

# CAMSIC: Content-aware Masked Image Modeling Transformer for Stereo Image Compression

Xinjie Zhang<sup>1,2\*</sup>, Shenyuan Gao<sup>1</sup>, Zening Liu<sup>1</sup>, Jiawei Shao<sup>1,3</sup>,  
Xingtong Ge<sup>2</sup>, Dailan He<sup>4</sup>, Tongda Xu<sup>5</sup>, Yan Wang<sup>5</sup>, Jun Zhang<sup>1†</sup>

<sup>1</sup>The Hong Kong University of Science and Technology

<sup>2</sup>SenseTime Research

<sup>3</sup>Institute of Artificial Intelligence (TeleAI), China Telecom

<sup>4</sup>The Chinese University of Hong Kong

<sup>5</sup>Institute for AI Industry Research (AIR), Tsinghua University

{xzhangga, sygao, zhening.liu}@connect.ust.hk, shaojw2@chinatelecom.cn, xingtong.ge@gmail.com, hedailan@link.cuhk.edu.hk, x.tongda@nyu.edu, wangyan202199@163.com, eejzhang@ust.hk

## Abstract

Existing learning-based stereo image codec adopt sophisticated transformation with simple entropy models derived from single image codecs to encode latent representations. However, those entropy models struggle to effectively capture the spatial-disparity characteristics inherent in stereo images, which leads to suboptimal rate-distortion results. In this paper, we propose a stereo image compression framework, named CAMSIC. CAMSIC independently transforms each image to latent representation and employs a powerful decoder-free Transformer entropy model to capture both spatial and disparity dependencies, by introducing a novel content-aware masked image modeling (MIM) technique. Our content-aware MIM facilitates efficient bidirectional interaction between prior information and estimated tokens, which naturally obviates the need for an extra Transformer decoder. Experiments show that our stereo image codec achieves state-of-the-art rate-distortion performance on two stereo image datasets Cityscapes and InStereo2K with fast encoding and decoding speed.

## Introduction

Stereo Image Codec (SIC) compresses a pair of stereoscopic images captured from distinct viewpoints by the same camera. SIC has attracted significant interest due to the growing demand for high-quality stereo image transmission and storage across applications such as autonomous driving (Yin et al. 2020), virtual reality (Fehn 2004) and video surveillance (Stepanov and Tishchenko 2016). By leveraging the inherent correlation between views, they achieve higher coding efficiency compared with single image codecs.

Over the past decades, classical multi-view image coding standards, such as H.264-based MVC (Vetro, Wiegand, and Sullivan 2011) and H.265-based MV-HEVC (Tech et al. 2015), have catalyzed the development of learning-based

\*This work was partially performed when Xinjie Zhang was an Intern at SenseTime.

†Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

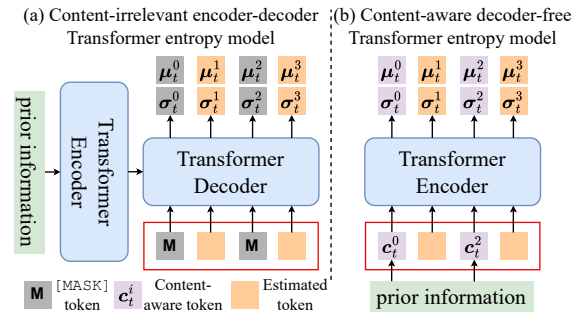


Figure 1: Comparison of two different mask image modeling Transformer entropy models. [MASK] token is a learnable parameter, which is irrelevant to specific image content. Tokens within the same red rectangle have bidirectional interactions. The vanilla content-irrelevant masked image modeling uses uniform [MASK] tokens to fill the unestimated positions and relies on an inefficient encoder-decoder Transformer to propagate the prior information. Differently, our content-aware masked image modeling replaces the uninformative [MASK] tokens with proposed content-aware tokens, which enables an efficient decoder-free Transformer architecture with more sufficient information interactions.

stereo image compression approaches (Liu, Wang, and Urtasun 2019; Deng et al. 2021; Wödlinger et al. 2022; Zhai et al. 2022; Deng et al. 2023). Given a reference view, these methods, rooted in the predictive coding paradigm, employ disparity-compensated prediction in either pixel or feature space to compress current view. Additionally, recent advancements have introduced a bidirectional coding framework (Lei et al. 2022; Wödlinger et al. 2024) that utilizes a cross-attention mechanism, further enhancing content correlation exploitation between stereo images. Albeit remarkable, the odyssey of stereo image compression research concentrates on producing compact representations through various information flows or network designs before entropy coding. However, an orthogonal direction, which reduces redundancy by devising a superior spatial-disparity entropy

model for entropy coding, is rarely explored.

