TRANSFORMER-BASED VARIABLE-RATE IMAGE COMPRESSION WITH REGION-OF-INTEREST CONTROL

Chia-Hao Kao[⋆]   *Ying-Chieh Weng*[⋆]   *Yi-Hsin Chen*   *Wei-Chen Chiu*   *Wen-Hsiao Peng*

National Yang Ming Chiao Tung University, Taiwan

ABSTRACT

This paper proposes a transformer-based learned image compression system. It is capable of achieving variable-rate compression with a single model while supporting the region-of-interest (ROI) functionality. Inspired by prompt tuning, we introduce prompt generation networks to condition the transformer-based autoencoder of compression. Our prompt generation networks generate content-adaptive tokens according to the input image, an ROI mask, and a rate parameter. The separation of the ROI mask and the rate parameter allows an intuitive way to achieve variable-rate and ROI coding simultaneously. Extensive experiments validate the effectiveness of our proposed method and confirm its superiority over the other competing methods.

Index Terms— Transformer-based image compression, variable-rate compression, region-of-interest, prompt tuning

1. INTRODUCTION

Transformers have recently emerged as an attractive alternative to convolutional neural networks (CNN) for constructing learned image compression systems [1, 2]. The attention-based convolution coupled with the shifted-windowing technique [3] offers both high compression performance and low computational cost. To make learned image codecs practical, much research has been devoted to the use of a single autoencoder for variable-rate compression. However, little work is done on transformer-based codecs.

One common approach to variable-rate compression with a single autoencoder is to adapt the feature distributions of the autoencoder. For example, Yang et al. [4] channel-wisely scale the feature maps of every convolutional layer in the autoencoder according to a rate parameter. In contrast, Cui et al. [5] perform channel-wise scaling for the image latents only. Specifically, they first optimize the scaling factors for a few distinctive rate points and then interpolate between the resulting scaling factors for continuous rate adaptation.

In a similar vein, Wang et al. [6] scale the value matrix channel-wisely in each self-attention layer for transformer-based codecs. In another direction, Song et al. [7] propose a spatially adaptive rate adaptation scheme by introducing spatial feature transform (SFT) [8] to convolutional layers for an element-wise affine transformation of the feature maps. In particular, the affine parameters are predicted by a conditioning network that takes as input a quality map reflecting the spatial importance of every image pixel. Notably, this quality map can be adapted for multiple uses, such as rate control, spatial bit allocation, region-of-interest (ROI) coding, and task-specific coding. However, how to determine the quality map is non-trivial and may involve time-consuming back-propagation at inference, especially when it is necessary to combine some of these tasks, e.g. ROI coding subject to a rate constraint.
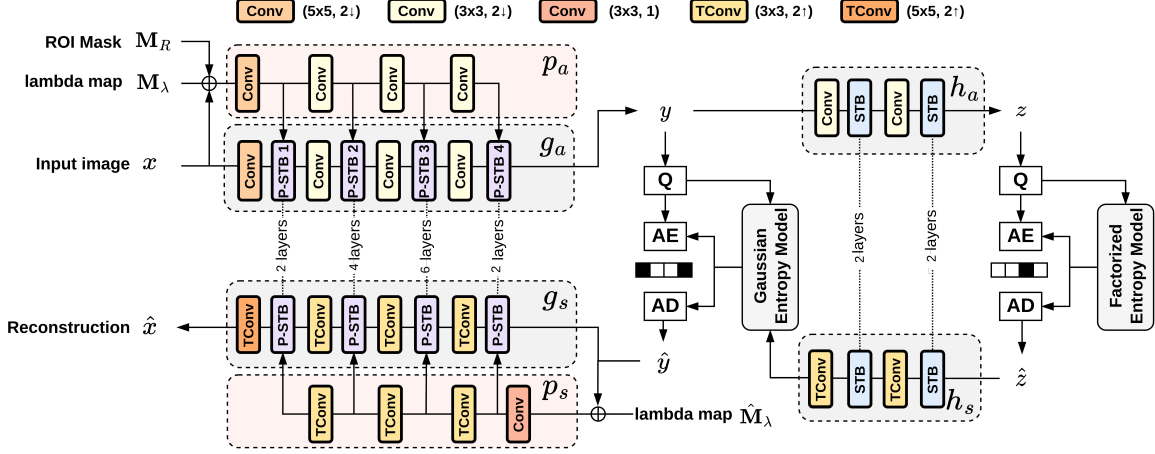
In this work, we propose a transformer-based image codec capable of achieving variable-rate compression while supporting ROI functionality. Inspired by prompting techniques [9, 10], we introduce prompt generation networks to condition our transformer-based codec. Unlike ordinary prompting [9, 10], which learns task-specific prompt tokens, our prompt generation networks generate content-adaptive tokens according to the input image, an ROI mask, and a rate parameter. The separation of the ROI mask and the rate parameter allows us to disentangle the rate and spatial quality controls. Our contributions are threefold. (1) To our best knowledge, this work is the first transformer-based image codec that leverages network-generated prompts to achieve variable-rate coding with ROI support. (2) Our scheme offers an intuitive way to specify the ROI and rate parameters. (3) Our scheme performs comparably to or better than the baselines while having lower computational complexity.

