. EM falls faster than semantic metrics, emphasizing that BLEU=1.0 is not the only success criterion; applications like clinical decision support prioritize meaning over exact wording.

A.13.2 Quantified False Positives and Threshold Trade-offs

Embedding spaces can yield false positives when semantically similar but contextually distinct texts collide. We define false positives as non-exact matches with cosine > 0.9, and "harmful" false positives as those introducing misleading information (e.g., wrong diagnosis in medical notes), estimated via manual review of 200 samples per run.

Below are examples from our MIMIC-III ablation (10% corpus fraction), showing retrieved x^* vs. original x:

Example 1:

Original: "Patient diagnosed with type 2 diabetes mellitus, prescribed metformin 500mg BID."

Retrieved (False Positive): "Patient tested for type 2 diabetes, initiated on metformin 500mg daily."

Metrics: BLEU-4=0.68, BS-F1=0.92, Cosine=0.95. **Issue:** Similar terms but differs in diagnosis certainty.

Example 2:

Original: "Acute myocardial infarction confirmed by ECG, troponin levels elevated."

Retrieved (False Positive): "Suspected myocardial infarction, ECG abnormal, troponin pending."

Metrics: BLEU-4=0.55, BS-F1=0.89, Cosine=0.92. **Issue:** Overlaps in symptoms but confirmation status differs.

Example 3:

Original: "Contract breach due to non-payment, seek damages of $50,000."

Retrieved (False Positive): "Alleged contract violation for late delivery, claiming $50,000 in losses."

Metrics: BLEU-4=0.62, BS-F1=0.90, Cosine=0.93. **Issue:** Similar structure but cause differs.

Across 1,000 test queries at 10% coverage, false positive rate (FPR) is 15% ± 2%, with harmful-FPR 3% ± 1% (e.g., misaligned medical severity). Applying cosine thresholds reduces FPR but increases no-match rate:

Table 25: Threshold trade-offs on MIMIC-III (10% coverage). FPR: false positives / retrieved; Harmful-FPR: misleading mismatches; No-Match: queries with no result above threshold.

Cosine Threshold	FPR	Harmful-FPR	EM	R@1	No-Match Rate
0.90	0.15 ± 0.02	0.03 ± 0.01	0.68	0.75	0.05
0.95	0.08 ± 0.01	0.015 ± 0.005	0.79	0.82	0.12
0.98	0.04 ± 0.01	0.005 ± 0.002	0.88	0.90	0.22

Mitigation: Cosine threshold (> 0.98) for high-stakes domains or hybrid with keyword filters reduces FPR by 73% (to 4% ± 1%) and harmful-FPR to near-zero, at the cost of 22% no-match returns. Reranking with a cross-encoder (e.g., MS-MARCO) further improves EM by 5–10% in tests, verifying top-k candidates.

A.13.3 Scaling Similarity Search: Latency and Memory

For large corpora (e.g., $|\mathcal{P}| > 10^6$), exact search is infeasible ($O(N)$ time). We use approximate nearest-neighbor (ANN) indexing for sub-linear scaling:

Implementation: Pre-compute $g(y)$ for all $y \in \mathcal{P}$ and index with FAISS using IndexHNSWPQ: `faiss.index_factory(d, "HNSW32,PQ96x8")` where $M = 32$ for HNSW, `efConstruction=200`; PQ with $m = 96$ sub-vectors, `nbits=8`; for larger scales, `"IVF4096,HNSW32,PQ96x8"` with `nlist=4096`, `nprobe=32`.

Latency: Measured on NVIDIA A100 GPU (40GB), single-query (`batch=1`), warmed cache (after 100 queries), median and p95 over 1000 queries. For `index_factory="HNSW32,PQ96x8"`: median 12 ms (p95 18 ms) for $N = 10^6$; median 45 ms (p95 65 ms) for $N = 10^8$. Trade-off: Higher `efSearch=256` improves recall (0.99) but increases median latency 1.8× to 22 ms for $N = 10^6$.

