# DEEP CROSS-MODAL STEGANOGRAPHY USING NEURAL REPRESENTATIONS

Gyojin Han, Dong-Jae Lee, Jiwan Hur, Jaehyun Choi, and Junmo Kim

School of Electrical Engineering, KAIST, South Korea

# ABSTRACT

Steganography is the process of embedding secret data into another message or data, in such a way that it is not easily noticeable. With the advancement of deep learning, Deep Neural Networks (DNNs) have recently been utilized in steganography. However, existing deep steganography techniques are limited in scope, as they focus on specific data types and are not effective for cross-modal steganography. Therefore, We propose a deep cross-modal steganography framework using Implicit Neural Representations (INRs) to hide secret data of various formats in cover images. The proposed framework employs INRs to represent the secret data, which can handle data of various modalities and resolutions. Experiments on various secret datasets of diverse types demonstrate that the proposed approach is expandable and capable of accommodating different modalities.

*Index Terms*— Deep Steganography, Implicit Neural Representation, Data Hiding

# 1. INTRODUCTION

Steganography is a technique that involves embedding secret data, such as binary messages, images, audio, or video, into another message or data in a way that makes it difficult to detect. The objective of steganography is to add an extra layer of security to communication, making it harder for unauthorized parties to detect the presence of hidden information.

Recently, there has been a surge of interest in utilizing deep neural networks (DNNs) in steganography pipelines. Many of these approaches, such as those outlined in [1, 2, 3, 4, 5], use DNNs as encoders and decoders for embedding and extracting hidden information. These deep steganography techniques, which involve end-to-end training of DNNs, have demonstrated advantages over traditional steganography methods in terms of capacity and security, while also being simpler to implement.

However, most of these methods are limited in scope and focus only on specific data types, so they do not provide a comprehensive solution for data hiding. The majority of deep steganography approaches concentrate on hiding binary messages [6, 2, 7] or natural images [1, 3]. While some efforts have been made to use deep steganography to hide video [8, 9] or audio [10] within the same data type, there is a short-



**Fig. 1**. We provide a comprehensive solution for deep crossmodal steganography dealing with secret data of various modalities by using INRs.

age of research on cross-modal steganography. This type of steganography involves hiding secret data in cover data that is in a different format than the secret data. This can be particularly challenging when the secret data has higher dimensions than the cover data, such as hiding a video in an image.

To address these limitations, we establish a new deep cross-modal steganography framework using Implicit Neural Representations (INRs) that can be generalized to tasks with secret data of various formats. We employed INRs to represent the secret data, as INRs have the ability to handle data of various modalities and resolutions as seen in Fig. 1. Under this framework, the objective is to send secret data expressed in INR through a container image with minimal distortion from a cover image. To achieve this goal, we make the following contributions:

- We enable the sender and recipient to share a base network and transmit a portion of the weights constituting that network in the form of an image.
- To reduce quantization error that occurs during the conversion process from the weights to the images, we propose a simplified quantization-aware training.

We demonstrate the expandability of our method and potential for various applications through experiments on datasets of various modalities.



Fig. 2. We train  $f_{\theta}$  to approximate secret data S with various modalities, with a constraint that a portion of weights  $\mathbf{W}_{v}$  to imitate the cover image C. Using a container image C', the only recipient who knows the remaining portion of the parameters  $\mathbf{W}_{p}$  can reveal the secret data S' by reconstructing the whole INR network.

## 2. METHOD

#### 2.1. Overview

Our goal is to solve the cross-modal steganography problem, specifically hiding secret data S of various modalities (such as a video, audio, or 3D shape) into a single cover image C. The container image C' including the information about Sshould be difficult to distinguish from the cover image C by human observation. In our framework, the sender and recipient share the architecture of an INR, referred to as the base network. The base network is composed of variable weights  $(\mathbf{W}_v)$  and pre-defined weights  $(\mathbf{W}_p)$ . The hiding stage is carried out by training the INR so that the container image C'converted from the  $\mathbf{W}_v$  is close to the cover image  $\mathbf{C}$ , while simultaneously making the revealed secret S' reconstructed by the INR to be close to the secret S. The decoding can be accomplished by reconstructing the entire INR using  $\mathbf{W}_{v}$ which is converted from C' and  $W_p$  which is pre-defined. This is achieved by treating each channel of the container image as a weight matrix and inserting the weight matrices with converted image channels in the base network.

