
On Perceptual Lossy Compression: The Cost of Perceptual Reconstruction and An Optimal Training Framework

Zeyu Yan¹ Fei Wen¹ Rendong Ying¹ Chao Ma¹ Peilin Liu¹

Abstract

Lossy compression algorithms are typically designed to achieve the lowest possible distortion at a given bit rate. However, recent studies show that pursuing high perceptual quality would lead to increase of the lowest achievable distortion (e.g., MSE). This paper provides nontrivial results theoretically revealing that, 1) the cost of achieving perfect perception quality is exactly a doubling of the lowest achievable MSE distortion, 2) an optimal encoder for the “classic” rate-distortion problem is also optimal for the perceptual compression problem, 3) distortion loss is unnecessary for training a perceptual decoder. Further, we propose a novel training framework to achieve the lowest MSE distortion under perfect perception constraint at a given bit rate. This framework uses a GAN with discriminator conditioned on an MSE-optimized encoder, which is superior over the traditional framework using distortion plus adversarial loss. Experiments are provided to verify the theoretical finding and demonstrate the superiority of the proposed training framework.

1. Introduction

Lossy compression is a fundamental problem in modern digital world for efficient transmission and storage of image, video and audio data. Recently, due in part to the roaring success of deep learning, the research of deep neural networks (DNNs) based compression has attracted much attention and shown promising results in image, audio, and video compression (Rippel & Bourdev, 2017; M Tschannen, 2018; Agustsson et al., 2019; Toderici et al., 2016; Toderici et al., 2017; Ballé et al., 2016; Ballé et al., 2017; 2018; Agustsson et al., 2017; Minnen et al., 2018; Li et al., 2018; Mentzer et al., 2018; Johnston et al., 2018; Galteri et al.,

2017; Santurkar et al., 2018; Shaham & Michaeli, 2018). For lossy compression, Shannon’s rate-distortion theory is a theoretical cornerstone, which characterizes the tradeoff between the bit rate of compressed representation and the distortion in reconstructing the data (Shannon, 1959; Cover & Thomas, 2006).

Typically, lossy compression algorithms are designed based on the rate-distortion theory to achieve the lowest possible distortion at a given bit rate. In this context, lower distortion is desired, e.g., lower mean square error (MSE), higher peak signal to noise ratio (PSNR), or higher structural similarity (SSIM) (Zhou Wang et al., 2004). However, recent studies have demonstrated that these distortion measures are not fully consistent with human’s perception (Johnson et al., 2016; Zhang et al., 2018; Blau & Michaeli, 2018; 2019; Agustsson et al., 2019; Galteri et al., 2017; Santurkar et al., 2018; Ma et al., 2017). Specifically, minimizing distortion alone does not necessarily lead to good perceptual quality. In fact, it has been shown that pursuing high perceptual quality would lead to increase of the lowest achievable distortion (e.g., MSE). Generally, there exist two methods to improve perceptual quality. The first is to incorporate a perceptual loss measuring the difference between deep features (Johnson et al., 2016; Simonyan & Zisserman, 2015; Gatys et al., 2016; Chen & Koltun, 2017; Dosovitskiy & Brox, 2016). The second is to incorporate an adversarial loss by using generative adversarial networks (GAN) (Goodfellow et al., 2014). Noteworthily, using an adversarial loss, the second method has shown remarkable effectiveness in achieving high perceptual quality (Rippel & Bourdev, 2017; Agustsson et al., 2019; Ledig et al., 2017; Wang et al., 2018; Wu et al., 2020; Iwai et al., 2020; Mentzer et al., 2020).

For perceptual reconstruction, new insights from recent studies have revealed that distortion and perceptual quality are at odds with each other. It can be well characterized by a perception-distortion tradeoff (Blau & Michaeli, 2018), in which perceptual quality is defined in terms of the deviation between the distributions of the source and the reconstructed data. More recently, this result has been extended to the lossy compression problem, resulting in a rate-distortion-perception tradeoff (Matsumoto, 2018a;b; Blau & Michaeli, 2019). The three-way tradeoff indicates that imposing a

¹School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University, Shanghai, China. Correspondence to: Fei Wen <wenfei@sjtu.edu.cn>.

high perceptual quality constraint on the lossy compression problem would lead to an elevation of the rate-distortion curve. Hence, a sacrifice in either rate or distortion is necessary to achieve high perceptual quality. The works (Blau & Michaeli, 2018; 2019) have paved the way for understanding the perception-distortion tradeoff.

While it has become increasingly accepted that high perceptual quality can be achieved with some increase of the lowest achievable distortion, there lacks quantitative analysis on such an increase of distortion. This work is motivated by the following two open questions:

1. Is it possible to quantitatively characterize the effect of perceptual constraint on the rate-distortion tradeoff?
2. How to build a framework for perfect perception reconstruction in lossy compression?

Toward answering these two questions, the main contributions of this work are as follows.

First, we derive a nontrivial result theoretically revealing that, at a given bit rate, the cost of achieving perfect perceptual quality is exactly a doubling of the lowest achievable MSE distortion. Our analysis also shows that an encoder optimized in terms of MSE is also an optimal encoder under perfect perception constraint. This result implies that the commonly used adversarial loss in state-of-the-art works (Blau & Michaeli, 2019; Agustsson et al., 2019; Iwai et al., 2020) is in fact unnecessary for optimizing the encoder.

Second, based on the analysis, we propose a training framework that can achieve the lowest MSE distortion under perfect perception constraint at a given bit rate. This framework uses a GAN with its discriminator conditioned on an encoder optimized in terms of MSE, which is superior over the traditional distortion plus adversarial loss (DAL) based framework.

