

Towards End-to-End Image Compression and Analysis with Transformers

Yuanchao Bai^{1,2*}, Xu Yang^{1*}, Xianming Liu^{1,2†}, Junjun Jiang^{1,2},
Yaowei Wang², Xiangyang Ji³, Wen Gao^{2,4}

¹Harbin Institute of Technology, ²Peng Cheng Laboratory, ³Tsinghua University, ⁴Peking University
yuanchao.bai@gmail.com, 20s103274@stu.hit.edu.cn, {csxm,jiangjunjun}@hit.edu.cn
wangyw@pcl.ac.cn, xyji@tsinghua.edu.cn, wgao@pku.edu.cn

Abstract

We propose an end-to-end image compression and analysis model with Transformers, targeting to the cloud-based image classification application. Instead of placing an existing Transformer-based image classification model directly after an image codec, we aim to redesign the Vision Transformer (ViT) model to perform image classification from the compressed features and facilitate image compression with the long-term information from the Transformer. Specifically, we first replace the patchify stem (*i.e.*, image splitting and embedding) of the ViT model with a lightweight image encoder modelled by a convolutional neural network. The compressed features generated by the image encoder are injected convolutional inductive bias and are fed to the Transformer for image classification bypassing image reconstruction. Meanwhile, we propose a feature aggregation module to fuse the compressed features with the selected intermediate features of the Transformer, and feed the aggregated features to a deconvolutional neural network for image reconstruction. The aggregated features can obtain the long-term information from the self-attention mechanism of the Transformer and improve the compression performance. The rate-distortion-accuracy optimization problem is finally solved by a two-step training strategy. Experimental results demonstrate the effectiveness of the proposed model in both the image compression and the classification tasks.

Introduction

Vision Transformer (ViT) (Dosovitskiy et al. 2021) and its variations (Touvron et al. 2021; Wu et al. 2021; Yuan et al. 2021; Chen et al. 2021b; Liu et al. 2021), inherited from Transformer architecture (Vaswani et al. 2017) in natural language processing (NLP), have recently demonstrated outstanding performance on a board range of image analysis tasks, such as image classification (Dosovitskiy et al. 2021), segmentation (Zheng et al. 2021) and object detection (Fang et al. 2021). With the self-attention mechanism, these models are capable of capturing long-range dependencies in the image data, but inevitably result in high computational cost. In practice, Transformer-based models are usually deployed in the cloud-based paradigm and executed remotely. For example, massive image data is acquired by the frontend de-

vices, such as mobile phones or surveillance cameras, and transmitted to the cloud (*i.e.*, data center) for further analysis, sharing and storage. Image compression serves as a fundamental infrastructure for data communication between the frontend and the cloud.

In the traditional paradigm of cloud-based applications, image compression is considered independent of image analysis, and adopts lossy image compression standards designed for *human vision*, such as JPEG (Wallace 1992). In particular, the raw images are first transformed to the frequency domain with Discrete Cosine Transform (DCT). The frequency coefficients are then quantized to discard high frequencies that are less sensitive to human eyes. The quantized coefficients are encoded to bitstreams with entropy encoding and are transmitted to the cloud. On the cloud side, the quantized coefficients are recovered from the received bitstreams, which are then inversely transformed to reconstruct images. The reconstruction distortions are minimized with respect to Peak Signal-to-Noise Ratio (PSNR). However, if the reconstructed images optimized by PSNR are fed into the downstream image analysis tasks, which are tailored to *machine vision* instead, the corresponding results may be inaccurate, because the principle of machine vision is different from human vision (Wang et al. 2020). Besides, the traditional image codecs are comprised of hand-crafted modules with complex dependencies. It is difficult to optimize the sophisticated compression frameworks together with subsequent machine analysis tasks.

Recently, learning-based image compression emerges as an active research area in computer vision community. A number of learning-based image codecs, such as (Toderici et al. 2016; Theis et al. 2017; Li et al. 2018; Ballé et al. 2018; Minnen, Ballé, and Toderici 2018; Cheng et al. 2020; Ma et al. 2020; Hu et al. 2021), have achieved comparable or even better perceptual performance than traditional image codecs for human vision. Besides, by replacing the hand-crafted modules with deep neural networks (DNNs), learning-based image compression can be integrated with high-level tasks and end-to-end optimized for machine vision (Torfason et al. 2018; Chamain et al. 2021; Le et al. 2021). However, compared with image compression for human vision, image compression for machine vision is still in its infancy, because it is challenging to achieve the best of both worlds for low-level and high-level tasks.

*Equal contribution. †Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we propose a novel paradigm that is friendly for both human vision and machine vision, which integrates learning-based image compression with Transformer-based image analysis. The derived end-to-end image compression and analysis model leads to the synergy effect of these two tasks. Instead of placing an existing Transformer-based image classification model directly after an image codec, we redesign the ViT model to perform image classification from the compressed features (Torfason et al. 2018) and facilitate image compression with the long-term information from the Transformer. Specifically, we replace the the *patchify stem* (*i.e.*, image splitting and embedding) of the ViT model with a lightweight image encoder modelled by a convolutional neural network (CNN). The compressed features generated by the image encoder are injected convolutional inductive bias and are more expressive than the features extracted by the patchify stem from the decoded images. When transmitted to the cloud, the compressed features are fed to the Transformer for image classification bypassing image reconstruction. We further propose a feature aggregation module to fuse the compressed features with the selected intermediate features of the Transformer, and feed the aggregated features to a deconvolutional neural network for image reconstruction. The aggregated features obtain the long-term information from the Transformer and effectively improve the compression performance. We interpret the corresponding *rate-distortion-accuracy* optimization problem based on variational auto-encoder (VAE) (Theis et al. 2017; Ballé et al. 2018) and information bottleneck (IB) (Tishby, Pereira, and Bialek 2000; Alemi et al. 2017), and finally solve it with a two-step training strategy.

The main contributions are summarized as follows:

- We propose an end-to-end image compression and analysis model, which performs image classification from the compressed features. We interpret the rate-distortion-accuracy optimization problem based on VAE and IB.
- We design the network by integrating learning-based image compression with ViT-based image analysis, which leads to the synergy between the two tasks.
- In terms of rate-distortion, the proposed model achieves PSNR performance close to BPG (Bellard 2014). In terms of rate-accuracy, the proposed model outperforms ResNet50 (He et al. 2016), DeiT-S (Touvron et al. 2021) and Swin-T (Liu et al. 2021) classification from the decoded images, while significantly reduces the computational cost under equivalent number of parameters.