2. PROPOSED METHOD

We propose a Swin-transformer-based image compression system. It is capable of performing variable-rate compression with a single model while offering spatially adaptive quality control for the ROI functionality. Fig. 1 illustrates our overall architecture. It is built upon TIC (**T**ransformer-based **I**mage **C**ompression [1]) but without the context model for entropy coding. The main autoencoder g_a, g_s and hyper-


[⋆] Equal contribution.

Fig. 1: The network architecture of the proposed transformer-based image codec. For simultaneous rate and spatial quality control, the prompt generation networks p_a, p_s produce prompt tokens for the encoder g_a and decoder g_s, respectively.

prior autoencoder h_a, h_s are comprised of Swin-transformer blocks (STB) interleaved with convolutional layers. The details of STB can be found in [1]. To encode an input image $x \in \mathbb{R}^{3 \times H \times W}$, the encoder takes two additional inputs, a lambda map $\mathbf{M}_\lambda \in \mathbb{R}^{1 \times H \times W}$ and a ROI mask $\mathbf{M}_R \in \mathbb{R}^{1 \times H \times W}$. The lambda map $\mathbf{M}_\lambda$ is a uniform map populated with the same rate parameter $m_\lambda \in [0,1]$ that controls the bit rate of the compressed bitstream. The ROI mask $\mathbf{M}_R$ specifies spatially the importance of individual pixels in the image. Each element in the ROI mask is a real value in $[0,1]$. Both inputs serve as the conditioning signal utilized to generate *prompt tokens* for adapting the main encoder g_a (Section 2.1). In a similar way, the decoder g_s is adapted by taking as inputs the quantized image latent $\hat{y}$ and a downscaled lambda map $\hat{\mathbf{M}}_\lambda \in \mathbb{R}^{1 \times \frac{H}{16} \times \frac{W}{16}}$ that matches the spatial resolution of the latent $\hat{y}$. Unlike [7], which relies on a (single) quality map for both rate and spatial quality control, our design has the striking feature of disentangling the rate (i.e. $\mathbf{M}_\lambda$) and spatial quality control (i.e. $\mathbf{M}_R$). In other words, it treats them as two independent dimensions, offering a more intuitive way to achieve simultaneous variable-rate and ROI coding.

2.1. Prompt-based Conditioning

Inspired by [10], we propose to use learned parameters, known as prompts, as the additional inputs to the Swin-transformer layers, in order to achieve variable-rate and ROI coding. The resulting STB is termed prompted Swin-transformer block (P-STB). As shown in Fig. 1, the learned prompts are produced by two generation networks p_a, p_s for conditioning the encoder g_a and decoder g_s, respectively. p_a consists of several convolutional layers that match those of the encoder g_a, and it takes as input the concatenation of the ROI mask $\mathbf{M}_R$, lambda map $\mathbf{M}_\lambda$, and image x. The feature maps of p_a are fed into the corresponding P-STBs to generate
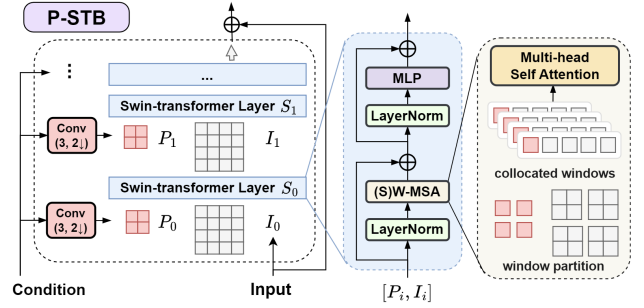
Fig. 2: Illustration of the prompted Swin-transformer block.