# 2.2. Deep Cross-Modal Hiding Framework

In this paper, we consider, but not limited to, a multi-layer perceptron (MLP) with L layers as an example of INR  $f_{\theta}$ where the parameter  $\theta$  can be defined as a set of matrices such that  $\theta = \{\mathbf{W}_l | \mathbf{W}_l \in \mathbb{R}^{C_{in} \times C_{out}}\}_{l=0}^{L-1}$  and  $C_{in}$  and  $C_{out}$  represent the number of units in each layer before and after  $\mathbf{W}_l$ . For the proposed method, the sender and recipient share the pre-defined weights  $\mathbf{W}_p$  and positions of variable weights. Formally, the variable weights  $\mathbf{W}_v$  refer to a subset of the weights constituting the base network that can be trained. Additionally, they share the hyperparameters  $w_{min}$  and  $w_{max}$ , which are required for scaling and quantization.

We create a container image C' by stacking the  $W_v$  and

scaling from the range of  $[w_{min}, w_{max}]$  to [0, 255] with the idea that weight matrix has the same matrix form as an image channel. Therefore, the  $\mathbf{W}_v$  are trained for two objects during the hiding phase: 1) the INR including the  $\mathbf{W}_v$  should represent secret data S well, and 2) the container image  $\mathbf{C}'$  converted from the stack of the  $\mathbf{W}_v$  should be close to the cover image  $\mathbf{C}$ . During the hiding phase, all other parts of the base network except for the  $\mathbf{W}_v$  are fixed.

For the first object, the INR  $f(x_i)_{\theta} : \mathbb{R}^k \to \mathbb{R}^m$  are trained to approximate the data point  $d_{i \in \mathcal{I}}$  of secret data **S** according to the input coordinate vectors  $x_{i \in \mathcal{I}}$ , where  $\mathcal{I}$  is a pre-defined set of input index. During the training, the  $\mathbf{W}_v$  are trained with a secret data reconstruction loss  $\mathcal{L}_s$ , where

$$\mathcal{L}_s(f_\theta) = \sum_{i \in \mathcal{I}} \|d_i - f_\theta(x_i)\|_2^2.$$
(1)

For the second object, we constrain a stack of the variable weight matrices  $\mathbf{W}_s \in \mathbb{R}^{N \times N \times 3}$  to resemble the cover image  $\mathbf{C} \in \mathbb{R}^{N \times N \times 3}$ . That is, we aim to directly utilize scaled  $\mathbf{W}_s$  as a container image  $\mathbf{C}'$  by minimizing a loss

$$\mathcal{L}_{c}(\mathbf{W}_{s}) = \|\mathbf{W}_{s} - (\mathbf{C}/255 \cdot (w_{max} - w_{min}) + w_{min})\|_{2}^{2}.$$
 (2)

However, we did not consider the quantization error caused by converting the weights to images in this section. Hence, it is imperative to address the quantization error during the hiding process.

# 2.3. Quantization-Aware Training for Weight-to-Image Conversion

It should be noted that converting weights stored in FP32 to an image stored in UINT8 can cause a significant perturbation, leading it to deviate from the converged state. This poses a major challenge and needs to be handled carefully during the process. To address this, we introduce simplified quantization-aware training (QAT) for weight-to-image conversion. Originally, QAT is a method that is widely used in network compression problems, which considers the quantization process during training. We use QAT with some modifications to solve the problem that the weights are shifted from the converged state during conversion.

We use a quantization module  $f_{\theta_Q}$  with quantized variable weights  $\mathbf{Q}_v$  and the same pre-defined weights, which simulates the quantization process during the training. For the simulation, the variable weights  $\mathbf{W}_v$  are quantized to  $\mathbf{Q}_v$  for every updates of  $\mathbf{W}_v$  as below formula:

$$\mathbf{Q}_{v} = \mathbf{ROUND}(255 \cdot (\mathbf{W}_{v} - w_{min}) / (w_{max} - w_{min}))$$
$$\mathbf{Q}_{v} \leftarrow \mathbf{Q}_{v} / 255 \cdot (w_{max} - w_{min}) + w_{min}.$$
(3)

The quantization module  $f_{\theta_Q}$  is used to calculate the gradient  $d\mathcal{L}/d\mathbf{Q}_v$ . As explained earlier, the total loss for the quantized network  $\mathcal{L}$  is:

$$\mathcal{L} = \mathcal{L}_s(f_{\theta_O}) + \beta \cdot \mathcal{L}_c(\mathbf{Q}_s) \tag{4}$$

where  $\mathbf{Q}_s$  is a stack of  $\mathbf{Q}_v$ . To back-propagate the gradient, we convert  $d\mathcal{L}/d\mathbf{Q}_v$  into  $d\mathcal{L}/d\mathbf{W}_v$  with straight through estimator (STE) [11]. Therefore,  $\mathbf{W}_v$  is finally updated as follows:

$$d\mathcal{L}/d\mathbf{W}_{v} \approx \mathbf{STE}(d\mathcal{L}/d\mathbf{Q}_{v})$$
  
$$\mathbf{W}_{v} \leftarrow \mathbf{W}_{v} - \alpha \cdot d\mathcal{L}/d\mathbf{W}_{v}.$$
 (5)

The provided Fig. 2 illustrates the proposed steganography framework as a whole.

