GHOSTFREAK: ENHANCED STEGANOGRAPHY IN PRINT-SCAN ENVIRONMENTS USING DEEP LEARNING

S. KOMAL KOUR[1], SOHAN GUNDOJU[2], SREE VAIBHAV DUVVURI[3], DR. T ADILAKSHMI[4], T JALAJA[5]

komalkour@staff.vce.ac.in[1], gundojusohan@gmail.com[3], sreevaibhav345@gmail.com[2], hodcse@staff.vce.ac.in[4], jalaja.t@staff.vce.ac.in[5]

[1,2,3,4,5] Department of Computer Science and Engineering, Vasavi College of Engineering, Hyderabad, India

ABSTRACT

GhostFreak, a novel deep steganography framework designed to address the challenges inherent in print-scan pipelines. Traditional steganographic approaches often struggle with the distortions introduced during the printing and scanning process, leading to significant degradation in the quality and recoverability of embedded data. To overcome these limitations, GhostFreak strategically leverages three distinct color spaces—RGB for digital displays, HSI for human visual perception, and CMYK for printing—to achieve robust and imperceptible data embedding while maintaining resilience against real-world distortions. Framework extends prior research on print-scan steganography by integrating a U-Net GAN architecture, which processes a concatenated tensor of multi-color representations alongside a secret code to generate a residual image. This approach ensures that the hidden information remains intact while preserving the quality of the Stego image. A key innovation of GhostFreak is its dynamically weighted loss function, which balances multiple loss components—secret loss, secret decay loss, LPIPS (perceptual similarity) loss, residual loss, and edge loss—to optimize the model's performance across different training phases. During initial training stages, the model prioritizes accurate encoding of the secret message, whereas in later stages, the emphasis shifts towards ensuring high-fidelity image reconstruction. We validate the effectiveness of Ghost Freak through comprehensive experiments on a large-scale image dataset, demonstrating that our method withstands the degradations introduced by print-scan operations. Comparative evaluations against prior steganographic approaches highlight significant improvements in terms of imperceptibility, robustness, and recoverability of the embedded data. The results indicate that GhostFreak is a promising advancement in deep steganography, offering a practical and resilient solution for secure data embedding in printed media.

Keywords: *Image Steganography, Deep Learning, Print-Scan, Color Spaces, U-Net GAN, Dynamic Loss Weighting, Error Correction.*

1. INTRODUCTION

Steganography is the science of hiding information within innocuous-looking media, primarily for secure communication. Over the years, this field has evolved from rudimentary practices such as microdots and invisible ink to advanced digital methods. In particular, image steganography has gained prominence due to the ubiquity of visual media and its capacity to carry embedded data without raising suspicion. Modern advances, especially those leveraging deep learning, have enabled embedding schemes that are both imperceptible to the human eye and resilient to a range of transformations including compression, resizing, and more recently, the print-scan cycle.

The goal of this study is to address a fundamental challenge in the field: maintaining the integrity of hidden information through real-world degradations introduced during the print-scan process, which many conventional steganographic models are unable to withstand. While most methods are tailored for digital-only media, they often fail when confronted with analog distortions such as ink bleeding, lighting variations, and scanning noise.

To overcome this, we propose GhostFreak, a novel deep steganographic framework that combines multiple color space representations (RGB, HSI, CMYK) and a U-Net-based GAN architecture to robustly embed and recover hidden data even after physical transformations. This approach is further strengthened by a dynamically weighted loss function that evolves over the course of training, balancing perceptual similarity, robustness, and fidelity. This work is primarily focused on short-message embedding (up to 100 bits) into natural images intended for printing and scanning. We assume that:

- The encoder has access to RGB images and can compute their HSI and CMYK representations.
- The decoder receives a digitally scanned image post-print, possibly affected by lighting noise, warping, and compression.
- The goal is not only to embed messages invisibly but to maintain message recoverability under real-world distortions.

