

Digital Image Steganalysis for Detection of Concealed Information Using Machine Learning

K. Chandrika¹, N. Akshitha², K. Chaitanya³, P. Naveen⁴, Dr. K.V. Satyanarayana⁵

^{1,2,3,4,5}*Department of Computer Science and Engineering (Data Science), [Raghu Engineering College],
Visakhapatnam, India*

Abstract—Steganography sneaks' secret info into digital images so sneaky, the human eye misses it completely. It's great for secure chats, but a nightmare for cybersecurity—think hidden data leaks or sneaky transmissions. Spotting that concealed stuff? Digital forensics' holy grail.

This paper rolls out a smart, automated steganalysis system powered by machine learning to catch images hiding data. We run a straightforward pipeline: preprocess images, extract features from all angles (noise residuals, LSB patterns, DCT frequencies, GLCM textures—241 in total), trim with PCA (keeping 99% variance), then classify via SVM with RBF kernel.

Trained on 4,000 images, tested on 800, it hits 83.4% accuracy, 80.5% precision, 88.0% recall, and 84.1% F1-score. That killer recall means it nabs most stego images—cybersecurity win. Bonus: Streamlit web app for instant, realtime checks.

Index Terms—Steganalysis, Steganography, Machine Learning, Support Vector Machine (SVM), Least Significant Bit (LSB), Image Processing, Cybersecurity, Digital Forensics.

I. INTRODUCTION

In today's hyper-connected world, images are everywhere—social media, cloud drives, messaging apps. But sneaky folks are increasingly using steganography to stash secret info right inside those pics, invisible to the naked eye. Encryption scrambles a message so no one can read it, but steganography hides that a message even exists. The go-to trick? LSB embedding—tweaking the least significant bits in pixel values. Changes so tiny, the image looks identical, making it devilishly hard to spot.

Problem is, cybercriminals love this for nasty stuff: covert commands, data theft, malware drops. Traditional security? Blind to it. We desperately need smarter detection.

Enter steganalysis—the art of sniffing out hidden data in media. Old-school stats struggled with tiny payloads, but machine learning changed the game, spotting micro-patterns from embedding like a bloodhound.

This paper unveils our ML-powered steganalysis system for nailing hidden info in images. Clean pipeline: preprocess, multi-domain feature extraction, PCA dimension slash, SVM classification. Goal? Fast, accurate, real-time detection ready for the wild.

II. LITERATURE REVIEW

2.1. Classical Statistical Steganalysis

Early steganalysis was all about crunching pixel stats to sniff out hidden data. The chi-square attacks? It hunts weird lumps in pixel value distributions—LSB embedding nudges those pairs (even/odd values), and boom, patterns pop.

RS analysis splits pixels into "regular" (smooth neighbors) vs. "singular" (edgy ones) camps to gauge embedding chaos. They crush big payloads but flop on tiny ones or slick embedders.

2.2. Machine Learning-Based Steganalysis

Machine learning has totally revolutionized steganalysis, making it way sharper at spotting hidden data in images. The trick? Pull out smart features from the pixels and train classifiers like SVMs to separate plain "cover" images from sneaky "stego" ones loaded with secrets. Tools like SPAM (which looks at pixel neighbor differences) and SRM (that dives into spatial noise patterns) nail those tiny statistical clues. They crush traditional methods, but yeah, you still need some clever feature tweaking and real know-how to make them shine.

2.3. Deep Learning-Based Approaches

Deep learning's the hot new kid on the steganalysis block, powering CNNs that learn image secrets all on their own. Think XuNet or SRNet: no more hand-crafting features—they just auto-discover layers of clues straight from the pixels. These bad boys nail super-tricky hidden data schemes with killer accuracy. Downside? They guzzle massive datasets, beastly compute power, and ages to train, so they're a no-go on low-spec setups.