#### 3. EXPERIMENTS

We experiment with the proposed method on three tasks: video-into-image steganography, audio-into-image steganography, and shape-into-image steganography. These tasks have not been explored well due to the difficulty that secret data has higher dimensions than cover data or contains temporal information. We demonstrate the performance and flexibility of our method by solving these challenging tasks for the first time.

## 3.1. Experimental Setting

**Datasets.** In common to all three tasks, cover images are randomly sampled from the ImageNet [12] by the number of data in the corresponding secret dataset, and all sampled cover images are resized to  $512 \times 512$ .

For video-into-image steganography, we use the Densely Annotated VIdeo Segmentation dataset (DAVIS) [13] as a secret dataset. We resize the secret videos to  $128 \times 128$ and extract 16 frames from the videos. For audio-into-image steganography, we hide audio data from GTZAN music genre dataset [14, 15]. Each hidden audio data is cropped to a length of 100,000 samples with a 22,050 sample rate (4.54 seconds).

Container Container Correction Container Correction Container Correction Container Correction Corre

Fig. 3. Qualitative results of video-into-image steganography.

For shape-into-image steganography, we selected two samples, Stanford Bunny and Dragon, as secret three-dimensional (3D) shapes from The Stanford 3D Scanning Repository [16].

**Details.** We use SIREN [17] with four hidden layers (total six layers) as the architecture of INRs to hide video and audio, and IGR [18] with six hidden layers (total eight layers) as the architecture of INRs to hide shapes. For SIREN and IGR, { $W_1$ ,  $W_2$ ,  $W_3$ } and { $W_1$ ,  $W_2$ ,  $W_4$ } are the variable weights and represent color channels of the container image, respectively. The remaining pre-defined weights  $W_p$  are fixed in an initialized state. The INRs are trained for 5,000 steps for video and shapes, and 20,000 steps for audio using Adam optimizer. In addition, the learning rate of the optimizer is set to  $1 \times 10^{-3}$  for video and audio and set to  $1 \times 10^{-4}$ for shapes.

# 3.2. Experimental Results

**Video-into-image steganography.** The Average Pixel Discrepancy (APD) of the cover and secret is calculated as the  $L_1$  distance between the cover and container and that between the frames of the secret video and revealed secret video, respectively. In addition to APD, we report Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Perceptual Similarity (LPIPS) in the same way in Table 1. The qualitative results for the video-into-image steganography can be confirmed in Fig. 3.

Errors	$APD\downarrow$	$PSNR\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Cover	20.32	19.77	0.5191	0.5262
Secret	5.79	29.98	0.9263	0.1333

Table 1. Performance of the video-into-image steganography.



Fig. 4. Qualitative results of audio-into-image steganography.



Fig. 5. Qualitative results of shape-into-image steganography.

**Audio-into-image steganography.** For evaluation of reconstructed audio, we report Absolute Error (AE) and Signalto-Noise Ratio (SNR) between the secret and revealed secret instead of APD and PSNR in Table 2. In addition, the audios of the secret and revealed secret are visualized as melspectrograms in Fig. 4.

Errors	APD/AE $\downarrow$	PSNR/SNR ↑	SSIM $\uparrow$	LPIPS $\downarrow$
Cover	13.73	23.33	0.7139	0.3948
Secret	2.84	17.92	0.9328	-

 Table 2. Performance of the audio-into-image steganography.

**Shape-into-image steganography.** We experiment with the proposed method on the shape-into-image steganography to demonstrate that the proposed method has a significantly broader scope of application than conventional deep steganography methods, as 3D shapes are fundamentally different from the data that previous deep steganography methods attempted to conceal. In Fig. 5, we present the reconstructed Stanford Bunny and Dragon from various viewpoints.

**Ablation study.** We measure the performance difference of video-into-image steganography depending on whether QAT is applied or not. The experimental results in Table 3 show that when QAT is not applied, the revealed secret is very



Fig. 6. Effectiveness of quantization-aware training.

altered from the original due to the weight shifts by quantization. We also present the visual effect of QAT through the sample included in Fig. 6.