Finally, experimental results on the MNIST dataset are provided to verify the theoretical finding and demonstrate the superiority of the proposed training framework. Particularly, extensive experimental results show that the effect of perception constraint on the rate-distortion tradeoff accords well with our theoretical result, i.e., a doubling of the MSE distortion in achieving high perceptual quality.

Though our result is restricted to the lossy compression problem, to the best of our knowledge, it is the first quantitative result on the lowest degradation of distortion in achieving perceptual reconstruction. Note that the work (Blau & Michaeli, 2019) has shown that perfect perceptual quality can be attained at a sacrifice of no more than 2-fold increase in MSE distortion. Our analysis is fundamentally different from (Blau & Michaeli, 2019) in three aspects. 1) We obtain a different result deterministically shows that the increase is exactly 2-fold. 2) The result in (Blau & Michaeli, 2019)

is derived through analyzing a constructed encoder-decoder pair that achieves perfect perceptual quality by concatenating a post-processing perceptual mapping. In contrast, our analysis follows a different line via directly analyzing the perception constrained lossy compression formulation itself. 3) Our analysis also provides a new insight that, to achieve high perceptual quality, perceptual loss is unnecessary for training an encoder. Besides, distortion loss, either on pixels or deep features, is unnecessary for training a decoder. We show that an adversarial loss, with the discriminator conditioned on an encoder optimized in terms of MSE distortion, is enough to achieve perfect perceptual quality. A more detailed comparison is provided in Section 4.2.

2. Preliminaries

2.1. Perceptual Quality

In image processing, peak signal to noise rate (PSNR), MSE and SSIM/MS-SSIM (Wang et al., 2003; Zhou Wang et al., 2004) are commonly used distortion measures. However, recent studies have shown that these measures are not fully consistent with human’s perception (Johnson et al., 2016; Zhang et al., 2018; Agustsson et al., 2019; Galteri et al., 2017; Santurkar et al., 2018). An image with lower distortion does not necessarily have better perceptual quality. In practice, it has been empirically shown that high perceptual property can be achieved at the cost of increased distortion. This behavior can be theoretically characterized as a perception-distortion tradeoff, which has been put forward in (Blau & Michaeli, 2018). It reveals that minimizing distortion would cause the distribution of reconstructed outputs deviating from that of the (ground-truth) source, which leads to worse perceptual quality.

The perceptual quality of a restored sample is the extent to which it looks like a natural sample from human’s perception, regardless its similarity to any reference sample. It can be conveniently defined in terms of the deviation of the distribution of restored outputs \hat{X} from that of natural samples X as (Blau & Michaeli, 2018)

$$d(p_X, p_{\hat{X}}), \tag{1}$$

which is some divergence (e.g., the Kullback-Leibler divergence or Wasserstein distance) satisfying $d(p, q) \geq 0$ and $d(p, q) = 0 \Leftrightarrow p = q$ for any distributions p and q . Such a definition conforms with the common practice of quantifying perceptual quality through real-versus-fake questionnaire studies (Zhang et al., 2018; Zhang et al., 2016; Salimans et al., 2016). Basically, this definition of perceptual quality is based on the deviation from natural sample statistics, which correlates well with human subjective score and has been widely used in designing no-reference image quality measures. (Mittal et al., 2013; Wang & Simoncelli, 2005)

2.2. Rate-Distortion-Perception Tradeoff

Shannon’s rate-distortion theory characterizes the fundamental tradeoff between the rate (bits per symbol) and the expected distortion. Specifically, the relation between the input $X \sim p_X$ of the encoder and the output $\hat{X} \sim p_{\hat{X}}$ of the decoder can be viewed as a mapping defined by a conditional distribution $p_{\hat{X}|X}$. The information rate-distortion function is defined as (Cover & Thomas, 2006)

$$\begin{aligned} R^{(I)}(D) &= \min_{p_{\hat{X}|X}} I(X; \hat{X}) \\ \text{s.t. } \mathbb{E}[\Delta(X, \hat{X})] &\leq D, \end{aligned} \quad (2)$$

where I stands for mutual information, and Δ is a distortion measure such as MSE or hamming distance. It has been well established that $R^{(I)}(D)$ is a convex and non-increasing function of D , which demonstrates the rate-distortion tradeoff.

The rate-distortion function (2) does not take into account the perceptual quality of reconstruction. To take the perceptual quality into consideration, the traditional rate-distortion tradeoff has been extended to a three-way tradeoff model including rate, distortion, and perception as (Blau & Michaeli, 2019)

$$\begin{aligned} R^{(I)}(D, P) &= \min_{p_{\hat{X}|X}} I(X; \hat{X}) \\ \text{s.t. } \mathbb{E}[\Delta(X, \hat{X})] &\leq D, d(p_X, p_{\hat{X}}) \leq P, \end{aligned} \quad (3)$$

where $d(p_X, p_{\hat{X}})$ is a divergence between the distributions of the source and the reconstruction output, typically measured by Kullback-Leibler divergence or Jensen-Shannon divergence. When $P = +\infty$, the perception constraint is invalid and (3) degenerates to the traditional rate-distortion function (2). When $P = 0$, the distributions of X and \hat{X} are constrained to be identical, which defines a rate-distortion function under perfect perceptual quality constraint. It has been shown in the work (Blau & Michaeli, 2019) that, the rate-distortion curve necessarily elevates when imposing the perfect perceptual quality constraint, but the elevation is bounded. Specifically, for the MSE distortion, the rate-distortion function under perfect perception constraint, i.e., $R^{(I)}(D, 0)$, is upper bounded by a scaled version of the rate-distortion curve under no perception constraint as

$$R^{(I)}(D, 0) \leq R^{(I)}\left(\frac{1}{2}D, +\infty\right). \quad (4)$$

It implies that, at a given bit rate, the increase of the MSE distortion incurred by perfect perceptual quality constraint is no more than 2-fold the MSE distortion in the case without any perception constraint. However, how much is the elevation of the rate-distortion curve incurred by the perfect perceptual quality constraint is still unclear.