The key novel contributions of this work include:

- A multi-color space encoding strategy that exploits perceptual and print-specific advantages of RGB, HSI, and CMYK.
- A dynamically weighted loss function that shifts emphasis from secret fidelity to visual quality throughout training.
- A robust print-scan resilient steganographic model, validated against leading methods such as StegaStamp.

By introducing a system that performs reliably across both digital and physical domains, GhostFreak sets a new benchmark for hybrid-media steganography.

2. RELATED WORK

Image steganography has progressed significantly, transitioning from elementary pixel-level manipulation techniques to sophisticated deep learning-driven frameworks. Below is a critical examination of key existing approaches categorized by their methodology, with an emphasis on their Plus, Minus, and Interesting (PMI) attributes.

2.1. Least Significant Bit (LSB) Method

LSB-based techniques are intuitive, computationally lightweight, and easy to implement. They manipulate the least significant bits of pixel values, ensuring minimal visual distortion [8].

Despite their simplicity, these methods are highly susceptible to image processing operations such as compression, scaling, and cropping. They also lack resistance to statistical steganalysis [6].

LSB remains a foundational technique for educational purposes and has been used as a baseline in comparative studies due to its deterministic behavior [9].

2.2. DCT-Based Methods

Discrete Cosine Transform (DCT)-based approaches operate in the frequency domain, making them more robust to compression and other lossy operations such as JPEG encoding [8].

These methods often compromise on payload capacity and require careful coefficient selection to maintain invisibility.

Techniques like DCT form the basis of widely-used image formats (e.g., JPEG), making them compatible with many real-world applications despite their limited embedding capacity.

2.3. Deep Learning (DL) Approaches:

DL-based methods, such as HiDDeN [10] and StegaStamp [1], have revolutionized the field by learning optimal embedding strategies from data, offering superior robustness and imperceptibility. Adversarial training setups further enhance their resilience against steganalysis [12].

Most of these models are trained and evaluated under digital-only distortions (e.g., cropping, JPEG compression), making them inadequate for scenarios involving analog noise such as print-scan degradation.

These approaches often use auxiliary loss functions (e.g., perceptual similarity metrics like LPIPS [12]) to ensure that encoded images remain visually plausible.

2.4. Print-Scan Method

A niche but growing area, print-scan steganography addresses physical media degradation, a challenge largely ignored by conventional techniques. For example, Light Field Messaging (LFM) incorporates camera and lighting distortions during training to simulate real-world capture scenarios [11].

LFM's approach requires extensive datasets of manually captured images, making it time-intensive and resource-heavy. Moreover, most models in this category fail to generalize well to variations in print media quality or scanning conditions.The use of physical transformations as part of the training loop represents a significant shift in steganographic design philosophy and sets the stage for hybrid-media embedding frameworks.

2.5. Summary and Research Gap

While numerous techniques aim to improve embedding capacity, visual quality, or resilience to digital attacks, few address the full print-scan pipeline with practical usability. Techniques like StegaStamp [1] offer partial robustness but leave noticeable perturbations in printed images. Others like LFM [11] improve robustness but lack scalability and efficiency.

GhostFreak addresses this gap by:

- Integrating multiple color spaces (RGB, HSI, CMYK) for enhanced perceptual alignment and print compatibility.
- Employing a dynamically weighted loss function that evolves during training to balance fidelity and robustness.
- Outperforming existing methods in both digital metrics and real-world (printed media) evaluations.

These contributions position GhostFreak as a practical and scalable solution for robust, high-fidelity steganography in hybrid media settings.

3. PROPOSED METHADOLOGY

Despite significant progress in deep learning-based image steganography, robust message embedding that survives the print-scan pipeline remains a critical and underexplored challenge. Most existing methods are optimized for digital distortions such as compression or cropping but fail when faced with analog distortions introduced by printing and scanning—such as color shifts, warping, uneven lighting, and scanning noise. These issues significantly degrade both the visual quality of stego images and the recoverability of hidden data. Moreover, conventional steganographic frameworks rely heavily on a single color space, typically RGB, which is ideal for digital displays but poorly aligned with human perception or print reproduction. These methods also employ static loss functions, leading to a rigid optimization process that fails to dynamically balance robustness and imperceptibility during training.