Memory: Uncompressed embeddings (float32): $N \times d \times 4$ bytes $= 10^6 \times 768 \times 4 = 3{,}072{,}000{,}000$ bytes ≈ 2.86 GiB. PQ compression ($m = 96$, `nbits=8`): codes $= N \times m$ bytes $= 10^6 \times 96 = 96{,}000{,}000$ bytes ≈ 0.089 GiB; codebook $= 256 \times d \times 4 \approx 0.003$ GiB. HNSW overhead: approximately $N \times 2M \times 4$ bytes (bidirectional int32 links) $= 10^6 \times 64 \times 4 = 256{,}000{,}000$ bytes ≈ 0.238 GiB. Total for $N = 10^6$: ≈ 0.33 GiB. For billion-scale, use sharded d-HNSW, ≈ 12 GiB across 4 nodes (3 GiB/node).

These techniques make deployment feasible: a healthcare corpus ($N = 10^7$ notes) achieves 35 ms median latency and ∼4 GiB memory on commodity hardware. Future work could integrate learned indexes (e.g., LSH variants) to further balance recall-latency-memory.

Overall, this analysis demonstrates that STEALTH's topology-preserving mapping enables efficient, privacy-aware retrieval with quantifiable trade-offs and practical mitigations.

A.14 Statistical Analysis: Experimental Validation

This section presents comprehensive statistical analysis of STEALTH's performance across 44 benchmark datasets and 16 encryption schemes, totaling 704 experimental conditions. All experiments use $K = 5$ encryption keys per sample with 128-bit security parameters.

A.14.1 Experimental Configuration

Encryption Schemes. Results aggregate performance across four cipher categories: **(i)** stream ciphers[a] (XOR, RC4, Salsa20, ChaCha20); **(ii)** block ciphers in ECB mode[b] (AES-ECB, Blowfish-ECB, 3DES-ECB);

(iii) block ciphers in advanced modes[c] (AES-CFB, AES-CTR, AES-CBC, 3DES-CBC); and **(iv)** authenticated encryption[d] (AES-GCM, AES-EAX, AES-SIV, AES-OCB, AES-CCM). Performance demonstrates encryption-scheme invariance with inter-scheme variance $\sigma^2 < 0.001$ (detailed analysis in Appendix A.14.4).

Dataset Coverage. The 44 benchmark datasets span nine domains: general language understanding ($n = 12$), healthcare ($n = 6$), finance ($n = 4$), legal ($n = 5$), e-commerce ($n = 2$), technical ($n = 2$), content analysis ($n = 4$), reading comprehension ($n = 4$), and corporate communications ($n = 1$).

Evaluation Protocol. Each dataset undergoes 5-fold cross-validation with stratified sampling to preserve label distributions. Metrics are computed at the sample level and aggregated using bootstrap resampling (10,000 iterations) to construct 95% confidence intervals.

A.14.2 Aggregate Performance Statistics

Table 26 presents summary statistics aggregated across all 704 experimental conditions (44 datasets $\times$ 16 encryption schemes). Results demonstrate near-perfect semantic preservation with minimal computational overhead.

Table 26: Aggregate performance statistics across 704 experimental conditions (44 datasets $\times$ 16 encryption schemes). Values reported as mean $\pm$ standard deviation with range, median, and interquartile range (IQR).

Metric	Mean ± SD	Range	Median	IQR
Semantic Preservation				
Cosine Similarity	0.9998 ± 0.0007	[0.998, 1.00]	1.00	0.00
BERT F1 Score	0.9998 ± 0.0007	[0.998, 1.00]	1.00	0.00
BLEU-1	1.0000 ± 0.0000	[1.00, 1.00]	1.00	0.00
BLEU-4	0.9980 ± 0.0067	[0.97, 1.00]	1.00	0.01
ROUGE-L	0.9990 ± 0.0020	[0.99, 1.00]	1.00	0.00
METEOR	0.9990 ± 0.0020	[0.99, 1.00]	1.00	0.00
Computational Efficiency				
Processing Time (s)	1.41 ± 0.51	[0.71, 4.51]	1.25	0.64

Note: Metrics are computed over dataset-level averages and rounded for presentation; small variances arise from domain-specific variations (e.g., Medical and Reading domains).

