

Distortion Agnostic Deep Watermarking

Xiyang Luo¹, Ruohan Zhan^{2*}, Huiwen Chang¹, Feng Yang¹, and Peyman Milanfar¹

¹Google Research
²Stanford University

Abstract

Watermarking is the process of embedding information into an image that can survive under distortions, while requiring the encoded image to have little or no perceptual difference from the original image. Recently, deep learning-based methods achieved impressive results in both visual quality and message payload under a wide variety of image distortions. However, these methods all require differentiable models for the image distortions at training time, and may perform poorly on unknown distortions. This is undesirable since the types of distortions applied to watermarked images are usually unknown and non-differentiable. In this paper, we propose a new framework for distortion-agnostic watermarking, where the image distortion is not explicitly modeled during training. Instead, the robustness of our system comes from two sources: adversarial training and channel coding. Compared to training on a fixed set of distortions and noise levels, our method achieves comparable or better results on distortions available during training and on unknown distortions.

1. Introduction

Digital watermarking [8] is the task of embedding information into an image in a visually imperceptible fashion, where the message can be reliably extracted under image distortions. There are two key factors to measure the performance of a digital watermarking system, *imperceptibility* and *robustness*. Given an image and a message, a good watermarking system produces an encoded image that is nearly identical to the original image, while carrying a message payload that will survive under a variety of distortions such as cropping, blurring, or JPEG compression. Traditional approaches found creative ways of hiding information in texture rich areas [5] or the frequency domain [16].

*Work done while an intern at Google.

| Watermarking Models | Known Distortions | | | | | Unknown Distortions | | | | |
|-----------------------|-------------------|-------------|----------------|-----------------|-----------------------|---------------------|-------------|----------------|-----------------|-----------------------|
| | Identity | JPEG (Q=50) | Crop (p=0.035) | Dropout (p=0.3) | Blur ($\sigma=1.0$) | Identity | JPEG (Q=50) | Crop (p=0.035) | Dropout (p=0.3) | Blur ($\sigma=1.0$) |
| Identity | 100.0 | 50.2 | 53.0 | 81.3 | 71.8 | 74.6 | 65.1 | 93.4 | 85.3 | 65.0 |
| JPEG (Q=50) | 99.0 | 80.8 | 85.7 | 87.0 | 86.1 | 87.7 | 71.0 | 95.2 | 78.7 | 90.6 |
| Crop (p=0.035) | 100.0 | 49.9 | 99.3 | 51.0 | 62.5 | 50.3 | 64.3 | 92.1 | 81.1 | 63.3 |
| Dropout (p=0.3) | 93.1 | 50.0 | 51.3 | 99.4 | 51.0 | 51.0 | 54.3 | 89.0 | 75.4 | 63.8 |
| Blur ($\sigma=1.0$) | 53.0 | 49.8 | 51.8 | 50.5 | 99.9 | 50.3 | 61.4 | 69.6 | 76.3 | 52.6 |
| Combined [40] | 100.0 | 77.0 | 99.1 | 98.7 | 99.1 | 93.5 | 70.8 | 94.2 | 84.9 | 88.6 |
| Our Model | 100.0 | 81.7 | 93.5 | 97.9 | 92.8 | 95.6 | 94.0 | 98.5 | 88.4 | 91.7 |

Figure 1: Bit accuracy of our model compared to models trained with explicit image distortions. Each column corresponds to a type of image distortion at test time, and each row corresponds to the image distortion used to train the watermarking model (with the exception of our model which requires no distortion model). The left half of the columns (separated by the black line) are *known distortions*, i.e., distortions included in training for the HiDDeN combined model [45], and the right half of the columns *unknown distortions*, i.e., a held-out set of commonly used distortions not used to train the HiDDeN combined model. See Section 4.1 for more details.



Figure 2: Example of original image, encoded image and difference between the two images from our model.