Related Work

Image Compression for Machine Vision. With the fast progress of artificial intelligence, an increasing amount of visual data is now not only viewed by humans but also analyzed by machines. Recently, image/video compression for machine vision has drawn significant interests in the computer vision community (Duan et al. 2020).

In order to optimize image compression with analysis, (Choi and Han 2020; Luo et al. 2021) and (Chamain, Cheung, and Ding 2019) proposed to optimize the quantization of the traditional codecs JPEG and JPEG2000 to improve

the performance of the following image classification. However, since the frameworks of traditional codecs are different from fully optimizable DNN and only the quantization is involved in the optimization, the improvement is limited. In contrast, learning-based image compression is more suitable to be jointly optimized with DNN-based image analysis. The related works can be divided into two categories: 1) *RGB inference*, such as (Chamain et al. 2021) and (Le et al. 2021), performs image analysis from RGB reconstructed images by placing image analysis methods directly after existing image codecs. 2) *Compressed inference*, such as (Torfason et al. 2018), performs image analysis directly from the compressed features bypassing image reconstruction.

In this paper, we propose an end-to-end image compression and analysis model with Transformers, inspired by (Torfason et al. 2018). Beyond (Torfason et al. 2018), we interpret the rate-distortion-accuracy optimization problem based on VAE and IB, and design the Transformer-based model leading to the synergy between the two tasks.

Transformers in Computer Vision. Nowadays, Transformers have shown their potential to be a viable alternative to CNNs in computer vision tasks. However, the ViT model (Dosovitskiy et al. 2021) without any human-defined inductive bias suffers from over-fitting when the training data is limited, and thus needs sophisticated data augmentation schemes (Touvron et al. 2021). In order to improve the performance and the robustness of Transformers, several works (Wu et al. 2021; Yuan et al. 2021; Chen et al. 2021b) incorporated CNNs into Transformers.

In this paper, we propose to replace the patchify stem of the ViT model with a CNN-based image encoder, which can enable image analysis from the compressed features and effectively improve the performance of image classification. The concurrent work (Xiao et al. 2021) also observes that early convolutions in Transformers can increase the optimization stability and improve the Top-1 accuracy. Our experimental results are consistent with the observation of (Xiao et al. 2021).

Proposed Method

Problem Formulation

We aim to perform image analysis from the compressed features. Given a raw image \mathbf{x} and its label \mathbf{y} , our goal is to learn a compressed representation $\hat{\mathbf{z}}$ that facilitates both image decoding (reconstruction) and analysis, as sketched in Fig. 1. Since the compressed representation $\hat{\mathbf{z}}$ is extracted from the image \mathbf{x} while not accessing the label \mathbf{y} , we assume that \mathbf{x} , \mathbf{y} , $\hat{\mathbf{z}}$ form a Markov chain $\mathbf{y} \leftrightarrow \mathbf{x} \leftrightarrow \hat{\mathbf{z}}$, leading to $p(\hat{\mathbf{z}}|\mathbf{x}, \mathbf{y}) = p(\hat{\mathbf{z}}|\mathbf{x})$.

Image Compression. We first formulate the lossy image compression model without taking image analysis into consideration. Following the standard framework of variational auto-encoder based image compression (Theis et al. 2017; Ballé et al. 2018), the latent representation \mathbf{z} is transformed from the raw image \mathbf{x} by an encoder and is quantized to the discrete-valued $\hat{\mathbf{z}}$. Then, $\hat{\mathbf{z}}$ is losslessly compressed with entropy encoding techniques (Witten, Neal, and Cleary 1987;




Figure 1: The theoretical framework of the proposed end-to-end image compression and analysis model.

Duda 2009) to form a bitstream. On the decoder side, $\hat{\mathbf{z}}$ is recovered from the bitstream and inversely transformed to a reconstructed image $\hat{\mathbf{x}}$. To optimize the performance of the compression model, it can be approximated by the minimization of the expectation of Kullback-Leibler (KL) divergence between the intractable true posterior $p(\hat{\mathbf{z}}|\mathbf{x})$ and a parametric inference model $q(\hat{\mathbf{z}}|\mathbf{x})$ over the data distribution $p(\mathbf{x})$ (Ballé et al. 2018):

$$\mathbb{E}_{p(\mathbf{x})} D_{kl}[q(\hat{\mathbf{z}}|\mathbf{x})||p(\hat{\mathbf{z}}|\mathbf{x})] = \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q(\hat{\mathbf{z}}|\mathbf{x})} [\log q(\hat{\mathbf{z}}|\mathbf{x}) - \log p(\mathbf{x}|\hat{\mathbf{z}}) - \log p(\hat{\mathbf{z}})] + \text{const} \quad (1)$$

where $D_{kl}[\cdot||\cdot]$ denotes KL divergence. Because the transform from \mathbf{x} to \mathbf{z} is deterministic and the quantization of \mathbf{z} is relaxed by adding noise from uniform distribution $\mathcal{U}(-\frac{1}{2}, \frac{1}{2})$, we have $q(\hat{\mathbf{z}}|\mathbf{x}) = \prod_i \mathcal{U}(z_i - \frac{1}{2}, z_i + \frac{1}{2})$ and thus the first term $\log q(\hat{\mathbf{z}}|\mathbf{x}) = 0$. The second term of (1) is interpreted as the expected distortion between \mathbf{x} and $\hat{\mathbf{x}}$, and the third term is interpreted as the cost of encoding $\hat{\mathbf{z}}$, leading to the rate-distortion trade-off (Shannon 1948).

