

InvisiCrypt: A Dual-Layer Secure Communication Framework Integrating AES-256 Encryption with Complexity-Adaptive LSB Steganography

Krunal Pandurang Bhadke¹, Yugal Deepak Bari², Prof.Rajni Hemade³

¹Department of Computer Science And Engineering ,Nutan College Of Engineering And Research

²Department of Computer Science And Engineering ,Nutan College Of Engineering And Research

³Department of Computer Science And Engineering ,Nutan College Of Engineering And Research

Abstract — Digital communication increasingly operates under conditions where encrypted traffic attracts forensic inspection. Traditional cryptographic techniques protect message contents but fail to conceal the existence of communication channels. InvisiCrypt addresses this gap by proposing a dual-layer security framework that fuses Advanced Encryption Standard (AES-256) cryptography with an adaptive Least Significant Bit (LSB) steganographic algorithm. Plaintext is first encrypted using AES-256 in Cipher Block Chaining (CBC) mode; the resulting ciphertext is then embedded invisibly within the pixel channels of a cover image using a complexity-adaptive LSB strategy that concentrates higher-order bit substitutions in textured image regions, exploiting the spatial masking properties of the human visual system. The stego-image is visually indistinguishable from the original, yielding a Peak Signal-to-Noise Ratio (PSNR) exceeding 52 dB and a Structural Similarity Index Measure (SSIM) of 0.9991 at half embedding capacity. Steganalysis resistance was evaluated against histogram analysis, chi-square attack, and RS analysis, achieving detection rates of only 3.8–6.1%, near the theoretical random-chance threshold. InvisiCrypt outperforms five state-of-the-art hybrid schemes across all imperceptibility and security metrics, confirming suitability for confidential enterprise communication, digital forensics, and covert intelligence operations.

Key Words: steganography, AES-256, LSB embedding, steganalysis, covert communication, information hiding.

1.INTRODUCTION

The explosion of digital communication across cloud platforms, mobile networks, and Internet of Things (IoT) ecosystems has made data security one of the defining engineering challenges of the twenty-first century. Two classical disciplines have independently addressed this challenge: cryptography, which scrambles message content so that unauthorized parties cannot read it, and steganography, which conceals the very existence of a message within an innocuous carrier medium [1]. Deployed alone, each carries well-documented weaknesses. A strongly encrypted file is computationally secure, yet its presence signals sensitivity and may invite targeted forensic inspection or traffic-analysis-based profiling. Steganography without encryption leaves readable plaintext inside the carrier; once the embedding scheme is discovered, the hidden data is fully exposed [2].

The convergence of both techniques into a unified pipeline offers compounded security: an adversary must first detect that a message is hidden, defeating steganography, and then decrypt its content, defeating cryptography. This layered philosophy

motivates InvisiCrypt. Crucially, AES-256 ciphertext is statistically near-uniform, resembling random noise, which means it blends naturally into high-texture image regions and makes steganalysis detectors less effective — a synergy purely single-layer systems cannot achieve. The specific contributions of this work are: (1) an end-to-end pipeline combining AES-256 encryption with adaptive LSB steganography in a single modular system; (2) a pixel-complexity-adaptive bit-plane selection algorithm that maximises embedding capacity while maintaining imperceptibility; (3) empirical evaluation against three categories of steganalysis attacks on a 200-image benchmark; and (4) a comparative analysis against five state-of-the-art hybrid schemes demonstrating superior PSNR, SSIM, and steganalysis resistance.

2. LITERATURE REVIEW

A. Classical Steganographic Approaches

Simmons [3] formalised the prisoner's problem, establishing the theoretical foundation for covert communication over observed channels. Chan and Cheng [4] introduced a pixel-value differencing (PVD) scheme that achieves higher capacity than naive LSB methods while preserving moderate image quality by exploiting the perceptual tolerance of edge pixels. Wu and Tsai [5] extended this with edge-region preferential embedding, raising capacity while limiting perceptual impact. Neither spatial-domain variant, however, encrypted the payload before hiding it, leaving hidden data unprotected once the embedding scheme is known. Frequency-domain techniques, including DCT-based approaches, offer improved robustness to JPEG compression at the cost of higher computational overhead [11].

B. Cryptographic Foundations

The Advanced Encryption Standard (AES), adopted by NIST in 2001, remains the gold standard for symmetric encryption [6]. AES-256 in Cipher Block Chaining (CBC) mode provides approximately 1.1×10^{77} possible keys. A fresh initialisation vector (IV) per session ensures semantic security — repeated encryption of identical plaintext yields distinct ciphertext blocks [7]. An important additional property relevant to steganography is that AES output bytes are statistically indistinguishable from a uniformly random sequence, so embedded ciphertext does not produce the structured patterns that statistical steganalysis tools are tuned to detect.