prompt tokens to be interacted with image tokens. p_s follows a similar architecture, replacing the convolutional layers with the transposed convolutional layers for upsampling.

Fig. 2 further details P-STB, where P_i, I_i denote the prompt and image tokens, respectively. They are fed into the i^{th} Swin-transformer layer S_i for window-based attention to arrive at I_{i+1}. Specifically, each window has its own image and prompt tokens. We divide spatially the prompt tokens in the same way as the image tokens. For multi-head self-attention, the key K and value V matrices, initially composed of only image tokens $\mathbf{X}_I \in \mathbb{R}^{S_I \times d}$, are augmented with the prompt tokens $\mathbf{X}_P \in \mathbb{R}^{S_P \times d}$, where S_I, S_P are the numbers of image and prompt tokens in a window, respectively, and d is the dimension of each token. In symbols, we have

$$\begin{aligned} Q &= \mathbf{X}_I \mathbf{W}_Q, \\ K &= [\mathbf{X}_I, \mathbf{X}_P] \mathbf{W}_K, \\ V &= [\mathbf{X}_I, \mathbf{X}_P] \mathbf{W}_V, \end{aligned} \quad (1)$$

where $[\cdot]$ indicates concatenation along the token dimension, $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d \times d}$ project their respective input matrices into query $Q \in \mathbb{R}^{S_I \times d}$, key $K \in \mathbb{R}^{(S_I + S_P) \times d}$, and value $V \in \mathbb{R}^{(S_I + S_P) \times d}$. Then, we have

$$\text{Attention}(Q, K, V) = \text{Softmax}(QK^\top / \sqrt{d} + B)V, \quad (2)$$
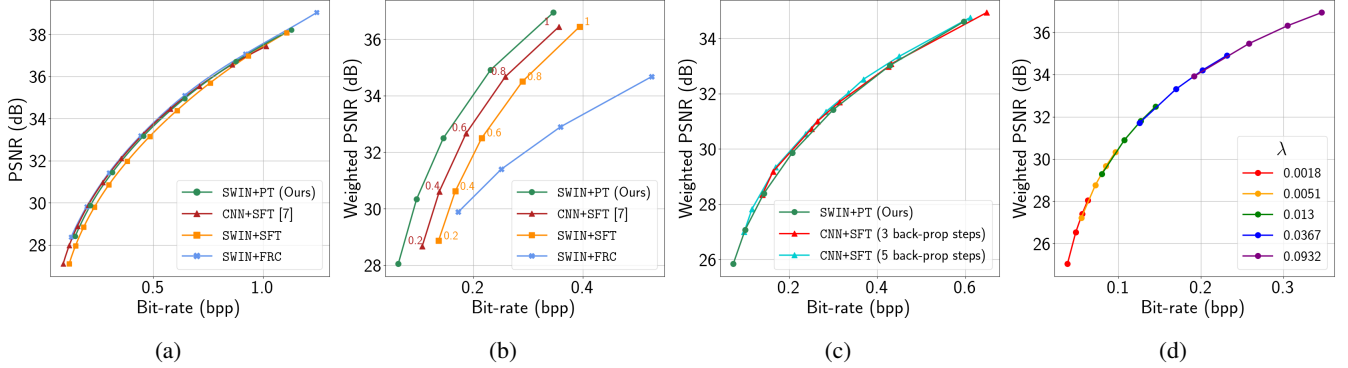
Fig. 3: (a) Variable-rate coding without ROI on Kodak. (b) Variable-rate coding with ROI on COCO. The annotated numbers indicate the quality values of ROI for `CNN+SFT` and `SWIN+SFT`. (c) Variable-rate coding optimized for a given ROI specification on 10 randomly selected images from COCO. (d) Rate-distortion plots by altering $\mathbf{M}_\lambda = \lambda$ and $\mathbf{M}_R$ on COCO.

