

LIDAR DEPTH MAP GUIDED IMAGE COMPRESSION MODEL

Alessandro Gnutti^{*}, Stefano Della Fiore[†], Mattia Savardi[§]
Yi-Hsin Chen[‡], Riccardo Leonardi^{*}, and Wen-Hsiao Peng[‡]

^{*} Department of Information Engineering, CNIT - University of Brescia, Italy

[†] Department of Informatics, University of Salerno, Italy

[§] DSMC, University of Brescia, Italy

[‡] Department of Computer Science, National Yang Ming Chiao Tung University, Hsinchu, Taiwan

ABSTRACT

The incorporation of LiDAR technology into some high-end smartphones has unlocked numerous possibilities across various applications, including photography, image restoration, augmented reality, and more. In this paper, we introduce a novel direction that harnesses LiDAR depth maps to enhance the compression of the corresponding RGB camera images. To the best of our knowledge, this represents the initial exploration in this particular research direction. Specifically, we propose a Transformer-based learned image compression system capable of achieving variable-rate compression using a single model while utilizing the LiDAR depth map as supplementary information for both the encoding and decoding processes. Experimental results demonstrate that integrating LiDAR yields an average PSNR gain of 0.83 dB and an average bitrate reduction of 16% as compared to its absence.

Index Terms— Learned image compression, LiDAR, depth map, prompts, transformer.

1. INTRODUCTION

LiDAR, an acronym for Light Detection and Ranging, is a technology that utilizes laser beams to measure distances and create detailed 3D maps of the surrounding environment [1]. Apple has recently introduced LiDAR technology into some of its products, including the iPad Pro, which was the first to incorporate this technology in March 2020, followed by the iPhone 12 Pro and iPhone 12 Pro Max (and later editions) [2]. The purpose of the LiDAR scanner is depth sensing. It typically emits laser pulses and measures the time they take to bounce back from objects in the environment, thereby generating a precise 3D map of the surroundings through specialized algorithms.

This study was carried out within the LICAM project – funded by European Union – Next Generation EU within the PRIN 2022 program (D.D. 104 - 02/02/2022 Ministero dell’Università e della Ricerca). This manuscript reflects only the authors’ views and opinions and the Ministry cannot be considered responsible for them.

The integration of LiDAR technology in Apple’s devices has opened up various possibilities. It significantly enhances augmented reality experiences, making them more immersive and accurate. In fact, the precise depth data obtained from LiDAR allows augmented reality applications to better understand the physical world and seamlessly overlay digital objects onto it. This technology has been employed in a variety of ARKit applications and games, which have seen improved performance and user experience. Furthermore, the LiDAR sensor has also improved photography. Its capabilities extend to enhancing autofocus, especially in low-light conditions, ensuring that users can capture sharp images even in challenging environments. It also contributes to better portrait mode effects and augmented reality photography. Beyond photography and augmented reality, the LiDAR sensor aids in object and scene recognition, with applications in accessibility features designed to assist users with visual impairments in understanding their surroundings.

Of all these applications, none has ever delved into the potential of using the LiDAR depth map to enhance the compression of RGB camera images. The field of image compression has seen continuous research for many decades, and its importance has grown substantially, especially with the widespread use of mobile devices for capturing and sharing images. Lossy image compression, in particular, plays a crucial role in this context, as it efficiently reduces the required storage space and transmission bandwidth, even though it comes at the cost of some degradation in the quality of reconstructed images. Deep learning-based techniques [3, 4, 5, 6, 7, 8, 9] have emerged as front-runners, surpassing traditional codecs such as JPEG [10], JPEG2000 [11], and BPG [12]. Many of these learning-based approaches leverage autoencoder networks to perform nonlinear transformations [13] and optimize the rate-distortion trade-off [14], making them promising candidates for the next generation of image compression.