Image Analysis. We then turn to consider image analysis. We propose to maximize the mutual information $I(\hat{\mathbf{z}}, \mathbf{y})$ between the compressed representation $\hat{\mathbf{z}}$ and the label \mathbf{y} , inspired by information bottleneck (Tishby, Pereira, and Bialek 2000; Alemi et al. 2017). The mutual information $I(\hat{\mathbf{z}}, \mathbf{y})$ is the reduction in the uncertainty of \mathbf{y} due to the knowledge of $\hat{\mathbf{z}}$:

$$\begin{aligned} I(\hat{\mathbf{z}}, \mathbf{y}) &= H(\mathbf{y}) - H(\mathbf{y}|\hat{\mathbf{z}}) \\ &= H(\mathbf{y}) + \sum_{\mathbf{y}, \hat{\mathbf{z}}} p(\mathbf{y}, \hat{\mathbf{z}}) \log p(\mathbf{y}|\hat{\mathbf{z}}) \end{aligned} \quad (2)$$

where $H(\cdot)$ denotes the entropy. Because the true posterior $p(\mathbf{y}|\hat{\mathbf{z}})$ is also intractable, we propose a variational approximation $q(\mathbf{y}|\hat{\mathbf{z}})$, which is the decoder for image analysis apart from the decoder for image reconstruction. Since $D_{kl}[p(\mathbf{y}|\hat{\mathbf{z}})||q(\mathbf{y}|\hat{\mathbf{z}})] \geq 0$, we have $\sum_{\mathbf{y}} p(\mathbf{y}|\hat{\mathbf{z}}) \log p(\mathbf{y}|\hat{\mathbf{z}}) \geq \sum_{\mathbf{y}} p(\mathbf{y}|\hat{\mathbf{z}}) \log q(\mathbf{y}|\hat{\mathbf{z}})$ and thus

$$I(\hat{\mathbf{z}}, \mathbf{y}) \geq H(\mathbf{y}) + \sum_{\mathbf{y}, \hat{\mathbf{z}}} p(\mathbf{y}, \hat{\mathbf{z}}) \log q(\mathbf{y}|\hat{\mathbf{z}}) \quad (3)$$

Because the entropy $H(\mathbf{y})$ is independent of $\hat{\mathbf{z}}$, we can maximize $\sum_{\mathbf{y}, \hat{\mathbf{z}}} p(\mathbf{y}, \hat{\mathbf{z}}) \log q(\mathbf{y}|\hat{\mathbf{z}})$ as a proxy for $I(\hat{\mathbf{z}}, \mathbf{y})$. Based on the Markov chain assumption, we replace $p(\mathbf{y}, \hat{\mathbf{z}})$ with $\sum_{\mathbf{x}} p(\mathbf{x}, \mathbf{y}, \hat{\mathbf{z}}) = \sum_{\mathbf{x}} p(\mathbf{x}, \mathbf{y}) p(\hat{\mathbf{z}}|\mathbf{x})$, and can rewrite $\sum_{\mathbf{y}, \hat{\mathbf{z}}} p(\mathbf{y}, \hat{\mathbf{z}}) \log q(\mathbf{y}|\hat{\mathbf{z}})$ as

$$\mathbb{E}_{p(\mathbf{x}, \mathbf{y})} \mathbb{E}_{p(\hat{\mathbf{z}}|\mathbf{x})} \log q(\mathbf{y}|\hat{\mathbf{z}}) \quad (4)$$




Figure 2: The network architecture of the encoder. We set $N = 128$ and $M = 192$ same as the low-rate setting of (Ballé et al. 2018). The encoder replaces the patchify stem of ViT and is relatively lightweight, which can be deployed at the frontend. **Q**: Quantization. **AE**: Arithmetic Encoder.

With $q(\mathbf{y}|\hat{\mathbf{z}})$, we can generate the estimated label $\hat{\mathbf{y}}$ from $\hat{\mathbf{z}}$.

Joint Optimization. With (1) and (4), we further formulate the joint optimization of both image compression and analysis. Since $p(\hat{\mathbf{z}}|\mathbf{x})$ in (4) is intractable, we share the inference model $q(\hat{\mathbf{z}}|\mathbf{x})$ in (1) as the approximation, and minimize the approximated negative (4) together with (1)¹:

$$\mathbb{E}_{p(\mathbf{x}, \mathbf{y})} \{-\alpha \mathbb{E}_{q(\hat{\mathbf{z}}|\mathbf{x})} \log q(\mathbf{y}|\hat{\mathbf{z}}) + D_{kl}[q(\hat{\mathbf{z}}|\mathbf{x})||p(\hat{\mathbf{z}}|\mathbf{x})]\} \quad (5)$$

where α is a trade-off parameter. Suppose that $p(\mathbf{x}|\hat{\mathbf{z}})$ is given by $\mathcal{N}(\mathbf{x}|\hat{\mathbf{x}}, (2\beta)^{-1}\mathbf{1})$, we can finally rewrite (5) to the objective function:

$$\mathbb{E}_{p(\mathbf{x}, \mathbf{y})} \mathbb{E}_{q(\hat{\mathbf{z}}|\mathbf{x})} [-\alpha \log q(\mathbf{y}|\hat{\mathbf{z}}) + \beta \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 - \log p(\hat{\mathbf{z}})] \quad (6)$$

The first term of (6) weighted by α can be interpreted as the cross-entropy loss for image analysis, such as image classification or segmentation, based on the types of label \mathbf{y} . In this work, we choose image classification as the target task. The second term weighted by β is the mean square error (MSE) distortion loss. The third term is the rate loss.

In contrast to the image compression models (Theis et al. 2017; Ballé et al. 2018), the compressed representation $\hat{\mathbf{z}}$ in (6) is also optimized for image analysis tasks. The complexity of $\hat{\mathbf{z}}$ in (6) is controlled by minimizing the cost of encoding $\hat{\mathbf{z}}$, rather than controlled by minimizing $I(\hat{\mathbf{z}}, \mathbf{x})$ in the information bottleneck models (Tishby, Pereira, and Bialek 2000; Alemi et al. 2017).

Transformer-based Network Architecture

We realize the theoretical framework in Fig. 1 by proposing an end-to-end image compression and analysis model with Transformers. The proposed model can promote the synergy between the two tasks.

