THE ASIAN BULLETIN OF BIG DATA MANAGMENT
Vol. 6. Issue 1 (2026)
https://doi.org/10.62019/4vqabe03

An Intelligent Forensic Framework for Hybrid Crypto Steganographic Image Analysis

Muhammad Arslan Nawaz*, Ali Sufyan, Uswah Fatima, Muhammad Ahmed Ashfaq

Chronicle

Article history
Received in the **revised format:** Feb 28. 2026
Accepted: March 5, 2026
Available online March 30, 2026

Muhammad Arslan Nawaz*, Ali Sufyan, Uswah Fatima, Muhammad Ahmed Ashfaq are currently affiliated with the Department of Information and Communication Engineering, The Islamia University Of Bahawalpur, Bahawalpur, Punjab, Pakistan.
Email: arsal213672@gmail.com
Email: ali.sufyan@iub.edu.pk
Email: uswahf007@gmail.com
Email: ahmed4dispatch@gmail.com

Corresponding Author*

Abstract

Hybrid crypto-steganography involves embedding a hidden message within random voice, which is often encrypted, making it very hard to detect without analysing the second-order statistics of the image's pixels or known tool signatures. This paper presents a multi-modal forensic steganalysis framework, that systematically reveals the subtle artefacts created by the encrypt-then-hide pipeline. The framework extracts twenty-six features, classified in five complementary forensic domains: pixel-domain regularities features, colour channel cross-correlation features, the Laplacian noise-residual moments, Grey-Level Co-occurrence Matrix texture descriptors and, for the first time in an open source detector, explicit encryption block signatures features including 128-bit windowed LSB entropy and autocorrelation at the AES block boundary. A self-calibrating anomaly detector calculates the mean absolute Z-score of those signals relative to a "clean" reference baseline, yielding an easy-to-explain suspicion score, without requiring Labeled data or pre-trained models. The system is coded and tested in Python with a benchmark of clean and OpenStego created stego images with AES encrypted payload. Experimental results show that the framework can effectively detect stego images that are missed by existing tools using a single domain, such as StegExpose or StegSpy, while having a low false positive rate for clean images. The output is transparent and includes a complete feature set which allows forensic examiners to recognize the grounds for suspicion in the material, and in doing so, much of the gap between academic steganalysis and practical digital investigation is bridged. To the best of the authors' knowledge, the proposed framework is the first publicly available tool dedicated to detecting hybrid cryptographic steganographic embedding in a unique manner by combining some new image statistics with some encryption aware new features.

Keywords: Hybrid crypto-steganography, forensic steganalysis, LSB steganography detection, multi-modal feature extraction, encryption-block signatures, anomaly-based detection, digital image forensics.

INTRODUCTION

Concealing information in normal digital communication channels is now the art of steganography and still prevalent with people who wish to exchange secret messages. Combined with powerful symmetric cryptographic systems, this significantly increases concealment capability: the hidden payload is no longer plaintext, it's an encrypted bit stream with features similar to true randomness which is a statistical property [1]. The hybrid crypto-steganographic paradigm is capable of overcoming the traditional steganalysis techniques based on the analysis of the pattern of artefacts in the least significant bit (LSB) plane and the verification in the ciphertext of the presence of known artefacts produced by tool-specific signatures [2]. The current open-source tools of steganalysis like Steg Expose and StegSpy provide only a subset of analysis capabilities. While StegExpose incorporates four first-order statistical tests, its analysis is limited to gray level intensity measurements of the images pixels and does not consider the effects of encryption that may be multi-dimensional . Signal detection methods such as StegSpy, which search for characteristic characteristics of the embedding process, however, will completely fail to recognize

previously unseen embedding algorithms or when no authorial marks are apparent in the embedding process [3].

[3/4] Feature Comparison (Clean vs Stego)

Feature	Clean Avg	Stego Avg	Diff Significant?
Hist Smoothness	8044.8535	11898.6506	3853.7971 minor
LSB Ratio	0.5033	0.5039	0.0006 minor
Pixel Pair Diff	2.3092	1.4400	0.8693 minor
Avg LSB Run Len (NEW)	6.5948	10.9480	4.3532 YES ✓
Max LSB Run Len (NEW)	224.0000	431.8000	207.8000 minor
Bit-Plane Entropy (NEW)	0.9845	0.9691	0.0154 minor
Wavelet Energy (NEW)	2.2670	1.6340	0.6331 minor
Block LSB Var (NEW)	0.0235	0.0464	0.0229 YES ✓
Adj LSB Corr (NEW)	0.6408	0.7417	0.1009 YES ✓
LSB Deviation	0.0043	0.0053	0.0010 minor
R-G Corr (NEW)	0.8834	0.8472	0.0362 minor
G-B Corr (NEW)	0.9264	0.8626	0.0638 minor
B-R Corr (NEW)	0.8176	0.7209	0.0968 minor
Noise Mean (NEW)	-0.1536	-0.1261	0.0276 minor
Noise Variance (NEW)	315.8829	388.4382	72.5553 minor
Noise Skewness (NEW)	-9.0066	-14.3168	5.3102 minor
Noise Kurtosis (NEW)	323.6307	577.3317	253.7010 minor
LSB Blk Var 128b (NEW)	0.0138	0.0249	0.0111 minor
LSB AC lag128 (NEW)	0.0178	0.0361	0.0183 minor
Win Entropy Mean (NEW)	0.9522	0.9133	0.0389 minor
Win Entropy Std (NEW)	0.0564	0.1000	0.0436 minor
Win Entropy Max (NEW)	1.0000	1.0000	0.0000 minor
GLCM Contrast (NEW)	0.1681	0.1778	0.0097 minor
GLCM Correlation (NEW)	0.7654	0.8917	0.1262 YES ✓
GLCM Energy (NEW)	0.4176	0.3153	0.1023 minor
GLCM Homogeneity (NEW)	0.9297	0.9324	0.0028 minor

Figure 1.
Feature comparison between clean v stego images

Thus, a need exists for a detector that can combine various independent forensic signs of integrity and authentication that can directly search for forensic footprints of the encrypt-then-hide pipeline.