To address these limitations, we introduce GhostFreak, a novel steganographic system that combines multi-color space encoding, a U-Net-based encoder-decoder architecture, and a dynamically weighted composite loss function. This integration enables reliable message embedding across both digital and physical domains.

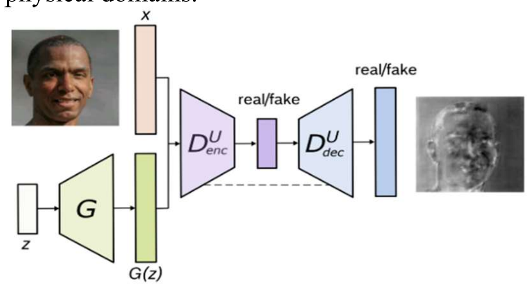

Figure 1. Architecture Of Unet GAN Model

3.1. Research Objectives

This study is guided by the following core objectives:

- O1: Design a deep learning-based steganographic model that performs reliably in both digital and physical media, with special attention to print-scan robustness.
- O2: Incorporate multi-color space encoding using RGB, HSI, and CMYK to enhance perceptual fidelity and media adaptability [3], [4], [5].
- O3: Introduce a dynamic loss-weighting strategy that adapts during training to balance secret recovery accuracy and visual imperceptibility [12].
- O4: Evaluate GhostFreak against state-of-the-art approaches such as StegaStamp [1], using both quantitative metrics (PSNR, SSIM, LPIPS) and qualitative performance (artifact visibility and decoding after print-scan).

3.3. Research Hypothesis

The study investigates the following hypotheses:

- H1: Multi-color space training improves visual similarity and decoding robustness compared to single-space models.
- H2: Dynamic loss weighting yields better trade-offs between imperceptibility and robustness than static approaches.

H3: GhostFreak will show lower visible artifacts and comparable or superior decoding accuracy in printed images compared to existing methods.

3.4. Architecture Overview

GhostFreak has two components, the Image Generator and the Discriminator. The Image Generator comprises encoder-decoder architecture.

We refer to these encoder decoder components as the GhostFreakEncoder and the GhostFreakDecoder.

- GhostFreakEncoder: The GhostFreakEncoder leverages a U-Net architecture, enabling the extraction of higher-order features and the generation of the required residual output, which matches the size of the input image (I_{RGB}). The GhostFreakEncoder receives a concatenated input ($I_{combined}$) formed by the RGB image, its HSI and CMYK representations, and a randomly generated secret tensor (S). It outputs a residual image (R) such that the encoded image (I_{enc}) is given by:

$$I_{combined} = I_{RGB} \oplus I_{HSI} \oplus I_{CMYK} \oplus S$$
$$R = GhostFreakEncoder(I_{combined})$$
$$I_{enc} = I_{RGB} + R$$

Note : ($\oplus$) represents concatenation and (+) represents element-wise addition

- GhostFreakDecoder: The GhostFreakDecoder processes a transformed version of (I_{enc}) (subjected to simulated print-scan distortions such as warping and uneven lighting) to recover the hidden secret (S).

The training process begins with the GhostFreakEncoder. First, a random RGB image is selected from the dataset. Using the formulas described earlier, an input ($I_{combined}$) is created and fed into the GhostFreakEncoder, which outputs a residual (R). Next, the encoded image (I_{enc}) is generated using the same formulas. This encoded image (I_{enc}) is then modified to introduce print-scan distortions, such as JPEG compression, warping, uneven lighting, and blurring. The modified encoded image ($\widehat{I_{enc}}$) is subsequently passed to the GhostFreakDecoder, which produces a secret tensor ($\hat{S}$). The Total Loss ($L_{total}(t)$), as detailed in Section 3.3, is computed and backpropagated. Simultaneously, the Discriminator attempts to distinguish between the original RGB image (I_{RGB}) and the encoded image (I_{enc}); a separate loss for the Discriminator is computed and backpropagated as well. This process is repeated for multiple iterations until the Total Loss ($L_{total}(t)$) converges. Figure 2. Represents the Training Procedure of the Ghost Freak Model.