Encoder. The network architecture of the proposed encoder is illustrated in Fig. 2. Similar to the setting of (Ballé et al. 2018), we employ four stride-2 5×5 convolutional layers to extract features with gradually reduced spatial resolution from the input image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$. We use LeakyReLU as the activation function instead of using the Generalized Divisive Normalization (GDN) (Ballé, Laparra, and Simoncelli 2016), because GDN results in convergence

¹ $\mathbb{E}_{p(\mathbf{x})} D_{kl}[q(\hat{\mathbf{z}}|\mathbf{x})||p(\hat{\mathbf{z}}|\mathbf{x})] = \mathbb{E}_{p(\mathbf{x}, \mathbf{y})} D_{kl}[q(\hat{\mathbf{z}}|\mathbf{x})||p(\hat{\mathbf{z}}|\mathbf{x})]$

problem when training with the Transformer blocks in the proposed model.

The resulting feature $\mathbf{z} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times M}$ is then quantized to the discrete-valued $\hat{\mathbf{z}}$. We employ the hyper-prior module (Ballé et al. 2018; Minnen, Ballé, and Toderici 2018) to estimate $p(\hat{\mathbf{z}}|\mathbf{h}_{\hat{\mathbf{z}}})$ for the entropy encoding of $\hat{\mathbf{z}}$, where $\mathbf{h}_{\hat{\mathbf{z}}}$ denotes the hyper-prior of $\hat{\mathbf{z}}$. We do not use the serial autoregressive module (Minnen, Ballé, and Toderici 2018), because its corresponding decoding time is too long on large-scale image classification datasets.

From the perspective of Transformer architecture design, the proposed encoder can be considered as the replacement of the patchify stem, *i.e.*, a stride-16 16×16 convolutional layer, applied to the input image in the ViT model (Dosovitskiy et al. 2021). Several concurrent works (Wu et al. 2021; Yuan et al. 2021; Chen et al. 2021b; Xiao et al. 2021) also replace the patchify stem with a stack of convolutional layers, in order to improve the performance of image classification. In the experiments, we observe that the proposed encoder is capable of extracting compressed features suitable for both image decoding reconstruction and classification through joint optimization. Note that the proposed encoder is relatively lightweight, thus can be deployed at the frontend, such as mobile phones or surveillance cameras.

Decoder-Classifier. The decoder receives the bitstream from the encoder and adopts the shared hyper-prior module $p(\hat{\mathbf{z}}|\mathbf{h}_{\hat{\mathbf{z}}})$ to recover $\hat{\mathbf{z}}$. For image classification, we directly feed $\hat{\mathbf{z}}$ to an inference network, instead of the common subsequent approach—first reconstruct the decoded image $\hat{\mathbf{x}}$ and then conduct inference on $\hat{\mathbf{x}}$.

Specifically, we adopt the standard Transformer blocks in the ViT model (Dosovitskiy et al. 2021) with the number of parameters equivalent to ResNet50 (He et al. 2016), as shown in Fig. 3. We expand the channel dimension of $\hat{\mathbf{z}}$ to C with a 1×1 convolutional layer, and reshape the resulting feature $\hat{\mathbf{z}}_0 \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times C}$ to a sequence $\hat{\mathbf{z}}_0 \in \mathbb{R}^{\frac{HW}{16^2} \times C}$. To maintain the spatial information of the feature $\hat{\mathbf{z}}_0$, we add learnable position embeddings \mathbf{p} to $\hat{\mathbf{z}}_0$ leading to $\tilde{\mathbf{z}}_0 = \hat{\mathbf{z}}_0 + \mathbf{p}$. Following (Dosovitskiy et al. 2021), we prepend a learnable class embedding $\tilde{\mathbf{c}}_0$, and feed the sequence $[\tilde{\mathbf{c}}_0; \tilde{\mathbf{z}}_0] \in \mathbb{R}^{(\frac{HW}{16^2} + 1) \times C}$ to the Transformer consisting of L Transformer blocks. The architecture of each Transformer block is illustrated in Fig. 3. The computation process can be formulated as

$$\begin{aligned} [\tilde{\mathbf{c}}'_i; \tilde{\mathbf{z}}'_i] &= \text{MSA}(\text{LN}([\tilde{\mathbf{c}}_{i-1}; \tilde{\mathbf{z}}_{i-1}])) + [\tilde{\mathbf{c}}_{i-1}; \tilde{\mathbf{z}}_{i-1}] \\ [\tilde{\mathbf{c}}_i; \tilde{\mathbf{z}}_i] &= \text{FFN}(\text{LN}([\tilde{\mathbf{c}}'_i; \tilde{\mathbf{z}}'_i])) + [\tilde{\mathbf{c}}'_i; \tilde{\mathbf{z}}'_i] \quad i = 1, \dots, L \end{aligned} \quad (7)$$

where $\text{MSA}(\cdot)$ denotes the multi-head self-attention module, $\text{FFN}(\cdot)$ denotes the feed forward network and $\text{LN}(\cdot)$ denotes the layer normalization (Ba, Kiros, and Hinton 2016), respectively.

With the self-attention mechanism, the class embedding $\tilde{\mathbf{c}}_i$ interacts with the image feature $\tilde{\mathbf{z}}_i$, and the final output $\tilde{\mathbf{c}}_L$ is used to compute $q(\mathbf{y}|\hat{\mathbf{z}})$ for image classification:

$$q(\mathbf{y}|\hat{\mathbf{z}}) = \text{Softmax}(\text{FFN}(\text{LN}(\tilde{\mathbf{c}}_L))) \quad (8)$$

where $\text{Softmax}(\cdot)$ denotes softmax operation. The $\text{FFN}(\cdot)$

is the classifier head mapping the embedding dimension from C to the number of classes.

Decoder-Reconstructor. Reconstructing image $\hat{\mathbf{x}}$ directly from $\hat{\mathbf{z}}$ (or $\hat{\mathbf{z}}_0$) ignores the global spatial correlations among the latent features. Recent image compression works (Qian et al. 2021; Guo et al. 2021) demonstrate that leveraging global context information during entropy coding can improve the compression performance. Transformers naturally capture the global spatial information among the latent features, which can also benefit low-level tasks, such as image processing (Chen et al. 2021a) and image generation (Jiang, Chang, and Wang 2021). Motivated by these works, we aim to extract the intermediate features $\tilde{\mathbf{z}}_i$'s of the Transformer and incorporate them into image reconstruction.