2.3. Adversarial Loss for Perceptual Reconstruction

To improve the perceptual quality of reconstruction, a natural way is to minimize the deviation from the distribution of natural samples. As GAN is very effective in aligning distributions, the GAN-based methods have shown remarkable improvement in perceptual quality (Zhang et al., 2018; Agustsson et al., 2019). For the lossy compression problem, a typical formulation incorporating an adversarial loss to optimize the encoder E and the decoder G through adversarial training is given by

$$\min_{E, G} \mathcal{L}_{rec} + \lambda \mathcal{L}_{adv} + \beta H, \quad (5)$$

where \mathcal{L}_{rec} is a distortion loss (e.g., MSE, L_1 norm, or distance between feature maps), \mathcal{L}_{adv} is an adversarial (perception) loss which is measured by a discriminator, and H is the entropy of the compressed representation. λ and β are positive parameters which balance the three terms. Although adversarial training helps to improve the perceptual quality, such a training framework has limitations. Specifically, if the value of λ is not large enough, the resulting perceptual quality would not be satisfactory. On the other hand, when λ is a relatively large value, the distortion of the system cannot be well optimized and would finally result in excessive increase in distortion, as will be shown later in experiments. Hence, it is difficult to achieve the lowest distortion under perfect perception constraint by simply balancing the distortion loss \mathcal{L}_{rec} and the adversarial loss \mathcal{L}_{adv} .

To address these limitations, we propose a training framework that can achieve the lowest distortion under perfect perception constraint. It is also based on GAN but avoids the balance between the distortion and adversarial losses.

3. Main Results and Proposed Training Framework

3.1. Analysis on the Rate-Distortion Tradeoff Under Perfect Perception Constraint

Suppose that X is a discrete source with a finite alphabet $\chi = \{x_i \in \mathbb{R}^N : 1 \leq i \leq m\}$, e.g., $m = 256^{3N}$ for 8-bit RGB image with N pixels. From (3), the information rate-distortion function under perfect perception constraint is

$$\begin{aligned} R^{(I)}(D, P) &= \min_{p_{\hat{X}|X}} I(X; \hat{X}) \\ \text{s.t. } \mathbb{E}[\Delta(X, \hat{X})] &\leq D, d(p_X, p_{\hat{X}}) \leq 0. \end{aligned} \quad (6)$$

Under the constraint that the distributions of X and \hat{X} are identical, χ is also the alphabet of \hat{X} .

Before proceeding to the analysis, we present some definitions. Define a joint distribution matrix $\mathbf{B} \in \mathbb{R}^{m \times m}$, of

which the elements are given by

$$b_{ij} := p(X = x_i, \hat{X} = x_j), 1 \leq i, j \leq m. \quad (7)$$

Meanwhile, define a distortion matrix $\mathbf{W} \in \mathbb{R}^{m \times m}$, of which the elements are given by

$$w_{ij} := \Delta(x_i, x_j), 1 \leq i, j \leq m. \quad (8)$$

Then the distortion constraint in (6) can be rewritten as $\langle \mathbf{W}, \mathbf{B} \rangle \leq D$. Moreover, let $G_{p_X}(\mathbf{B})$ denote the objective function of (6), then under perfect perception constraint, it follows that

$$\begin{aligned} G_{p_X}(\mathbf{B}) &:= I(X; \hat{X}) \\ &= H(X) + H(\hat{X}) - H(X, \hat{X}) \\ &= 2H(X) + \sum_{i=1}^m \sum_{j=1}^m b_{ij} \log b_{ij}, \end{aligned} \quad (9)$$

where the last equality follows from the fact that X and \hat{X} have the same distribution under the constraint $d(p_X, p_{\hat{X}}) \leq 0$. For fixed p_X , $p_{\hat{X}|X}$ can be equivalently represented by \mathbf{B} . Thus, the formulation (6) can be rewritten as

$$\begin{aligned} R^{(I)}(D, 0) &= \min_{\mathbf{B}} G_{p_X}(\mathbf{B}) \\ \text{s.t. } &\langle \mathbf{W}, \mathbf{B} \rangle \leq D, \\ &\sum_{i=1}^m b_{ij} = \sum_{i=1}^m b_{ji} = p(X = x_j), 1 \leq j \leq m. \end{aligned} \quad (10)$$

Note that the perception constraint $d(p_X, p_{\hat{X}}) \leq 0$ in (6), which means $p_X = p_{\hat{X}}$, has been rewritten as the equivalence constraint between the row summation and column summation of \mathbf{B} , which are the distributions of X and \hat{X} , respectively. Based on (10), we have the following result (proof is given in the supplementary material).