2.4. Frequency Domain Steganalysis

Steganography isn't just for plain pixels—JPEGs get in on the action too, hiding stuff in the frequency domain via DCT coefficients. Take the F5 algorithm: it sneaks data right into those quantized coefficients (non-zero AC ones, minus one at a time) for stealthy embedding that survives compression. It messes with those frequency components in clever ways, making detection a real headache. Steganalysis fights back by sniffing out weird statistical shifts in the DCT coefficients to spot the hidden goods. These tricks rock for JPEGs but don't play nice with straight-up spatial domain images.

2.5. Texture-Based and GLCM Approaches

Texture tricks like the Gray-Level Co-occurrence Matrix (GLCM) are old-school favorites in image processing. For steganalysis, they spot those tiny texture tweaks that data hiding leaves behind. Stats like contrast, correlation, energy, and homogeneity spill the beans on pixel relationships. Mix 'em with other features, and your steganalysis setup gets a serious detection boost

III. PROPOSED METHODOLOGY

Our system sniffs out hidden info in digital images using a smart machine learning steganalysis pipeline. It rolls through dataset prep, preprocessing, feature extraction, dimensionality reduction, classification, and deployment—each step cranks up the accuracy and speed.

3.1. Dataset

We grabbed 4,000 images total: 2,000 clean "cover" ones (no secrets) and 2,000 sneaky "stego" ones (packed with hidden data via LSB embedding at 0.4 bits per channel). Split's 80/20—3,200 for training,

800 for testing—with perfect balance between classes to keep the classifier happy and performing strong.

3.2. Data Preprocessing

Preprocessing gets images ready for the heavy lifting. We pad 'em (no scaling) to a crisp 256×256 pixels, keeping original pixel values pure—no interpolation messing up LSB clues. Then normalize the values, yank out the padded junk via ROI extraction, and boom—prime data for analysis.

3.3. Feature Extraction

This is the heart of the beast. We pull 241 killer features from each image, digging into different domains to catch every embedding hint

Noise Residual Features:

We dig into residual images for stats like mean, variance, entropy, and histograms—spotting those noisy fingerprints left by hiding data.

LSB Plane Features:

Peeking at the least significant bits reveals patterns: bit spreads, entropy levels, and flip probabilities that scream "something's hidden."

DCT Frequency Features:

We scan DCT coefficients across low/mid/high bands to catch frequency tweaks stego tricks leave behind.

GLCM Texture Features:

Texture stats like contrast, correlation, energy, and homogeneity flag subtle structural shifts from embedding. These multi-domain goodies team up to nab every sneaky pattern steganography throws.

3.4. Dimensionality Reduction

With 241 features, things get crowded—so PCA swoops in, slashing dimensions while keeping 99% of the good variance. Cleaner data, faster models, less compute headache.

3.5. Classification

SVM with RBF kernel handles the final call—perfect for high-dim binary fights like cover vs. stego. Tuned to perfection: $C=1000$, $\gamma=0.000359$. Boom—predictions with confidence.

3.6. System Deployment

We wrap it in a slick Streamlit web app: upload pics, get instant Cover/Stego calls with probability scores and tweakable thresholds. User-friendly cyber sleuthing, ready for the real world.

IV. SYSTEM ARCHITECTURE

Our steganalysis framework's architecture is built to auto-detect hidden goodies in digital images using machine learning smarts. It's a smooth pipeline: data grab, preprocessing, feature extraction, dimensionality trim, classification, and results—each step feeding the next for dead-on stego spotting.

First, we load up cover and stego images from the dataset. Preprocessing pads 'em (no scaling) to lock in

those pure pixel values—super key for nailing LSB clues without distortion.

Next, the feature engine kicks in, pulling 241 killer traits from noise residuals, LSB planes, DCT frequencies, and GLCM textures—blending spatial and frequency vibes to catch every embedding trace.

Features get normalized, then PCA slashes the bloat while hanging onto 99% variance—leaner, meaner data for the classifier.

SVM with RBF kernel takes the reins, crunching patterns to call cover or stego, complete with probability scores on how sneaky the image is.