Specifically, we select $\hat{\mathbf{z}}_0$ and $[\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2, \tilde{\mathbf{z}}_3]$, and propose a feature aggregation module to fuse these features, similar to (Zheng et al. 2021). Selecting $[\tilde{\mathbf{z}}_1, \tilde{\mathbf{z}}_2, \tilde{\mathbf{z}}_3]$ means that the first three Transformer blocks are also involved in the image reconstruction process. Since image reconstruction may work independently of image classification, we avoid using $\{\tilde{\mathbf{z}}_i (i > 3)\}$ that involve too many Transformer blocks in the image reconstruction, in order to reduce the computational complexity. The feature aggregation module is illustrated in Fig. 3. The computational process can be formulated as

$$\begin{aligned} \hat{\mathbf{z}}''_0 &= \text{Conv}_1^0(\hat{\mathbf{z}}_0), \quad \tilde{\mathbf{z}}''_1 = \text{Conv}_1^1(\tilde{\mathbf{z}}_1) \\ \tilde{\mathbf{z}}''_2 &= \text{Conv}_1^2(\tilde{\mathbf{z}}_2), \quad \tilde{\mathbf{z}}''_3 = \text{Conv}_1^3(\tilde{\mathbf{z}}_3) \\ \hat{\mathbf{z}}_f &= \text{Conv}_2([\hat{\mathbf{z}}''_0; \tilde{\mathbf{z}}''_1; \tilde{\mathbf{z}}''_2; \tilde{\mathbf{z}}''_3]) \end{aligned} \quad (9)$$

where $\text{Conv}_1^i(\cdot)$ denotes a 1×1 convolutional layer reducing the channel dimension of the input to $\frac{C}{4}$. $\text{Conv}_2(\cdot)$ is another 1×1 convolutional layer with C channels fusing the four concatenated input features. $\hat{\mathbf{z}}_f$ is the fused feature.

Finally, we input the fused feature $\hat{\mathbf{z}}_f$ to four stride-2 5×5 deconvolutional layers gradually increasing the spatial resolution, leading to the reconstructed RGB image $\hat{\mathbf{x}}$.

Training Strategy

We observe that the one-step training strategy, *i.e.*, minimizing (6) to train the encoder and decoder from scratch, leads to convergence problem in the experiments. Instead, we employ a two-step training strategy:

- 1) We pretrain the proposed model without considering the quantization of \mathbf{z} and the hyper-prior module of $\hat{\mathbf{z}}$. We remove the rate loss in (6) temporarily, and minimize the cross-entropy loss together with the MSE loss. Because the value of the cross-entropy loss is much smaller than that of the MSE loss, we set $\alpha = 1$ and $\beta = 0.001$ in (6) to balance the contributions of the two losses.
- 2) We load the pretrained parameters and minimize (6) to train the entire network including the quantization of \mathbf{z} and the hyper-prior module of $\hat{\mathbf{z}}$. The α and β in (6) are tuned with fixed $\frac{\alpha}{\beta}$ to achieve different bit rates.

Experiments

Experimental Settings

Datasets. We perform extensive experiments on the ImageNet dataset (Deng et al. 2009) and iNaturalist19 (INat19)

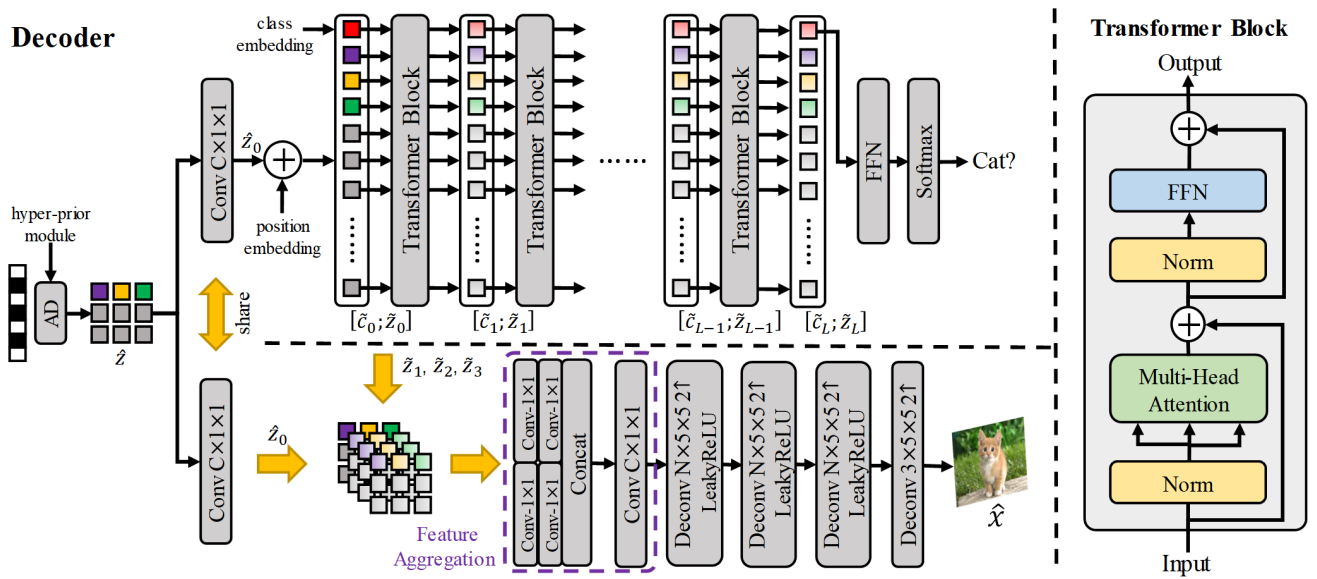


Figure 3: The network architecture of the decoder. We set $C = 384$, $L = 12$ and $N = 128$. The number of parameters is equivalent to ResNet50 (He et al. 2016). The decoder is deployed on the cloud to classify or reconstruct images from the received bitstreams. **AD**: Arithmetic Decoder.

dataset (Horn and Aodha 2019). ImageNet is well-known image classification dataset containing 1000 object classes with 1, 281, 167 training images and 50, 000 validation images. INat19 is a fine-grained classification dataset containing 1010 species of plants and animals with 265, 213 training images and 3030 validation images.