Lemma 1. *Suppose that Δ is symmetric, then any optimal solution \mathbf{B}^* to (10) is a symmetric matrix.*

Lemma 1 shows that when the distortion measure is symmetric (e.g., squared-error), for any optimal encoder-decoder pair to (10) (equivalently (6)), the joint distribution of X and \hat{X} is also symmetric as

$$p(X = x_i, \hat{X} = x_j) = p(X = x_j, \hat{X} = x_i). \quad (11)$$

Since the achievability of the information rate-distortion-perception function $R^{(I)}(D, P)$ has not been proved yet, we analyze the relation between the rate-distortion functions $R(D, \infty)$ and $R(D, 0)$. Assume a discrete memoryless stationary source X , let $Y = (X_1, X_2, \dots, X_t)$ be a source sequence with length t . Consider a compression task $Y \rightarrow Z \rightarrow \hat{Y}$, where Z is the output of the encoder and \hat{Y} is the output of the decoder. As defined in (Cover &

Thomas, 2006), $R(D, \infty)$ is the infimum of rates R for a given distortion D , such that (R, D) is achievable. Hence, $R(D, \infty)$ can be expressed as

$$R(D, \infty) = \inf_t F_t(D, \infty), \quad (12)$$

where

$$\begin{aligned} F_t(D, \infty) &:= \min_{p_{Z|Y}, p_{\hat{Y}|Z}} \frac{1}{t} H(Z) \\ \text{s.t. } &\frac{1}{t} \mathbb{E}[\Delta(Y, \hat{Y})] \leq D. \end{aligned} \quad (13)$$

Considering the condition of perfect perception constraint, $R(D, 0)$ can be similarly expressed as (Matsumoto, 2018b)

$$R(D, 0) = \inf_t F_t(D, 0), \quad (14)$$

where

$$\begin{aligned} F_t(D, 0) &:= \min_{p_{Z|Y}, p_{\hat{Y}|Z}} \frac{1}{t} H(Z) \\ \text{s.t. } &\frac{1}{t} \mathbb{E}[\Delta(Y, \hat{Y})] \leq D, d(p_Y, p_{\hat{Y}}) \leq 0. \end{aligned} \quad (15)$$

For the MSE distortion measure, we have the following result (proof is given in the supplementary material).

Theorem 1. *Let X be a discrete memoryless stationary source and $Y = (X_1, X_2, \dots, X_t)$ be a source sequence with length t . For a compression task $Y \rightarrow Z \rightarrow \hat{Y}$, where Z and \hat{Y} are the outputs of the encoder and decoder, respectively. Then, to any t , if Δ is the squared-error distortion, there exists an optimal encoder-decoder pair to (15) satisfying*

$$p_{Y|Z} = p_{\hat{Y}|Z}. \quad (16)$$

Theorem 1 indicates that when the distortion measure is MSE, replacing the perfect perception constraint $d(p_X, p_{\hat{X}}) \leq 0$ in (15) with the conditional distribution constraint (16) does not change the optimal objective value. This is because under the condition that the data source is memoryless and stationary, it follows from theorem 1 that there exists an optimal solution to (15) satisfying (16), under which the distributions of X and \hat{X} are the same. Therefore, analyzing the relation between $F_t(D, \infty)$ and $F_t(D, 0)$ is equivalent to analyzing the relation between $F_t(D, \infty)$ and

$$\begin{aligned} F_t(D, 0) &= \min_{p_{Z|Y}, p_{\hat{Y}|Z}} \frac{1}{t} H(Z) \\ \text{s.t. } &\frac{1}{t} \mathbb{E}[\Delta(Y, \hat{Y})] \leq D, p_{Y|Z} = p_{\hat{Y}|Z}. \end{aligned} \quad (17)$$

For any encoder-decoder pair satisfying (16), the MSE distortion between Y and \hat{Y} can be expressed as (proof is given in the supplementary material)

$$\frac{1}{t} \mathbb{E} \left[\left\| Y - \hat{Y} \right\|^2 \right] = \frac{2}{t} \mathbb{E} \left[\left\| Y - \mathbb{E}[Y|Z] \right\|^2 \middle| Z \right], \quad (18)$$

which substituted into (17) yields

$$F_t(D, 0) = \min_{p_{Z|Y}, p_{\hat{Y}|Z}} \frac{1}{t} H(Z) \quad (19)$$

s.t. $\frac{2}{t} \mathbb{E} \left[\|Y - \mathbb{E}[Y|Z]\|^2 | Z \right] \leq D, p_{Y|Z} = p_{\hat{Y}|Z}.$

It can be seen that the objective value of (19) is independent on the decoder $p_{\hat{Y}|Z}$. Hence, with MSE distortion, the encoder of (17) can be optimized separately as

$$F_t(D, 0) = \min_{p_{Z|Y}} \frac{1}{t} H(Z) \quad (20)$$

s.t. $\frac{2}{t} \mathbb{E} \left[\|Y - \mathbb{E}[Y|Z]\|^2 | Z \right] \leq D.$

In (20), the constraint $p_{Y|Z} = p_{\hat{Y}|Z}$ is removed since it does not affect the objective under the constraint $\frac{2}{t} \mathbb{E} \left[\|Y - \mathbb{E}[Y|Z]\|^2 | Z \right] \leq D.$

Next, we consider the lossy compression task without perception constraint. Given a compressed representation z and when the distortion measure is MSE, the output of an optimal decoder is $\mathbb{E}[Y|z]$. Accordingly, the expected distortion can be expressed as

$$\mathbb{E}[\Delta(Y, \mathbb{E}[Y|Z])] = \frac{1}{t} \mathbb{E} \left[\|Y - \mathbb{E}[Y|Z]\|^2 | Z \right], \quad (21)$$

with which formulation (13) can be rewritten as

$$F_t(D, \infty) = \min_{p_{Z|Y}} \frac{1}{t} H(Z) \quad (22)$$

s.t. $\frac{1}{t} \mathbb{E} \left[\|Y - \mathbb{E}[Y|Z]\|^2 | Z \right] \leq D.$

Then, it follows from (20) and (22) that $F_t(D, 0) = F_t(D/2, \infty)$ for any t , which implies

$$\inf_t F_t(D, 0) = \inf_t F_t(D/2, \infty). \quad (23)$$

Thus, we have the following result.