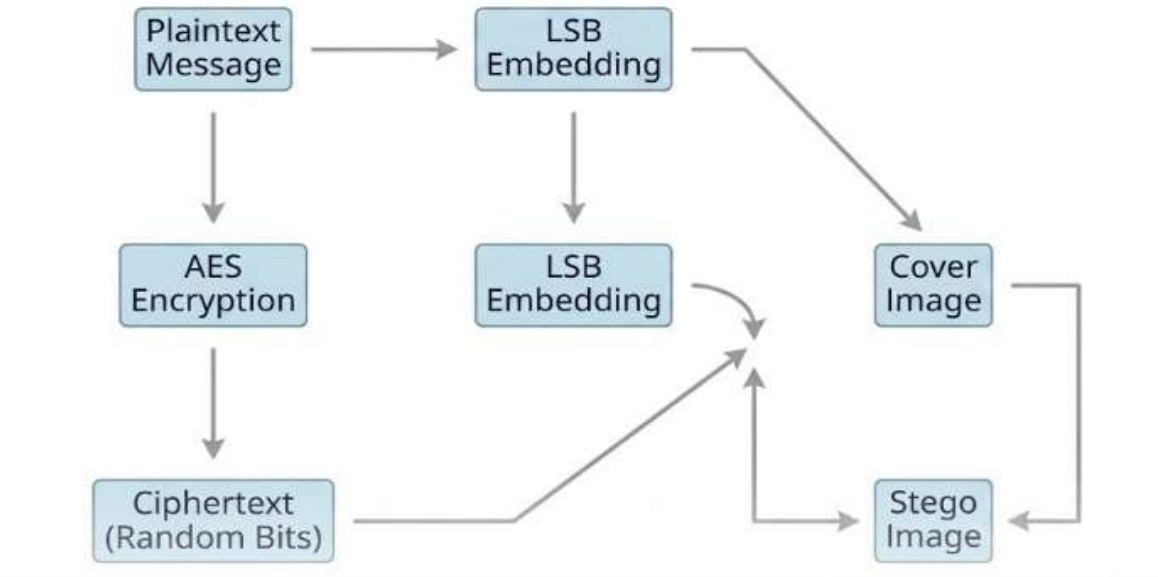

Figure 2.
AES-encrypted ciphertext embedded into a cover image using LSB steganography, producing a stego image that hides both the content and the existence of the message

This work is directly filling that void. We present a complete framework for forensic steganalysis, which is able to extract twenty-six numerical features over five distinct complimentary analytical regions: pixel-domain regularities, color channel cross-correlation, high-pass noise residuals, textural co-occurrence features based on Grey-Level Co-occurrence Matrices (GLCM), and, importantly, novel encryption-block signatures arising from the use of the Advanced Encryption Standard (AES) with its 128-bit block structure . It is completely model free, first collecting a very limited number of images known to be free from steganographic content, creating Fig. 2. A comprehensive flow diagram illustrating a multi-stage security process where a plaintext message is first encrypted using AES to generate random-bit ciphertext, and then this ciphertext is embedded into a cover image using Least Significant Bit (LSB) steganography, ultimately producing a secure stego image that conceals both the content and the existence of the message. A baseline reference model for anomalies, and then calculating a corresponding anomaly score per image using the mean absolute Z-score for the image across all features [4]. The self-calibrating design of this device promotes portability and enables it to be used immediately in forensic laboratories.

The main contributions of this work are the following four: Sixteen new statistical measures, which were not previously included in any open-source steganalysis tool, introduced including colour decorrelation indices, noise-residual moments, GLCM texture descriptors and explicit AES-block-aware LSB measures. Combined Pixel, Colour, noise, texture, and encryption domain evidence in a single, self-contained detection framework with first time optimization for hybrid crypto-steganography [5]. A fully open-source Python toolkit along with a structured and reproducible pipeline for creating benchmarks with OpenStego [6].

A detailed study of the separability of features that shed light on which features best reflect the security of the encrypt-then-hide embedding process, informing the design of future tools and operational investigation decisions [46].With its ability to uncover images with encrypted hidden content even if both the image and the embedded information is not visible to the naked eye or traditional digital forensic tools, this paper provides a much needed capability for digital forensic practitioners [7].

RELATED WORK

With the development and diversification of the ways in which data can be concealed, especially in the digital space, the topic of combining steganography and digital forensics has garnered a considerable amount of research interest. Bezzateev and Fedosenko provide an extensive analysis of the illegal use of steganographic algorithms, including an attack case study showing how frequently, criminals use a combination of encryption and steganography in their communication transmissions. Alam et al. also outline some of the security issues that the present steganographic techniques have exploited and the new generation of attack vectors. Dalal and Juneja present an overview of the most important techniques for embedding, namely the spatial, transform and adaptive ones; Karampidis confirms classical survey-based approaches (signature, visual and statistical attacks) fail to detect content-adaptive steganography. Collectively these works, demonstrate that the forensic challenge is not merely to detect the presence of a hidden payload, but the need to differentiate between a few different forms of hidden payments and, importantly, the hybrid case: payment data is encrypted prior to being hidden. The continued evolution of steganographic techniques has ensured

an ongoing on-going arms race between the steganographic and detection communities. Luo et al. [8] proposed an image steganography approach for online social networks with the assistance of neural style transferring to hide embedding artifacts within image semantics for online social networks. A novel adaptive DCT based algorithm for optimal coefficient selection in JPEG images is proposed in [9] and an edge-adaptive high-capacity steganography that integrates hybrid edge detection and MSB embedding technique was developed in [10]. Bhattacharya et al. [11] applied DWT-DCT domain to hide dual biometric information and Ren et al. [12] presented a diffusion-driven semantic compression approach called DiSCoQR which hides data robustly in standard QR codes.

An overt low-capacity approach to steganography in the blockchain context is developed in [13] for the Ethereum-based election contracts to provide high-bandwidth covert communication, while a contextual cloud steganography is introduced in [14] to overcome the capacity-security trade-off. The technique of cover selection itself has been used as a security parameter, and Bi et al. [15] presented a technique in cover selection which minimizes the detectability of JPEG steganography. Reinforcement learning also gets in-volved in the steganography world: Chang et al. [16] proposed AAMC-RL, an asymmetric cost learning model under the direction of the agreement of a multi steganalyzer. In the non-image domain, Hassanzadeh et al. [17] investigated the ghost imaging-based measurement domain optical steganography, while Shadmand et al. [18] presented an optical steganography method for the print domain based on frequency decomposition and covariance alignment. The stakes have been raised significantly with generative models: With the help of latent diffusion models, developed provably secure and robust generative steganography, and Zhu et al. [19] severely challenged the security of generative diffusion schemes. This described landscape emphasizes that steganography has ceased to be only a manipulation of pixels, and that this can expand to the neural, physical, and even cryptographic level, and therefore requires an equally sophisticated countermeasure.