Previous learning-based SIC approaches often adapt existing single image compression solutions, such as hyperprior (Ballé et al. 2018) and auto-regressive prior (Minnen, Ballé, and Toderici 2018), into conditional formats. These adapted entropy models, typically based on convolutional neural network (CNN) architectures, aim to capture spatial-disparity correlations. However, CNNs are inherently limited by their local receptive fields, which hampers their ability to model long-range dependencies. This limitation often leads to a rough estimation of probability distribution and inferior compression performance, especially when the disparity distance between stereo images is large. Consequently, we argue that better coding gains can be achieved as long as we find a more effective spatial-disparity entropy model.

In this paper, we present a stereo image compression framework that centers on a powerful Transformer entropy model to leverage spatial-disparity dependencies between current and previously decoded views. This framework incorporates a straightforward image encoder-decoder pair from He et al. (2022a) to extract latent representations for each view. As illustrated in Fig. 1 (a), we introduce the vanilla masked image modeling (MIM) style (Chang et al. 2022) to construct a content-irrelevant encoder-decoder Transformer entropy model. This model captures spatial information from the context of decoded tokens and disparity information from previous decoded view. Since the bi-directional self-attention in Transformer allows [MASK] tokens to use context from both directions, it supports a non-sequential auto-regressive decoding process and decodes all tokens in the images in a few steps, achieving a good performance-speed trade-off (Chang et al. 2022).

Nevertheless, the vanilla MIM-based Transformer entropy model still exhibits two significant shortcomings. Firstly, it uses uniform [MASK] tokens to indicate the positions of the tokens whose probability distributions are yet to be estimated, which are irrelevant to the image content and uninformative to entropy coding. As a result, it wastes much computation and leads to inferior coding gains. Secondly, it needs an extra Transformer decoder to propagate the prior information via cross-attention. Such a complicated encoder-decoder Transformer architecture only permits an inefficient unidirectional propagation from the prior information to the input tokens of the Transformer decoder.

To overcome these challenges, we develop a novel **Content-Aware Masked image modeling Transformer entropy model for Stereo Image Compression (CAMSIC)** shown in Fig. 1 (b). Specifically, we introduce content-aware tokens generated by prior information to replace the uninformative [MASK] tokens, forming a new content-aware masked image modeling style, which brings three major benefits. First, the interactions between the proposed content-aware tokens and the estimated tokens provide useful information about each specific position, preventing distraction of previous uninformative [MASK] tokens when updating the token features. Second, it allows a bidirectional information flow between prior information and estimated tokens via self-attention. Compared with static prior information in the vanilla MIM, content-aware tokens, carrying

specific prior information, can also be iteratively updated. Thus, it can effectively capture the context information from already estimated tokens and reduce the estimation uncertainty of remaining tokens. Third, since our design naturally embeds prior information propagation into each self-attention operation, we discard the whole Transformer decoder and devise an encoder-only Transformer architecture with low computation overhead. In summary, our main contributions are three-fold:

- We introduce a learning-based stereo image compression framework with a simple image encoder-decoder pair, which uses an elegantly neat but powerful Transformer entropy model based on masked image modeling to exploit the relationship between the left and right images.
- We present a unique content-aware masked image modeling style for Transformer entropy model. This innovation enables more effective and extensive interactions between prior information and estimated tokens, while also facilitating an efficient yet strong decoder-free Transformer entropy model architecture design.
- Experimental results show that our proposed method with lower encoding and decoding latency significantly outperforms existing learning-based stereo image compression methods. Detailed ablation studies and analyses validate the effectiveness of each proposed component.