where B denotes the relative position bias. Due to the use of an additional strided convolution in P-STB, the number of the prompt tokens is only one fourth of that of the image tokens. This helps reduce the complexity. Our design differs from [10], where all the non-overlapping windows in a Swin-transformer layer share the same learned prompts. We argue that this is not optimal for spatially adaptive quality control such as ROI coding. With our design, an output image token in I_{i+1} aggregates information from the input image I_i and prompt P_i tokens in the same window.

2.2. Loss Function

We design our loss function in such a way that the model will respond to both the ROI mask $\mathbf{M}_{R_i}$ and the rate parameter m_λ properly. Specifically, it is given by a weighted sum of the masked distortion and the bit rate:

$$\mathcal{L}_{rd}(x) = \lambda \cdot \sum_{i=1}^{N} \frac{\mathbf{M}_{Ri} \cdot (x_i - x'_i)^2}{N} + \mathcal{R}, \qquad (3)$$

where x_i, x'_i are the i-th pixel in the original and compressed images, respectively, N is the total number of pixels in the image, $\mathcal{R}$ denotes the bit rate in bits-per-pixel, and $\lambda = f(m_\lambda) = \exp((\log \lambda_{max} - \log \lambda_{min}) \cdot m_\lambda + \log \lambda_{min})$ is a Lagrange multiplier, which is a function of the rate parameter m_λ, with λ_{max} and λ_{min} being the highest and lowest λ, respectively. In Eq. (3), the squared error of each pixel is weighted by ROI mask $\mathbf{M}_R$ to reflect its spatial importance, and λ trades off the distortion against the bit rate.

3. EXPERIMENTS

Training Details. We train our model in three stages using Flicker2W [11] and COCO 2017 [12] training sets. In each training iteration, the input images are randomly cropped to 256×256. We first pre-train a base codec (i.e. g_a, g_s, h_a, h_s) without the prompt generation networks for the highest rate point. We then train the whole model jointly for variable-rate coding by sampling λ uniformly from $\lambda_{min} = 0.0018$ to

$\lambda_{max} = 0.0932$. In this stage, a uniform ROI mask filled with 1's (i.e. every pixel is equally important) is applied. Lastly, we follow [7] to produce 4 types of random ROI masks and fine-tune the model for spatial quality control.

Evaluation. We evaluate our model for variable-rate coding *without* and *with* ROI. In the absence of ROI, we evaluate our model on Kodak [13] by having $\mathbf{M}_R = \mathbf{1}$. In the presence of ROI, we adopt COCO 2017 [12] validation set for testing. The image reconstruction quality is measured in terms of weighted PSNR, for which the weighted mean-squared error is evaluated by $(\alpha \text{MSE}_{\text{ROI}} + \beta \text{MSE}_{\text{NROI}})/(\alpha N_{\text{ROI}} + \beta N_{\text{NROI}})$, where α, β are the weighting factors for the ROI and non-ROI regions, respectively, MSE_{ROI} (respectively, MSE_{NROI}) is the sum of the squared errors over ROI (respectively, non-ROI), and N_{ROI} (respectively, N_{NROI}) is the number of ROI (respectively, non-ROI) pixels. Under this setting, the ROI is specified by the union of all the foreground objects in the ground-truth segmentation mask.

Baselines. For comparison, the baseline methods include (1) training separate models for variable-rate coding without p_a, p_s (i.e. separate models for different rates), denoted as `SWIN+FRC`, and (2) the spatial feature transform (SFT) method in [7] (i.e. `CNN+SFT`). To validate the effectiveness of prompt tuning, we further construct (3) a model, denoted as `SWIN+SFT`, that introduces a SFT layer after every STB in our transformer-based codec for variable-rate ROI coding.