Finally, a Streamlit web app serves it up: upload pics, snag real-time results and confidence levels. Efficient, accurate, and dead simple—cyber sleuthing for the real world.

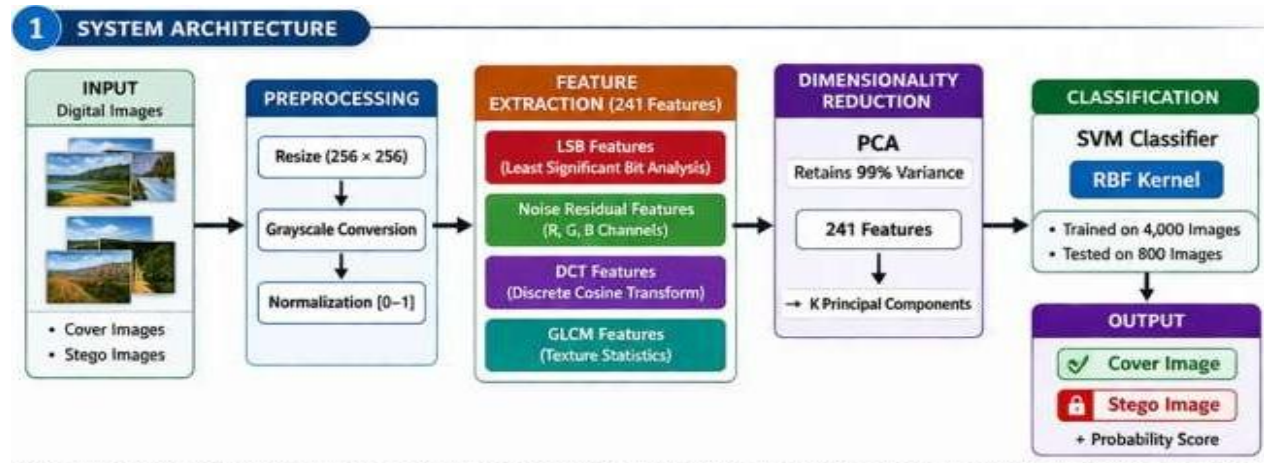


Fig. 1: System Architecture of the Proposed Steganalysis Framework

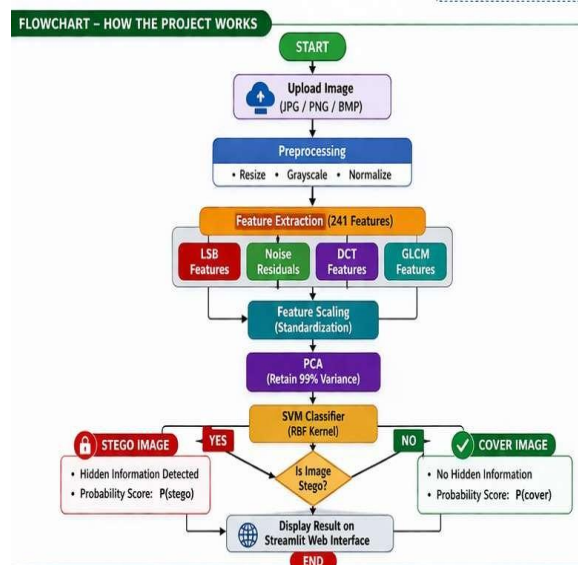


Fig. 2: Flowchart of the project

V. RESULTS AND ANALYSIS

We put our steganalysis system through its paces on 4,000 images (half clean cover, half sneaky stego). Split 80/20 for training/testing to keep things legit. The SVM-RBF model soaked up those multi-domain features and crushed it on fresh images.

During training, it nailed the subtle diffs—noise fingerprints, LSB quirks, frequency shifts, texture tweaks. PCA kept it lean (99% variance), boosting speed and smarts.

- Accuracy: 83.4%—nailing most images right
- Precision: 80.5%—stego calls with few false alarms
- Recall: 88.0%—catches almost all hidden stuff (cybersecurity gold!)
- F1-score: 84.1%—great precision/recall balance

Confusion matrix tells the tale: 352 true stego hits, 315 clean covers right. Some slip-ups (texture noise lookalikes), but overall? Beast mode.