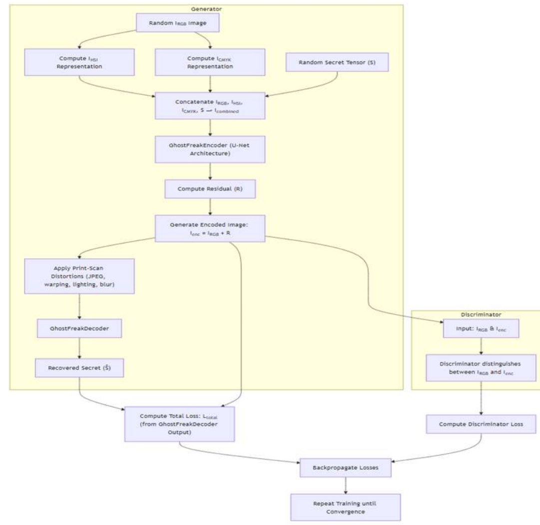

Figure 2: Overview Of Training Process

3.5. Multi-Color Space Encoding

GhostFreak leverages the complementary properties of three fundamental color spaces to ensure that our encoded images maintain high visual fidelity across both digital and print media. Rich Color Space (RGB): As baseline image data capturing image color essential elements. By separating bright channels from chromatic components, the Hue, Saturation, Intensity (HSI) color space is closer to human visual perception, and also allows applications to adjust perceptually meaningful aspects. You are being use to outputs in a single color but then finally, CMYK colorspace must be used to reproduce graphics on print outs with precision control and amount of ink for harmonious colors.

To convert RGB into HSI color space, we use the following conversion formulas:

$$I = \frac{R + G + B}{3}$$
$$S = 1 - \frac{min\,(R, G, B)}{I}$$
$$H = \left(\frac{1}{2}\frac{(R - G) + (R - B)}{\sqrt{(R - G)^2 + (R - G)(G - B)}}\right)$$

Note: If $I = 0$, *then* $S = 0$ *and if* $B > G$ *, then* $H = 360° - H$.

To convert RGB to CMYK color space, we first normalize the RGB values, where the RGB values are in the range [0,1] and the use the below conversion formulas:

$$K = 1 - max(R, G, B)$$

$$C = \frac{1-R-K}{1-K}$$
$$M = \frac{1-G-K}{1-K}$$
$$Y = \frac{1-B-K}{1-K}$$

Note: If $K = 1$, *then* $C = 0,\ M = 0,\ Y = 0.$

3.6. Loss Functions and Dynamic Weighting

Our training objective is formulated as a weighted sum of several loss components that each target a specific aspect of the steganographic process. This dynamic loss formulation is essential to balance the competing goals of effective secret embedding, maintaining visual fidelity, and ensuring robust decoding despite print-scan distortions. The total loss at time(t) is given by the below formula where each (λ) is dynamically adjusted over training:

$$L_{total}(t) = \lambda_{secret}(t)L_{secret} + \lambda_{secret-decay}(t)L_{secret-decay} + \lambda_{LPIPS}(t)L_{LPIPS} + \lambda_{residual}(t)L_{residual} + \lambda_{edge}(t)L_{edge}$$

3.6.1. Secret Loss(L_{secret})

The secret loss is the primary objective during the initial training phase. It is computed as the cross-entropy between the true secret (S) and the predicted secret $(\hat{S})$:

$$L_{secret} = -\sum_{j} S_j \hat{S}_j$$

This loss ensures that the encoder-decoder pair focuses on embedding and accurately recovering the secret message. However, relying solely on this loss can lead the model to exploit shortcuts, which brings us to the concept of unwanted residuals.

3.6.2. Secret Decay Loss$(L_{secret-decay})$

In early training, the model may converge to a suboptimal solution where the encoder produces a residual that, while sufficient to reduce the secret loss, does so by embedding patterns that are not robust or imperceptible. These patterns are termed *unwanted residuals*—they represent encoding artifacts that can be easily detected or may interfere with the visual quality of the image (as shown in Figure 3, Section 9).