Pretraining w/o Compression. As aforementioned, we pretrain the proposed model without the quantization of \mathbf{z} and the hyper-prior module of $\hat{\mathbf{z}}$. We remove the rate loss, and minimize the cross entropy loss together with the MSE loss. We set $\alpha = 1$ and $\beta = 0.001$, respectively. We set the input size to 224×224 and adopt the same data augmentation as DeiT (Touvron et al. 2021), except for the Exponential Moving Average (EMA) (Polyak and Juditsky 1992), which do not enhance the performance of the proposed model. The input images are normalized with ImageNet default mean and standard deviation, and are denormalized during image reconstruction. We observe that random erasing (Zhong et al. 2020), mixup (Zhang et al. 2017) and cutmix (Yun et al. 2019) designed for the training of image classification are also compatible with the training of image reconstruction in our experiments.

On the ImageNet dataset, we train the proposed network from scratch. We use AdamW optimizer (Loshchilov and Hutter 2019) for 300 epochs with minibatches of size 1024. We set the initial learning rate to 0.001 and use a cosine decay learning rate scheduler with 5 epochs warm-up.

On the INat19 dataset, we initialize the network with the pretrained parameters on the ImageNet dataset. The classifier head is adjusted to the class number of INat19. We use AdamW optimizer for 100 epochs with minibatches of size 512. We set the initial learning rate to 0.0005 and use a cosine decay learning rate scheduler with 2 epochs warm-up.

Training w/ Compression. We load the pretrained parameters on the ImageNet and INat19 datasets, respectively. We recover the quantization of \mathbf{z} and the hyper-prior module of $\hat{\mathbf{z}}$. We fix $\frac{\alpha}{\beta} = 100$ and set $\alpha \in \{0.1, 0.3, 0.6\}$. We observe that the hyper-prior module of $\hat{\mathbf{z}}$ is sensitive to data augmentation, and thus we only employ RandomResizedCropAndInterpolation and RandomHorizontalFlip during training.

On the ImageNet dataset, we load the corresponding pretrained parameters, and use Adam optimizer (Kingma and Ba 2015) with a initial learning rate of 0.0001, following (Ballé et al. 2018). We train the proposed network for 300 epochs with minibatches of size 1024, and use a cosine decay learning rate scheduler with 5 epochs warm-up.

On the INat19 dataset, we load the corresponding pretrained parameters, and also use Adam optimizer with a initial learning rate of 0.0001. We train the proposed network for 300 epochs with minibatches of size 512, and use a cosine decay learning rate scheduler with 2 epochs warm-up.

Experimental Results

Pretrained Model. Table 1(a) reports the experimental results of our pretrained model without compression on ImageNet. We compare with the existing image classification models including CNN-based models, such as ResNet50 (He et al. 2016) and RegNetY-4G (Radosavovic et al. 2020), and Transformer-based models, such as ViT-B (Dosovitskiy et al. 2021), DeiT-S (Touvron et al. 2021), CvT-13 (Wu et al. 2021), CeiT-S (Yuan et al. 2021), Visformer-S (Chen et al. 2021b), Swin-T (Liu et al. 2021) and ViT_C-4GF (Xiao et al. 2021). We select the specific settings of the models with the number of parameters closest to ResNet50.

Table 1(b) reports the experimental results of our pretrained model on INat19 dataset, compared with ResNet50,

(a) Results on ImageNet Dataset				
Model	input size	Params (M)	Top-1 (%)	PSNR (dB)
ResNet50	224	25.6	75.9	—
RegNetY-4G	224	20.6	80.0	—
ViT-B	224	86.5	77.9	—
DeiT-S	224	22.1	79.9	—
CvT-13	224	20.0	81.6	—
CeiT-S	224	24.2	82.0	—
Visformer-S	224	40.2	82.3	—
Swin-T	224	28.3	81.2	—
ViT _C -4GF	224	17.8	81.4	—
Ours	224	25.6	81.7	31.7

(b) Results on INat19 Dataset				
Model	input size	Params (M)	Top-1 (%)	PSNR (dB)
ResNet50	224	25.6	71.8	—
DeiT-S	224	22.1	76.2	—
Swin-T	224	28.3	77.9	—
Ours	224	25.6	78.0	31.5

Table 1: Results of our pretrained model without compression on the ImageNet/INat19 datasets, compared with the existing image classification models.


DeiT-S and Swin-T. ResNet50 is finetuned based on (He et al. 2016). DeiT-S and Swin-T are finetuned using the same setting as our pretrained model.

From Table 1, our pretrained model can achieve comparable or better Top-1 accuracies than the existing image classification models, which demonstrates the efficacy of the image encoder for the Transformer-based image classification. In terms of image reconstruction evaluated by PSNR, our pretrained model achieves 31.7 dB and 31.5 dB, respectively. All these results demonstrate that the pretrained model can provide a satisfactory initialization for the following training with compression.


Full Model. In Fig. 4, we report the rate-distortion and rate-accuracy results of our full model on ImageNet and INat19. We compare with the existing image codecs and the image classification models applied to the decoded images.

We select the traditional image codecs including JPEG (Wallace 1992), JPEG2000 (Skodras, Christopoulos, and Ebrahimi 2001), BPG (Bellard 2014), and the learning-based image codecs including *bmsbj* (Ballé et al. 2018) and *mbt-m* (Minnen, Ballé, and Toderici 2018). The *mbt-m* removes the serial autoregressive module, avoiding long decoding time on the large-scale datasets. The sophisticated learning-based image codecs with complex entropy models and network architectures, such as (Hu, Yang, and Liu 2020; Cheng et al. 2020; Qian et al. 2021; Guo et al. 2021), are too time-consuming to be evaluated on the large-scale datasets, despite their good compression performance.

We select the decoded images of the best performed traditional and learning-based image codecs in our experiments, *i.e.*, BPG and *mbt-m*, and adopt the image classification models ResNet50, DeiT-S and Swin-T to compute the Top-1 accuracies in comparison with the proposed model. More-



(a) Rate-distortion and rate-accuracy results on ImageNet.



(b) Rate-distortion and rate-accuracy results on INat19.

Figure 4: Rate-distortion and rate-accuracy results of the proposed model, compared with the existing image codecs and the image classification methods applied to the reconstructed RGB images. JFT means joint finetune.

over, we jointly finetune *mbt-m* together with ResNet50, DeiT-S and Swin-T by minimizing (6) with $\frac{\alpha}{\beta} = 100$, same as the proposed model.