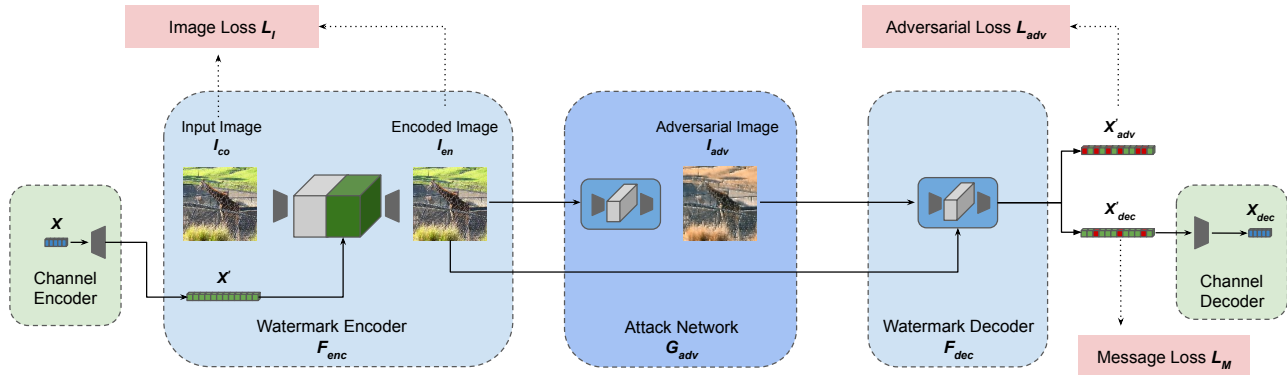


Figure 3: Overview of proposed architecture. The input message X is first fed through the channel encoder to produce a redundant message X' , which is then combined with the input image I_{co} to generate the encoded image I_{en} by the watermark encoder F_{enc} . The decoder F_{dec} produces a decoded message X'_{dec} , where it is further processed by the channel decoder to produce the final message X_{dec} . The attack network generates adversarial examples I_{adv} , which are fed to the image decoder to obtain X'_{adv} . F_{enc} , F_{dec} , are trained on a combination of the image loss L_I which includes both proximity to the cover image I_{co} and perceptual quality as in Equation 1, the message loss L_M as in Equation 2, and the message loss on the decoded adversarial message X'_{adv} as in Equation 4. The attack network G_{adv} is trained to minimize the adversarial loss L_{adv} as in Equation 3. The training updates G_{adv} and the F_{enc} , F_{dec} in an alternating fashion.

More recently, convolutional neural networks (CNNs) have been used to provide an end-to-end solution to the watermarking problem. Zhu *et al.* [45] proposed HiDDeN, a CNN-based framework for image watermarking. These learning-based methods [45, 1, 40] improve their model’s robustness by including a diverse combination of distortions during training. While effective, there are several issues with this approach. First of all, the distortions included in training must be differentiable in order to train the encoder-decoder jointly. However, many practical distortions, such as GIF compression, projector-camera distortion in Light Field Messaging [40], etc., are non-differentiable and challenging to make differentiable. Secondly, in most cases, distortions are not available during training. Selecting a diverse and well-balanced set of distortions for training is non-trivial and requires careful tuning.

To this end, we propose a framework for adding robustness to a watermarking system *without* any prior knowledge on the type of image distortions during training. We achieve this by applying adversarial training with CNN generated perturbations, and using channel coding to inject redundancy in the encoded message. To our knowledge, this is the first paper explore distortion agnostic methods for deep watermarking. Our model achieves comparable or better performance on a wide variety of distortions, without the need to explicitly model the distortions at training time.

Our main contributions are the following.

- We apply adversarial training to improve model robustness in a distortion agnostic fashion. In particular, our CNN generated adversarial examples implicitly incorporates a rich collection of image distortions that

co-adapt with training.

- We propose augmenting the watermarking system with channel coding, adding an additional layer of robustness through channel redundancy.
- We combine the two ideas above and achieve comparable results to models trained with explicit distortion, and better performance to unknown distortions.

2. Related Work

There are three main areas of research relevant to this work: watermarking, adversarial training, and channel coding. We give a brief review for each topic in the subsections below.

2.1. Watermarking

Digital watermarking [8, 15, 5, 14, 24, 31, 35, 12] has been an active research area with many important applications such as content copyright protection. More recently, deep learning based approaches have been applied to train an end-to-end watermarking system [45, 23, 1, 21, 43] with impressive results. HiDDeN [45] was one of the first deep learning solutions for image watermarking. RedMark [1] introduced residual connections with a strength factor for embedding binary images in the transform domain. Deep watermarking has since been generalized to video [39, 44] as well. Modeling more complex and realistic image distortions also broadened the scope in terms of application [40, 36].

There are several works that applied attacks to the encoded image when training the watermarking system. Mun

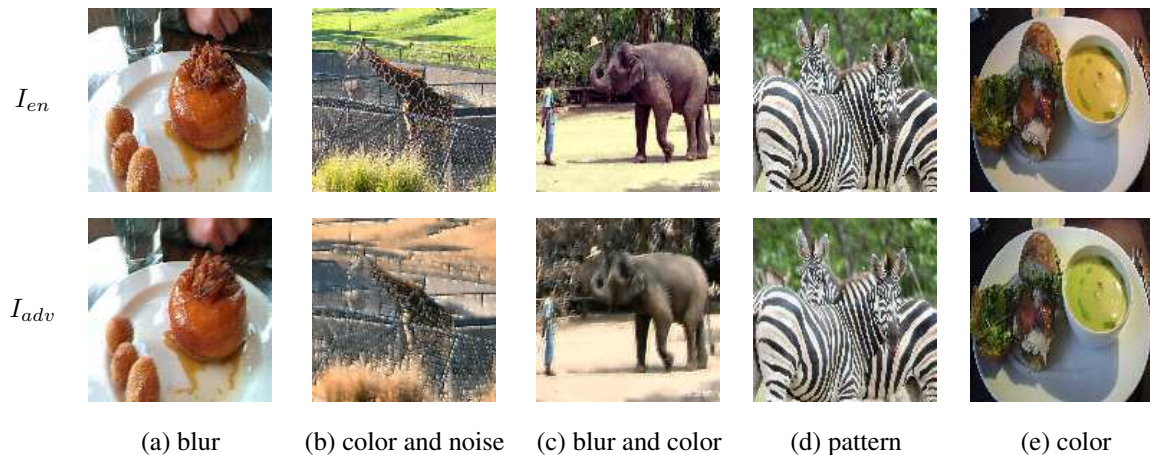


Figure 4: Visualization of adversarial examples generated by the attack network G_{adv} . **Top:** Encoded image I_{en} . **Bottom:** Adversarial examples I_{adv} generated from the attack network G_{adv} . We observe a diverse set of image manipulations generated from the attack network, consisting of a combination of blur, color change, and other types of distortions.