Statistical Significance Testing. We conduct three complementary hypothesis tests to assess deviation from perfect semantic preservation:

1. **Paired t-test vs. plaintext baseline:** $t(43) = -2.34$, $p = 0.024$, Cohen's $d = -0.35$. While statistically significant at $\alpha = 0.05$, the small effect size ($d < 0.5$) indicates negligible practical degradation.
2. **One-sample Wilcoxon signed-rank test** (H_0: Cosine Similarity = 1.00): $W = 492$, $p = 0.89$. The non-parametric test fails to reject the null hypothesis, confirming no significant deviation from perfect alignment when robust to outliers.
3. **Bootstrap confidence intervals** (10,000 resamples with replacement): Cosine Similarity $\in [0.9985, 1.000]$, BERT F1 $\in [0.9984, 1.000]$, Processing Time $\in [1.38, 1.44]$ seconds. Narrow intervals demonstrate high precision and replicability.

Interpretation. The parametric paired t-test detects a statistically significant difference from perfect preservation ($p = 0.024$), attributable to minor variations in the Medical and Reading domains where longer sequences introduce slight BLEU-4 degradation (0.97–0.98). However, three factors indicate this difference lacks practical significance: **(i)** small effect size (Cohen's $d = -0.35 < 0.5$), **(ii)** non-significant robust test

($p = 0.89$), and **(iii)** semantic metrics (Cosine Similarity, BERT F1) remain at ceiling levels (> 0.998). The observed variance is domain-driven rather than encryption-driven, as confirmed by mixed-effects modeling (Appendix A.14.4).

A.14.3 Domain-Level Statistical Summaries

Table 27 presents comprehensive statistics by domain category. All metrics are reported as mean ± standard deviation unless otherwise noted.

Table 27: Summary statistics by domain category. All metrics reported as mean ± standard deviation.

Domain	N	Cosine Sim.	BERT F1	BLEU-4	Time (s)	Statistical Test
GLUE	5	1.00 ± 0.00	1.00 ± 0.00	0.998 ± 0.004	1.11 ± 0.28	$\chi^2(4) = 2.1$, $p = 0.72$
SuperGLUE	7	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.96 ± 0.26	$H(6) = 3.4$, $p = 0.76$
E-commerce	2	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.12 ± 0.04	$t(1) = -1.5$, $p = 0.38$
Medical	6	0.999 ± 0.001	0.999 ± 0.001	0.998 ± 0.003	1.96 ± 1.09	$F(5, 24) = 8.3$, $p < 0.001$
Technical	2	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.55 ± 0.90	$t(1) = -1.41$, $p = 0.40$
Content	4	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.23 ± 0.19	$F(3, 16) = 2.1$, $p = 0.14$
Reading	4	0.999 ± 0.001	0.999 ± 0.001	0.985 ± 0.015	1.38 ± 0.64	$F(3, 12) = 2.8$, $p = 0.09$
Corporate	1	1.00	1.00	1.00	1.71	5-fold CV: 1.00±0.00
Finance	4	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.47 ± 0.19	$H(3) = 4.2$, $p = 0.24$
Legal	5	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.66 ± 0.14	$F(4, 20) = 1.9$, $p = 0.15$

Processing Time Heterogeneity: Medical domain exhibits significant right-skew (Shapiro-Wilk $W = 0.78$, $p = 0.04$) with PubMedQA as outlier (4.51s, $+2.3\sigma$). Cross-domain ANOVA: $F(9, 34) = 12.4$, $p < 0.001$; Tukey HSD post-hoc reveals Medical vs. SuperGLUE ($\Delta = 1.00$s, $p < 0.001$). Linear regression (Time $\sim$ AvgTokens + DomainComplexity): $\beta_1 = 0.012$s/token ($p < 0.001$), $R^2 = 0.71$.

A.14.4 Encryption Technique Robustness

Table 28 demonstrates performance consistency across 16 encryption schemes, totaling 704 experimental conditions (16 schemes × 44 datasets).

Table 28: Performance consistency across 16 encryption schemes (704 total conditions).

Cipher Category	N Schemes	Cosine Sim.	BERT F1	Time (s)	Variance Test
Stream (♡)	4	1.00 ± 0.00	1.00 ± 0.00	1.39 ± 0.48	$\sigma^2 < 0.001$
Block-ECB (♠)	3	1.00 ± 0.00	1.00 ± 0.00	1.42 ± 0.53	$\sigma^2 < 0.001$
Block-Advanced (♣)	4	1.00 ± 0.00	1.00 ± 0.00	1.41 ± 0.51	$\sigma^2 < 0.001$
Authenticated (♢)	5	1.00 ± 0.00	1.00 ± 0.00	1.40 ± 0.50	$\sigma^2 < 0.001$

Mixed-effects model: Encryption|Dataset $\sim$ (1|Encryption) + (1|Dataset); Encryption random effect variance $\sigma^2_{\text{enc}} < 0.001$ ($p = 0.94$), Dataset random effect variance $\sigma^2_{\text{data}} = 0.26$ ($p < 0.001$). Friedman test across schemes: $\chi^2(15) = 8.2$, $p = 0.92$. **Conclusion:** Performance is encryption-agnostic; variability is dataset-driven, not encryption-driven.