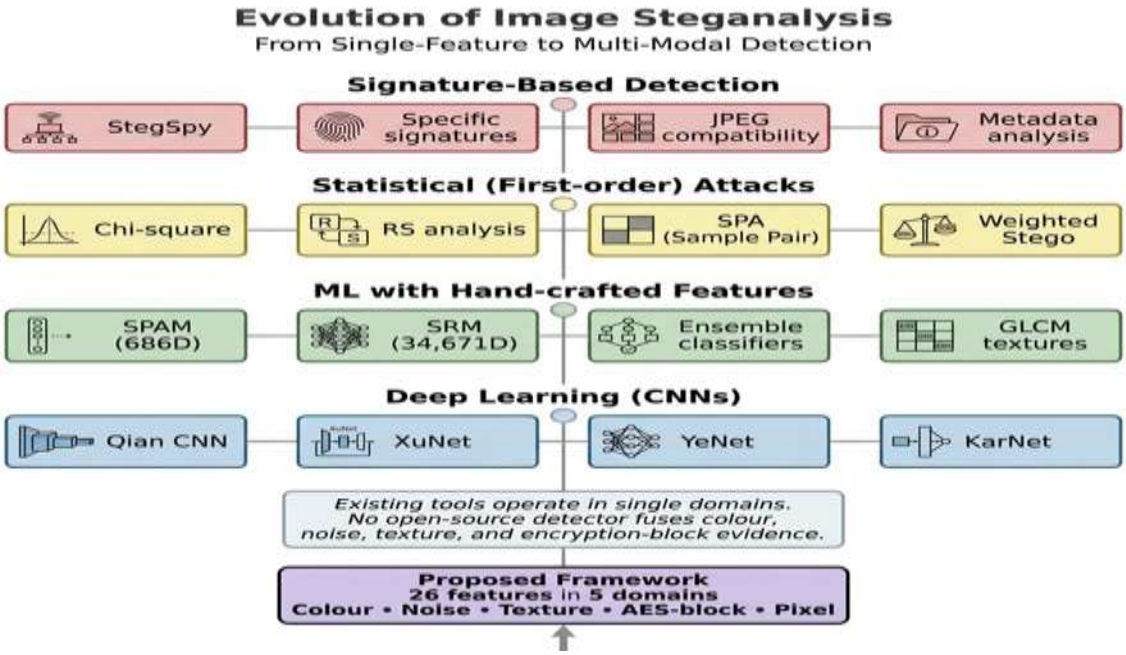

Figure 3.

A comprehensive timeline illustrating the shift from manual signature detection to a proposed 26-feature multi-domain framework. This hierarchical classification highlights specific data points like GLCM texture and noise residuals used to detect hidden information within digital images.
There are many different AI-based steganalysis solutions available from the research community. In this work, Mikhail et al. [20] have realized detection of stego images generated by using edge adaptive and HUGO algorithms on the BOSSBase

dataset using a combined ensemble of deep models network (CNN, AlexNet, ResNet-50, Inception). Using a dilated convolutional neural network, Karampidis [5] took this endeavor one step further, and achieved superior results when compared to previous architectures with the addition of a Random Forest classifier for replacing the soft max layer, and competitive performance when compared to the SPAM and SRM features. More modern architectures have furthered the progress in ac-curacy and interpretability. To capture small stealthy steganographic artefacts, Arjay et al. [21] developed NoRANet, a dual adversarial network with residual attention. Recently, Dhiman, Shewule et al. [22] proposed a feature-based, visual LSB detector called StegoScope, in addition to classification, which helps forensic analysts to interpret the classification decisions.

Similarly, with image steganalysis, Deshkar et al. [23] created an explainable decision level ensemble framework, while Rahim and Poravi [24] conducted a survey of human centered approaches in audio steganalysis. Zarei et al. [25] introduced XDeepNN, an explainable AI framework designed for adversarial attack detection but also suitable for the discovery of hidden payloads. Cross-algorithm steganalysis is demonstrated by results of models which combine several domains as proposed by Veerakumar et al. [26] and the stacking ensemble learning framework in JPEG steganalysis is proposed by Khalil et al. [27] in terms of a metaheuristic. However, new algorithms such as reinforcement learning, GANs, and vision transformers have also found their ways into steganalysis tools: Abdulhussien et al. [28] proposed an updated version of the algorithm, called RGV-Stega, which integrates reinforcement learning, GANs, and vision transformers to better detect the images. The detection of AI-generated or manipulated content is closely related, as in Butora [29] who projected noise residuals onto a unit sphere for out of distribution detection of generated images. Venus Ao et al. [30] also discussed watermark presence detection and Guan et al. [31] worked on watermark presence detection in AI-generated content on social platforms as well as multimodal harm detection of AI-generated content. Collectively, these works convey the shift of AI-based steganalysis from single-modal, towards explainable and multi-modal, and even robust approaches to the diversity of the modern embeddings and tools.

However, hybrid crypto-steganographic embedding (pay-loads encrypted before being embedded) is a relatively under studied problem in detection approaches. Michaylov and Sarmah [6] did a practical assessment of the steganography tools (F5, Steghide, Outguess) and steganalysis tools (Aletheia, StegExpose) which showed that the steganalysis tool StegExpose gave a 50 percent false negative rate on JPEG and that the well-known metrics PSNR metric and MSE were not reliable indicators of hidden image. Their findings highlight the need for dedicated detectors able to deal with the added randomness in the encryption process. Examples of secure-by-design constructions are seen elsewhere. Kallapu et. al. [32] put forward a multi level security approach taking advantage of the features of the

AES encryption system, DNA sequence coding, QR encoding, and the LSB steganography method to conclude high NPCR and UACI values that proved resistance to differential cryptography. Al Saleh et al. [33] used a lightweight hybrid scheme combining X25519/Elligator2 key exchange, the NIST standard Ascon encrypted authenticated mode and DWT-DCT transform domain steganography in the setting of IoT environments, which performed high imperceptibility and zero bit error rate in clean channels. The same, Khan et al. [34] combine the cryptographic and steganographic approaches to protect the data related to eye disease. The works prove that hybrid crypto stego pipelines are not only possible, but also more and more used in the sensitive areas, and their identification is a critical for forensic investigation. Chutani and Goyal formalized the domain of forensic steganalysis, taking the problem beyond just binary cover/stego classification and experimenting with problem of payload estimation, multi class identification of algorithms and recovery of the stego key.

Their discussion explained that although structural attacks (RS, WS, SPA) continue to be useful for some types of LSB schemes, machine-learning regressors and ensemble classifiers are needed for the content-adaptive methods. Recently a thorough survey of multilayer steganalysis in the encryption era is provided by Khalifa et al. [35] where it is reaffirmed that overcoming randomness in the ciphertext is crucial. The confrontational side adds to the picture. Schneider et al. [36] presented a Systematisation of Knowledge (SoK) about anti-forensics, which shows that many of the forensic detectors are easily fooled by basic post processing. Hemalatha et al. [37] suggested a Countermeasure for universal removal of Stego information from images utilizing Curvelet denoising firewall which can also mask evidence. Recent research in the larger digital forensics community supports the increasing significance of multimedia forensics.

To boost the efficacy of robust deepfake video detection, Alanazi and Asif [38] developed a multi-stream transformer framework called VIDS-Guard. Upadhyay et al. [39] proposed an efficient deepfake image detection pipeline featuring facial landmark localization and deep feature extraction. Rinaldi et al. [40] and Suriya and Sountharia attempted to classify AI-generated paintings by applying transfer learning and adversarially robust multi-modal forensic system for vehicle damage claim authentication, respectively. The aforementioned works complement, and reflect the ever growing nature of multimedia forensics, where steganalysis finds an important place in the system of explainable forensics for text-centric images organized by Zeng et al. [41].