C. Hybrid Cryptography-Steganography Systems

Bhattacharyya et al. [8] proposed embedding RSA-encrypted data within DCT coefficients, achieving PSNR near 44 dB. Al-Dmour and Al-Ani [9] used DES encryption followed by LSB embedding; however, DES is now cryptographically broken due to its 56-bit key length. Thiyagarajan et al. [10] combined Triple-DES with frequency-domain steganography, reporting

improved robustness but reduced embedding capacity. Deep-learning steganalysis frameworks such as SRNet [11] have demonstrated that naive LSB schemes are vulnerable to convolutional neural network (CNN) detectors, underscoring the need for adaptive, key-dependent embedding strategies.

D. Identified Research Gaps

A survey of the literature reveals three persistent gaps: most hybrid schemes rely on deprecated cryptographic primitives such as DES or Triple-DES; few evaluate resistance beyond simple histogram inspection; and embedding capacity and imperceptibility are rarely optimised simultaneously. InvisiCrypt is designed to close all three gaps by adopting AES-256, evaluating against three steganalysis categories, and employing complexity-adaptive embedding.

3. METHODOLOGY

A. System Architecture Overview

The InvisiCrypt pipeline consists of two primary modules: the Encryption Engine and the Steganographic Embedding Engine. A corresponding receiver-side pair handles extraction and decryption. Both the AES key and the pseudo-random traversal permutation share a single secret key K , so possession of K is both necessary and sufficient for authorised communication. Figure 1 illustrates the end-to-end architecture.




Fig. 1: InvisiCrypt End-to-End System Architecture

B. Encryption Engine: AES-256 CBC

The plaintext message M is padded using PKCS7 to align with the 128-bit AES block boundary. PKCS7 always adds at least one byte, ensuring unambiguous pad removal by the receiver. A 256-bit key K is generated via a CSPRNG drawing entropy from the OS hardware entropy pool. A fresh 128-bit IV is generated per session. Encryption proceeds as:

$$C_i = E_k(P_i \oplus C_{i-1}), \quad C_0 = IV \quad \dots (1)$$

CBC mode chains each plaintext block to the previous ciphertext block via XOR, preventing identical plaintext blocks from producing identical ciphertext blocks, a critical weakness of ECB mode [7]. The full ciphertext blob $C = \{IV \parallel C_1 \parallel \dots \parallel C_n\}$ is serialised to a byte stream and passed to the embedding engine. Prepending the IV to the payload allows the receiver to recover it without any additional out-of-band channel; the IV is not secret and its inclusion does not weaken AES-256 security.

C. Adaptive LSB Steganographic Embedding

The ciphertext byte stream is embedded within a 24-bit true-colour PNG cover image I of dimensions $W \times H$ pixels. Unlike static 1-bit LSB substitution, InvisiCrypt assigns 1, 2, or 3 bits per channel based on local pixel complexity.

1) Pixel Complexity Measurement: For each pixel $P(x, y)$, the local complexity score $\phi(x, y)$ is the standard deviation of pixel values within a 3×3 neighbourhood:

$$\phi(x, y) = \sigma\{P(x', y') : |x' - x| \leq 1, |y' - y| \leq 1\} \quad \dots (2)$$

Standard deviation is computationally inexpensive and correlates reliably with perceptual roughness, making it well-suited for real-time complexity estimation [13].

2) Bit Allocation Policy: Three tiers are defined: $\phi < 10 \rightarrow 1$ bit per channel (smooth regions); $10 \leq \phi < 30 \rightarrow 2$ bits per channel (moderate complexity); $\phi \geq 30 \rightarrow 3$ bits per channel (high-complexity edge regions). This directly exploits the spatial contrast masking property of the human visual system, whereby distortions in visually complex areas require significantly larger magnitudes to become perceptible.

3) Total Embedding Capacity:

$$T = 3 \times (1 \cdot N_1 + 2 \cdot N_2 + 3 \cdot N_3) \quad \dots (3)$$

where N_1, N_2, N_3 are pixel counts in each complexity tier. The embedding traversal order is determined by a pseudo-random permutation seeded with a hash of key K , ensuring that blind extraction is infeasible without K .

D. Extraction and Decryption

The receiver uses the shared key K and stego-image S to reconstruct the keyed traversal order, re-evaluate pixel complexity tiers, and reassemble the ciphertext byte stream. Decryption reverses Eq. (1):

$$P_i = D_k(C_i) \oplus C_{i-1} \quad \dots (4)$$

PKCS7 de-padding yields the original plaintext. Figure 2 illustrates the full workflow.




Fig. 2: InvisiCrypt Encryption and Embedding Workflow

E. Security Analysis

InvisiCrypt's security rests on two independent hard problems: (1) breaking AES-256 CBC, which has no known polynomial-time attack; and (2) locating the payload without key K , equivalent to inverting the keyed pseudo-random permutation. An adversary defeating the steganographic layer recovers only ciphertext; one who breaks concealment still faces AES-256 [12]. The combination ensures that defeating one layer provides essentially no advantage in attacking the other.