In this paper, we propose a novel approach where the LiDAR depth map is leveraged to enhance the compression of the corresponding RGB camera image. To the best of our

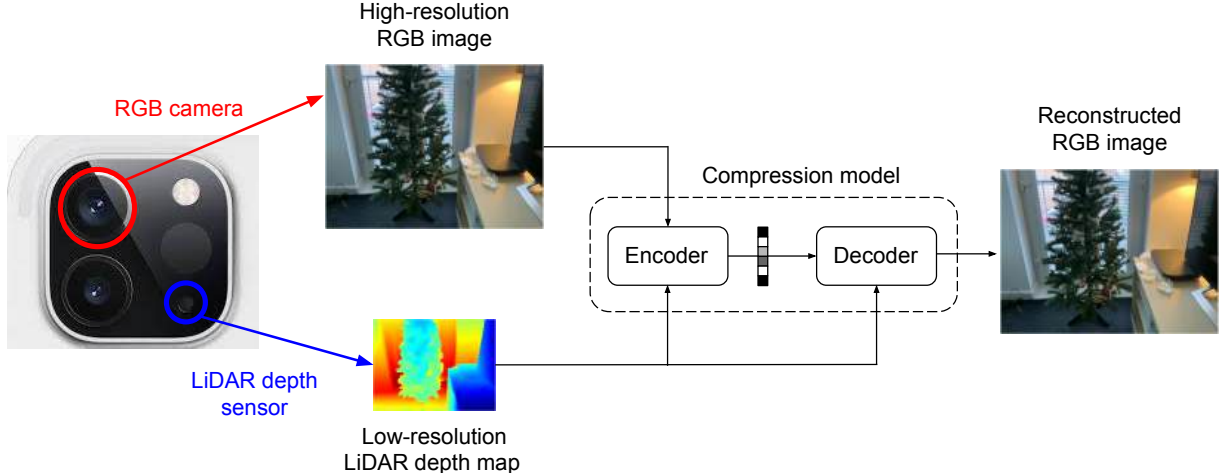


Fig. 1: High-level architectural framework. Our LiDAR Depth Map Guided Image Compression Model takes input from both the RGB image and the LiDAR depth data, using the LiDAR map to enhance the compression process of the RGB image. In the considered scenario, the LiDAR map is available at both the encoder and the decoder.

knowledge, this marks the initial exploration of this particular direction. Given the uncharted nature of this field, we opt to begin by considering the simple scenario in which LiDAR information is available both at the encoder and the decoder. To evaluate our approach, we choose to adopt the Swin-Transformer-based image compression system proposed in [15] as the reference architecture for our experiments. In fact, in the realm of learned image compression systems, Transformers have emerged as a compelling alternative to convolutional neural networks (CNNs). Their attention-based convolution, in combination with the shifted-windowing technique, offers a balance of superior compression performance and reduced computational requirements. Thus, we opt to incorporate LiDAR information into this architecture, which stands as one of the current state-of-the-art methods for learned image compression. Experimental results show that the integration of LiDAR technology results in an average improvement of 0.83 dB in PSNR-RGB performance and an average reduction of 16% in bit rate when compared to its absence.

The rest of the paper is structured as follows. Sec. 2 depicts the specific scenario under investigation. Sec. 3 introduces the architecture proposed for incorporating LiDAR information, and Sec. 4 shows the experimental results. Finally, Sec. 5 provides the concluding remarks.

2. EXAMINED SCENARIO

Let us consider a simple scenario to assess the effectiveness of using a LiDAR map in compressing an RGB image, that is, we assume that the LiDAR map is available at both the encoder and the decoder, as illustrated in Fig. 1. In light of this scenario, we aim to investigate two distinct cases.

As a first case, we consider to have the LiDAR map available in its original **uncompressed** format. Thus, we examine the situation where the RGB image and the corresponding LiDAR map are simultaneously acquired from the device, and the uncompressed LiDAR map is stored without compression on the acquisition device, as it will be utilized for the various imaging manipulations previously mentioned. Then, the uncompressed LiDAR data will assist in both encoding the RGB image, which will be then locally stored on the device, and decoding the reconstructed version. In essence, this approach capitalizes on the presence of the LiDAR map on the device to enhance the compression of the RGB image. In this context, we will not consider the rate cost of the LiDAR map, as its presence on the device is independent of the RGB image compression goal. Of course, this initial case is confined to an encoding and decoding process performed locally on the device.