Theorem 2. *Suppose that Δ is the squared-error distortion, then $R(D, P)$ satisfies*

$$R(D, 0) = R\left(\frac{1}{2}D, +\infty\right). \quad (24)$$

Theorem 2 indicates that, when the distortion is measured by MSE, and for fixed bit rate, the lowest achievable distortion under perfect perception constraint is 2-fold that under no perception constraint. In other words, the cost of attaining perfect perceptual quality is exactly a doubling of the lowest achievable MSE distortion.

Moreover, (20) and (22) also imply that, with distortion-constrained by $D/2$, an optimal encoder under no perception

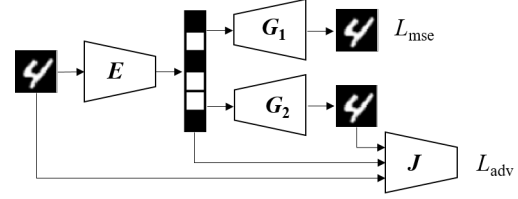


Figure 1. **The proposed training framework.** First, the encoder-decoder pair (E, G_1) is trained by MSE loss. Then G_2 is trained by adversarial loss conditioned on E .

constraint is also an optimal encoder under perfect perception constraint with distortion constrained by D .

Theorem 3. *When the distortion is measured by MSE, an optimal encoder to (13) (without perception constraint) with distortion constrained by $D/2$ is also an optimal encoder to (15) (with perfect perception constraint) with distortion constrained by D .*

Theorem 3 implies that the rate-distortion function with perfect perception constraint is achievable. Besides, it sheds some light on how to optimize an encoder-decoder pair under perfect perception constraint. Specifically, the encoder can be independently optimized without considering perception constraint. In light of this, we can first optimize an encoder with only MSE constraint, and then fix the encoder and optimize the decoder to satisfy (16), which would double the MSE distortion and has been proven to be an optimal decoder that can achieve the bound in Theorem 2. Such a framework is detailed in the next subsection.

3.2. A Framework for Perfect Perceptual Reconstruction in Lossy Compression

Based on the above results, we further propose a framework for training an encoder-decoder pair to achieve perfect perceptual reconstruction in lossy compression. The overall architecture is shown in Figure 1, which includes an encoder E , two decoders G_1, G_2 , and a discriminator J . The desired encoder-decoder pair is (E, G_2) , which is obtained by a training procedure with two steps:

- i) *Encoder optimization:* (E, G_1) are optimized only in terms of MSE distortion, e.g. by minimizing MSE loss without considering the perception constraint. From Theorem 3, such an optimized encoder E is also an optimal encoder under perfect perception constraint.
- ii) *Decoder optimization:* Fixing the optimized encoder E in the first step, the decoder G_2 is optimized via adversarial training, e.g. iteratively optimize G_2 and the discriminator J . In this step, the goal is to minimize the divergence between $p_{Y|Z}$ and $p_{\hat{Y}|Z}$. To achieve this, we

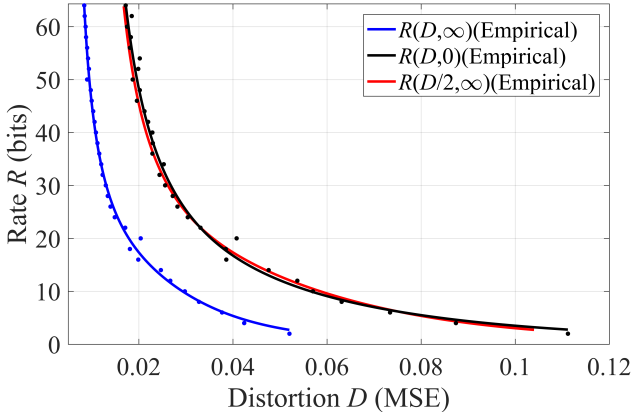


Figure 2. Empirically fitted rate-distortion functions in the cases with or without perception constraint.

use a conditional discriminator (Mirza & Osindero, 2014) which takes both the data (Y or \hat{Y}) and the bit stream Z as the input.

Our proposed framework is different from traditional DAL framework as follows. As discussed in Section 2.3, DAL involves a balance between the distortion and adversarial losses, and is difficult to achieve the lowest distortion under perfect perception constraint. Unlike the discriminator in DAL method being used to discriminate whether its input is real, the discriminator in our framework not only discriminate whether input data is real but also discriminate whether the input data is consistent with the bit stream Z . By this, though without distortion constraint, G_2 is trained to decode data consistent with the input rather than randomly generate a realistic output. This results in an advantage that the proposed framework avoids the balance between the distortion and adversarial losses in DAL, and can achieve the lowest MSE distortion under perfect perception constraint.

To optimize the decoder to satisfy (16), we add a random noise input in the decoder G_2 . This is a commonly used trick to achieve high perceptual quality (Blau & Michaeli, 2019; Agustsson et al., 2019), In our framework, (16) holds when (G_2, J) are optimal, which means the output of the decoder would have perfect perceptual quality.

4. Experimental Illustration

We demonstrate the theoretical finding and the effectiveness of the proposed framework on the MNIST dataset (Lecun et al., 1998). Note that the rate-distortion-perception function does not have an analytical expression for a general data distribution, here we empirically demonstrate the 2-fold relation in Theorem 2 by experiment. As the compression task on the MNIST data is relatively simple, and the pro-

posed framework can theoretically achieve the lowest MSE distortion with perfect perceptual quality as discussed in Section 3.2, an implementation using DNN can be expected to closely approach the rate-distortion curve under perception constraint. In this context, we can compare the empirical rate-distortion curve of the proposed framework with that under no perception constraint to demonstrate the derived 2-fold relation. Moreover, we demonstrate the superiority of the proposed framework by comparing it with the DAL method. The code is available online¹.