To counteract this, we introduce the secret decay loss, which penalizes the model for generating residuals that resemble these unwanted patterns. This is achieved by computing an exponential decay based on the LPIPS [12] distance between the produced residual (R) and a reference $(R_{unwanted})$:

$$L_{secret-decay} = -k \cdot exp\,(LPIPS(R,\ R_{unwanted}))$$

where (k) is a scaling constant. This loss forces the network to abandon shortcuts that only satisfy the secret loss, encouraging a more robust and visually imperceptible embedding.

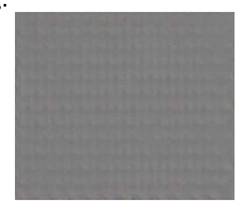

Figure 3 Example of an unwanted residual

3.6.3. LPIPS Loss(L_{LPIPS})

Also, it is important to keep the visual fidelity of the encoded image. For this we use the perceptually based metric LPIPS (Learned Perceptual Image Patch Similarity) [12], which measures perceptual similarity between images. We compute the LPIPS loss as the average of three contributions corresponding to different color spaces in our framework:

$$L_{LPIPS} = \frac{1}{3}\left[LPIPS\left(I_{RGB},\ I_{enc,RGB}\right)\right] + \frac{1}{3}\left[LPIPS\left(I_{HSI},\ I_{enc,HSI}\right)\right] + \frac{1}{3}\left[LPIPS(I_{CMYK},\ I_{enc,CMYK})\right]$$

- $LPIPS\left(I_{RGB},\ I_{enc,RGB}\right)$measures the perceptual difference in the RGB domain.
- $LPIPS\left(I_{HSI},\ I_{enc,HSI}\right)$ensures that the human perception-aligned HSI representation remains close to the original.
- $LPIPS(I_{CMYK},\ I_{enc,CMYK})$ guarantees that the print-relevant CMYK features are preserved.

This multi-space averaging is critical to maintain image quality across all media, ensuring that the alterations remain imperceptible to both human observers and printing devices.

Journal of Theoretical and Applied Information Technology
31st July 2025. Vol.103. No.14
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

3.6.4. Residual Loss($L_{residual}$)

A potential trivial solution for the encoder is to generate a near-zero residual (R),which minimizes the LPIPS loss without effectively embedding any information. To prevent this, we include a residual loss defined as:

$$L_{residual} = exp\,(-\beta \cdot L_{LPIPS})$$

with βas a scaling parameter. This loss discourages the network from taking shortcuts by ensuring that the residual carries meaningful encoded information rather than converging to a zero tensor.

3.6.5. Edge Loss(L_{edge})

Human vision is particularly sensitive to distortions in smooth regions. The edge loss is designed to guide the model to embed data preferentially in textured areas, where alterations are less perceptible. This is achieved by generating an edge mask (G) via a Canny filter and subsequent Gaussian blurring:

$$E = Canny(I, \sigma)$$

$$G = GaussianBlur(E, kernel\ size, \sigma)$$

and then computing:

$$L_{edge} = \frac{1}{N}\sum_{i=1}^{N} [(1 - G_i)R_i]^2$$

By weighting the residual loss more heavily in smooth regions (where G is low), this loss ensures that the hidden information is embedded in parts of the image where it will remain visually inconspicuous.

3.6.6. Dynamic Weighting Schedule

The dynamic weighting of these loss components is a crucial innovation. During the initial 10% of the training steps ($t < 0.1T$), the focus is solely on minimizing the secret and secret decay losses to guarantee that the secret is being accurately embedded. As training progresses ($t \geq 0.1T$), the weights for LPIPS, residual, and edge losses are gradually increased. This shift directs the model to refine the encoding, enhancing visual fidelity and suppressing unwanted residuals while still maintaining decoding accuracy.

This carefully orchestrated balance ensures that GhostFreak not only embeds the secret robustly but also produces an encoded image that is nearly indistinguishable from the original—a critical requirement for real-world print-scan applications.