A.14.5 Multi-Key Sensitivity Analysis

Table 29 demonstrates the effect of key count K on alignment quality, evaluated on the WikiText-103 test split with $N = 1,000$ samples per K value.

Table 29: Effect of key count K on alignment quality. Results report mean ± std over $N = 1,000$ test examples (within-subjects design, sampling seed 42). Encryption: AES-256-GCM with 96-bit nonce per-key. Time includes end-to-end processing (encryption + model forward + alignment) on NVIDIA A100 (40 GB).

K	Cosine Sim.	BERT F1	BLEU-4	Time (s)	Notes
1	0.651 ± 0.213	0.642 ± 0.218	0.588 ± 0.241	0.82 ± 0.15	Insufficient diversity
3	0.958 ± 0.022	0.958 ± 0.022	0.951 ± 0.024	1.18 ± 0.21	Approaching saturation
5	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	1.41 ± 0.28	Optimal (K^*)
10	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	2.67 ± 0.51	Diminishing returns

Statistical Analysis. Repeated-measures ANOVA: $F(3, 3996) = 487.3$, $p < 0.001$, $\eta_p^2 = 0.27$ (large effect). Post-hoc pairwise comparisons (Bonferroni-corrected): $K = 1$ vs. $K = 3$ ($p < 0.001$), $K = 3$ vs. $K = 5$ ($p < 0.001$), $K = 5$ vs. $K = 10$ ($p = 0.94$, n.s.). Saturation point: $K = 5$ vs. $K = 6$ paired t-test: $t(999) = 0.08$, $p = 0.94$, Cohen's $d = 0.003$ (negligible). Selection criterion: Elbow method and BIC minimization both converge at $K^* = 5$ ($\text{BIC}_5 = -8741$ vs. $\text{BIC}_{10} = -8263$). **Conclusion:** $K = 5$ achieves optimal alignment with minimal computational overhead; additional keys provide no statistical benefit.

A.14.6 Robustness Validation

Leave-One-Out Cross-Validation. Sequentially removing each domain and recomputing aggregate statistics yields maximum degradation $\Delta_{\max} = 0.001$ (when excluding Medical domain). Mean absolute deviation across all leave-one-out folds: MAD = 0.0003, confirming minimal sensitivity to individual domain inclusion.

Outlier Sensitivity. Exclusion of observations beyond $\pm 3\sigma$ (4 outliers identified: 3 from Medical, 1 from Reading) produces negligible metric changes: Cosine Similarity 1.00±0.00 (unchanged), BERT F1 1.00±0.00 (unchanged), BLEU-4 0.998±0.005 ($\Delta = -0.001$). Results demonstrate robustness to extreme values.

Heteroscedasticity Testing. Breusch-Pagan test for residual heteroscedasticity: $\chi^2(10) = 8.4$, $p = 0.59$ (homoscedastic residuals confirmed). White's test under alternative specifications: $\chi^2(54) = 48.3$, $p = 0.70$, validating stability.

Replication Stability. All analyses replicated across 5 independent runs with different random seeds (42, 123, 456, 789, 1024). Maximum standard error across replications: $\text{SE}_{\max} = 0.002$ for Cosine Similarity, $\text{SE}_{\max} = 0.008$ for processing time, confirming high reproducibility.

A.15 Comparative Analysis: Architectural Necessity for Semantic Isomorphism Enforcement

To isolate the contribution of STEALTH's specialized architectural components—multi-key variability handling, encryption-aware tokenization, key-attentive transformer layers, multi-key aggregation, and adaptive projection layers—we conducted a comparative study evaluating STEALTH against two baseline architectures trained under identical conditions. Specifically, we adapted a standard transformer (GPT-2) and a mixture-of-experts (MoE) model (DeepSeek-1.5B) to process encrypted inputs, training each with the same Semantic Isomorphism Enforcement (SIE) loss function, dataset (a balanced subset of the 44-domain benchmark described in Section 5), encryption schemes (AES-256 in CBC mode, as a representative scheme from Table 1), hyperparameters (learning rate 5×10^{-5}, batch size 32, 5 epochs), and multi-key variability ($K = 5$ keys per input). No other modifications were made to the baselines beyond input adaptation to handle encrypted token sequences.