## Related Works

**Single Image Compression.** Traditional standard image codecs, such as JPEG (Wallace 1991), JPEG2000 (Skodras, Christopoulos, and Ebrahimi 2001), BPG (Bellard 2014), and VVC-intra (Bross et al. 2021), typically employ a three-step process involving transformation, quantization, and entropy coding. Regrettably, these steps are optimized independently, which may compromise overall compression performance. In contrast, recent learning-based image compression methods (Ballé, Laparra, and Simoncelli 2017; Ballé et al. 2018; Minnen, Ballé, and Toderici 2018; Cheng et al. 2020; Chen et al. 2021; Zhu, Yang, and Cohen 2021; Chen, Xu, and Wang 2022; Zou, Song, and Zhang 2022; He et al. 2022a; Liu, Sun, and Katto 2023) have demonstrated remarkable success, outperforming traditional methods in terms of rate-distortion (RD) performance. These modern approaches utilize nonlinear transformations (e.g., generalized divisive normalization (Ballé, Laparra, and Simoncelli 2015), attention mechanism (Cheng et al. 2020; Chen et al. 2021), and stacks of residual bottleneck blocks (He et al. 2022a)) to generate concise latent representations. Moreover, a series of advanced entropy models, such as factorized prior (Ballé, Laparra, and Simoncelli 2017), hyperprior (Ballé et al. 2018), and auto-regressive context prior (Minnen, Ballé, and Toderici 2018), are employed to accurately approximate the probability distribution of quantized latent representations. These existing works serve as foundational elements for our stereo image compression solution.

**Stereo Image Compression.** Traditional multi-view image codecs, such as H.264-based MVC (Vetro, Wiegand, and Sullivan 2011) and H.265-based MV-HEVC (Tech et al. 2015), typically compress stereo images using disparity

compensation prediction, which enjoys high compression efficiency but heavily depends on prior knowledge to design hand-crafted modules. Recently, learning-based stereo image compression methods have shown substantial improvements in compression performance. These existing works can be roughly classified into two categories: (1) Unidirectional coding (Liu, Wang, and Urtasun 2019; Deng et al. 2021; Wödlinger et al. 2022; Zhai et al. 2022; Deng et al. 2023) begins with the prediction of a disparity-compensated view in either pixel or feature space, followed by the compression of residuals between the current and predicted views. (2) Bidirectional coding (Lei et al. 2022; Wödlinger et al. 2024; Liu et al. 2025) incorporates a cross-attention mechanism within both the encoder and decoder to exploit the mutual information across stereo images. Apart from the above mentioned joint encoding methods, some recent works (Zhang, Shao, and Zhang 2022; Mital et al. 2023) have ventured into distributed multi-view image coding by using cross-attention alignment to achieve competitive RD performance. Contrary to the prevailing focus on intricate nonlinear transformations by exploring different network structures or information flows, we investigate the application of the Transformer architecture in entropy coding and propose a gracefully simple but powerful Transformer entropy model for stereo image compression. It effectively exploits spatial-disparity correlations, which eliminates the need for complex feature extraction or warping operations, thereby streamlining the compression process and establishing a neat yet potent coding framework.

**Masked Image Modeling.** Inspired by the masked language modeling (Kenton and Toutanova 2019) in natural language processing, masked image modeling (MIM) has been recently applied to representation learning for vision tasks (He et al. 2022b; Xie et al. 2022; Liu et al. 2022). It has shown great success on various downstream tasks, including image generation (Chang et al. 2022; Liang et al. 2022), action recognition (Tong et al. 2022; Feichtenhofer et al. 2022), video compression (Xiang, Tian, and Zhang 2023) and video prediction (Gupta et al. 2023). In this paper, we apply the MIM technique to stereo image compression. Note that the vanilla MIM style uses a uniform [MASK] token to take the places of the tokens whose probability distribution are not yet estimated, which wastes much computation on the content-irrelevant [MASK] tokens, thus leading to sub-optimal coding gains. In addition, it resorts to an extra Transformer decoder to propagate the prior information in a unidirectional manner. To address these limitations, we develop a novel content-aware MIM style, which replaces the uninformative [MASK] tokens with content-aware tokens. It not only enables a bidirectional interaction between the prior information and estimated tokens, but also allows us to design an efficient decoder-free Transformer entropy model.