et al. [23] iteratively simulated attacks to the watermarking system. RedMark [1] introduced an attack layer which consists of random combinations of a fixed set of distortions. However, these attacks are not adversarial since they do not adapt with the watermarking model during training. Recently, ROMark [38] applied a simple form of adversarial training where the distortion type and distortion strength are adaptively selected to minimize the decoding accuracy.

One key distinction of our method from the above is that we do *not* generate our attacks from a *fixed* pool of common distortions. Instead, the adversarial examples are generated from a trained CNN. This also has the benefit that the watermarking training is end-to-end differentiable, which is not true for ROMark [38].

2.2. Adversarial Training

Deep neural networks are susceptible to certain tiny perturbations in the input space. Since the discovery of adversarial examples by Szegedy *et al.* [34], a variety of methods have been proposed for both adversarial attack [3, 17, 25], and adversarial defense [11, 13, 27, 42]. One of the earliest and most effective defense mechanism against adversarial attacks is adversarial training [11], but is computationally expensive on large datasets. Many attempts have since been made to reduce the cost of adversarial training, *e.g.*, using approximations to the optimization step [27], or using generative models in place of iterative optimization [4, 18].

2.3. Channel coding

Channel coding is a mechanism for detecting and correcting errors during signal transmission [6]. Shannon’s capacity theorem [28] gives the theoretical limit to transfer data through a noisy channel, and channel coding is designed to approach this limit. In implementation, var-

ious classical methods such as the Reed-Solomon (RS) codes [41], low-density parity-check (LDPC) codes [26], turbo codes [32], and polar codes [37], have been widely applied in the field of telecommunication. More recently, learning based solutions have gained attention in this field as well [2, 7, 9].

3. Proposed Method

3.1. Motivation

In designing a general purpose watermarking model, the distortions at test time could be any image manipulation that still preserves some image content. A typical solution would involve identifying a set of representative distortions, and applying a carefully tuned combination of distortions during training.

Motivated by the recent success of using CNNs to perform various image manipulation tasks, *e.g.*, style transfer [10], HDRNet [19], we propose automating the distortion tuning process by *training* a CNN to generate distortions that exploits the weakest link in the current watermarking model. Figure 4 shows some samples of distorted images generated by our attack CNN, which contain a rich and complex combination of distortions.

The use of channel coding is motivated by the idea of injecting extra redundancy to the system. Shannon’s capacity theorem tells us that redundancy is necessary in order to achieve robustness. In the HiDDeN architecture, spatially repeating the input message is an example of adding redundancy. Channel coding simply provides another alternative on top of the current methods.

3.2. Method Overview

Figure 3 gives an overview of our overall architecture. Our method adds two key components on top of the watermarking encoder / decoder networks F_{enc} and F_{dec} in [45]:

- We replace the distorted image with I_{adv} , where I_{adv} is an adversarial example generated from a convolutional neural network trained to *maximize* the message loss.
- We replace the input message with a longer binary message X' generated from channel coding.

3.3. Adversarial Training

Adversarial training generates distortions that co-adapt with the training of our watermarking model, actively strengthening the weakest point of the current model. Adversarial training was first introduced by Goodfellow *et al.* [11] as a method to defend against adversarial attacks. In our context, adversarial training equates to minimizing the message loss given the worst-case distortion in an ϵ -ball. This is expressed as the following min-max problem,

$$\min_{\Theta_{enc}, \Theta_{dec}} \max_{\|\delta\| \leq \epsilon} \{L_M(F_{dec}(F_{enc}(I_{co}; \Theta_{enc}) + \delta; \Theta_{dec}), X)\}, \quad (1)$$

where Θ_{enc} , Θ_{dec} are the model parameters for watermarking encoder/decoder networks F_{enc} , F_{dec} , and X is the input message. Here we consider the L_2 norm $\|\cdot\|_2$ to constrain δ , the perturbation to the encoded image $I_{en} = F_{enc}(I_{co}; \Theta_{enc})$. But more semantically meaningful measures such as L_2 distance on VGG [30] activations could also be used.