Secondly, we examine the case in which the RGB image must be transmitted to a second device. As transmitting the uncompressed LiDAR map solely for image replication on a remote device would be impractical due to the high transmission cost, we explore the impact of employing a **compressed** LiDAR map within the considered image compression framework and its associated rate cost. Thus, in this context, the LiDAR map is first compressed independently, and then the compressed version is used to assist the encoding on the first device and the decoding on the second device of the RGB image. It is important to highlight that, for the image encoding process, we would have the option to employ the uncompressed LiDAR map, assuming it has been stored on the acquisition device without compression (or at a very high quality). However, we opt to use the compressed LiDAR map for assisting the image encoding as well, ensuring

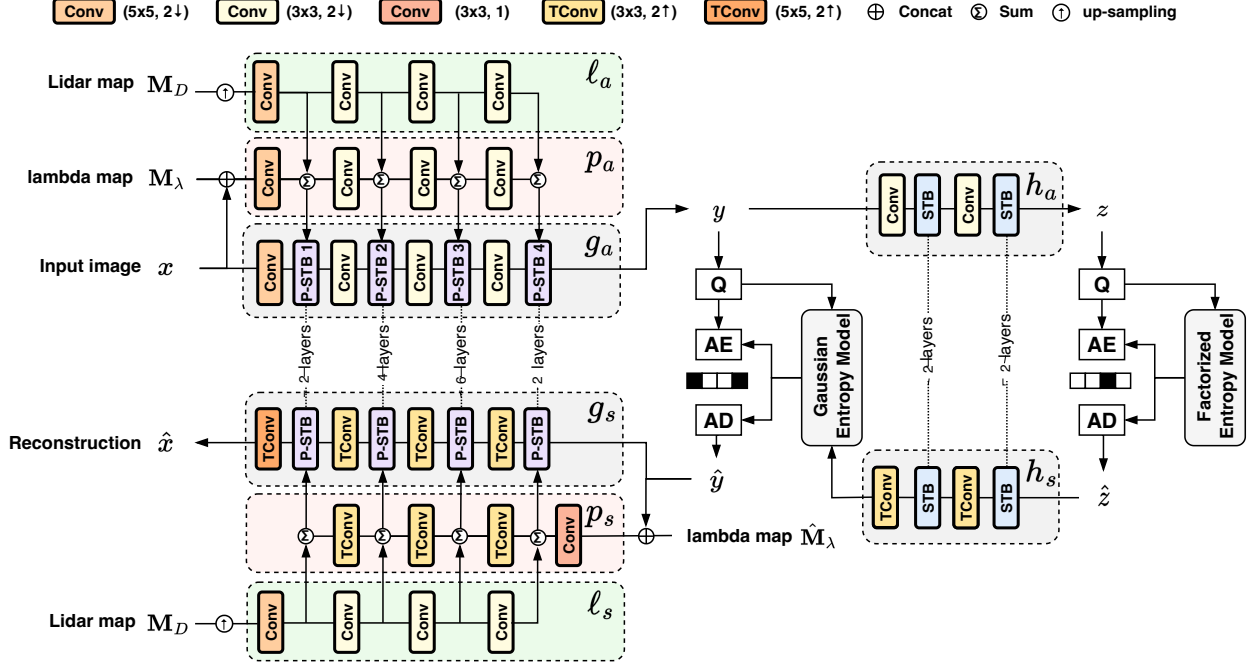


Fig. 2: The network architecture of the proposed LiDAR depth map guided image compression model. For variable rate control, the prompt generation networks p_a and p_s produce prompt tokens for the encoder g_a and decoder g_s , respectively. Two additional prompt generation networks, denoted as l_a and l_s , receive the LiDAR map as input and generate prompt tokens, which are subsequently combined at the convolutional layers of the prompt networks p_a and p_s , respectively.

uniformity in the LiDAR data used for both image encoding and decoding. It is also important to note that incorporating a compressed depth map into the bitstream typically results in a marginal increase in coding cost, even due to the considerably lower spatial resolution of the LiDAR map compared to the RGB image (see Sec. 4.1). A similar framework is discussed in [16], where indeed a semantic/instance label map of the original image is utilized as side information at both the encoder and decoder for extreme low rate coding.