**Transformer Architecture.** Recent advancements in Transformer architectures have primarily employed two configurations: (i) The encoder-decoder Transformer is tailored for sequence-to-sequence tasks such as machine translation and question answering. Despite its flexibility in generating varying output lengths, it is computationally intensive and requires substantial resources for training due to its dual-




Figure 2: Our overall stereo image compression framework. Given an input image  $x_v$ , we first encode the input view to produce the latent representation  $y_v$  that is quantized to  $\hat{y}_v$ . To transmit  $\hat{y}_v$  with fewer bits, an MIM-based Transformer entropy model is introduced to predict a probability distribution for  $\hat{y}_v$  given hyperprior  $\hat{z}_v$ , and previously stored representations  $\hat{y}_{v-1}$ . Finally, the quantized representation  $\hat{y}_v$  is decoded to the reconstructed view  $\hat{x}_v$ . When compressing the first view, a learned embedding  $\hat{y}_{-1}$  is used. The arithmetic encoder and decoder after the hyperprior encoder are omitted for brevity.

module structure. (ii) The decoder-free Transformer is suitable for input understanding tasks (e.g., sentence matching and feature extraction). This structure usually offers quicker training and lower computational demands. Given the effectiveness of Transformer in information fusion and interaction, recent compression practices like VCT (Mentzer et al. 2022) and MIMT (Xiang, Tian, and Zhang 2023) have incorporated encoder-decoder Transformer architectures into entropy models, significantly advancing compression performance. By contrast, we leverage a decoder-free Transformer to achieve a better speed-performance trade-off.

## Method

### Overview of CAMSIC

Our stereo image compression framework CAMSIC is shown in Fig. 2. Let  $\mathcal{X} = \{x_1, x_2\}$  denotes a pair of stereo images, where  $x_v \in \mathbb{R}^{H \times W \times 3}$  is the image at view  $v$  ( $v = 1, 2$ ) with height  $H$  and width  $W$ . An image encoder  $E$  is used to compress each individual view  $x_v$  to a latent representation  $y_v \in \mathbb{R}^{h \times w \times d}$ , where  $h = H/16$ ,  $w = W/16$  and  $d$  is the latent dimension.  $y_v$  is then quantized to  $\hat{y}_v$  by the quantizer  $Q$ . Finally, we feed  $\hat{y}_v$  to an image decoder  $D$  to reconstruct the view  $\hat{x}_v$ . Formally, the overall coding procedure can be formulated as:

$$y_v = E(x_v; \phi), \hat{y}_v = Q(y_v), \hat{x}_v = D(\hat{y}_v; \theta) \quad (1)$$

where  $\phi$  and  $\theta$  are learnable parameters of the image encoder and decoder, respectively. To further reduce the bits to be transmitted, the quantized representation  $\hat{y}_v$  is losslessly compressed by entropy coding, where the probability distribution of  $\hat{y}_v$  is estimated by a conditional entropy model. Given the spatial context  $s_v$ , hyperprior  $\hat{z}_v$ , and disparity prior  $\hat{y}_{v-1}$ , we model the probability of quantized representation  $\hat{y}_v$  as a Gaussian distribution:

$$p_{\hat{y}_v}(\hat{y}_v | \hat{y}_{v-1}, \hat{z}_v, s_v) = \prod_i (\mathcal{N}(\mu_v, \sigma_v^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}))(\hat{y}_v^i) \quad (2)$$