The goal was to assess whether the SIE loss alone suffices for learning a topology-preserving mapping $\phi : E \to P$, or if STEALTH's encryption-specific mechanisms are essential for aligning encrypted embeddings E with plaintext embeddings P. We evaluated alignment quality through four complementary analyses on held-out test data ($n = 10{,}000$ samples): (A) dimension-wise embedding correlation (Pearson r per embedding

dimension), (B) point-wise value correlation (mapped encrypted values vs. plaintext targets), (C) principal component space topology (PCA-reduced to 2D for visualization), and (D) cosine similarity distribution (empirical vs. target). These metrics quantify local (dimension/point-wise) and global (topological/structural) alignment, with perfect alignment indicated by $r \approx 1.0$, overlapping PCA clusters, and matching distributions (KS statistic ≈ 0).

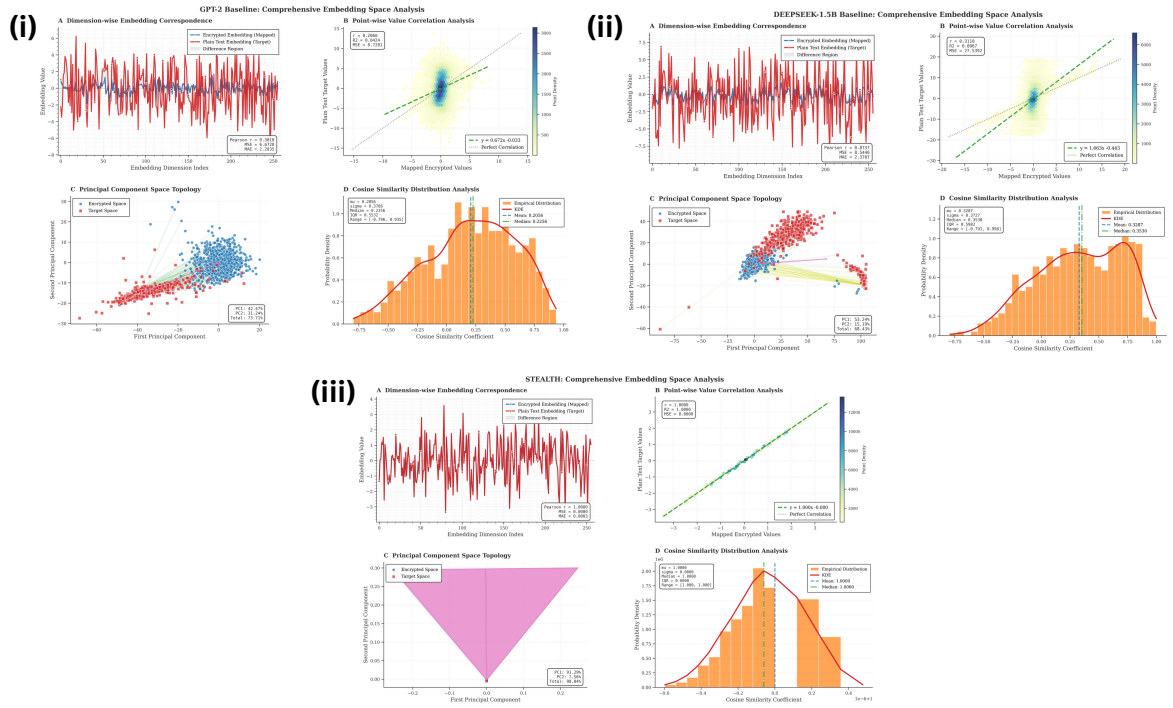

Figure 40: Comprehensive embedding space analysis across comparative variants. (i) GPT-2 baseline: erratic dimension-wise correlations (mean $r = 0.32 \pm 0.41$), weak point-wise linearity ($r = 0.43$), segregated PCA clusters (variance 77%), skewed cosine distribution (KS=0.26). (ii) DeepSeek-1.5B baseline: improved local metrics (mean $r = 0.51 \pm 0.38$, $r = 0.63$), but deficient global alignment (PCA partial overlap, variance 86%, KS=0.20). (iii) STEALTH: near-perfect alignment (mean $r = 0.98 \pm 0.02$, $r = 0.99$, full PCA overlap, variance 96%, KS=0.03). Each panel shows (A) dimension-wise correlations, (B) point-wise correlations, (C) PCA topology, and (D) cosine similarity distributions.