In the proposed framework, E , G_1 , G_2 and J are convolutional neural networks (CNN). The encoder E maps a 32×32 image into a $d \times 1$ vector with each element be a quantized integer. To preserve gradient for backpropagation, the elements are quantized by

$$\hat{z} = z + \text{sg}[Q(z) - z], \quad (25)$$

where $Q(\cdot)$ is a quantization operator and $\text{sg}[\cdot]$ is the stop-gradient operator (van den Oord et al., 2017; Razavi et al., 2019). The bit rate is controlled by limiting the dimension and quantization level of encoder output. The output is quantized into binary values and correspondingly the bit rate is bounded by d .

Our model is trained by two steps as presented in Section 3.2, in which (E, G_1) is firstly optimized as

$$\min_{E, G_1} \mathbb{E} \|X - G_1(E(X))\|^2. \quad (26)$$

Then, with E fixed, (G_2, J) is optimized as

$$\min_{J \in \mathcal{F}} \max_{G_2} \mathbb{E}[J(G_2(E(X)), E(X))] - \mathbb{E}[J(X, E(X))],$$

where \mathcal{F} denotes the bounded 1-Lipschitz functions. G_2 is trained to generate images of which the distribution is the same as that of X . The input of G_2 is a 100×1 vector, in which the first d ($d < 100$) elements are the output of E and the other elements are Gaussian noise. When $d = 0$, (G_2, J) degenerates to a pure generative adversarial network. Extra experiments show that, with different dimensions of the noise vector, the MSE curves converge to almost the same value. To stabilize the process of adversarial training, WGAN-gp (Arjovsky et al., 2017; Ishaan et al., 2017) is employed, in which a gradient penalty is added into the loss function of J as

$$\mathbb{E}[J(G_2(E(X)), E(X))] - \mathbb{E}[J(X, E(X))] + \lambda_{gp} L_{gp},$$

where L_{gp} is the gradient penalty term and λ_{gp} is a parameter set to 10. We use a pre-training scheme to address the degeneration problem in adversarial training, which is detailed in the supplementary material.

¹<https://github.com/ZeyuYan/Perceptual-Lossy-Compression>

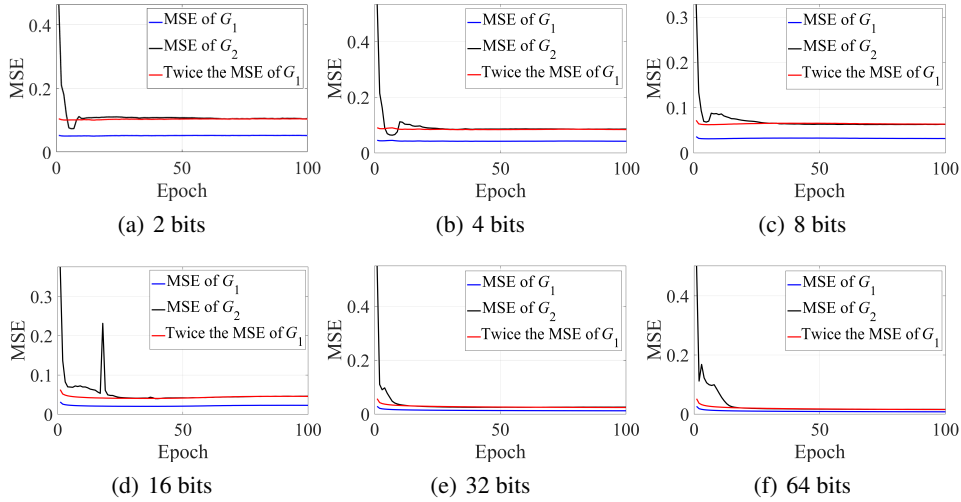


Figure 3. MSE loss versus training epoch for different bit-rates.

4.1. Rate-Distortion Tradeoff with or without Perception Constraint

For 32 different bit rates, $R \in \{2, 4, 6, \dots, 64\}$, we train 32 encoder-decoder pairs (E, G_1) by minimizing MSE-only loss and another 32 pairs (E, G_2) by our proposed framework. Figure 2 shows the fitted empirical rate-distortion curves of the encoder-decoder pairs. In Figure 2, $R(D, \infty)$ and $R(D, 0)$ are the fitting results, whilst $R(D/2, \infty)$ is a scaled version of the fitted $R(D, \infty)$ curve. Clearly, the empirical rate-distortion curves are monotonically non-increasing and convex, which is consistent with the theoretical properties of $R(D, P)$ (Blau & Michaeli, 2019). Interestingly, $R(D/2, \infty)$ closely approaches $R(D, 0)$, which empirically demonstrates the result given by Theorem 2, i.e., the 2-fold relation between the lowest achievable MSE distortion under perfect perception constraint and that under no perception constraint.

Figure 3 presents the MSE loss in training G_1 and G_2 versus training epoch for six rate cases $R \in \{2, 4, 8, 16, 32, 64\}$. For each case, the curve which is twice the amplitude of the MSE of G_1 is also plotted for reference. It can be seen that, without perception constraint, the MSE loss of G_1 converges rapidly in all bit-rate cases. In comparison, the MSE loss of G_2 converges slower, which is due to the fact that G_2 is jointly trained with J in an alternating manner. Although G_2 is trained only using an adversarial loss and without using MSE loss, it closely converges to the 2-fold MSE curve of G_1 in each case, which accords well with Theorem 2.