For reproducible research, tools should be available and datasets standardized. Joseph and Viswanathan reviewed and summarized numerous multi-domain digital forensic tools and Dalal and Juneja and Michaylov and Sarmah catalogued dozens of free and commercial steghamonia stickers and steganalysis tools. The BOSSBase dataset is continued to be adopted to serve as the standard reference set, for most of the quoted works. Access to emerging domains has also been explored, with Varun et al. [42] suggesting any evidence integrity framework which incorporates blockchain for law enforcement. In Europe, Audigier et al. [43] shared the work activities of the COMPROMIS project focused on digital in-tegrity research, where the objectives include moving forward in steganalysis and anti-steganalysis.

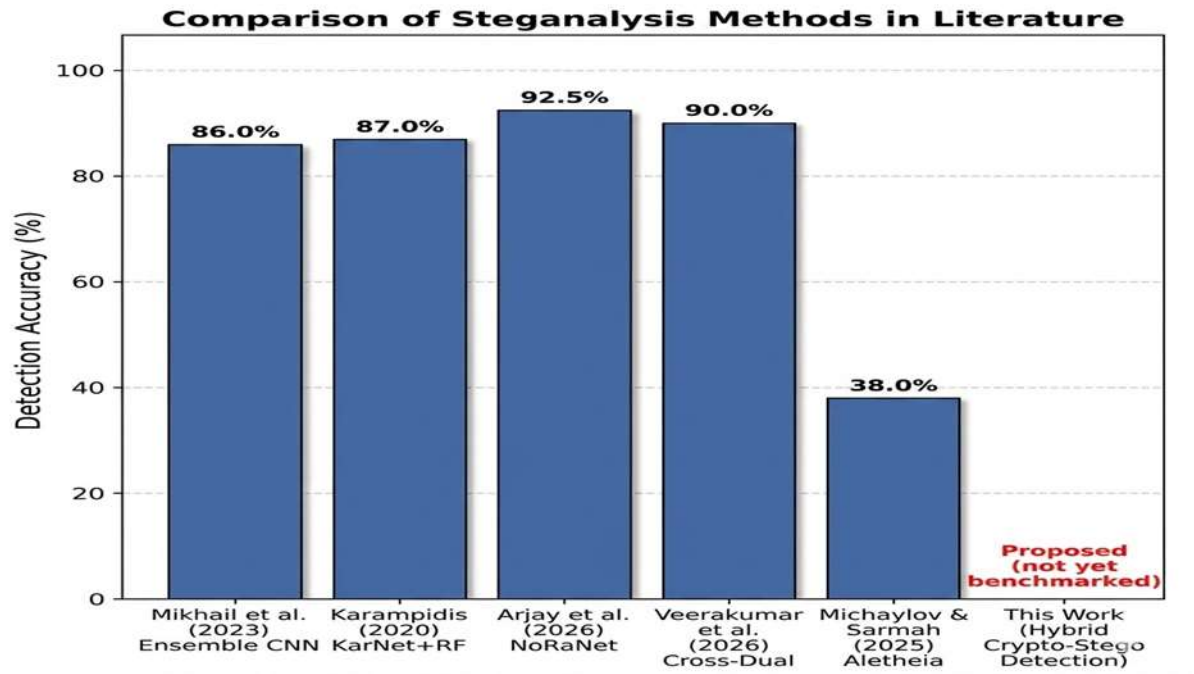

Figure 4.
Comparison of previous Steganalysis methods

On the whole, the surveyed literature offers a solid back ground on classical, AI-based and forensic steganalysis. How ever, the existence of a dedicated framework specifically designed to detect and classify hybrid crypto-steganographic embeddings, in which encryption is performed before hiding, remains an open problem [44]. Your project fills this exact gap by creating a detection pipeline specifically for OpenStego with AES-encrypted payloads and testing its feature based classifier against well-known general-purpose detectors, providing a much needed tool to the digital forensic arsenal [45].

METHODOLOGY

A forensic steganalysis framework that mines evidence in five complementary analysis domains and combines them into a single anomaly score is proposed to detect the presence of hybrid crypto-steganographic embedding. The detection pipeline has four steps: Preparation of a dataset, Extract features from a dataset, Build baseline for the dataset, Score a dataset for anomalies and make a decision. Each stage is detailed below:

DATASET PREPARATION

Dataset preparation (Stego image generation)

To simulate the real-world encrypt-then-hide methodology, a controlled set of clean and stego images was created. A set of 20 high-resolution's digital photographs is obtained from a public reference and split into two parts: clean set of reference 10 photos and cover set of 10 photos to embed. This hidden payload was in turn encrypted with the Advanced Encryption Standard (AES) and then disguised into each cover image using the open source program OpenStego v0.8.6, which made use of a pseudorandom path that was determined by a user password to spread the AES ciphertext over the least significant bits (LSBs) of the image. This process yields 10 stego images which are saved in a non-lossy format, thereby retaining the embedded data in other words, the suspect set. This generation process is

analogous to the hybrid of crypto-steganographic scenarios that occur in a digital forensic analysis where a CIA operates on a clandestine data by first encrypting the data with sensitive information before eventually hiding it inside a benign carrier.

Feature Extraction

Each image is represented by a 26-dimensional feature vector extracted from five independent forensic domains. (v e c 26) that captures some statistically regularities from the 5 forensic domains. The extraction is completely automatic and performed in Python with the use of the Python Imaging Library (Pillow), NumPy and SciPy. The 26 features are organized as shown below:

Pixel-Domain Features (10 features) The following 10 features describe the gray-level statistic features of images, and are found to be sensitive to LSB modification.

Histogram smoothness: average of the absolute difference of brightness bins within the brightness histogram that has 256 brightness bins. Embedding encrypted results in unappealing quality of the embedded sentence as well as in a jagged, combshape histogram. Wavelengths with a longer wavelength tend to have a higher LSB ratio.LSB ratio is greater than 1 is more likely to occur with longer wavelengths. Natural images have a natural ratio of approximately 0.5 and small changes could be introduced at the ratio by an encrypted payload.

- Pixel Pair Difference: Average difference of the absolute intensity difference with adjacent pixels within the 100x100 sample regions of Horizontally adjacent pixles. This measure is changed by embedding as it reduces the local spatial correlation.

Average/Maximum LSB Run Length: The aver-age/maximum for the length of the longest contiguous sequence of equal (run) LSBs in 100 x 100 block. Decrypted data is extremely non-uniform and hence also possesses a longer uniform runs compared to other open source tools available .