3.1. Variable-rate Compression

Fig. 3a compares the competing methods for variable-rate compression without ROI. We see that our method with prompt tuning (`SWIN+PT`) performs very close to the baseline method with training separate models (`SWIN+FRC`). In contrast, the variant `SWIN+SFT` incurs a slight rate-distortion loss. These results suggest that prompt tuning is a more effective approach to variable-rate compression than SFT for our transform-based codec. We also see that `SWIN+PT` performs comparably to `CNN+SFT` while having lower computational complexity (Section 3.2).
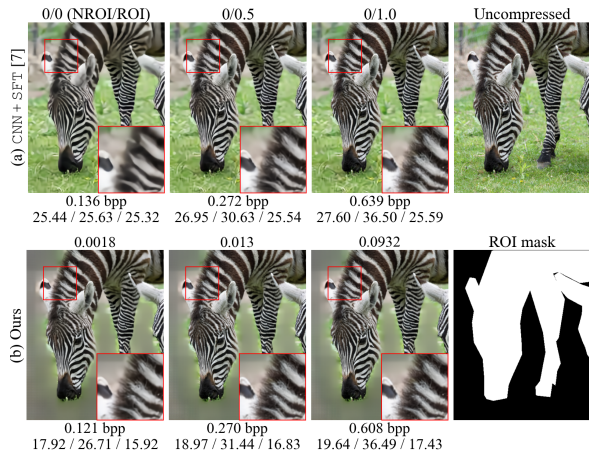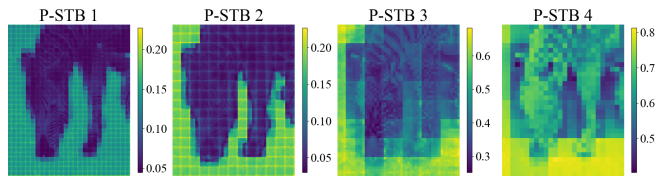
Fig. 4: Subjective quality comparison of our method and CNN+SFT [7]. The corresponding rate and PSNR (full image / ROI / non-ROI) are presented below each image.

Fig. 5: Visualization of the attention map in S_0 of each P-STB in g_a with a fixed binary $\mathbf{M}_R$ at $\lambda = 0.013$.

Table 1: Comparison of the kMACs/pixel and model size.

	kMACs/pixel	Params (M)
SWIN+FRC	718.50	17.66
CNN+SFT [7]	1480.16	27.56
SWIN+SFT	1915.80	21.71
Ours	1070.68	32.7

Fig. 3b shows how these methods perform in terms of variable-rate coding with ROI. For this experiment, we set $\alpha = 1, \beta = 0$ in evaluating the weighted PSNR; that is, we focus only on the quality of the ROI region. Recall that our scheme (SWIN+PT) separates the rate control M_λ from the spatial quality control M_R. For the present task, M_R is chosen to be the binary ground-truth ROI mask while several distinct M_λ values are used for variable-rate compression. Without the disentanglement of the rate and spatial quality control, both CNN+SFT and SWIN+SFT have to rely on adjusting a quality map. Specifically, we fix the quality value of the non-ROI region at 0 and adjust that of the ROI region for rate control. In Fig. 3b, our scheme (SWIN+PT) consistently shows higher weighted PSNR than the other baselines. Fig. 4 further demonstrates that as compared to CNN+SFT [7], our method (SWIN+PT) is more effective in blurring the background for better foreground coding.

Taking one step further, Fig. 3c compares our SWIN+PT with CNN+SFT [7] under a more general setting, where the ROI and non-ROI quality is weighted by 0.8 and 0.2, respectively. For SWIN+PT, we simply set $\mathbf{M}_R$ to be 0.8 and 0.2 for the ROI and non-ROI regions, respectively, while adjusting $\mathbf{M}_\lambda$ for variable-rate compression. To achieve the same effect with CNN+SFT, we back-propagate an input-specific rate-distortion loss, $\lambda \times (0.8\mathrm{MSE}_{\mathrm{ROI}} + 0.2\mathrm{MSE}_{\mathrm{NROI}}) + R$, to the input quality map because there is no straightforward way to determine its values for simultaneous rate and spatial quality control. In Fig. 3c, our feed-forward-based approach achieves nearly identical rate-distortion performance to the more complicated back-prop-based optimization for CNN+SFT (with 3 and 5 update steps).