In terms of the rate-distortion performance, the proposed model significantly outperforms the traditional image codecs JPEG and JPEG2000. It is comparable to BPG at relatively low bit-rates but is surpassed by BPG at high bit-rates. The proposed model outperforms the learning-based image codecs *bmsbj* and *mbt-m*. The rate-distortion performance of the jointly finetuned *mbt-m* is similar to the original *mbt-m*. The proposed model also outperforms the jointly finetuned *mbt-m*.

In terms of the rate-accuracy performance, the proposed model outperforms ResNet50, DeiT-S and Swin-T applied to the decoded images of BPG and *mbt-m*, because the image classification models trained on the original datasets are not robust to the decoded images at low bit-rates. Although Swin-T outperforms DeiT-S on the original datasets (Table. 1), DeiT-S outperforms Swin-T on the decoded images at low bit-rates. After jointly finetuned, Swin-T surpasses DeiT-S on the decoded images at the corresponding bit-rates. The proposed model also outperforms the jointly




Figure 5: (a) JPEG+DeiT-S, 29.3dB, 0.60bpp, maraca (×). (b) JPEG2000+DeiT-S, 29.17dB, 0.41bpp, can opener (×). (c) BPG+DeiT-S, 29.41dB, 0.18bpp, can opener (×). (d) mbt-m+DeiT-S, 29.79dB, 0.18bpp, can opener (×). (e) mbt-m+Swin-T(JFT), 29.75dB, 0.19bpp, reel (×). (f) Ours, 29.52dB, 0.19bpp, teddy bear (✓).

finetuned ResNet50, Swin-T and DeiT-S, because the image reconstruction constrained by MSE loss damages the high-frequency information effective for the image classification. The proposed model bypasses the image reconstruction process and directly do inference from the compressed features, which can utilize these high-frequency information.

In Fig. 5, we show an illustrative example of the experimental results. Under similar PSNRs, our model achieves comparable or less bit-per-pixel (bpp) and the correct classification result, compared with the other methods.

Ablation Studies

Rate-Distortion-Accuracy Trade-off. When the bit-rates of \hat{z} is constrained, minimizing (6) with different $\frac{\alpha}{\beta}$ leads to bit allocation between image classification and reconstruction. We empirically set $\frac{\alpha}{\beta} \in \{50, 100, 200\}$ and test the rate-distortion-accuracy trade-off as shown in Fig. 6. We can observe that larger $\frac{\alpha}{\beta}$ leads to better Top-1 accuracy but sacrifices PSNR. In contrast, smaller $\frac{\alpha}{\beta}$ leads to better PSNR but lower Top-1 accuracy.

Feature Aggregation. In Fig. 6, we compare with the proposed model without the feature aggregation module, *i.e.*, directly feeding \hat{z}_0 to the deconvolutional neural network for image reconstruction. Although the feature aggregation module is only applied to the image reconstruction process, it can benefit both the image classification and reconstruction through rate-distortion-accuracy optimization. The reduced bit-rates from the image reconstruction are potentially allocated to the image classification, leading to the improvement of both tasks. Although the improvement of Top-1 accuracy is more obvious than PSNR in Fig. 6, the decrease




Figure 6: Ablation studies of rate-distortion-accuracy trade-off and the feature aggregation (FA) module.




Figure 7: Ablation study of the computational cost.

of the cross entropy loss $-\alpha \log q(\mathbf{y}|\hat{\mathbf{z}})$ is actually similar to that of the MSE loss $\beta \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$ in (6) in the experiments.

Computational Cost. In Fig. 7, we compare the computational cost of the proposed model with the concatenation of the learning-based image codec *mbt-m* and the image classification methods including ResNet50, DeiT-S and Swin-T. The architecture of the proposed encoder is similar to the low-rate setting of *mbt-m*, thus their computational costs of image encoding are similar. In terms of image reconstruction on the decoder side, our image reconstructor needs more FLOPs than *mbt-m* due to the feature aggregation module. In terms of image classification on the decoder side, our image classifier directly performs inference from the compressed features without the image reconstruction process, and thus needs far less computational cost compared with the inference from reconstructed RGB images.

Conclusion

In this paper, we learn an end-to-end image compression and analysis model with Transformers, targeting to the cloud-based image classification application. At the frontend, a CNN-based image encoder extracts compressed features from raw images and transmits them to the cloud. On the cloud, the compressed features injected convolutional inductive bias are directly fed to the Transformer for image classification bypassing image reconstruction. Meanwhile, the intermediate features of the Transformer capturing global information are aggregated with the compressed features for image reconstruction. Experimental results demonstrate the effectiveness of the proposed model in both rate-distortion and rate-accuracy performance.

Acknowledgements

This work was supported by National Key Research and Development Project under Grant 2019YFE0109600, National Natural Science Foundation of China under Grants 61922027, 61827804, 6207115, 61971165 and U20B2052, PCNL Key Project under Grant PCL2021A07, China Postdoctoral Science Foundation under Grant 2020M682826.