A direct optimization of Equation 1 is both computationally expensive and overly restrictive for the watermarking model. Instead, we relax Equation 1 by restricting the set of distortions δ to be generated from some class of convolutional neural network $G_{adv}(I; \Theta_{adv})$.

$$\min_{\{\Theta_{enc, dec}\}} \max_{\{\|G_{adv}(I_{en}) - I_{en}\| \leq \epsilon\}} \{L_M(F_{dec}(G_{adv}(I_{en})); X)\}. \quad (2)$$

Using CNN generated adversarial examples have the benefit of retaining the ability to generate a diverse set of image distortions, as shown in Figure 4. An alternative is to generate the adversarial samples via the Fast Gradient Sign Method (FGSM) as in [11]. But we found this yielded less diverse examples compared to CNN generated examples, and resulted in poorer overall robustness against distortion.

To train the attack network G_{adv} , we minimize the following adversarial training loss:

$$L_{adv} = \alpha_1^{adv} \|I_{adv} - I_{en}\|^2 - \alpha_2^{adv} L_M(F_{dec}(I_{adv}); X), \quad (3)$$

where $I_{adv} = G_{adv}(I_{en})$ is the adversarial example, L_M is the message loss which we set as the L_2 loss in this paper,

and α_1^{adv} , α_2^{adv} are the scalar weights. α_1^{adv} controls the strength of the distortion generated by the attack network G_{adv} , while α_2^{adv} controls the strength of the message loss for G_{adv} .

For the network G_{adv} , we use a two-layer CNN,

$$G_{adv}(I) = \text{Conv}_3 \circ \text{Leaky ReLU} \circ \text{Conv}_{16}(I). \quad (4)$$

In general, we find that finding the right balance of attack strength, controlled by the complexity of G_{adv} and the ratio between α_1^{adv} and α_2^{adv} , is important for training. An overly strong attack results in slow training and a failure of the watermarking network to adapt to the adversarial examples, while an overly simple attack results in less robustness of the trained model. A detailed analysis can be found in Section 4.3.

3.4. Channel coding




Figure 5: Illustration of channel coding. Given an input message X , the channel encoder produces a redundant message X' of longer length. The redundant message X' is transmitted through a noisy channel and received by the decoder as X'_{no} . Finally the decoder recovers the input X from the corrupted message X'_{no} .

Channel coding provides an additional layer of robustness through injecting redundancy to the system. Given a binary message $X \in \{0, 1\}^D$ of length D , a channel encoder produces a redundant message $X' \in \{0, 1\}^N$ of length $N > D$, which can be used to recover X through the channel decoder given reasonable amounts of channel distortion to X' , as shown in Figure 5.

In this paper, we generate a channel code X' from the input message X , before passing X' to the watermarking encoder as shown in Figure 3. The channel distortions in this context are the errors from the watermarking model, between X' and X'_{dec} . Given that we do not explicitly model the image distortions, it is impossible to know the true channel distortion model. Instead, we use a binary symmetric channel (BSC) to approximate the channel distortion. BSC is a standard channel model which assumes each bit is independently and randomly flipped with probability p . Even though this assumption is not strictly satisfied in our case, we find using BSC works well in this application.

Conceptually, any standard error correcting code such as low-density parity-check (LDPC) codes [26] can be used to generate X' . However, traditional codes such as LDPC require the decoder to have an estimate of the channel noise strength, which is impractical in our application since the




Figure 6: Channel noise strength versus decoder bit accuracy for various redundant message lengths. The input message length is fixed at $D = 30$, where the redundant message length N is varied from 90 to 150. All models are trained on random binary input with BSC noise. The training noise level is uniformly sampled from $[0, 0.3]$, and $[0, 0.4]$ at test time.

noise strength can vary greatly from image to image. Therefore, we use NECST [7], a learning based solution for joint source and channel coding to cover a broad range of channel distortion strengths. We use BSC for training the channel model, where the input message X is randomly sampled, and the channel noise strength is chosen from the interval $[0, \text{maxstrength}]$ uniformly at random. Figure 6 shows the bit accuracy of the NECST model on a range of BSC channel noise.

We emphasize here that the channel coding model is *not* jointly trained with the rest of the watermarking model. This decoupling prevents the channel models from co-adapting with the image models during training, which results in overfitting and less robustness across a wide spectrum of image distortions.

3.5. Watermarking Training and Losses

We give a detailed description of the algorithms for training the watermarking models. We first define the training losses, using the same notations as in Figure 3.

Image loss

$$L_I = \alpha_1^I \|I_{co} - I_{en}\|^2 + \alpha_2^I L_G(I_{en}) \quad (5)$$

Message loss

$$L_M = \alpha^M \|X'_{dec} - X'\|^2 \quad (6)$$

Attack network training loss

$$L_{adv} = \alpha_1^{adv} \|I_{adv} - I_{enc}\|^2 - \alpha_2^{adv} \|X'_{adv} - X'\|^2 \quad (7)$$

Watermarking training loss

$$L_W = L_I + L_M + \alpha_W^{adv} \|X'_{adv} - X'\|^2 \quad (8)$$

The image loss in Equation 5 consists of an L_2 loss, and a GAN loss L_G with spectral normalization [22] to control

the perceptual quality of the encoded image. This is similar to the *adversarial loss* defined in the HiDDeN network [45]. For the message loss L_M , we use the L_2 loss between the decoded message and input. Equation 7 defines the loss used to train the attack network G_{adv} . Finally, Equation 8 defines the overall loss for training F_{enc} and F_{dec} . The various α s are the weights for each loss. Training alternates between updating the attack network G_{adv} and the watermarking networks F_{enc}, F_{dec} , detailed in Algorithm 1.