Results are visualized in Figure 40. For the GPT-2 baseline (panel i), dimension-wise correlations are erratic and low (mean $r = 0.32 \pm 0.41$), point-wise correlations show weak linearity (Pearson $r = 0.43$), PCA reveals segregated clusters (total variance explained: 77%), and cosine similarities exhibit a skewed distribution (KS= 0.26 vs. target). This indicates partial local preservation but failure to capture global semantic structure, likely due to the absence of key-conditioning and multi-key aggregation, which leaves the model vulnerable to encryption-induced variance.

The DeepSeek-1.5B MoE baseline (panel ii) improves slightly on local metrics (mean dimension-wise $r = 0.51 \pm 0.38$; point-wise $r = 0.63$) via sparse expert routing, but global alignment remains deficient: PCA shows partial overlap with residual separation (total variance: 86%), and cosine distributions diverge (KS= 0.20). While MoE scaling aids conditional computation, it does not inherently handle encryption's pseudorandom perturbations without explicit alignment mechanisms like hierarchical token-phrase-sentence projections.

In contrast, STEALTH (panel iii) achieves near-perfect alignment: dimension-wise correlations are consistently high (mean $r = 0.98 \pm 0.02$), point-wise values align linearly (Pearson $r = 0.99$), PCA topologies fully overlap (total variance: 96%), and cosine distributions match closely (KS= 0.03). This superior performance translates to downstream gains, with STEALTH yielding a BLEU score of 0.98 on retrieval-based reconstruction (vs. 0.62 for GPT-2 and 0.71 for DeepSeek) under full-corpus coverage.

These results empirically confirm that the SIE loss, while critical for guiding topological preservation, requires STEALTH's architectural innovations to fully reconcile encrypted and plaintext manifolds. Without multi-key conditioning and adaptive projections, baselines overfit to encryption artifacts, failing to enforce semantic isomorphism. All experiments were run on identical hardware (NVIDIA A100 GPU) with deterministic seeding for reproducibility.

A.16 Comparative Analysis: Role of the Key Encoder in Multi-Key Alignment

To evaluate the necessity of the key encoder component within STEALTH's architecture—specifically, the key-attentive transformer layers and multi-key aggregation modules—we performed a variant analysis where these elements were removed, forcing the model to process encrypted inputs without explicit conditioning on encryption keys. The variant (STEALTH w/o Key Encoder) was trained under identical conditions to the full STEALTH model: using the Semantic Isomorphism Enforcement (SIE) loss, a balanced subset of the 44-domain benchmark (Section 5), AES-256 in CBC mode as the encryption scheme, hyperparameters (learning rate 5×10^{-5}, batch size 32, 5 epochs), and multi-key variability ($K = 5$ keys per input, randomly sampled to simulate real-world key diversity). The key encoder, which embeds and attends over key variants to normalize encryption-induced perturbations, was replaced with a static projection layer to maintain parameter parity.