We would like to remark that in the case of transmitting the LiDAR map alongside, an effective solution might involve jointly compressing both the RGB image and the LiDAR depth map. This approach would entail designing a model that takes both the RGB image and LiDAR map as inputs, encoding both of them, and producing both their reconstructed versions as outputs. Alternatively, one could envision a model that uses the LiDAR map exclusively at the encoder side, without involving it in the decoding process. In fact, despite the established principle derived from Berger’s result, indicating that side information, when solely available at the encoder and not influencing the distortion measure, does not offer benefits [17], one could explore whether incorporating LiDAR data at the encoder could improve the convergence of an end-to-end learning model and potentially result in a more efficient data compression model. We leave these potential investigations for future works.

3. PROPOSED METHOD

Following the approach presented in [15], we propose to adopt their Swin-Transformer-based image compression system as our reference architecture. The structure of our LiDAR depth map guided image compression model is shown in Fig. 2.

3.1. Background: reference architecture

The system introduced in [15] is an evolution of TIC (Transformer-based Image Compression [18, 19]), but it omits the inclusion of a context model for entropy coding. The core components of this system consist of a main autoencoder, denoted as g_a and g_s , and a hyperprior autoencoder, represented as h_a and h_s , which incorporate Swin-Transformer blocks (STB) interspersed with convolutional layers. More details about STBs are presented in [19].

As usually performed in learned end-to-end image compression frameworks, the analysis transform g_a encodes an image $x \in \mathbb{R}^{3 \times H \times W}$ into its latent representation $y \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times 192}$. This latent representation y undergoes uniform quantization as \hat{y} and is further encoded into a bitstream using a learned prior distribution. On the decoder side, the bitstream is entropy decoded to reconstruct $\hat{x} \in \mathbb{R}^{3 \times H \times W}$ via the synthesis transform g_s . The prior distribution significantly impacts the required number of bits to signal the

quantized latent \hat{y} . Therefore, it is modeled in a content-adaptive manner by a hyperprior autoencoder, consisting of a hyperprior analysis transform h_a and a hyperprior synthesis transform h_s . As illustrated in Fig. 2, h_a converts the latent y into $z \in \mathbb{R}^{\frac{H}{64} \times \frac{W}{64} \times 192}$, typically representing a small portion of the compressed bitstream. Its quantized version \hat{z} is ultimately decoded from the bitstream through h_s . A factorized entropy method is applied for \hat{z} , while a gaussian entropy model is used to characterize \hat{y} for probability estimation (see [4]).

In [15], to encode the input image x , the encoder requires two additional inputs: a lambda map $M_\lambda \in \mathbb{R}^{1 \times H \times W}$ and an ROI mask. The lambda map M_λ is a uniform map with a parameter $m_\lambda \in [0, 1]$, which regulates the bit rate of the compressed bitstream. The ROI mask specifies the spatial importance of individual pixels in the image. Both of these inputs serve as conditioning signals that are used to generate prompt tokens for adapting the primary encoder g_a . Likewise, the decoder g_s is adapted by receiving inputs in the form of the quantized latent \hat{y} and a downsampled lambda map $\hat{M}_\lambda \in \mathbb{R}^{1 \times \frac{H}{16} \times \frac{W}{16}}$, which matches the spatial resolution of \hat{y} .