Figure 4 shows some samples of the reconstructed images by G_1 and G_2 under different bit rates. Obviously, the output of G_2 is clearer and sharper than that of G_1 . As the

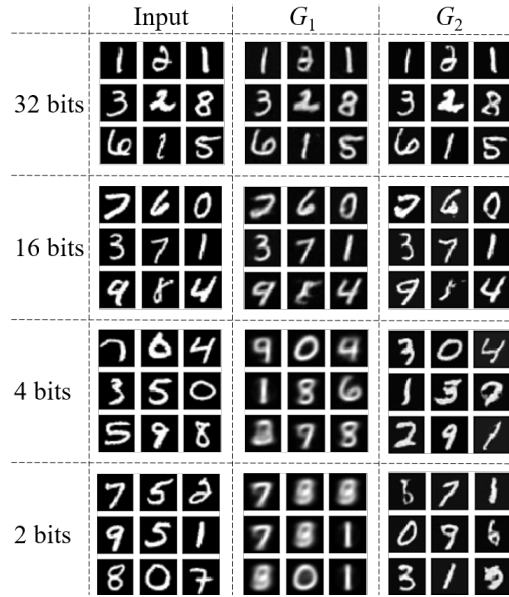


Figure 4. Visual comparison between the two cases with or without perception constraint.

decrease of bit rate, the output of G_1 becomes more blurry and unrecognizable, especially when bit rate is less than 4, which is due to the information lost in the encoding step. In comparison, the output of G_2 does not suffer from such deteriorating problem, which is clear and recognizable even when the bit rate is 2. However, for fixed bit rate, better perceptual quality would lead to larger distortion. Hence clear output images do not necessarily mean correct numbers. In principle, G_2 additionally takes a noise input to generate details, which would change the typeface, inclination,



















Input	G_1	Distortion plus adversarial loss			G_2
		$\lambda=0.1$	$\lambda=1$	$\lambda=5$	
	MSE: 0.0424	MSE: 0.0594	MSE: 0.0891	MSE: 0.1140	MSE: 0.0874
  	  	  	  	  	  

Figure 5. Performance comparison between the traditional distortion plus adversarial loss based method and the proposed method.

thickness of the original input and may even reconstruct a completely different but clear number.

Next, we compare our method with the DAL method, of which the parameter λ balances between the MSE distortion loss and the adversarial loss. Based on DAL for different values of λ , we train a number of encoder-decoder pairs for bit rate $R = 4$. Figure 5 compares the results of G_1 , G_2 and the DAL method with selected $\lambda \in \{0.1, 1, 5\}$, including the MSE and output image samples. It can be seen that, for the DAL method, as the increase of λ , the perceptual quality improves but in the meantime the distortion increases. Our method has better perceptual quality than the DAL method even when the MSE of DAL is larger than that of G_2 . As discussed in Section 2.3 and 3.2, while our method can achieve the lowest distortion under perfect perception constraint, the DAL method cannot.

4.2. Related Works

Our work is closely related to (Blau & Michaeli, 2019) but different from it in the following aspects. First, the analysis in (Blau & Michaeli, 2019) is based on constructing an encoder-decoder pair that can attain perfect perception quality, e.g., by concatenating a post-processing perceptual mapping after an optimal encoder-decoder pair under MSE measure. Since the constructed compression system can attain perfect perception quality with a doubling of the MSE distortion, in theory the lowest achievable MSE distortion under perfect perception constraint should be no more than twice that under no perception constraint. In contrast, we derive the rate-distortion bound under perfect perception constraint through a completely different line of analysis, via analyzing the lossy compression formulation itself as presented in Section 3.1. Second, our result deterministically shows that the lowest achievable distortion under perfect perception constraint is exactly 2-fold that under no perception constraint. Third, we show that, to achieve the lowest MSE distortion under perfect perception constraint, perceptual loss is unnecessary for training the encoder whilst distortion loss, either on pixel or deep features, is unnecessary for training the decoder. Fourth, we propose a training

framework that can achieve the lowest MSE distortion under perfect perception constraint at a given bit rate.

Like (Zhang et al., 2018; Ledig et al., 2017), our method is also based on GAN. However, unlike the formers using a distortion plus adversarial loss to train the encoder and decoder in an end-to-end manner, our method first uses a MSE loss to train the encoder and, then, uses an adversarial loss to train the decoder conditioned on the so obtained encoder. The superiority of the new method is discussed in Section 2.3 and 3.2. Recently, a conditioned discriminator similar to that in our framework has been employed in (Mentzer et al., 2020) to achieve high-fidelity image compression. Our work provides a theoretical foundation of the conditioned discriminator. The work (Mentzer et al., 2020) still uses a distortion plus adversarial loss, while we show that the decoder can be trained to achieve perfect perceptual reconstruction using adversarial-only loss and adversarial loss is unnecessary when training encoder.

We leave the application of the proposed framework to color image compression to future work, as it has a high requirement on computation hardware and poses a challenge in tuning a big GAN model.

5. Conclusion

We analyzed the effect of perception constraint on the rate-distortion function in lossy compression. We proved that, for fixed bit rate, the cost of imposing a perfect perception constraint is exactly a doubling of the lowest achievable MSE. The analysis also provided new insights on how to build a training framework for perfect perception reconstruction in lossy compression. Accordingly, we proposed a framework for training an encoder-decoder pair to achieve the lowest MSE under perfect perception constraint. Experimental results well verified the theoretical finding and demonstrated the superiority of the new framework over the traditional distortion plus adversarial loss based framework.