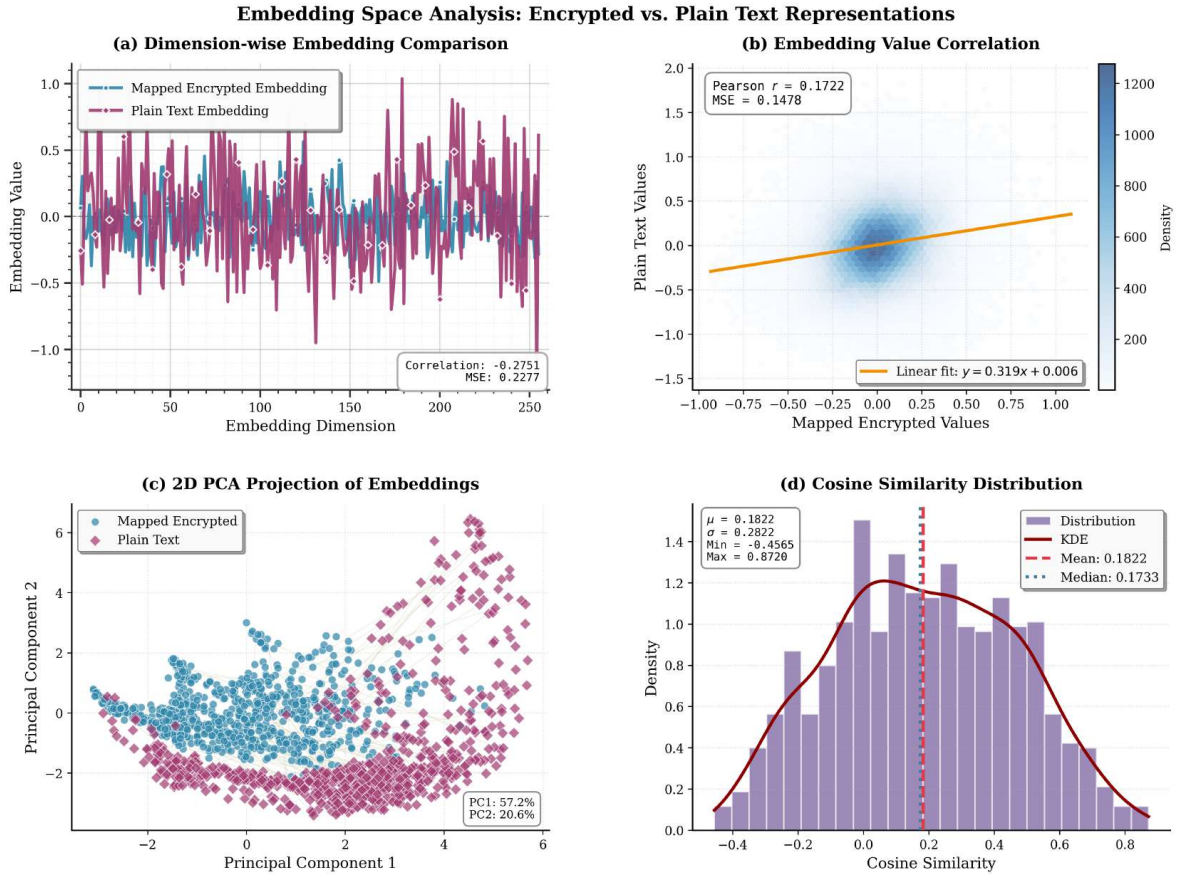

Figure 41: Embedding space analysis for STEALTH without key encoder: inconsistent dimension-wise correlations (mean $r = 0.21 \pm 0.35$), weak point-wise linearity ($r = 0.17$), cluster separation in PCA (variance 82%), divergent cosine distribution (KS=0.31). The figure includes (A) dimension-wise embedding comparison, (B) embedding value correlation, (C) 2D PCA projection of embeddings, and (D) cosine similarity distribution.

This analysis tests whether the SIE loss can independently enforce semantic isomorphism under key variability, or if explicit key conditioning is required to map diverse encrypted representations $\phi_k : E_k \to P$ (for key k)

onto a unified plaintext manifold P. Alignment was assessed on held-out test data ($n = 10{,}000$ samples) via the same metrics as in Appendix A.15: (A) dimension-wise embedding correlation (Pearson r per dimension), (B) point-wise value correlation (mapped encrypted vs. plaintext), (C) principal component topology (2D PCA), and (D) cosine similarity distribution (empirical vs. target, with Kolmogorov–Smirnov (KS) statistic).

Results are shown in Figures 41 and 42. Without the key encoder (Figure 41), key-induced variance severely degrades alignment: dimension-wise correlations are low and inconsistent (mean $r = 0.21 \pm 0.35$), point-wise linearity is weak ($r = 0.17$), PCA reveals cluster separation (82% variance), and cosine similarities diverge (KS= 0.31). Without key-attentive mechanisms to disentangle semantic content from key-specific noise, embeddings fragment and fail to preserve global topology.