These prompts enable variable-rate and ROI coding and draw inspiration from [20]. The outcome is a modified Swin-Transformer block termed prompted Swin-Transformer block (P-STB). The generation of these learned prompts involves two networks: p_a and p_s , conditioning the encoder g_a and decoder g_s , respectively. The p_a network comprises several convolutional layers matching those of the encoder g_a and takes as input the concatenation of the ROI mask, the lambda map M_λ , and the image x . The feature maps generated by p_a are channeled into the respective P-STBs to produce prompt tokens, which subsequently interact with image tokens. On the other hand, p_s follows a similar architectural pattern, replacing the convolutional layers with transposed convolutional layers for upsampling. More details can be found in [15].

3.2. LiDAR Depth Map Guided Image Compression Model

Fig. 2 depicts our image compression model guided by LiDAR depth maps. It draws inspiration from [15]. However, two significant distinctions set our approach apart from [15]. Firstly, our method does not incorporate the ROI mask, as it plays no role in our model, in contrast to the approach outlined in [15], which proposes a ROI-based coding. Secondly, the distinctive feature of our model lies in the inclusion of two supplementary prompt generation networks, referred to as l_a and l_s situated at the encoder and decoder sides, respectively. Both of them receive as input the LiDAR depth map $M_D \in \mathbb{R}^{1 \times H \times W}$ and generate prompt tokens, which are then added to the prompt tokens of p_a and p_s , respectively. Note that the LiDAR depth map M_D undergoes an initial upsampling process to align its size with that of the RGB image.

This is necessary because the resolution of LiDAR sensors is typically smaller (256×192 for Apple’s devices), compared to the RGB camera.

We would like to remark that in [15], the ROI map information is concatenated with the lambda map M_λ (only at the encoder side). However, we found that concatenating the LiDAR depth map with M_λ does not allow the model to significantly exploit the depth information. Instead, incorporating the specific prompts l_s and l_a at both the encoder and decoder sides, which receive both M_D , enables the system to effectively leverage the LiDAR depth information (as we will show in the next section).

In training, the loss function is formulated to make the model responsive to the rate parameter m_λ , without incorporating any LiDAR information into the loss function itself. Thus, it is defined as a weighted combination of both distortion and bit rate components:

$$J_{rad}(x) = \lambda \sum_{i=1}^N \frac{(x_i - \hat{x}_i)^2}{N} + R, \quad (1)$$

where x_i and \hat{x}_i are the i -th pixel in the original and reconstructed images, respectively, N is the total number of pixels in the image, R is the bit rate and denotes the number of bits-per-pixel needed to represent the latent representations \hat{y} and \hat{z} , and $\lambda = f(m_\lambda) = (\lambda_{\max})^{m_\lambda} (\lambda_{\min})^{1-m_\lambda}$ 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.

4. EXPERIMENTS

4.1. Dataset

The choice of datasets is critical to ensure they accurately represent the smartphone-based sensor fusion scenario. This poses a challenge because smartphones equipped with both cameras and LiDAR sensors have only become available recently, and traditional RGB-D datasets were primarily designed for dedicated devices like the Microsoft Kinect. Among the existing datasets, ARKitScenes [21], constructed using the built-in cameras and LiDAR sensors in Apple smartphones and tablets, is the only one that aligns with our scope.

The dataset consists of indoor scene images captured using the 2020 iPad Pro, which was utilized to collect both RGB images from the RGB camera and dense depth maps from the LiDAR scanner. The dataset comprises a total of 44,599 acquisitions, each one including both the low-resolution and high-resolution RGB images (256×192 and 1920×1440 pixels, respectively), and the corresponding low-resolution depth maps (256×192 pixels). To work with high-resolution data, we performed upsampling on the low-resolution depth maps, increasing their resolution by a factor of 7.5. In line with the existing division in the ARKitScenes dataset, we partitioned




Fig. 3: Comparison of the model performance with and without LiDAR in terms of average RD curves, evaluating PSNR (a) and SSIM (b) as distortion metrics. The blue curve corresponds to the model’s performance when given the uncompressed LiDAR map, while the red curve represents the model’s performance when supplied with the compressed LiDAR map, which includes the cost of transmitting LiDAR information. The results are evaluated on the 5,999 images in the test set.

the dataset into a training set, which includes 39,000 images, and a test set, consisting of 5,599 images.