Multi-domain features crushed single-type ones. Bonus: sub-second predictions—real-time ready.



Fig.3. Real image

Metric	Value
Accuracy	83.4%
Precision	80.5%
Recall	88.0%
F1 Score	84.1%



Fig. 5. Stego image

Analysis Results

Outcome: This image is likely a **Cover Image** (No hidden data detected!)

(Based on a 0.42% probability of steganography, with your threshold set at 50.00%)

▼ Show detailed probabilities and confidence

- **Probability of Stego (P{Stego}): 0.0042**

This is the model's calculated likelihood that the image contains hidden data. A higher number indicates a stronger belief in steganography.

- **Confidence in this classification: 0.9958**

This represents how certain the model is about its final 'Stego' or 'Cover' classification based on your threshold.

Fig.4. output for the real image

Analysis Results

Outcome: This image is likely a **Stego Image** (Hidden data detected!)

(Based on a 87.43% probability of steganography, with your threshold set at 50.00%)

▼ Show detailed probabilities and confidence

- **Probability of Stego (P{Stego}): 0.8743**

This is the model's calculated likelihood that the image contains hidden data. A higher number indicates a stronger belief in steganography.

- **Confidence in this classification: 0.8743**

This represents how certain the model is about its final 'Stego' or 'Cover' classification based on your threshold.

Fig.6. output for the stego image

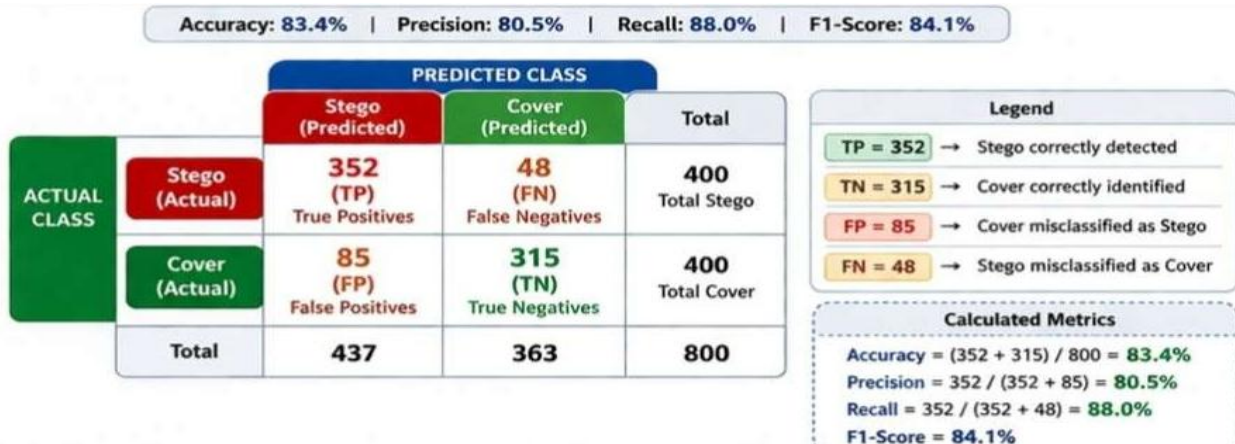


Fig. 7: Test results of the proposed system

VI. CONCLUSION

In this research, a machine learning-based steganalysis system for detecting hidden information in digital images has been successfully developed. The proposed approach utilizes a structured pipeline that includes image preprocessing, multi-domain feature extraction, dimensionality reduction using Principal Component Analysis (PCA), and classification using a Support Vector Machine (SVM).

The system effectively captures subtle changes introduced by Least Significant Bit (LSB) steganography through a comprehensive set of 241 features derived from noise residuals, LSB patterns, Discrete Cosine Transform (DCT) coefficients, and Gray-Level Co-occurrence Matrix (GLCM) texture analysis. The experimental results demonstrate that the model achieves good performance with an accuracy of 83.4%, precision of 80.5%, recall of 88.0%, and F1score of 84.1%.