Algorithm 1 Watermarking Training

1: procedure WATERMARKING TRAIN

Input: $I_{co}, X' \sim \text{Unif}(\{0, 1\}^N)$.

Output: Trained networks $G_{adv}, F_{enc}, F_{dec}$.

Training Variables: $\Theta_{adv}, \Theta_{enc}, \Theta_{dec}$.

2: while Step < max_steps do

3: Compute $I_{en} = F_{enc}(I_{co}, X)$

4: for $i = 1$ to num_iter do

5: Compute $I_{adv} = G_{adv}(I_{en})$

6: Update $\Theta_{adv} = \Theta_{adv} + lr \times \text{Adam}(L_{adv})$

7: Update $\Theta_{dec} = \Theta_{dec} + lr \times \text{Adam}(L_W)$

8: Update $\Theta_{enc} = \Theta_{enc} + lr \times \text{Adam}(L_W)$

4. Experiments

For comparison, we train two versions of HiDDeN [45] as the baseline, one without image distortion which we name the *identity* model, and another trained on a combination of standard image distortions which we name the *combined* model. We note here that our methodology is agnostic to the specific architecture of the watermarking networks. We use the original HiDDeN architecture throughout the experiments since it is a well studied model and a commonly used benchmark, but other architectures such as RedMark [1] could also be used as well.

We compare the bit accuracy on distortions seen during training and those that have not, and also report the peak signal-to-noise ratio (PSNR) of the encoded images. All models are trained and evaluated on the MS COCO dataset [20] resized to 128×128 , where a random selection of 3000 images are used for evaluation. Unless otherwise stated, we use $D = 30$ for the encoded message size, and $N = 120$ for the redundant message size. For the watermarking networks, we use the same architecture as used in HiDDeN, with the exception that the embedded message size is 120 instead of 30 due to the increased message length from channel coding. Detailed training parameters can be found in the supplementary materials.

4.1. Comparison with HiDDeN

We compare our method with both the HiDDeN identity and combined models. For the combined model, we

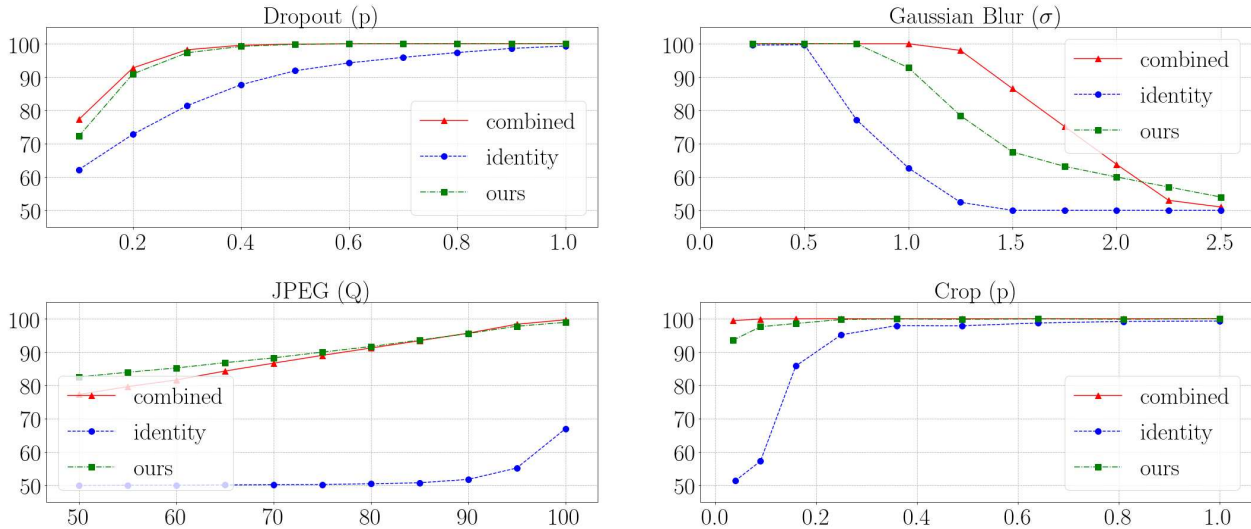


Figure 7: Comparison of our model with HiDDeN identity model and combined noise model for different types of image distortions.

| Model | Identity | Combined | Ours |
|-----------|----------|----------|------|
| RGB- PSNR | 40.3 | 32.3 | 33.7 |
| Y- PSNR | 47.5 | 34.2 | 35.7 |
| U- PSNR | 44.5 | 39.7 | 40.7 |
| V- PSNR | 43.1 | 39.5 | 40.1 |

Table 1: Comparison of encoded image quality. The PSNR values in both RGB and YUV are reported for our model, as well as the HiDDeN identity and combined distortion model.

use JPEG ($q = 50$), dropout ($q = 0.3$), crop ($p = 0.3$), Gaussian blur ($\sigma = 1.0$), where the JPEG distortion is approximated by the differentiable JPEG function [29]. We also compare with *specialized* models trained only on a single type of distortion with the noise levels in Figure 1. For a fair comparison, we adjust α_1^I to obtain a slightly higher PSNR compared to the combined model, as shown in Table 1.