4. EXPERIMENTAL RESULTS

A. Experimental Setup

Experiments ran on a workstation with an Intel Core i7-12700K, 32 GB DDR5 RAM, and NVIDIA RTX 3060 GPU. The implementation used Python 3.11 with PyCryptodome for AES and OpenCV 4.8 for image processing. The benchmark dataset comprised 200 uncompressed 24-bit PNG images from the USC-SIPI image database at 512×512 resolution, spanning natural scenes, portraits, textures, and synthetic patterns to ensure generalisation across the full range of complexity distributions.

B. Imperceptibility Metrics

Imperceptibility was measured using PSNR and SSIM [13]. Table I summarises results across payload sizes ranging from 5% to 99% of maximum embedding capacity.

TABLE - I: Imperceptibility Metrics Across Payload Sizes

Payload (KB)	PSNR (dB)	SSIM	Cap. (%)
10	58.42 ± 0.3	0.9997	5.2%
50	55.17 ± 0.4	0.9994	26.0%
100	52.83 ± 0.5	0.9991	52.1%
150	49.61 ± 0.7	0.9985	78.1%
190	47.29 ± 0.9	0.9978	99.0%

PSNR remains above 52 dB at half capacity, well above the 40 dB threshold accepted as visually imperceptible [9]. SSIM remains above 0.997 across all tested payload sizes, confirming structural preservation. The small standard deviations indicate consistent performance across diverse image types.




Fig. 3: PSNR vs. Payload Fill Rate (LSB Adaptive vs. Static)

C. Steganalysis Resistance

Three steganalysis attacks were applied to 50 randomly selected stego-images carrying 100 KB payloads. Table II compares InvisiCrypt against naive LSB and DCT-based schemes.

TABLE - II: Steganalysis Detection Rates (Lower = Better)

Attack Method	Naive LSB	DCT-Based	InvisiCrypt
Histogram Analysis	82.0%	38.4%	4.2%
Chi-Square Attack	94.6%	41.2%	3.8%
RS Steganalysis	96.3%	44.7%	6.1%

InvisiCrypt's 3.8% chi-square detection rate falls within the theoretical 5% significance-level bound, meaning stego-images are statistically consistent with natural images under this attack. The keyed pseudo-random traversal disrupts the pixel-pairing symmetry that RS analysis exploits, yielding a 6.1% detection rate near the random-chance baseline.

D. Comparative Performance

Table III benchmarks InvisiCrypt against five representative hybrid schemes from recent literature, evaluated on identical 100 KB payloads in 512×512 PNG images.

TABLE - III: Comparative Performance (100 KB Payload, 512×512 PNG)

Scheme	PSNR (dB)	SSIM	RS Det. Rate
Bhattacharyya et al. [8]	44.1	0.9942	51.3%
Al-Dmour & Al-Ani [9]	46.8	0.9961	38.9%
Thiyagarajan et al. [10]	43.7	0.9934	44.6%
AES + Naive LSB (Baseline)	48.2	0.9972	87.4%
AES + DCT Hybrid	50.1	0.9981	42.1%
InvisiCrypt (Proposed)	52.8	0.9991	6.1%

InvisiCrypt achieves the highest PSNR and SSIM scores while recording the lowest steganalysis detection rate — a simultaneous improvement across all three critical dimensions without sacrificing one metric for another. The AES + Naive LSB baseline achieves reasonable PSNR (48.2 dB) but catastrophically high RS detection (87.4%), confirming that encryption alone cannot compensate for a poor embedding strategy.

E. Processing Performance

Encryption and embedding of a 100 KB message into a 512×512 PNG image requires a median 0.43 seconds; extraction and decryption requires 0.39 seconds — acceptable latency for non-real-time secure communication. InvisiCrypt is

23% faster than the DCT-based alternative, attributable to pixel-domain rather than frequency-domain embedding, which avoids the overhead of forward and inverse discrete cosine transforms.

5. DISCUSSION

A. Interpretation of Results

The PSNR results confirm that complexity-adaptive bit allocation offers a measurable advantage over static LSB approaches. By concentrating higher-order bit substitutions in textured regions, InvisiCrypt exploits the spatial contrast masking property of the human visual system [13]. The near-random RS and chi-square detection rates validate the effectiveness of keyed pseudo-random traversal as a defence against classical steganalysis. Crucially, these rates fall within confidence intervals consistent with the null hypothesis of no embedding — a qualitatively stronger result than simply reporting low detection values, as it formally indicates steganalytic indistinguishability under the tested attack models.