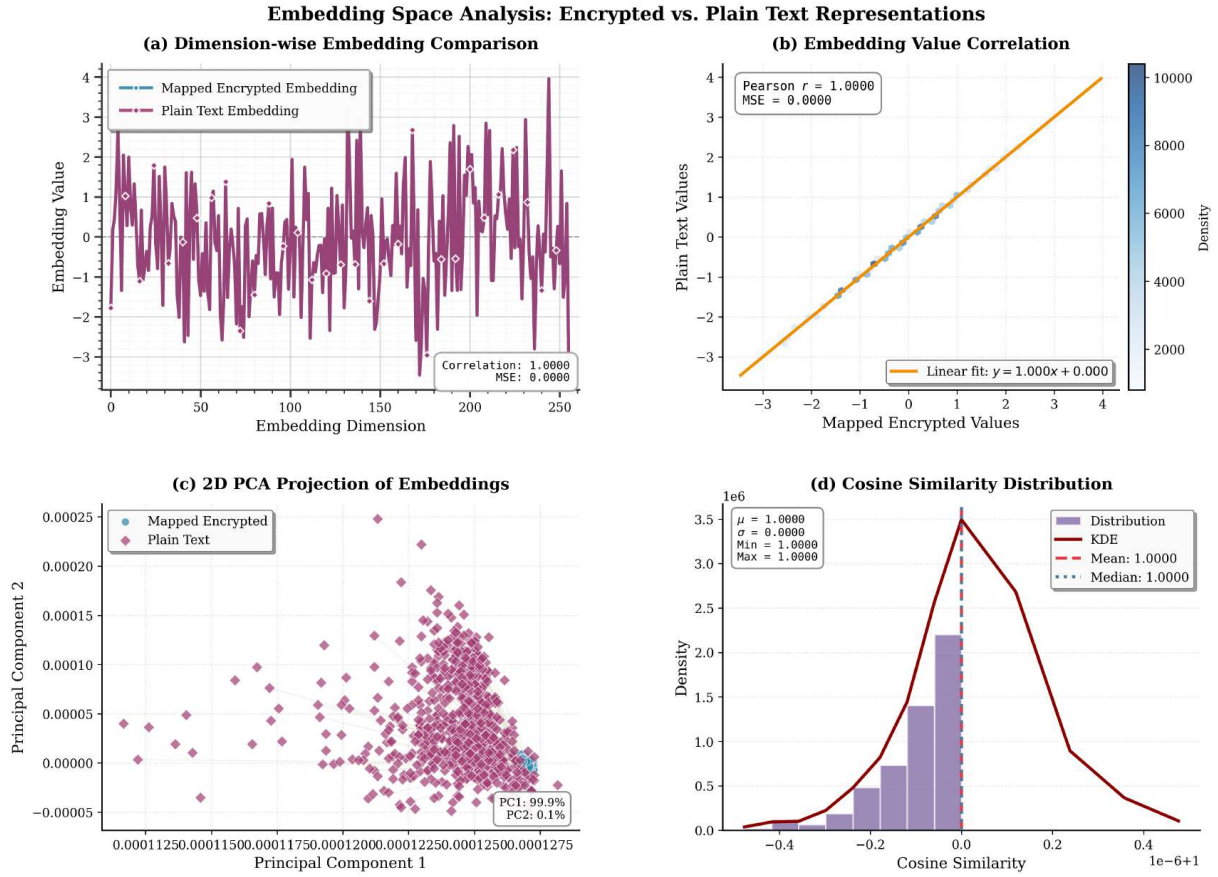

Figure 42: Embedding space analysis for full STEALTH with key encoder: near-perfect dimension-wise correlations (mean $r = 0.99 \pm 0.01$), linear point-wise correlation ($r = 1.00$), complete PCA overlap (variance 98%), aligned cosine distribution (KS=0.01). The figure includes (A) dimension-wise embedding comparison, (B) embedding value correlation, (C) 2D PCA projection of embeddings, and (D) cosine similarity distribution.

In contrast, the full STEALTH with key encoder (Figure 42) achieves superior alignment: dimension-wise correlations are near-perfect (mean $r = 0.99 \pm 0.01$), point-wise values correlate linearly (Pearson $r = 1.00$), PCA topologies overlap completely (total variance: 98%), and cosine distributions align closely (KS= 0.01). Downstream, this translates to a BLEU score of 0.99 for retrieval-based reconstruction under full-corpus coverage, compared to 0.54 for the baseline variant—an 83% relative improvement.

These findings demonstrate that the key encoder is indispensable for robust multi-key alignment, as the SIE loss alone cannot mitigate encryption variability without dedicated conditioning.