4.2. Results

We conduct a single-stage training of our model using the ARKitScenes dataset’s training set. During each training iteration, input images are randomly cropped to a size of 256×256 and the uncompressed LiDAR maps undergo cropping in corresponding positions. Our entire model is trained jointly for variable rate coding by uniformly sampling m_λ from 0 to 1. To ensure fairness in comparison, we adhere to the same training procedure for the model without LiDAR prompts; that is, the model presented in [15] when considering uniform ROI maps.

Fig. 3 shows the rate-distortion (RD) curves, analyzing both PSNR-RGB (used in the training loss function) and SSIM, shown separately in Fig. 3a and Fig. 3b. The blue curve represents the model adopting the uncompressed LiDAR for both the encoding and decoding stages. As it can be observed, this model consistently outperforms the model without LiDAR for both metrics. In terms of Bjøntegaard’s metric, integrating uncompressed LiDAR results in an average PSNR gain of 0.83 dB and an average bit rate reduction of 16% compared to its absence.

Furthermore, the red curve corresponds to the performance of our model when considering the case in which the uncompressed LiDAR data is unavailable, such as during transmission. In this context, the LiDAR information needs to be compressed separately. In our experiments, it undergoes

separate compression using our model, where we specifically provide zero matrices as input to the prompts l_a and l_s and set $m_\lambda = 1$ to maximize the quality of the compressed LiDAR map. Then, the compressed LiDAR map is used for assisting the RGB image encoding and decoding. Of course, the final bitstream will consist of both the latent spaces of the RGB image and the LiDAR map. Interestingly, the resulting performance is comparable to that achieved by the model using the original LiDAR map, both in terms of PSNR and SSIM. In fact, the increase in coding cost when encoding the LiDAR depth map is only marginal. On average, it accounts for about 9% of the total cost, which includes both the RGB image and the LiDAR map. This is due to the significantly lower spatial resolution of the LiDAR map compared to the RGB image. Moreover, using the compressed LiDAR map in place of the original one does not significantly affect the compression process. Specifically, in terms of Bjøntegaard’s metric, integrating compressed LiDAR results in an average PSNR gain of 0.57 dB and an average bit rate reduction of 10% compared to its absence (see Tab. 1 for recap). Notably, we remark that we have not retrained the model using the compressed LiDAR maps.

To better appreciate the visual quality, Fig. 4 displays the reconstructed images for two samples from the test set (Figs. 4a and 4b). Specifically, it showcases the results using our model with LiDAR input (Figs. 4c and 4d) and the model without LiDAR (Figs. 4e and 4f). Additionally, to highlight the effectiveness of LiDAR integration within our model, Figs. 4g and 4h display the reconstructed images from our model when a random map serves as the side information.




Fig. 4: Visual results for two images from the test set are displayed in (a)-(b). The reconstructed images using LiDAR as side information are depicted in (c)-(d), while those reconstructed without LiDAR are shown in (e)-(f). The reconstructed images using a random map as side information are seen in (g)-(h). Corresponding LiDAR maps are represented in (i)-(j).

	Uncompressed LiDAR	Compressed LiDAR
Δ PSNR [dB]	0.83	0.57
Δ rate [%]	-16	-10

Table 1: *BD* rate associated to Fig. 3a. The baseline performance is set with the model that does not use the LiDAR map (no LiDAR).

Evidently, there is a notable decrease in visual quality in these instances, indicating the significant impact of LiDAR data on the compression process.

5. CONCLUSION

In this paper, we have shown that integrating the LiDAR depth map within the compression process of an RGB camera image results in substantial performance improvements. As the use of LiDAR in mobile devices continues to grow, our work can serve as a starting point for a new research direction. There are some aspects that deserve further exploration in our study: (i) extending the experiments to include additional and more variegated RGB-LiDAR datasets, which are expected to increase considering the growing prevalence of LiDAR technology; (ii) investigating the universality of the method by integrating the LiDAR map into various image compression frameworks and assessing its efficiency; (iii) while our findings demonstrate that utilizing a decoded LiDAR map at the decoder side does not compromise the performance compared to the model using the original map, we believe that employing a combined compression approach for both the image and

LiDAR map could potentially yield even greater enhancements. We leave this possibility for exploration in future research endeavors.