Figure 1 shows the bit accuracy of our model and those trained with explicit image distortion. Each row corresponds to a different watermarking model, and each column a specific type of distortion applied at evaluation time. The top five rows (specialized models) clearly show a diagonal pattern, indicating poor generalization to other types of image distortions. From the bottom two rows, we see that both the combined model and our adversarially trained model are robust to distortions used to train the combined model (first five columns).

Figure 7 gives a more comprehensive comparison across a range of distortion levels. Our model reaches comparable performance on crop and dropout, outperforms the combined model on JPEG, and underperforms on Gaussian blur.

For small distortion strengths, our accuracy is nearly identical to the combined model. On all noise levels and distortion types, we outperform the identity model by a wide margin.

In terms of visual quality, our model is less prone to small artifacts in flat regions of an image. A qualitative comparison can be found in Figure 9.

4.2. Generalization to Unknown Distortions

A practical watermarking system must be robust to a wide range of image distortions, not just the distortions seen during training. Therefore, we compare the performance of our model and the combined model on a held-out set of commonly used image distortions. We choose six types of distortions, *i.e.*, saturation, hue, resize, Gaussian noise, salt and pepper noise, GIF encoding from four broad categories: color adjustment, pixel-wise noise, geometric transformations, and compression. For each type of distortion, we evaluate the models on three different values of distortion strengths. Figure 8 gives a visualization of the additional distortions. We choose the range of distortion strength strong enough to differentiate the performance between different models, but in a regime where the distorted image still resembles the original. We also test our method on the StirMark benchmark [33] with the default parameters.

Table 2 reports the bit accuracy of our model on these additional distortions. Overall, our model performs better on the unknown distortions, especially on the category of color change. We also note that the overall variance of bit accuracy across distortions is less compared to both the identity and combined model, indicating a more stable performance across different types of distortions.




Figure 8: Visualization of additional image distortions.

| Method | Identity | Combined | Ours |
|--------------------------|-------------|-------------|--------------|
| Gaussian Noise (0.06) | 74.6 | 93.5 | 95.6 |
| Gaussian Noise (0.08) | 67.7 | 87.2 | 93.5 |
| Gaussian Noise (0.10) | 63.2 | 80.4 | 89.5 |
| Salt and Pepper (0.05) | 99.1 | 97.2 | 95.7 |
| Salt and Pepper (0.10) | 93.1 | 89.4 | 85.0 |
| Salt and Pepper (0.15) | 83.4 | 79.6 | 77.1 |
| Adjust Hue (0.2) | 65.1 | 70.8 | 94.0 |
| Adjust Hue (0.4) | 34.0 | 45.3 | 70.7 |
| Adjust Hue (0.6) | 18.1 | 28.8 | 42.4 |
| Adjust Saturation (5.0) | 96.3 | 98.1 | 99.9 |
| Adjust Saturation (10.0) | 94.8 | 96.0 | 99.6 |
| Adjust Saturation (15.0) | 93.4 | 94.2 | 98.5 |
| GIF (64) | 87.1 | 96.5 | 97.6 |
| GIF (32) | 76.8 | 93.4 | 95.7 |
| GIF (16) | 65.0 | 88.6 | 91.7 |
| Resize Width (0.9) | 99.3 | 99.7 | 99.9 |
| Resize Width (0.7) | 85.3 | 84.9 | 88.4 |
| Resize Width (0.5) | 66.5 | 67.3 | 67.1 |
| StirMark [33] | 69.6 | 82.1 | 83.5 |
| Average | 77.91 | 84.15 | 88.05 |

Table 2: Comparison of our model with HiDDeN identity model and combined noise model on additional image distortions. We report the bit accuracy of our model, the HiDDeN combined and identity model. When computing the average, results lower than 50% are truncated to 50% since they are no better than random chance.

4.3. Detailed Analysis

4.3.1 Ablation Study

Table 3 reports the individual effect of channel coding and adversarial training. We see that adversarial training contributes to a large portion of the model robustness, while channel coding further boosts performance in terms of accuracy. Table 3 also shows that channel coding alone does not provide enough robustness without a robust watermarking model. However, combined with adversarial training channel coding further boosts the performance of the wa-

termarking system especially if the bit accuracy is already high.

| | JPEG (Q=50) | Crop (p=0.09) | Blur ($\sigma=1.0$) | Dropout (p=0.3) |
|----------|----------------|------------------|--------------------------|--------------------|
| Identity | 50.2 | 53.0 | 59.6 | 81.3 |
| Channel | 51.3 | 60.5 | 50.2 | 90.3 |
| Adv. | 85.0 | 90.6 | 86.2 | 95.0 |
| Both | 81.7 | 93.5 | 92.8 | 97.9 |

Table 3: Model ablation study. We report the bit accuracy for models trained with only channel coding, only adversarial training, both, and the identity model. For models trained with only adversarial training, the input message length to the watermarking model is 30 instead of 120.