4. EXPERIMENTAL SETUP

4.1. Dataset and Training Protocol

We utilized the MIRFLICKR-1M [2] dataset for its diverse real-world imagery. The model was trained for 140,000 iterations on images resized to 400×400 pixels. Each image carried a secret payload of 7 characters (56 bits) or 100 bits when incorporating BCH-based error correction [13]. Data augmentation simulated print-scan distortions (e.g., warping, blurring, uneven lighting) to ensure robust performance. Figure 4- represents the ways in which the input image is split into different models when passed to GhostFreakEncoder Architecture. The images shown are symbolic representations intended for conceptual understanding and do not depict actual data. Figure 3 represents the output of an unwanted residual.

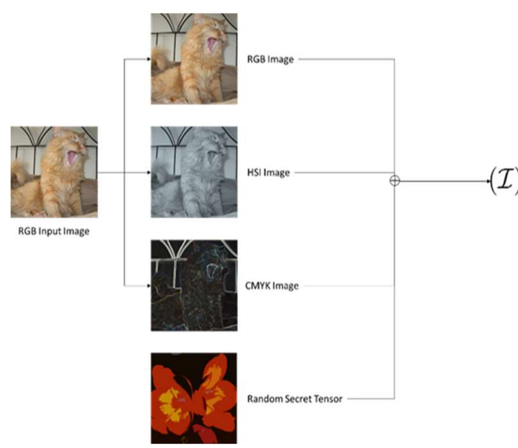

Figure 4: Input Tensor ($I_{combined}$) to GhostFreakEncoder Architecture.

The images shown are symbolic representations intended for conceptual understanding and do not depict actual data.

4.2. Evaluation Metrics

To quantitatively assess performance, we employed the following metrics:

- PSNR (Peak Signal-to-Noise Ratio) [14]:

$$PSNR = 10 \cdot (\frac{MAX^2}{MSE})$$

where MAX is the maximum pixel value and MSE is the mean squared error.

- SSIM (Structural Similarity Index Measure) [15]:

$$SSIM(x,y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}$$

which evaluates perceptual similarity between images.

- LPIPS (Learned Perceptual Image Patch Similarity):
As defined earlier, LPIPS quantifies the perceptual difference between the original and encoded images.

5. RESULTS

We compared GhostFreak against the established StegaStamp method on a standard test set of 400×400 images with a 100-bit payload (including BCH Error Correcting Codes). The performance of our model is summarized in Table 1.

Table 1: Comparison Of PSNR, SSIM, And LPIPS Metrics Between Stegastamp And Ghostfreak, Averaged Over 600 Images.

Metric	**StegaStamp**	**GhostFreak**	**Percentage Difference**
PSNR (dB) (↑)	28.50	**30.88**	**+ 8.35%**
SSIM (↑)	**0.905**	0.898	- 1.32%
LPIPS (↓)	0.101	**0.098**	**+ 2.97%**

While the numerical metrics are similar between the two models, the real advantage of GhostFreak is observed in practical applications. In printed form, GhostFreak exhibits significantly lesser visible perturbations compared to StegaStamp, while still maintaining sufficient decoding accuracy—often matching or surpassing that of StegaStamp. The residual images generated by the two models differ greatly; GhostFreak's residual is far less perceptible, making the encoding much less visible compared to the more nogticeable patterns produced by StegaStamp.

5.1 Quantitative Evaluation

Three metrics were used for numerical comparison:

- PSNR**:** GhostFreak outperforms StegaStamp with an 8.35% higher PSNR, indicating better image fidelity after encoding.
- SSIM**:** A slight drop (−1.32%) in SSIM suggests marginally less structural similarity—likely due to changes introduced by multi-color space transformations.
- LPIPS**:** GhostFreak yields a lower LPIPS, suggesting improved perceptual quality, despite the drop in SSIM.

5.2 Qualitative Evaluation

In real-world settings, steganographic performance hinges not only on numerical metrics but on how well the encoded data survives physical transformations. We compared printed and re-scanned outputs for both models.