References

- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2017. Deep variational information bottleneck. In *International Conference on Learning Representations*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450*.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2016. Density modeling of images using a generalized normalization transformation. In *International Conference on Learning Representations*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*.
- Bellard, F. 2014. BPG Image format. <https://bellard.org/bpg/>. Accessed: 2020-09-30.
- Chamain, L. D.; Cheung, S.-c. S.; and Ding, Z. 2019. Quantnet: Joint Image Compression and Classification Over Channels with Limited Bandwidth. In *IEEE International Conference on Multimedia and Expo (ICME)*, 338–343.
- Chamain, L. D.; Racapé, F.; Begaint, J.; Pushparaja, A.; and Feltman, S. 2021. End-to-End optimized image compression for machines, a study. In *Data Compression Conference (DCC)*, 163–172.
- Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; and Gao, W. 2021a. Pre-trained image processing transformer. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Chen, Z.; Xie, L.; Niu, J.; Liu, X.; Wei, L.; and Tian, Q. 2021b. Visformer: The vision-friendly transformer. *arXiv preprint arXiv:2104.12533*.
- Cheng, Z.; Sun, H.; Takeuchi, M.; and Katto, J. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *IEEE Conference on Computer Vision and Pattern Recognition*, 7939–7948.
- Choi, J.; and Han, B. 2020. Task-Aware Quantization Network for JPEG Image Compression. In *European Conference on Computer Vision*, 309–324. Cham: Springer International Publishing. ISBN 978-3-030-58565-5.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; and Gelly, S. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- Duan, L.-Y.; Liu, J.; Yang, W.; Huang, T.; and Gao, W. 2020. Video Coding for Machines: A Paradigm of Collaborative Compression and Intelligent Analytics. *IEEE Transactions on Image Processing*, 29: 8680–8695.
- Duda, J. 2009. Asymmetric numeral systems. *arXiv preprint arXiv:0902.0271*.
- Fang, Y.; Liao, B.; Wang, X.; Fang, J.; Qi, J.; Wu, R.; Niu, J.; and Liu, W. 2021. You Only Look at One Sequence: Rethinking Transformer in Vision through Object Detection. *arXiv preprint arXiv:2106.00666*.
- Guo, Z.; Zhang, Z.; Feng, R.; and Chen, Z. 2021. Causal Contextual Prediction for Learned Image Compression. *IEEE Transactions on Circuits and Systems for Video Technology*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Horn, G. V.; and Aodha, O. M. 2019. iNaturalist2019. <https://www.kaggle.com/c/inaturalist-2019-fgvc6>. Accessed: 2021-05-30.
- Hu, Y.; Yang, W.; and Liu, J. 2020. Coarse-to-fine hyperprior modeling for learned image compression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11013–11020.
- Hu, Y.; Yang, W.; Ma, Z.; and Liu, J. 2021. Learning end-to-end lossy image compression: A benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Jiang, Y.; Chang, S.; and Wang, Z. 2021. TransGAN: Two Transformers Can Make One Strong GAN. *arXiv preprint arXiv:2102.07074*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Le, N.; Zhang, H.; Cricri, F.; Youvalari, R. G.; and Rahtu, E. 2021. Image Coding For Machines: an End-To-End Learned Approach. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1590–1594.
- Li, M.; Zuo, W.; Gu, S.; Zhao, D.; and Zhang, D. 2018. Learning Convolutional Networks for Content-Weighted Image Compression. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3214–3223. ISBN 2575-7075.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *arXiv preprint arXiv:2103.14030*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Luo, X.; Talebi, H.; Yang, F.; Elad, M.; and Milanfar, P. 2021. The Rate-Distortion-Accuracy Tradeoff: JPEG Case Study. In *Data Compression Conference (DCC)*, 354–354.
- Ma, H.; Liu, D.; Yan, N.; Li, H.; and Wu, F. 2020. End-to-End Optimized Versatile Image Compression With Wavelet-Like Transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.

- Minnen, D.; Ballé, J.; and Toderici, G. D. 2018. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, 10771–10780.
- Polyak, B. T.; and Juditsky, A. B. 1992. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4): 838–855.
- Qian, Y.; Tan, Z.; Sun, X.; Lin, M.; Li, D.; Sun, Z.; Li, H.; and Jin, R. 2021. Learning Accurate Entropy Model with Global Reference for Image Compression. In *International Conference on Learning Representations*.
- Radosavovic, I.; Kosaraju, R. P.; Girshick, R. B.; He, K.; and Dollár, P. 2020. Designing Network Design Spaces. *IEEE Conference on Computer Vision and Pattern Recognition*, 10425–10433.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- Skodras, A.; Christopoulos, C.; and Ebrahimi, T. 2001. The jpeg 2000 still image compression standard. *IEEE Signal Processing Magazine*, 18(5): 36–58.
- Theis, L.; Shi, W.; Cunningham, A.; and Huszár, F. 2017. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations*.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Toderici, G.; O’Malley, S. M.; Hwang, S. J.; Vincent, D.; Minnen, D.; Baluja, S.; Covell, M.; and Sukthankar, R. 2016. Variable Rate Image Compression with Recurrent Neural Networks. In *International Conference on Learning Representations*.
- Torfason, R.; Mentzer, F.; Agustsson, E.; Tschannen, M.; Timofte, R.; and Gool, L. 2018. Towards Image Understanding from Deep Compression without Decoding. In *International Conference on Learning Representations*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wallace, G. K. 1992. The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1): xviii–xxxiv.
- Wang, H.; Wu, X.; Huang, Z.; and Xing, E. P. 2020. High-frequency component helps explain the generalization of convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 8684–8694.
- Witten, I. H.; Neal, R. M.; and Cleary, J. G. 1987. Arithmetic coding for data compression. *Communications of the ACM*, 30(6): 520–540.
- Wu, H.; Xiao, B.; Codella, N.; Liu, M.; Dai, X.; Yuan, L.; and Zhang, L. 2021. Cvt: Introducing convolutions to vision transformers. *arXiv preprint arXiv:2103.15808*.
- Xiao, T.; Singh, M.; Mintun, E.; Darrell, T.; Dollár, P.; and Girshick, R. 2021. Early Convolutions Help Transformers See Better. *arXiv preprint arXiv:2106.14881*.
- Yuan, K.; Guo, S.; Liu, Z.; Zhou, A.; Yu, F.; and Wu, W. 2019. Incorporating Convolution Designs into Visual Transformers. *arXiv preprint arXiv:2103.11816*.
- Yun, S.; Han, D.; Oh, S. J.; Chun, S.; Choe, J.; and Yoo, Y. 2019. CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. *International Conference on Computer Vision*, 6022–6031.
- Zhang, H.; Cissé, M.; Dauphin, Y.; and Lopez-Paz, D. 2017. mixup: Beyond Empirical Risk Minimization. *arXiv preprint arXiv:1710.09412*.
- Zheng, S.; Lu, J.; Zhao, H.; Zhu, X.; Luo, Z.; Wang, Y.; Fu, Y.; Feng, J.; Xiang, T.; and Torr, P. H. 2021. Rethinking Semantic Segmentation from a Sequence-to-Sequence Perspective with Transformers. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Zhong, Z.; Zheng, L.; Kang, G.; Li, S.; and Yang, Y. 2020. Random Erasing Data Augmentation. In *AAAI Conference on Artificial Intelligence*.