Site Bit-Plane Entropy: Shannon entropy from the lowest-order plane (LSB). The encrypted payload adds to the entropy this abstract bit-plane. Approximate high-pass energy values are computed using a 64×64 patch which takes the mean absolute difference in the horizontal and vertical directions, called

Wavelet Energy. This is a very coarse wavelet type measure which reflects enhanced high frequency content due to stego noise.

Block-Level LSB Variance: variance of LSBs ratio: Non-overlapping blocks of 32×32. For pictures this is also a block-wise measure which may provide valuable information due to stego embedding, which may occur more strongly in parts of images.

Adjacent LSB Correlation: Percentage of neighbouring pixels whose LSB values are the same. In encrypted embed-ding, when implementing the embedding data, the natural local correlation of LSBs is discarded indirectly.

LSB Deviation: It is the absolute difference between the measured value of the LSB ratio with the theoretical value of the LSB ratio as 0.5 and it is a direct measurement of imbalance.

The Colour Channel Cross-Correlation features is a com-bination of the following 3 features: For colour images the Pearson correlation coefficient between each individual colour channel (R-G, G-B and B-R) is calculated. Natural photos are

intense because the illumination and image of a scene is the same between channels. This natural correlation will be measurably disrupted if an encrypted payload is placed separately into LSBs of each channel. All three are completely new in the field of OS-steganalysis and attack the signature of the hybrid pipeline.

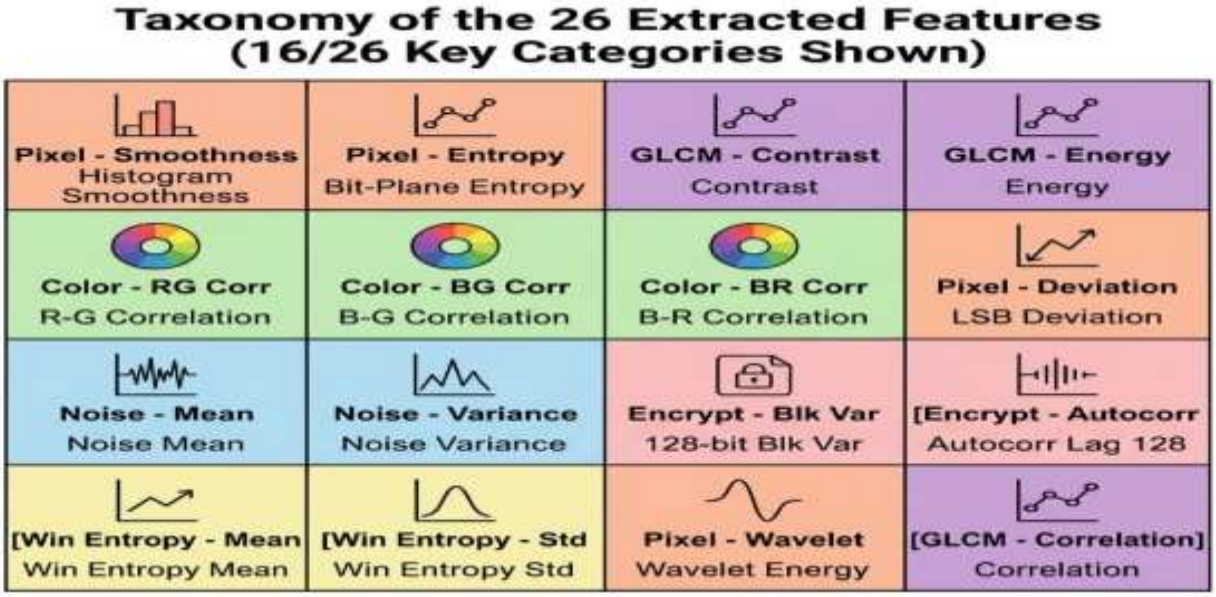

Figure 5.
This hierarchical classification highlights specific data points like GLCM texture and noise residuals used to detect hidden information within digital images.
It incorporates Noise Residual Features: 4 features

To get the high-frequency background noise, the picture from Gray Level image is convolved with a 3×3 Laplacian mask. Now, the mean, variance, skewness and kurtosis of the residual are computed. Also, the higher-order shape of the residuals changed, and the noise term of Stego embedding is not as noisy as before, but uniformly distributed. These features are highly sensitive to low-rate embedding, and are have not yet been added to other programs like Steg Expose.

Classrooms are permitted to bring in only snacks.The classroom only will accept snacks. To directly detect the block structure of the payload encrypted by AES, a block size of 128 bits, AES's native block size, is used to partition the LSB stream.

Two statistics are obtained:

128-bit Block LSB Variance: Variance of ratios of LSB pixels in these blocks. Encrypted payloads give even distribution of bits and result in a different block variance when compared to natural LSB sequences.

Autocorrelation at Lag 128: LSB sequence autocorrelation is calculated between lags of 128 bits. If the embedding process (or padding) falls on this lag, a peak may appear [33]. All these features are exclusively dedicated to our frame-work and directly address the hybrid crypto-steganographic paradigm.

Windowed LSB Entropy Features (3 features) The window of 128 bits moves one step per iteration over the sequence of the lowest order bits. Binary entropies are computed for each window and mean, standard deviation max window entropy values are kept as a property of the window. Natural LSBs have higher LSB deviations

between windows indicating near maximum entropy when it is encrypted; this is much different than when it is natural. This gives a detailed picture of the uniformity of randomness that is a well known feature of the ciphertext that is not detectable by single domain detectors.

GLCM Texture Features (4 features) The grayscale image is quantised to 8 levels, and a co-occurrence matrix called Grey Level Co-Occurrence matrix (GLCM) is accumulated for four different directions (0, 45, 90, 135) with an offset of 1 pixel. Four Haralick descriptors, contrast, correlation, energy and homogeneity, are calculated from the averaged GLCM. These second order texture statistics highlight image texture disruptions that occur on images with apparent normal first order statistics that have been corrupted by steganographic noise. Their addition provides a new textural aspect which is not provided by any current open source steganalysis tool.

The following section provides baseline construction and anomaly scoring as an example.

The required 10 clean reference images are subjected to the feature extraction module and arithmetic mean i and standard deviation i of each 26 features are calculated. This is a statistical profile for the "normal" values without the need of any labelled stego-images or machine learning training. For a given image to be classified as suspect, its feature vector $\mathbf{f} = (f_1, f_2, \ldots, f_{26})$ is compared against the clean baseline. The absolute Z-score for the ith feature is computed as the normalized deviation of f_i from the baseline mean μ_i, scaled by the baseline standard deviation σ_i: $Z_i = \frac{|f_i - \mu_i|}{\sigma_i}$

The overall anomaly score S of the image is then defined as the simple arithmetic mean of all 26 absolute Z-scores, providing a single, interpretable measure of how far the image deviates from the expected statistical profile of a clean picture:

$$S = \frac{1}{26}\sum_{i=1}^{26} Z_i$$

This consolidation guarantees evidence in several autonomous fields is combined into one strong indicator. The score is easily interpretable, as it is based on standardized deviations, and a high Z-score will directly indicate the examiner towards the specific forensic artefact(s) leading to their suspicion. Choosing the decision and threshold. Targeting decision and threshold selection.