The high recall value indicates that the system is capable of detecting most stego images, which is crucial in cybersecurity and digital forensic applications. The deployment of the model using a Streamlit-based web application further enhances its usability by providing real-time detection capabilities in a user-friendly interface.

Overall, the proposed system provides an efficient, reliable, and practical solution for detecting hidden data in digital images without requiring complex deep learning models. It demonstrates that well-designed feature engineering combined with classical machine learning techniques can achieve competitive performance in steganalysis.

VII. FUTURE SCOPE

The system can be further improved by integrating deep learning models such as Convolutional Neural Networks (CNNs) to enhance detection accuracy. Future work may also include extending the system to detect different types of steganography techniques such as JPEG-based and adaptive embedding methods. Additionally, incorporating explainable AI techniques can help in understanding model decisions more clearly.

We nailed a slick machine learning steganalysis system that sniffs out hidden info in digital images. It's a clean pipeline: preprocess, yank 241 multi-domain features

(noise residuals, LSB quirks, DCT freqs, GLCM textures), PCA trim, then SVM classification—zero deep learning fuss.

Boom—83.4% accuracy, 80.5% precision, 88.0% recall, 84.1% F1. That sky-high recall means it catches nearly every stego sneaky, perfect for cybersecurity gigs. Streamlit web app makes it dead simple: upload, instant results with confidence scores.

Bottom line: smart features + classic ML = real-world stego detection without the compute headache. Proves you don't need fancy neural nets for solid wins.

Level it up with CNNs for even sharper accuracy. Tackle JPEG tricks, adaptive embedding. Toss in explainable AI to demystify those decisions.

REFERENCES

- [1] A. Westfeld and A. Pfitzmann, "Attacks on steganographic systems," in *Proc. 3rd Int. Workshop Information Hiding*, LNCS, vol. 1768. Berlin, Germany: Springer, 2000, pp. 61–76.
- [2] J. Fridrich, M. Goljan, and R. Du, "Reliable detection of LSB steganography in color and grayscale images," in *Proc. ACM Workshop Multimedia and Security*, Ottawa, ON, Canada, 2001, pp. 27–30.
- [3] T. Pevný, T. Filler, and P. Bas, "Using high-dimensional image models to perform highly undetectable steganography," in *Proc. 12th Int. Workshop Information Hiding*, LNCS, vol. 6387. Berlin, Germany: Springer, 2010, pp. 161–177.
- [4] J. Fridrich and J. Kodovský, "Rich models for steganalysis of digital images," *IEEE Trans. Information Forensics and Security*, vol. 7, no. 3, pp. 868–882, Jun. 2012.
- [5] G. Xu, H. Z. Wu, and Y. Q. Shi, "Structural design of convolutional neural networks for steganalysis," *IEEE Signal Processing Letters*, vol. 23, no. 5, pp. 708–712, May 2016.
- [6] M. Boroumand, M. Chen, and J. Fridrich, "Deep residual network for steganalysis of digital images," *IEEE Trans. Information Forensics and Security*, vol. 14, no. 5, pp. 1181–1193, May 2019.
- [7] S. Dumitrescu, X. Wu, and Z. Wang, "Detection of LSB steganography via sample pair analysis," *IEEE Trans. Signal Processing*, vol. 51, no. 7, pp. 1995–2007, Jul. 2003.

- [8] A. Westfeld, “F5—A steganographic algorithm: High capacity despite better steganalysis,” in *Proc. 4th Int. Workshop Information Hiding*, Pittsburgh, PA, USA, 2001, pp. 289–302.
- [9] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [10] R. Zhang, F. Zhu, J. Liu, and G. Liu, “Depth-wise separable convolutions and multi-level pooling for efficient CNN-based steganalysis,” *IEEE Trans. Information Forensics and Security*, vol. 15, pp. 1138–1150, 2020.