References

- Agustsson, E., Mentzer, F., Tschannen, M., Cavigelli, L., Timofte, R., Benini, L., and Gool, L. V. Soft-to-hard vector quantization for end-to-end learning compressible representations. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 1141–1151, 2017.
- Agustsson, E., Tschannen, M., Mentzer, F., Timofte, R., and Van Gool, L. Generative adversarial networks for extreme learned image compression. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 221–231, 2019.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- Ballé, J., Laparra, V., and Simoncelli, E. P. End-to-end optimized image compression. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Ballé, J., Minnen, D., Singh, S., Hwang, S. J., and Johnston, N. Variational image compression with a scale hyperprior. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Ballé, J., Laparra, V., and Simoncelli, E. P. End-to-end optimization of nonlinear transform codes for perceptual quality. In *Proceedings of the Picture Coding Symposium (PCS)*, pp. 1–5, 2016.
- Blau, Y. and Michaeli, T. The perception-distortion tradeoff. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6228–6237, 2018.
- Blau, Y. and Michaeli, T. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 675–685, 2019.
- Chen, Q. and Koltun, V. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 1520–1529, 2017.
- Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. John Wiley and Sons, 2nd edition, 2006.
- Dosovitskiy, A. and Brox, T. Generating images with perceptual similarity metrics based on deep networks. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 658–666, 2016.
- Galteri, L., Seidenari, L., Bertini, M., and Bimbo, A. D. Deep generative adversarial compression artifact removal. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 4836–4845, 2017.
- Gatys, L. A., Ecker, A. S., and Bethge, M. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2414–2423, 2016.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 2672–2680, 2014.
- Ishaan, G., Faruk, A., Martin, A., Vincent, D., and Courville, A. C. Improved training of wasserstein gans. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 5767–5777, 2017.
- Iwai, S., Miyazaki, T., Sugaya, Y., and Omachi, S. Fidelity-controllable extreme image compression with generative adversarial networks. *CoRR*, abs/2008.10314, 2020.
- Johnson, J., Alahi, A., and Fei-Fei, L. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 694–711, 2016.
- Johnston, N., Vincent, D., Minnen, D., Covell, M., Singh, S., Chinen, T., Jin Hwang, S., Shor, J., and Toderici, G. Improved lossy image compression with priming and spatially adaptive bit rates for recurrent networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4385–4393, 2018.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., and Shi, W. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 105–114, 2017.
- Li, M., Zuo, W., Gu, S., Zhao, D., and Zhang, D. Learning convolutional networks for content-weighted image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3214–3223, 2018.
- M Tschannen, E Agustsson, M. L. Deep generative models for distribution-preserving lossy compression. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 5933–5944, 2018.
- Ma, C., Yang, C., Yang, X., and Yang, M. Learning a no-reference quality metric for single-image super-resolution. *Comput. Vis. Image Underst.*, 158:1–16, 2017.

- Matsumoto, R. Introducing the perception-distortion trade-off into the rate-distortion theory of general information sources. *IEICE Communications Express*, 7(11):427–431, 2018a.
- Matsumoto, R. Rate-distortion-perception tradeoff of variable-length source coding for general information sources. *IEICE Communications Express*, 2018b.
- Mentzer, F., Agustsson, E., Tschannen, M., Timofte, R., and Gool, L. V. Conditional probability models for deep image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4394–4402, 2018.
- Mentzer, F., Toderici, G., Tschannen, M., and Agustsson, E. High-fidelity generative image compression. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Minnen, D., Ballé, J., and Toderici, G. Joint autoregressive and hierarchical priors for learned image compression. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 10794–10803, 2018.
- Mirza, M. and Osindero, S. Conditional generative adversarial nets. *CoRR*, abs/1411.1784, 2014.
- Mittal, A., Soundararajan, R., and Bovik, A. C. Making a completely blind image quality analyzer. *IEEE Signal Process. Lett.*, 20(3):209–212, 2013.
- Razavi, A., van den Oord, A., and Vinyals, O. Generating diverse high-fidelity images with VQ-VAE-2. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 14837–14847, 2019.
- Rippel, O. and Bourdev, L. Real-time adaptive image compression. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 2922–2930, 2017.
- Salimans, T., Goodfellow, I. J., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 2226–2234, 2016.
- Santurkar, S., Budden, D. M., and Shavit, N. Generative compression. In *Proceedings of the Picture Coding Symposium (PCS)*, pp. 258–262, 2018.
- Shaham, T. R. and Michaeli, T. Deformation aware image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2453–2462, 2018.
- Shannon, C. Coding theorems for a discrete source with a fidelity criteria. *IRE Nat. Conv. Rec.*, 4, 01 1959.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- Toderici, G., O’Malley, S. M., Hwang, S. J., Vincent, D., Minnen, D., Baluja, S., Covell, M., and Sukthankar, R. Variable rate image compression with recurrent neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- Toderici, G., Vincent, D., Johnston, N., Hwang, S. J., Minnen, D., Shor, J., and Covell, M. Full resolution image compression with recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5435–5443, 2017.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, pp. 6306–6315, 2017.
- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., and Loy, C. C. ESRGAN: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, pp. 63–79, 2018.
- Wang, Z. and Simoncelli, E. P. Reduced-reference image quality assessment using a wavelet-domain natural image statistic model. In Rogowitz, B. E., Pappas, T. N., and Daly, S. J. (eds.), *Human Vision and Electronic Imaging X*, volume 5666, pp. 149–159, 2005.
- Wang, Z., Simoncelli, E. P., and Bovik, A. C. Multiscale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems Computers*, 2003.
- Wu, L., Huang, K., and Shen, H. A gan-based tunable image compression system. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 2323–2331, 2020.
- Zhang, R., Isola, P., and Efros, A. A. Colorful image colorization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.

Zhou Wang, Bovik, A. C., Sheikh, H. R., and Simoncelli,
E. P. Image quality assessment: from error visibility
to structural similarity. *IEEE Transactions on Image
Processing*, 13(4):600–612, 2004.