3.2. Analyses of Prompt-based Conditioning

Rate and Spatial Quality Disentanglement. Fig. 3d demonstrates the effectiveness of our disentanglement of the rate and spatial quality control. The rate-distortion segments of different colors correspond to different choices of $\mathbf{M}_\lambda$. The rate-distortion points of the same colored segment are obtained by setting $\mathbf{M}_R$ in the ROI region to 0.25, 0.5, 0.75, and 1 (and to 0 otherwise). We see that $\mathbf{M}_\lambda$ determines where the major rate point is while $\mathbf{M}_R$ contributes to local bit-rate variations.

Attention Maps. Fig. 5 visualizes the attention maps for image tokens in different P-STB blocks of the encoder. Each attention map reveals how the prompt tokens attend collectively to every image token. That is, for every image token, we visualize the weighting factors summed over all prompt tokens in the same window. We see that the prompt tokens contribute more distinctively to the ROI and non-ROI regions in the P-STB blocks closer to the input image. In the deeper layers, the distinction between the ROI and non-ROI regions becomes less obvious.

Complexity Comparison. Table 1 compares the multiply-accumulate-operations per pixel (kMACs/pixel) and model size of different methods. Even though our model is larger than CNN+SFT [7] and SWIN+SFT due to the higher number of channels in p_a, p_s, it has lower kMACs/pixel. This is because we process the conditioning signal in lower-resolution feature maps due to a smaller number of prompt tokens, while [7] needs to generate affine parameters in the same resolution as the input image. Note that our variable-rate model can be more cost-effective than SWIN+FRC, for which separate models must be trained for different bit rates. That is, its effective model size is a multiple of 17.66M depending on the number of supported rate points.

4. CONCLUSION

This work proposes a transformer-based image compression system. It features prompt generation networks to adapt the autoencoder for simultaneous variable-rate and ROI coding. The major finding is that our content-adaptive prompt tuning is more effective than spatial feature transform (SFT) in terms of adapting the transformer-based autoencoder. It also incurs lower kMACs/pixel than SFT.

5. REFERENCES

[1] Ming Lu, Peiyao Guo, Huiqing Shi, Chuntong Cao, and Zhan Ma, "Transformer-based image compression," in *Proceedings of the Data Compression Conference (DCC)*. IEEE, 2022.

[2] Yinhao Zhu, Yang Yang, and Taco Cohen, "Transformer-based transform coding," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022.

[3] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, 2021.

[4] Fei Yang, Luis Herranz, Joost van de Weijer, Jose A. Iglesias Guitian, Antonio M. Lopez, and Mikhail G. Mozerov, "Variable rate deep image compression with modulated autoencoder," *IEEE Signal Processing Letters*, vol. 27, pp. 331–335, 2020.

[5] Ze Cui, Jing Wang, Shangyin Gao, Tiansheng Guo, Yihui Feng, and Bo Bai, "Asymmetric gained deep image compression with continuous rate adaptation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

[6] Xining Wang, Ming Lu, and Zhan Ma, "Block-level rate control for learnt image coding," in *Proceedings of the Picture Coding Symposium (PCS)*. IEEE, 2022.

[7] Myungseo Song, Jinyoung Choi, and Bohyung Han, "Variable-rate deep image compression through spatially-adaptive feature transform," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

[8] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy, "Recovering realistic texture in image super-resolution by deep spatial feature transform," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.

[9] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023.

[10] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim, "Visual prompt tuning," in *Proceedings of European Conference on Computer Vision (ECCV)*, 2022.

[11] Jiaheng Liu, Guo Lu, Zhihao Hu, and Dong Xu, "A unified end-to-end framework for efficient deep image compression," *arXiv preprint arXiv:2002.03370*, 2020.

[12] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick, "Microsoft coco: Common objects in context," in *Proceedings of European Conference on Computer Vision (ECCV)*, 2014.

[13] Eastman Kodak, "Kodak lossless true color image suite (PhotoCD PCD0992)," .