- GhostFreak shows noticeably reduced visual artifacts after printing, producing stego images that appear more natural to the human eye.
- Both models experience degradation in extreme lighting conditions, but GhostFreak maintains more stable decoding.
The residuals generated by GhostFreak are less visually apparent, thanks to its dynamic edge-aware loss and color-space blending, making it better suited for physical steganography.These results are summarized in Table 2.

Table 2: Qualitative Evaluation Of Visual Stealth And Print-Scan Robustness Ghostfreak And Stegastamp Trained With The Same Hyperparameters

Aspect	**Stegastamp**	**GhostFreak**
Visible perturbations (print)	Moderate	Minimal
Secret recoverability	Consistent	Consistent
Residual visibility	High	Low

6. DISCUSSIONS AND MODEL LIMITATIONS

While GhostFreak demonstrates promising improvements in robustness and fidelity, certain limitations remain:

- Payload Capacity:The current design supports embedding a 7-character (56-bit) message per image. Attempts to increase this capacity result in degraded encoding performance and higher training costs.
- Computational Complexity: Training the model for 140,000 iterations on high-resolution images necessitates significant computational resources.
- Color Space Integration Complexity:Incorporating RGB, HSI, and CMYK into a single framework introduces additional design complexity, which may hinder scalability to higher resolutions or larger payloads.
- Sensitivity to Extreme Distortions:Although our dynamic loss weighting strategy improves robustness, extreme print-scan artifacts may still impact decoding accuracy.

Future research should focus on enhancing payload capacity, optimizing the training process, and refining the dynamic loss adjustments for varied operational conditions.

7. CONCLUSION

GhostFreak introduces a significant advancement in print-scan steganography by fusing the strengths of multiple perceptual and print-relevant color spaces (RGB, HSI, CMYK) with a dynamically weighted composite loss function. The proposed U-Net GAN-based architecture, trained under an evolving objective regime, enables the model to embed secret messages with high fidelity while adapting to the complex degradations introduced during physical printing and scanning. This dynamic weighting strategy—shifting emphasis from secret accuracy to image realism over time—proves critical for maintaining robustness across diverse distortion conditions.

Comparative evaluations with StegaStamp [1] reveal that although both models perform similarly on standard metrics like PSNR, SSIM, and LPIPS, GhostFreak demonstrates superior real-world resilience as demonstrated and summarized in table 3, which shows the visual difference between the traditional model and our model. Specifically, it produces significantly fewer visible perturbations in printed images and maintains equal or improved decoding accuracy. Additionally, the residuals generated by GhostFreak are less visually detectable, enhancing the steganographic covertness—an essential quality for practical deployments in document security, watermarking, and covert communication.
Overall, this study establishes a reproducible and scalable methodology for robust steganographic encoding in hybrid media. It also lays the groundwork for future research into adaptive encoding techniques, higher payload capacity, and domain-specific optimization that can extend deep steganography to broader real-world applications.

REFERENCES

[1] M. Tancik, B. Mildenhall, and R. Ng, "StegaStamp: Invisible Hyperlinks in Physical Photographs," arXiv preprint, 2020. [Online]. Available: https://arxiv.org/abs/1904.05343

[2] M. J. Huiskes and M. S. Lew, "The MIR Flickr Retrieval Evaluation," in *Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval (MIR '08)*, ACM, 2008.

[3] J. Maule, A. E. Skelton, and A. Franklin, "The Development of Color Perception and Cognition," *Annual Review of Psychology*, vol. 74, pp. 87-111, 2023. doi: https://doi.org/10.1146/annurev-psych-032720-040512

[4] E. Blotta, A. Bouchet, V. Ballarin, and J. Pastore, "Enhancement of medical images in HSI color space," *Journal of Physics: Conference Series*, vol. 332, 2011. doi: 10.1088/1742-6596/332/1/012041

[5] V. ŽIljak, K. Pap, and I. ŽIljak, "Infrared hidden CMYK graphics," *The Imaging Science Journal*, vol. 58, no. 1, pp. 20-27, 2010. doi: https://doi.org/10.1179/136821909X12520525092882