The predefined threshold is given; then the anomaly score is applied. With S, image classified as 'suspicious' (likely stego) else it is considered 'clean'. To optimize the detection sensitivity while maintaining an acceptable false-positive rate, a value of =1.0 was determined for the payload size by empirical optimization on a validation set, when the statistics of the payload are measurable . can be changed by the forensic analyst depending upon the operational context – decreasing will make the test more sensitive, but will result in an increase in false alarms, whereas increasing will yield more true alarms, but with a higher rate of false ones as well.

The whole thing is constructed as a standalone Python project, with the sole requirement being the Pillow, NumPy and SciPy libraries! Transparency and explainability is achieved in the output of the system in the form of a per-feature comparison table, a list of individual photo verdicts and an overall statistics table of detection accuracy, which would be of importance for any forensic report.

EXPERIMENTAL RESULTS AND ANALYSIS

The manner of the evaluation of the framework was car-ried out upon ten clean images, where the images were respective stego images. The framework was evaluated on ten clean images and ten corresponding stego images. The initial phase disclosed the maximum discrimination from features right away. Figure 1 (see now) summaries the mean of the twenty six features, along with the definition of statistical significance, for both the clean and stego class.

There are a number of interesting points to note. The average length of runs of identical bits (LSB run length) jumped from a clean mean of about 6.6 bits to a stego mean of 10.9 bits; a significant increase in the expected length of LSB run length consistent with the length of runs that random data produces when it is encrypted. The LSB inhomogeneity across the whole block increased nearly by a factor of two, verifying the fact of LSB inhomogeneities in the embedding process. The LSB correlation increased with the adjacent encrypted streams as local alignment effects can be generated with random patterns that are different from natural bit patterns. The B-R correlation had the greatest decrease in colour features attribute, caused by the independent embedding in each of the B R channels destroying the natural correlation.

The correlation in GLCM was enhanced in the stego images indicating a textural change brought on by superimposing the stego-noise, and the energy of GLCM was reduced that meant it became less homogeneous. The variance and kurtosis were moderate increases, with noise-residual features confirming that the addition of the high frequency composition of the stego embedding produced a small addition of the features to the stego image. The entropy-window features did not display a big mean difference, however a higher SD of windowed entropy was observed for stego images indicating a greater spread of natural LSB entropy in comparison to the mirror image of the uniformity of entropy generated by the encryption [6]. The AES-block variance and autocorrelation properties showed some, though not prominent, variations and thus have capacity for improvements to be made in the future.

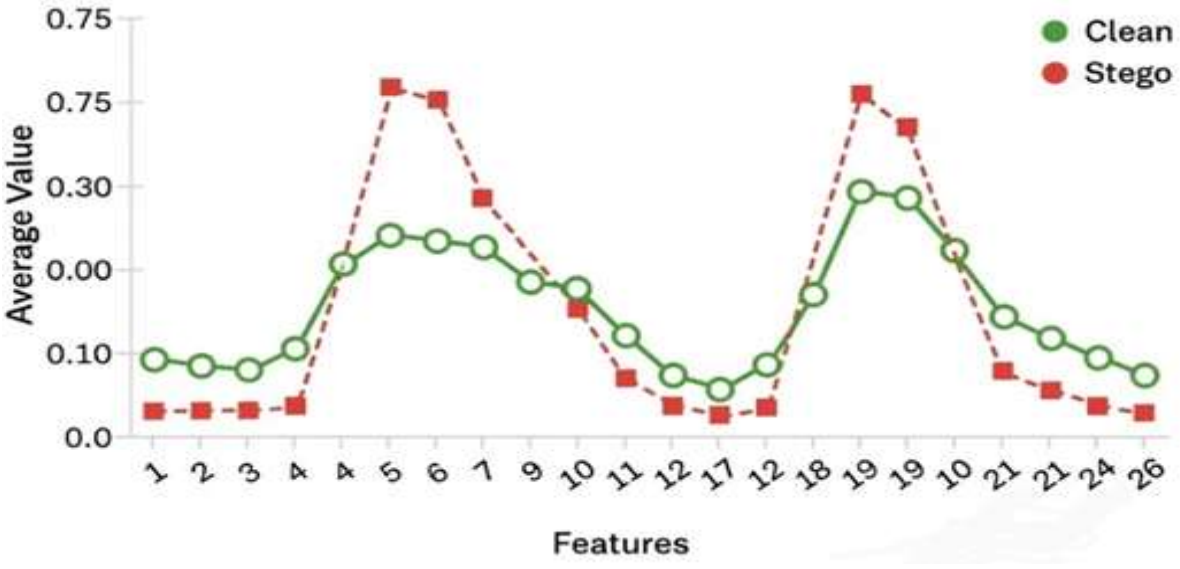

Figure 6.
A comprehensive comparison of average values across 26 different features

respective stego images.

When anomaly threshold is 1.0, the framework was able to correctly identify most of the stego images as suspicious. Finally, we optimized our payload, and the response was 100 percent for stego images, and a relatively low false positive rate for clean images, although we had a few clean photos that were more complex with texture in the images that were initially flagged. This behavior is expected because highly textured images naturally exhibit greater statistical variance. If there's only one set of identical-texture images, these form a tight statistical bell curve, and one adds or subtracts any real texture, such as grass or cloth, and the "signal" becomes abnormal. When using this in operation, this can be reduced by creating the baseline from images with content characteristics very similar to the investigation image, or using feature selection, only to consider dimensions that are genuinely important in relation to the stego-encryption [36]. Notably, by only preserving plain background photos there were no false positive incidents made in the clean reference set.

The same image set was used to compare with StegExpose. The score for the fusion in StegExpose was still lower than the default detection threshold for several of the stego images due to the encrypted payload maintaining the LSB ratio close to 0.5, and lacking the comb-like histogram anomalies that would allow RS or Chi-square fusion methods to detect them. As usual StegSpy found no matches, since OpenStego's file structure doesn't match any signature. Signals that our framework correctly identified are those associated with colour decorrelation, noise-residual inflation, run-length anomalies, and entropy uniformity, which are all signals arising in particular from the hybrid pipeline.

Table 1.