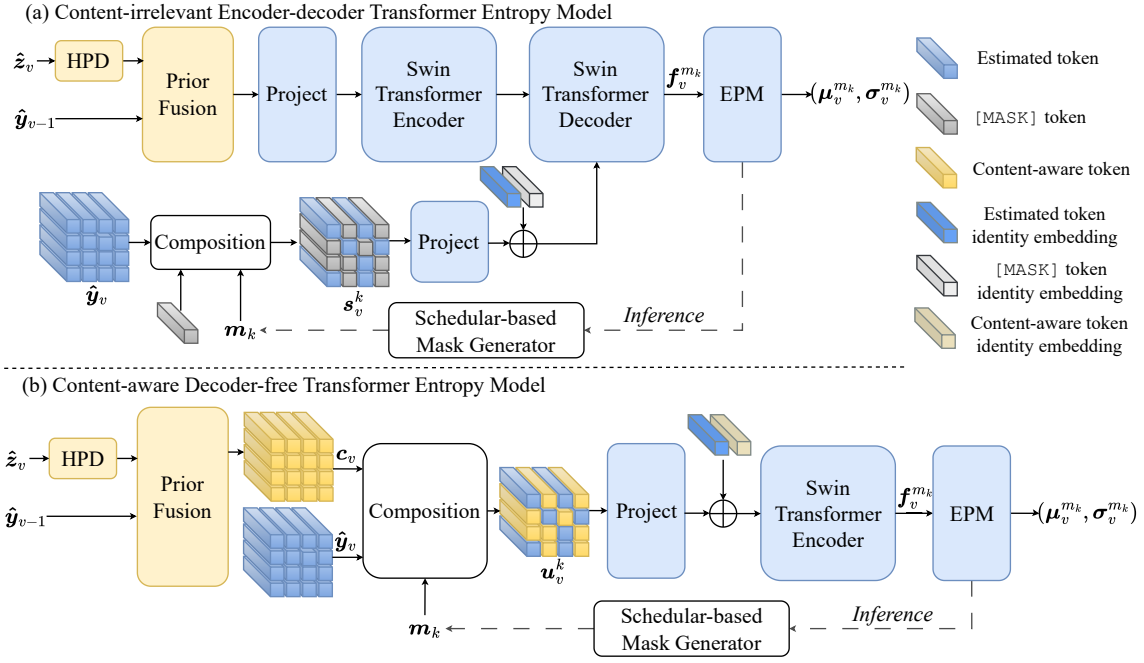


Figure 3: Structures of content-irrelevant encoder-decoder Transformer entropy model and our content-aware decoder-free Transformer entropy model. HPD indicates hyperprior decoder. EPM means entropy parameter module. The mask  $m_k$  starts at zero and is then updated iteratively by a scheduler-based generator during inference. In each iteration, a sinusoidal scheduler computes the number of tokens  $n_k$  to be encoded/decoded. The mask generator then sets the values of the  $n_k$  tokens with the lowest estimated bitrates and the previously encoded/decoded tokens as 1 to update  $m_k$ . (a) By using  $m_k$  to mix the [MASK] token and  $\hat{y}_v$ , a composite context  $s_v^k$  is synthesized and sent to a projection module, where two identity embeddings are added to further different token types. Both the prior information and the composite context are input into a Transformer decoder to estimate the probability distribution parameters  $(\mu_v^{m_k}, \sigma_v^{m_k})$  for the masked tokens. (b) The [MASK] token is replaced by the context-aware tokens  $c_v$  to form the composite context  $u_v^k$ . This context is then input into a projection module with two identity embeddings and a Transformer encoder to output the estimated Gaussian distribution parameters  $(\mu_v^{m_k}, \sigma_v^{m_k})$ .

where  $\mathcal{U}(\cdot)$  is uniform noise,  $*$  denotes convolution,  $\mu_v$  and  $\sigma_v^2$  are the estimated means and variances of the Gaussian distributions. When encoding the first view, we expand a learned embedding with the dimension as  $\mathbb{R}^{1 \times 1 \times d}$  to get  $\hat{y}_{-1} \in \mathbb{R}^{h \times w \times d}$ . In this paper, we refer to prior information as the combination of disparity prior  $\hat{y}_{v-1}$  and hyperprior  $\hat{z}_v$  to distinguish the spatial context from other priors. By leveraging complementary priors, the conditional entropy model can estimate more accurate probability distributions.

### Content-irrelevant Encoder-decoder Transformer Entropy Model

**Content-irrelevant Masked Image Modeling.** The input of the Transformer is a sequence of tokens. With some abuse of notation, we still refer to these tokens as  $\hat{y}_v$ , which is flattened from the latent representation. Given the quantized latent tokens  $\hat{y}_v$ , the content-irrelevant MIM defines the spatial context  $s_v$  as

$$s_v = \hat{y}_v \odot m + [\text{MASK}] \odot (1 - m) \quad (3)$$

where  $\odot$  denotes element-wise multiplication and  $m$  is a Boolean mask whose value is assigned as

$$m_i = \begin{cases} 1 & \text{if position } i \text{ is already estimated} \\ 0 & \text{if position } i \text{ is to be estimated} \end{cases} \quad (4)$$

Based on this assignment, [MASK] tokens take up the positions of the tokens whose probability distributions are to be estimated.

**Encoder-decoder Transformer.** Since these [MASK] tokens are irrelevant to the image content, a Transformer decoder is required to propagate the content-related prior information to the [MASK] tokens through cross-attention. Fig. 3 (a) shows the architecture of an encoder-decoder Transformer entropy model based on content-irrelevant MIM. Specifically, we use a prior fusion network and a Transformer encoder to integrate the hyperprior  $\hat{z}_v$  and disparity prior  $\hat{y}_{v-1}$  for producing prior features. During training, we randomly generate a Boolean mask to composite the spatial context  $s_v$ . The prior features and unmasked spatial tokens convey their information to the [MASK] tokens via cross-attention and self-attention in the Transformer decoder, respectively. Finally, we feed the feature  $f_v^m$  produced by the Transformer decoder into an entropy parameter module to estimate the probability distribution parameters  $(\mu_v^m, \sigma_v^m)$ .