B. Limitations

The current implementation does not address JPEG compression: quantisation in the frequency domain destroys spatial-domain LSB information, limiting InvisiCrypt to lossless PNG or BMP carriers. This excludes a substantial portion of common digital communication channels such as social media platforms and email clients that apply automatic compression. The shared-key distribution problem inherent in symmetric cryptography is not resolved within this framework; a pre-established secure channel for key exchange — such as Diffie-Hellman or RSA key encapsulation — is assumed but not provided. Resistance against deep-learning steganalysis frameworks such as SRNet [11], which learn statistical signatures beyond the reach of classical analytical methods, was not evaluated and constitutes an urgent direction for future work.

C. Practical Applications

Three direct deployment contexts are identified: (1) Confidential enterprise communication, where organisations transmit proprietary data over public channels without creating a detectable pattern of encrypted traffic; (2) Digital forensics and law enforcement, embedding tamper-evident authentication watermarks in chain-of-custody evidence images without visually degrading them; and (3) Academic and cybersecurity research, providing a transparent modular benchmark for studying the capacity-imperceptibility-security trade-off.

6. CONCLUSIONS

This paper presented InvisiCrypt, a dual-layer secure communication framework unifying AES-256 CBC cryptography with a complexity-adaptive LSB steganographic pipeline. Evaluation on a 200-image benchmark confirmed PSNR of 52.83 dB and SSIM of 0.9991 at 52% embedding capacity, outperforming five state-of-the-art hybrid schemes. Steganalysis detection rates of 3.8–6.1% across histogram, chi-square, and RS methods confirm that the embedded payload is statistically indistinguishable from natural image noise.

The primary limitation — vulnerability to JPEG quantisation — motivates extending embedding to the DCT frequency

domain for lossless-equivalent operation on JPEG carriers. Incorporating public-key encapsulation for key distribution and evaluating resistance against deep convolutional steganalysis networks are further planned enhancements. InvisiCrypt demonstrates that strong cryptography and perceptually adaptive steganography are mutually reinforcing: encryption randomises the payload into noise-like ciphertext, while adaptive embedding distributes that noise in visually complex regions where it is imperceptible — a synergy that is the core architectural principle of the framework.

ACKNOWLEDGEMENT

The authors sincerely thank the Department of Computer Science and Engineering, Nutan College of Engineering and Research, Pune, for providing the computational infrastructure and research guidance that supported this work. Gratitude is also extended to the anonymous peer reviewers whose constructive feedback considerably strengthened this manuscript.

REFERENCES

- [1] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Information hiding — a survey," *Proc. IEEE*, vol. 87, no. 7, pp. 1062–1078, Jul. 1999.
- [2] N. Provos and P. Honeyman, "Hide and seek: An introduction to steganography," *IEEE Security & Privacy*, vol. 1, no. 3, pp. 32–44, May–Jun. 2003.
- [3] G. J. Simmons, "The prisoners' problem and the subliminal channel," in *Proc. CRYPTO 1983, LNCS*, vol. 196, pp. 51–67, Springer, 1984.
- [4] C.-K. Chan and L. M. Cheng, "Hiding data in images by simple LSB substitution," *Pattern Recognition*, vol. 37, no. 3, pp. 469–474, Mar. 2004.
- [5] D.-C. Wu and W.-H. Tsai, "A steganographic method for images by pixel-value differencing," *Pattern Recognition Letters*, vol. 24, no. 9–10, pp. 1613–1626, Jun. 2003.
- [6] National Institute of Standards and Technology, "Advanced Encryption Standard (AES)," *FIPS Pub. 197*, Nov. 2001. [Online]. Available: <https://doi.org/10.6028/NIST.FIPS.197>
- [7] B. Schneier, *Applied Cryptography*, 2nd ed. New York, NY: Wiley, 1996.
- [8] S. Bhattacharyya, I. Banerjee, and G. Sanyal, "A novel approach of data hiding using pixel mapping method," *Int. J. Comput. Sci. Inf. Security*, vol. 8, no. 2, pp. 146–154, 2010.
- [9] H. Al-Dmour and A. Al-Ani, "A steganography embedding method based on edge identification and XOR coding," *Expert Syst. Appl.*, vol. 46, pp. 293–306, Mar. 2016.
- [10] P. Thiyagarajan, G. Aghila, and V. P. Venkatesan, "Reversible dynamic secure steganography for medical image using genetic algorithm," *IET Image Process.*, vol. 7, no. 4, pp. 342–351, Jun. 2013.
- [11] B. Xu, J. Wang, T. Chen, and S. Pan, "Structural design of convolutional neural networks for steganalysis," *IEEE Signal Process. Lett.*, vol. 23, no. 5, pp. 708–712, May 2016.
- [12] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *Handbook of Applied Cryptography*. Boca Raton, FL: CRC Press, 1997.
- [13] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.