Errors	$\text{APD}\downarrow$	$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Cover w/o. QAT	19.17	20.37	0.5700	0.4785
Secret w/o. QAT	27.76	17.73	0.5094	0.4279
Cover w. QAT	20.32	19.77	0.5191	0.5262
Secret w. QAT	5.79	29.98	0.9263	0.1333

Table 3. Ablation study results on QAT.

#### 3.3. Discussion and Future Work

Through our experiments, we show that it is possible to hide secret data of various modalities with only subtle changes to the cover image. However, there is a limitation that the format of cover data is fixed as an image, and there is some loss in high-frequency information for revealed secret. Since this paper introduces a steganography solution using INRs for the first time, we believe that there is still room for improvement and various applications. We also look forward to future work to improve the robustness of the proposed approach against container distortions such as blur or JPEG compression.

# 4. CONCLUSION

We proposed a novel approach to deep cross-modal steganography utilizing INRs, which enables the hiding of data of diverse modalities within images. We believe that our proposed method will stimulate further investigation, as it introduces a fresh research direction for deep steganography.

# 5. ACKNOWLEDGEMENT

This work was supported by Artificial intelligence industrial convergence cluster development project funded by the Ministry of Science and ICT(MSIT, Korea)&Gwangju Metropolitan City.

# 6. REFERENCES

- Shumeet Baluja, "Hiding images in plain sight: Deep steganography," in *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. 2017, vol. 30, Curran Associates, Inc.
- [2] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei, "Hidden: Hiding data with deep networks," in Proceedings of the European Conference on Computer Vision (ECCV), September 2018.
- [3] Chaoning Zhang, Philipp Benz, Adil Karjauv, Geng Sun, and In So Kweon, "Udh: Universal deep hiding for steganography, watermarking, and light field messaging," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, Eds. 2020, vol. 33, pp. 10223– 10234, Curran Associates, Inc.
- [4] Shao-Ping Lu, Rong Wang, Tao Zhong, and Paul L. Rosin, "Large-capacity image steganography based on invertible neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 10816–10825.
- [5] Zhenyu Guan, Junpeng Jing, Xin Deng, Mai Xu, Lai Jiang, Zhou Zhang, and Yipeng Li, "Deepmih: Deep invertible network for multiple image hiding," *IEEE Transactions on Pattern Analysis and Machine Intelli*gence, vol. 45, no. 1, pp. 372–390, 2023.
- [6] Jamie Hayes and George Danezis, "Generating steganographic images via adversarial training," in Advances in Neural Information Processing Systems, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. 2017, vol. 30, Curran Associates, Inc.
- [7] Matthew Tancik, Ben Mildenhall, and Ren Ng, "Stegastamp: Invisible hyperlinks in physical photographs," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- [8] Xinyu Weng, Yongzhi Li, Lu Chi, and Yadong Mu, "High-capacity convolutional video steganography with temporal residual modeling," in *Proceedings of the 2019* on International Conference on Multimedia Retrieval, New York, NY, USA, 2019, ICMR '19, p. 87–95, Association for Computing Machinery.
- [9] Aayush Mishra, Suraj Kumar, Aditya Nigam, and Saiful Islam, "Vstegnet: Video steganography network using spatio-temporal features and micro-bottleneck.," in *BMVC*, 2019, vol. 274.

- [10] Felix Kreuk, Yossi Adi, Bhiksha Raj, Rita Singh, and Joseph Keshet, "Hide and speak: Deep neural networks for speech steganography," *arXiv preprint arXiv:1902.03083*, 2019.
- [11] Yoshua Bengio, Nicholas Léonard, and Aaron Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," arXiv preprint arXiv:1308.3432, 2013.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, "Imagenet: A large-scale hierarchical image database," in 2009 IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 248– 255.
- [13] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung, "A benchmark dataset and evaluation methodology for video object segmentation," in *The IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2016.
- [14] George Tzanetakis, *Manipulation, analysis and retrieval systems for audio signals*, Princeton University, 2002.
- [15] George Tzanetakis and Perry Cook, "Musical genre classification of audio signals," *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293– 302, 2002.
- [16] Stanford University Computer Graphics Laboratory, "The stanford 3d scanning repository," http://graphics.stanford.edu/data/3Dscanrep/.
- [17] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein, "Implicit neural representations with periodic activation functions," in Advances in Neural Information Processing Systems, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, Eds. 2020, vol. 33, pp. 7462–7473, Curran Associates, Inc.
- [18] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman, "Implicit geometric regularization for learning shapes," in *Proceedings of the 37th International Conference on Machine Learning*, Hal Daumé III and Aarti Singh, Eds. 13–18 Jul 2020, vol. 119 of *Proceedings of Machine Learning Research*, pp. 3789– 3799, PMLR.