### Content-aware Decoder-free Transformer Entropy Model

**Content-aware Masked Image Modeling.** The vanilla MIM uses a content-irrelevant [MASK] token to occupy the

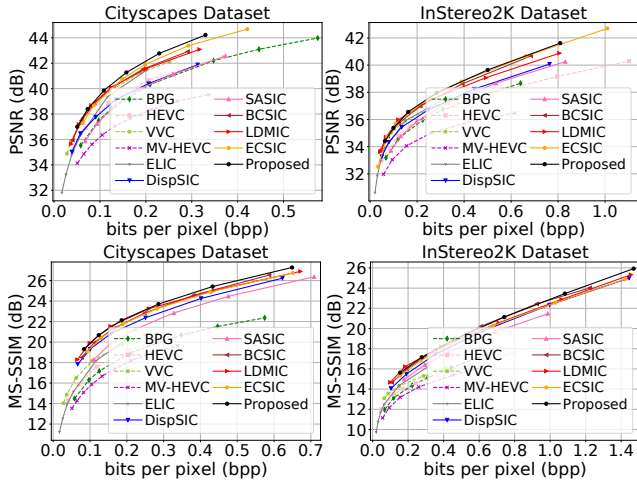


Figure 4: Rate-distortion curves of our approach and different baselines on Cityscapes and InStereo2K datasets.

places of the tokens whose probability distribution are not yet estimated, which is not the most efficient and effective choice for entropy coding. On the one hand, interacting with the uninformative [MASK] tokens may distract the specific features of the estimated tokens during the repetitive self-attention operations. On the other hand, the encoder-decoder Transformer architecture is complicated but only allows a unidirectional propagation from prior information to the input tokens of the Transformer decoder. The prior information is fixed once extracted and cannot utilize the updated spatial context information from the already estimated tokens. Consequently, it is difficult for the entropy model to make full use of the power of Transformer architecture.

To overcome these limitations, we propose content-aware tokens to unleash the potential of the Transformer entropy model. Specifically, the pre-acquired prior information, including the disparity prior  $\hat{y}_{v-1}$  and hyperprior  $\hat{z}_v$ , is used to generate the content-aware tokens  $c_v$ . Then, a composite context  $u_v$  is created by filling the unestimated positions with content-aware tokens:

$$u_v = \hat{y}_v \odot m + c_v \odot (1 - m) \quad (5)$$

This novel content-aware MIM style brings multiple advantages. First, since our proposed content-aware tokens in place have already carried specific prior information about each position, they can provide a more informative context to update the features of the estimated tokens. Second, performing self-attention over the composite context can bridge a bidirectional information flow between the prior information and estimated tokens. It allows the prior information to be iteratively updated by the features of the previously estimated token features. Thus, the composite context becomes more and more concrete as the number of the estimated tokens gradually increases, thereby effectively reducing the estimation uncertainty of the remaining tokens. Third, as the prior information propagation is seamlessly embedded into each self-attention operation, it facilitates our design of a new decoder-free Transformer entropy model.

**Decoder-free Transformer.** Since the prior information has sufficiently interacted with the estimated tokens in the Transformer encoder, an explicit Transformer decoder is no longer needed. Thus, we discard the whole Transformer decoder and build an encoder-only Transformer architecture. As shown in Fig. 3 (b), it comprises two parts, i.e., a prior generation network that produces content-aware tokens  $c_v$  based on different priors and a decoder-free Transformer that estimates the probability distributions of tokens. Note that we enable relative position embedding (Liu et al. 2021) in the Transformer and assign two identity embeddings to the token sequence to further distinguish different types of tokens, which effectively promote the token interactions.