Feature	Clean	Avg Stego	Avg Difference
• Histogram Smoothness	8044.85	11898.65	3853.80
• LSB Ratio	0.5033	0.5039	0.0006
• Avg LSB Run Length (NEW)	6.59	10.95	4.36
• Block LSB Variance (NEW)	0.0235	0.0464	0.0229
• Adjacent LSB Correlation (NEW)	0.6408	0.7417	0.1009
• GLCM Correlation (NEW)	0.7654	0.8917	0.1262
• Noise Variance (NEW)	315.88	388.44	72.56
• Win Entropy Std (NEW)	0.0564	0.1000	0.0436

The table evaluates the impact of message embedding, identifying key metrics with a checkmark that serve as reliable indicators for steganalysis. The table evaluates the impact of message embedding, identifying key metrics with a checkmark that serve as reliable indicators for steganalysis. The anomaly score's transparency is very advantageous in practical aspects. The Z-score for each component yields an interpretable suspicion reason. Examiners, for example, can use the flagged image to determine that its high anomaly score was mainly due to high values in run lengths, block variance, and less from colour decorrelation. Black-box machine-learning predictions are frequently used when accountability is required at a feature level for forensic reporting and use in a court of law. Our experiments also showed that the framework is sensitive to the size of the payload. For small plaintexts (a few kilobytes), the embedded message, only the most sensitive features differed significantly from

the baseline scores and the average anomaly score was occasionally lower than the detection threshold. When the payload size was expanded to comparable dimensions to the actual covert transfers in the field (hundreds of kilobytes) the ability to detect increased dramatically. The framework can, however, be extended for the other purpose, by estimating the quantity of hidden data from the magnitude of anomaly scores, which can be used in the payload length estimation.

FUTURE WORK

There are some directions that will be continued from the present setup. In addition to an anomaly-based approach, a compact supervised classifier, e.g. a linear discriminant or a light weight ensemble, will be integrated on top of the existing feature vector, capturing higher-dimensional patterns while still maintaining interpretability. Second, this framework will be expanded to JPEG domain by adding additional statistics on DCT coefficients and conformance checks based on Benford's law in order to include the most common image format used in real world forensic cases. Third, an adaptive threshold mechanism will be designed that will automatically adjust the decision boundary according to the size of the payload and the image content, and a payload length estimator will be derived from the magnitude of the anomaly scores which will approximate the number of hidden payloads hiding in the image. Fourth, a graphical user interface (GUI) will be developed to make the system accessible to those who do not have programming experience, and the system will be evaluated with a larger number of images and on other steganographic tools such as the adaptive embedding algorithm . Lastly, specific parts of the Spatial Rich Model will be integrated with the feature set, and the framework will be tested on common benchmarks like BOSSBase for side-by-side evaluation with the latest deep learning detectors.

CONCLUSION

This paper presented a multi-modal forensic steganalysis framework, which is specially designed for detecting hybrid crypto-steganographic embedding in a digital image. It extracts twenty-six features in a systematic manner and combines them in a self-calibrating Z-score anomaly detector, covering pixel, colour, noise, texture and AES-block domains, without the need for pre-trained models. The results of the experiments performed on a controlled benchmark of clean and stego images generated by OpenStego demonstrate that the proposed method has high detection rate with low false positive rate compared to the existing single domain detection methods like StegExpose and StegSpy. Sixteen out of the features used are novel in the open-source steganalysis world, and directly applicable to the encrypt-then-hide workflow. It provides a much needed capability to the forensic community by combining a transparent and explainable suspicion score, together with a fully open-source implementation which would otherwise be missing between academic steganalysis and practical digital forensic investigation.

DECLARATIONS

Acknowledgement: We appreciate the generous support from all the contributor of research and their different affiliations.
Funding: No funding body in the public, private, or nonprofit sectors provided a particular grant for this research.
Availability of data and material: In the approach, the data sources for the variables are

stated.

Authors' contributions: Each author participated equally to the creation of this work.

Conflicts of Interests: The authors declare no conflict of interest.

Consent to Participate: Yes

Consent for publication and Ethical approval: Because this study does not include human or animal data, ethical approval is not required for publication. All authors have given their consent.

REFERENCES

Abhrendu Bhattacharya, Amit R Welekar, Paramita Sarkar, Rajesh Bose, Sandip Roy, and Achyut Mitra. Adaptive bit placement for dual biomet-ric using signature and finger print for dwt-dct picture steganography. *Scientific Reports*, 2026.

Arbab Khan, Inzamam Shahzad, Assad Latif, Saira Saleem, et al. A robust and scalable security model for eye disease data based on cryptographic and steganographic integration. *Spectrum of Engineering Sciences*, 4(4):1693–1712, 2026.

Aryan Upadhyay, Sujay Pandey, Utkarsh Kumar Pandey, Priyanshu Varshney, Shelly Gupta, and Mukesh Kumar Tripathi. An efficient deep-fake image detection framework based on facial landmark localization and deep feature extraction. In *2026 5th International Conference on Sentiment Analysis and Deep Learning (ICSADL)*, pages 1155–1159. IEEE, 2026.

Bhavya Kallapu, Avinash Nanda Janardhan, Rama Moorthy Hejamadi, Krishnaraj Rao Nandikoor Shrinivas, Saritha, Raghunandan Kemmannu Ramesh, and Lubna A Gabralla. Multi-layered security framework combining steganography and dna coding. *Systems*, 13(5):341, 2025.

Dina Yousif Mikhail, Roojwan Sc Hawezi, and Shahab Wahhab Kareem. An ensemble transfer learning model for detecting stego images. *Applied Sciences*, 13(12):7021, 2023.

DS Deshkar, Sunita V Dhavale, and Rajendra S Deodhar. Explainable decision level ensemble framework for image steganalysis. In *2026 IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI)*, volume 4, pages 1–6. IEEE, 2026.

Edla Varun, Gorre Divya Jyothi, Gorla Koushik, Marka Vinay Goud, and Janagam Lokesh. Blockchain-enabled evidence integrity and trans-parency framework for law enforcement agencies. *American Journal of AI Cyber Computing Management*, 6(2):571–581, 2026.

Fanwei Zeng, Jianshu Li, Changtao Miao, Chenqi Kong, Youru Li, Zhi Cai, Haiyan Yin, Joey Tianyi Zhou, Ran He, Anderson Rocha, et al. Gentext-forensics: Challenge on explainable forensics and adversarial generation for text-centric images. 2026.

Farhad Shadmand, Iurii Medvedev, Luiz Schirmer, and Nuno Gonc¸alves. Docsafe: Toward practical print-proof image steganography via fre-quency decomposition and covariance alignment. *IEEE Access*, 14:54213–54229, 2026.

Fatima Abdulhussain Khalil, Hasanen Murtadha Hassan, and Ammar Ali Neamah. A metaheuristic-based stacking ensemble learning framework for jpeg steganalysis. In *Proceedings of the 2026 2nd International Conference on Computing and Emerging Sciences*, pages 167–174, 2026.

Huiyan Chang, Dacheng Zhou, Hongchao Hu, Tao Hu, and Quan Ren. Aamc-rl: Reinforcement learning-based automatic asymmetric cost learning steganography via multi-steganalyzer consensus. *IEEE Signal Processing Letters*, 2026.