6. REFERENCES

- [1] Pinliang Dong and Qi Chen, *LiDAR remote sensing and applications*, CRC Press, 2017.
- [2] “Apple unveils new ipad pro with breakthrough lidar scanner and brings trackpad support to ipados,” <https://www.apple.com/newsroom/2020/03/apple-unveils-new-ipad-pro-with-lidar-scanner-and-trackpad-support-in-ipados/>, [Online; accessed 07-November-2023].
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli, “End-to-end optimized image compression,” *arXiv preprint arXiv:1611.01704*, 2016.
- [4] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston, “Variational image compression with a scale hyperprior,” *arXiv preprint arXiv:1802.01436*, 2018.
- [5] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee, “Variable rate deep image compression with a conditional autoencoder,” in *Proc. of the IEEE/CVF ICCV*, 2019, pp. 3146–3154.
- [6] Tong Chen, Haojie Liu, Zhan Ma, Qiu Shen, Xun Cao, and Yao Wang, “End-to-end learnt image compression

- via non-local attention optimization and improved context modeling,” *IEEE Transactions on Image Processing*, vol. 30, pp. 3179–3191, 2021.
- [7] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool, “Conditional probability models for deep image compression,” in *Proc. of the IEEE Conference on CVPR*, 2018, pp. 4394–4402.
- [8] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson, “High-fidelity generative image compression,” *Advances in Neural Inf. Proc. Syst.*, vol. 33, pp. 11913–11924, 2020.
- [9] David Minnen and Saurabh Singh, “Channel-wise autoregressive entropy models for learned image compression,” in *2020 IEEE ICIP*. IEEE, 2020, pp. 3339–3343.
- [10] William B Pennebaker and Joan L Mitchell, *JPEG: Still image data compression standard*, Springer Science & Business Media, 1992.
- [11] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi, “The jpeg 2000 still image compression standard,” *IEEE Signal processing magazine*, vol. 18, no. 5, pp. 36–58, 2001.
- [12] F Bellard, “Bpg image format.,” 2014.
- [13] Johannes Ballé, Philip A Chou, David Minnen, Saurabh Singh, Nick Johnston, Eirikur Agustsson, Sung Jin Hwang, and George Toderici, “Nonlinear transform coding,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 2, pp. 339–353, 2020.
- [14] Lee D Davisson, “Rate-distortion theory and application,” *Proc. of the IEEE*, vol. 60, no. 7, pp. 800–808, 1972.
- [15] Chia-Hao Kao, Ying-Chieh Weng, Yi-Hsin Chen, Wei-Chen Chiu, and Wen-Hsiao Peng, “Transformer-based variable-rate image compression with region-of-interest control,” *arXiv preprint arXiv:2305.10807*, 2023.
- [16] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool, “Generative adversarial networks for extreme learned image compression,” in *Proc. of the IEEE/CVF ICCV*, 2019, pp. 221–231.
- [17] T. Berger, *Rate Distortion Theory: A Mathematical Basis For Data Compression*, Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [18] Ming Lu, Peiyao Guo, Huiqing Shi, Chuntong Cao, and Zhan Ma, “Transformer-based image compression,” in *2022 Data Compression Conference (DCC)*. IEEE, 2022, pp. 469–469.
- [19] Ming Lu, Peiyao Guo, Huiqing Shi, Chuntong Cao, and Zhan Ma, “Transformer-based image compression,” *arXiv preprint arXiv:2111.06707*, 2021.
- [20] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim, “Visual prompt tuning,” in *ECCV*. Springer, 2022, pp. 709–727.
- [21] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al., “Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data,” *arXiv preprint arXiv:2111.08897*, 2021.