[6] J. Fridrich, T. Pevný, and J. Kodovský, "Statistically Undetectable JPEG Steganography: Dead Ends, Challenges, and Opportunities," in *Proceedings of the 9th Workshop on Multimedia & Security*, 2007.
[7] T. Pevný, T. Filler, and P. Bas, "Using High-Dimensional Image Models to Perform Highly Undetectable Steganography," in *International Workshop on Information Hiding*, 2010.
[8] A. Cheddad, J. Condell, K. Curran, and P. McKevitt, "Digital Image Steganography: Survey and Analysis of Current Methods," *Signal Processing*, vol. 90, no. 3, 2010.
[9] I. Cox, M. Miller, J. Bloom, J. Fridrich, and T. Kalker, *Digital Watermarking and Steganography*, Morgan Kaufmann, 2007.
[10] J. Zhu, R. Kaplan, J. Johnson, and L. Fei-Fei, "HIDDEN: Hiding Data with Deep Networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
[11] E. Wengrowski and K. Dana, "Light Field Messaging with Deep Photographic Steganography," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1515–1524, 2019.
[12] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
[13] R. C. Bose and D. K. Ray-Chaudhuri, "On a Class of Error Correcting Binary Group Codes," *Information and Control*, 1960.
[14] C. Wang, Y. Yang, and J. Shen, "PSNR Estimate for JPEG Compression," in *Advances in Multimedia Information Processing – PCM 2017*, Lecture Notes in Computer Science, vol. 10736, Springer, Cham, 2018. doi: https://doi.org/10.1007/978-3-319-77383-4_68
[15] G. P. Renieblas, A. T. Nogués, A. M. González, N. G. León, and E. G. del Castillo, "Structural Similarity Index Family for Image Quality Assessment in Radiological Images," *Journal of Medical Imaging*, vol. 4, no. 3, p. 035501, July 2017.

	StegaStamp	GhostFreak
Image Input		
Secret (7-character text)	Meow!!!	Meow!!!
Residual	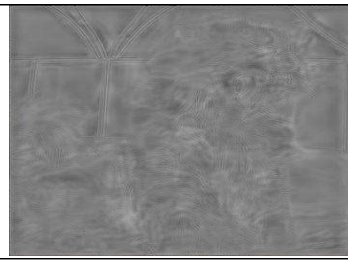	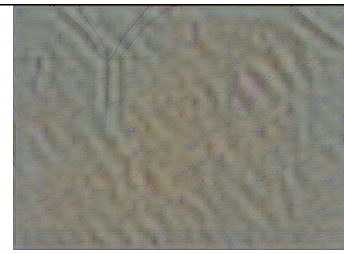
Encoded Image		
Printed and Rescanned Encoded Image	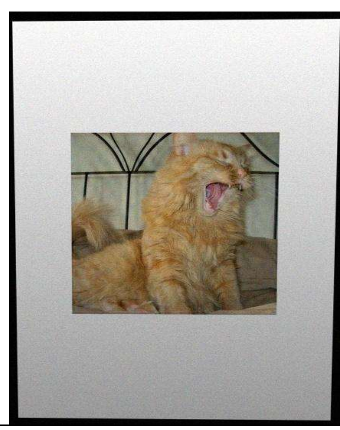	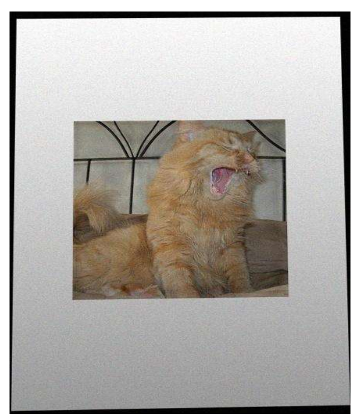
Extracted Secret from Printed Encoded Image	Meow!!!	Meow!!!

Table 3: Visual Comparison Of Stegastamp And Ghostfreak Trained With The Same Hyper Parameters