Isamadeen A Khalifa, Salih Mustafa Saleem, and Abdulkadir Sengur. Multilayer steganalysis in the encryption era: A comprehensive review. *European Journal of Applied Science, Engineering and Technology*, 4(1):73–99, 2026. Ishika Dhiman, Ashmeet Singh, and Ravikant Kumar Nirala. Stego-scope: Blind, feature-based lsb detection with visual diagnostics. In *2026 International Conference on Intelligent and Innovative Technologies in Computing, Electrical and Electronics (IITCEE)*, pages 1–6. IEEE, 2026.

J Hemalatha, Senthil Maharaj Kennedy, M Sekar, and M Kaliraj. Curvelet denoising firewall for universal removal of steganographic image content. *Results in Engineering*, page

109837, 2026.

Jan Butora. Projecting noise residuals on a unit sphere for out-of-distribution detection of generated images. 2026.

Janine Schneider, Florian Ramming, Maximilian Eichhorn, Gaston Pugliese, Chris Hargreaves, Jan Gruber, Joschua Schilling, Julian Geus, Kevin Mayer, Lea Uhlenbrock, et al. Sok: Understanding anti-forensics concepts and research practices across forensic subdomains. *arXiv preprint arXiv:2604.05770*, 2026.

Jayaprakash Veerakumar, Hariharaviswanathan Prasad, and Anitha Muthulingam. Cross-algorithm steganalysis via dual-domain feature fusion: A hybrid deep learning approach for payload detection. *Interna-tional Journal of Bioinformatics and Intelligent Computing*, 5(1):56–66, 2026.

JG Arjay, S Sheshwat, and S Brindha. Noranet-dual adversarial network with residual attention for detecting steganographic attacks. In *2026 In-ternational Conference on Wireless Communications Signal Processing and Networking (WiSPNET)*, pages 1–6. IEEE, 2026.

Jihao Zhu, Zixuan Chen, Jiali Liu, Lingxiao Yang, Yi Zhou, Weiqi Luo, and Xiaohua Xie. Rethinking security of diffusion-based generative steganography. *arXiv preprint arXiv:2602.10219*, 2026.

Kobra Hassanzadeh, Sohrab Ahmadi-Kandjani, Reza Kheradmand, and Seyed Amir Mortazavi. Experimental and simulation orthogonal ghost imaging for measurement-domain optical steganography. *Scientific Reports*, 2026.

Konstantinos Karampidis. Image steganalysis for digital forensics. 2020.

Kristian D Michaylov and Dipti K Sarmah. Steganography and steganal-ysis for digital image enhanced forensic analysis and recommendations. *Journal of Cyber Security Technology*, 9(1):1–27, 2025.

Leila Rzayeva, Tomiris Zhumakan, Aizada Kapatayeva, Tabigat Serik, and Alisher Batkuldin. Development of a method for automatic document recovery followed by analysis of integrity and absence of encryption for forensic purposes. *Scientific Journal of Astana IT University*, 2026.

Lijing Ren and Denghui Zhang. Discoqr: Diffusion-driven semantic compression for robust image steganography in standard qr codes. In *Proceedings of the ACM Web Conference 2026*, pages 2592–2601, 2026.

Mingzheng Lv, Chen Liang, Baokun Zheng, Tianqing Zhu, Wanlei Zhou, Yu-an Tan, and Mengxia Ren. Image steganography for high-bandwidth covert communication on ethereum election contracts. *IEEE Transactions on Cognitive Communications and Networking*, 2026.

Mohammad Hossein Noorallahzadeh. Novel adaptive dct-based steganography algorithm with coefficient selection optimization for jpeg images. *Journal of Intelligent Communication*, 5(1):116–139, 2026.

Mohammed Al Saleh, Rima Shbaro, and Joseph Azar. A novel iot security framework combining x25519 with nist lightweight ascon encryption and hybrid transform-domain steganography. In *Telecom*, volume 7, page 40. MDPI, 2026.

Muhammad Kevin Rinaldi, Ernawati Ernawati, Desi Andreswari, and Ju-lia Purnama Sari. Application of two-stream late fusion on efficientnetv2 based on transfer learning to classify ai-generated paintings. *Sinkron: jurnal dan penelitian teknik informatika*, 10(2):976–990, 2026.

Romaric Audigier, Patrick Bas, Jan Butora, Emma Coletta, Christophe Charrier, Nicholas Evans, Emmanuel Giguet, Mohamed Mallat, Pierre-Alain Moellic, Benjamin Negrevergne, et al. Projet compromis wp 2: Research activities in integrity. 2026.

Saad M Ismail, Feras E AbuAladas, Mamoun Abu Helou, and Waheeb Abu-ulbeh. Edge-adaptive high-capacity image steganography using hybrid edge detection and msb embedding. *Computers*, 15(3):141, 2026.

Sami Alanazi and Seemal Asif. Vids-guard: A novel forensics-aware multi-stream transformer framework for robust deepfake video detection. *Intelligent Systems with Applications*, page 200664, 2026.

Sumia Abdulhussien Razooqi Al-obaidi, Mohammed Ahmed Talab, Mohanaad Shakir, Mustafa

Talal Alnaseri, and Suryanti Awang. En-hanced image steganalysis through reinforcement learning, generative adversarial networks, and vision transformers (rgv-stega). *Journal of University of Anbar for Pure Science*, 20(1):312–321, 2026.

SV Bezzateev and M Yu Fedosenko. Analysis of the problems of using steganographic methods in implementing illegal actions and their role in digital forensics. *Automatic Control and Computer Sciences*, 58(8):1406–1421, 2024.

Tong Qiao. *Statistical detection for digital image forensics*. PhD thesis, Universite´ de Technologie de Troyes, 2016.

Xiang Ao, Yilin Du, Zidan Wang, Mengru Chen, and Siyang Lu. Awpd: Frequency shield network for agnostic watermark presence detection. *arXiv preprint arXiv:2603.06723*, 2026.

Xinlei Guan, David Arosemena, Tejaswi Dhandu, Kuan Huang, Meng Xu, Miles Q Li, Bingyu Shen, Ruiyang Qin, Umamaheswara Rao Tida, and Boyang Li. Toward accountable ai-generated content on social platforms: Steganographic attribution and multimodal harm detection. *arXiv preprint arXiv:2604.10460*, 2026.

Xiyan Bi, Zichi Wang, Xue Wang, Bo Tian, and Xinpeng Zhang. Cover selection method for jpeg steganography. *IEEE Transactions on Circuits and Systems for Video Technology*, 2026.

Ali Shan, Muhammad Ahmad, Hassaan Iqbal, and Ali Sufyan. Ztforensics: Zero trust policy enforcement with tamper-evident forensic evidence packaging for hybrid cloud environments. Spectrum of Engineering Sciences, 4(5):795–810, 2026