4.3.2 Attack Complexity

We study the effect of varying the complexity and architecture of the attack network G_{adv} . On top of adjusting the network size and depth, we also consider two variants of the attack network: the *residual* network (Res) where we add a skip connection from the input, and a *capped* network (Capped) where we limit the maximum pixel difference by setting $G_{adv}(I) = I + \epsilon \tanh(\text{CNN}(I))$. We also report the results from the fast gradient sign method (FGSM) for completeness.

| | JPEG (Q=50) | Dropout (p=0.3) | Blur ($\sigma=1.0$) | Acc. (adv.) |
|----------------|----------------|--------------------|--------------------------|----------------|
| Conv (3,16) | 81.7 | 97.9 | 92.8 | 90.6 |
| Conv (3,32) | 80.5 | 98.0 | 84.9 | 78.7 |
| Conv (3,32,32) | 75.0 | 95.3 | 81.5 | 72.0 |
| Res (3,16) | 84.5 | 96.3 | 86.3 | 95.0 |
| Capped (0.03) | 57.3 | 93.9 | 77.3 | 96.5 |
| Capped (0.06) | 53.2 | 94.6 | 78.1 | 99.5 |
| FGSM | 50.1 | 86.2 | 50.1 | 98.0 |

Table 4: Performance when varying attack network complexity. Each row corresponds to models trained with a different configuration of attack network. The first three columns show the bit accuracy on various image distortions. The last column shows the bit accuracy on adversarial message X'_{adv} .

From Table 4, we observe that the bit accuracy on the adversarial example decreases as the attack network complexity increases, causing a slight degradation in the final result. Capping the attack network yielded poor results on JPEG and Gaussian blur, indicating that this approach over-restricts the attack network. The residual network yielded very similar performance to the regular convolutional model, slightly underperforming on Gaussian blur.




Figure 9: Samples of encoded and cover images for the watermarking algorithm. **First row:** Cover image with no embedded message. **Second row:** Encoded image from HiDDeN combined distortion model. **Third row:** Encoded images from our model. **Fourth row:** Normalized difference of the encoded image and cover image for the HiDDeN combined model. **Fifth row:** Normalized difference for our model.

5. Conclusion

We propose a distortion agnostic watermarking method that does not explicitly model the image distortion at training time. Our method consists of two core components, adversarial training and channel coding, to improve the robustness of our system. Compared with conventional methods of improving model robustness, our methods do not require the explicit modeling of the image distortions at training time. Through empirical evaluations, we validate that our model reaches comparable performance to the combined distortion model on distortions seen during training, and better generalization to unseen distortions. In future work, we would like to improve upon our current methodology to further increase model robustness, and explore deeper

the connections between watermarking and adversarial attacks.

References

- [1] Mahdi Ahmadi, Alireza Norouzi, Nader Karimi, Shadrokh Samavi, and Ali Emami. Redmark: Framework for residual diffusion watermarking based on deep networks. *Expert Systems with Applications*, 146:113157, 2020. [2](#), [3](#), [5](#)
- [2] Fayçal Ait Aoudia and Jakob Hoydis. Model-free training of end-to-end communication systems. *IEEE Journal on Selected Areas in Communications*, 37(11):2503–2516, 2019. [3](#)
- [3] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing

- defenses to adversarial examples. *International Conference on Machine Learning*, 2018. 3
- [4] Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. *AAAI Conference on Artificial Intelligence*, 2017. 3
- [5] Walter Bender, Daniel Gruhl, Norishige Morimoto, and Anthony Lu. Techniques for data hiding. *IBM systems journal*, 35(3&4):313–336, 1996. 1, 2
- [6] Martin Bossert. *Channel coding for telecommunications*. John Wiley & Sons, Inc., 1999. 3
- [7] Kristy Choi, Kedar Tatwawadi, Tsachy Weissman, and Stefano Ermon. Necst: Neural joint source-channel coding. *International Conference on Machine Learning*, 2018. 3, 5
- [8] Ingemar J Cox, Matthew L Miller, Jeffrey Adam Bloom, and Chris Honsinger. *Digital watermarking*, volume 53. Springer, 2002. 1, 2
- [9] Rick Fritschek, Rafael F Schaefer, and Gerhard Wunder. Deep learning for channel coding via neural mutual information estimation. In *2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications*, pages 1–5. IEEE, 2019. 3
- [10] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016. 3
- [11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations*, 2014. 3, 4
- [12] Mohamed Hamidi, Mohamed El Haziti, Hocine Cherifi, and Driss Aboutajdine. A blind robust image watermarking approach exploiting the dft magnitude. In *IEEE/ACS 12th International Conference of Computer Systems and Applications*, pages 1–6. IEEE, 2015. 2
- [13] Hossein Hosseini, Yize Chen, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. Blocking transferability of adversarial examples in black-box learning systems. *arXiv preprint arXiv:1703.04318*, 2017. 3
- [14] Mei Jiansheng, Li Sukang, and Tan Xiaomei. A digital watermarking algorithm based on DCT and DWT. In *The International Symposium on Web Information Systems and Applications*, page 104, 2009. 2
- [15] S Katzenbeisser and FAP Petitcolas. *Digital Watermarking*. Springer, 2000. 2
- [16] Deepa Kundur and Dimitrios Hatzinakos. Digital watermarking using multiresolution wavelet decomposition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 5, pages 2969–2972. IEEE, 1998. 1
- [17] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *International Conference on Learning Representations*, 2016. 3
- [18] Hyeungill Lee, Sungyeob Han, and Jungwoo Lee. Generative adversarial trainer: Defense to adversarial perturbations with gan. *arXiv preprint arXiv:1705.03387*, 2017. 3
- [19] Jinghui Li and Peiyu Fang. Hdrnet: Single-image-based hdr reconstruction using channel attention cnn. In *Proceedings of the International Conference on Multimedia Systems and Signal Processing*, pages 119–124. ACM, 2019. 3
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 5
- [21] Yang Liu, Mengxi Guo, Jian Zhang, Yuesheng Zhu, and Xiaodong Xie. A novel two-stage separable deep learning framework for practical blind watermarking. In *Proceedings of the ACM International Conference on Multimedia*, pages 1509–1517. ACM, 2019. 2
- [22] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *International Conference on Learning Representations*, 2018. 5
- [23] Seung-Min Mun, Seung-Hun Nam, Han-UI Jang, Dongkyu Kim, and Heung-Kyu Lee. A robust blind watermarking using convolutional neural network. *arXiv preprint arXiv:1704.03248*, 2017. 2, 3
- [24] Joseph JK O’Ruanaidh and Thierry Pun. Rotation, scale and translation invariant digital image watermarking. In *Proceedings of International Conference on Image Processing*, volume 1, pages 536–539. IEEE, 1997. 2
- [25] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the ACM on Asia Conference on Computer and Communications Security*, pages 506–519. ACM, 2017. 3
- [26] William E Ryan et al. *An introduction to LDPC codes*. CRC Press Boca Raton, FL, 2004. 3, 4
- [27] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *Conference on Neural Information Processing Systems*, 2019. 3
- [28] Claude Elwood Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948. 3
- [29] Richard Shin and Dawn Song. Jpeg-resistant adversarial images. In *Conference on Neural Information Processing Systems 2017 Workshop on Machine Learning and Computer Security*, 2017. 6
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2015. 4
- [31] Prabhishkek Singh and RS Chadha. A survey of digital watermarking techniques, applications and attacks. *International Journal of Engineering and Innovative Technology*, 2(9):165–175, 2013. 2
- [32] Bernard Sklar. A primer on turbo code concepts. *IEEE communications Magazine*, 35(12):94–102, 1997. 3
- [33] Martin Steinebach, Fabien AP Petitcolas, Frederic Raynal, Jana Dittmann, Caroline Fontaine, S Seibel, Nazim Fates, and Lucilla Croce Ferri. Stirmark benchmark: audio watermarking attacks. In *Proceedings International Conference on Information Technology: Coding and Computing*, pages 49–54. IEEE, 2001. 6, 7
- [34] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. In-

- triguing properties of neural networks. *International Conference on Learning Representations*, 2013. 3
- [35] Kiyoshi Tanaka, Yasuhiro Nakamura, and Kineo Matsui. Embedding secret information into a dithered multi-level image. pages 216–220. IEEE, 1990. 2
- [36] Matthew Tancik, Ben Mildenhall, and Ren Ng. Stegastamp: Invisible hyperlinks in physical photographs. *arXiv preprint arXiv:1904.05343*, 2019. 2
- [37] Peter Trifonov. Efficient design and decoding of polar codes. *IEEE Transactions on Communications*, 60(11):3221–3227, 2012. 3
- [38] Bingyang Wen and Sergul Aydore. Romark: A robust watermarking system using adversarial training. *arXiv preprint arXiv:1910.01221*, 2019. 3
- [39] 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*, pages 87–95. ACM, 2019. 2
- [40] Eric Wengrowski and Kristin Dana. Light field messaging with deep photographic steganography. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1515–1524, 2019. 2
- [41] Stephen B Wicker and Vijay K Bhargava. *Reed-Solomon codes and their applications*. John Wiley & Sons, 1999. 3
- [42] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. *arXiv preprint arXiv:1711.01991*, 2017. 3
- [43] Kevin Alex Zhang, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Steganogan: Pushing the limits of image steganography. *arXiv preprint arXiv:1901.03892*, 2019. 2
- [44] Kevin Alex Zhang, Lei Xu, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Robust invisible video watermarking with attention. *arXiv preprint arXiv:1909.01285*, 2019. 2
- [45] 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)*, pages 657–672, 2018. 1, 2, 4, 5