## Training and Inference

**Bitrate Estimation.** The training phase involves a two-step prediction procedure to estimate the probability distributions of all tokens. Taking the content-aware MIM method as an example, a Boolean mask  $m$  is randomly generated to simulate various possible masking scenarios. Initially, purely based on the content-aware tokens  $c_v$ , we estimate the bitrate  $R_{c_v}(\hat{y}_v)$  of the initial tokens where the masked values are 1. Subsequently, a composite context  $u_v$  is created to estimate the bitrate  $R_{u_v}(\hat{y}_v)$  of the remaining tokens. Hence, the bitrate estimation based on the content-aware MIM is expressed as:

$$\begin{aligned} R(\hat{y}_v) &= R_{c_v}(\hat{y}_v) + R_{u_v}(\hat{y}_v) \\ &= \mathbb{E}[-\log_2 p_{\hat{y}_v}(\hat{y}_v|c_v)] + \mathbb{E}[-\log_2 p_{\hat{y}_v}(\hat{y}_v|u_v)] \end{aligned} \quad (6)$$

Similarly, for the content-irrelevant MIM, we replace the above content-aware token with a [MASK] token to estimate the bitrate during training.

During inference, the probability distributions of the tokens are iteratively estimated following a sinusoidal decoded ratio scheduler (Chang et al. 2022). Let  $R(\hat{y}_v^{i,k})$  denotes the estimated bitrate of  $i$ -th token in  $\hat{y}_v$  during  $k$ -th iteration. As shown in Fig. 3, the mask  $m_0$  is initially zero during inference. In each iteration, we use the composite context  $u_v^k$  to estimate the bitrate  $R(\hat{y}_v^{i,k})$ . The mask generator employs a scheduling function to calculate the number of tokens  $n_k$  to be encoded/decoded, and selects those  $n_k$  tokens that consume the fewest bits to update the mask  $m_{k+1}$ , setting the values in the  $n_k$  tokens and previous encoded/decoded tokens to 1. Finally,  $u_v^{k+1}$  integrates  $\hat{y}_v$  and  $c_v$  using  $m_{k+1}$  for next iteration. This iterative process ensures consistency between encoder and decoder, producing identical  $(\mu_v^{m_k}, \sigma_v^{m_k})$ . More details about iterative decoding process are available in the appendix.

**Quantization.** During training, since the quantization operation is not differentiable, we leverage a mixed quantizer to allow end-to-end optimization. To be specific, it adds  $y_v$  and a uniform noise  $\mathcal{U}(-\frac{1}{2}, \frac{1}{2})$  together as the input to the entropy model, but rounds  $y_v$  as the input to the decoder based on straight-through estimation (Bengio, Léonard, and Courville 2013). We refer readers to Minnen and Singh (2020) for more details.

**Loss Function.** The objective of our proposed method is to optimize the rate-distortion trade-off. Thus, a training loss

made up of two metrics is utilized:

$$\mathcal{L} = \lambda \sum_{v=1}^2 d(\mathbf{x}_v, \hat{\mathbf{x}}_v) + \sum_{v=1}^2 (R(\hat{\mathbf{y}}_v) + R(\hat{\mathbf{z}}_v)) \quad (7)$$

where  $d(\mathbf{x}_v, \hat{\mathbf{x}}_v)$  represents the distortion between the original image  $\mathbf{x}_v$  and reconstructed image  $\hat{\mathbf{x}}_v$  under a given metric such as MSE or MS-SSIM (Wang, Simoncelli, and Bovik 2003).  $R(\hat{\mathbf{y}}_v)$  denotes the estimated bitrate of the quantized representation  $\hat{\mathbf{y}}_v$  and  $R(\hat{\mathbf{z}}_v)$  denotes the estimated bitrate of the hyper latent  $\hat{\mathbf{z}}_v$ .  $\lambda$  is a hyper parameter that is used to balance the compression rate and distortion.

## Experiments

### Experimental Setup

**Dataset.** We evaluate the coding efficiency of our proposed CAMSIC on two public stereo image datasets: (1) Cityscapes (Cordts et al. 2016): A dataset of urban outdoor scenes with distant views. It includes 5000 image pairs at a resolution of  $2048 \times 1024$  pixels. We divide these into 2975 pairs for training, 500 pairs for validation, and the remaining 1525 pairs for testing. (2) InStereo2K (Bao et al. 2020): A dataset of indoor scenes with close views. It contains 2060 image pairs at a resolution of  $1080 \times 860$  pixels. We allocate 2010 and 50 stereo image pairs for training and testing, respectively. We randomly crop the images into  $256 \times 256$  patches during training, and evaluate different codecs on full-resolution images. When the resolution of the source image is not supported, we apply the replicate padding to the right and bottom sides of the image.

**Evaluation Metrics.** We use bit per pixel (bpp) to measure the number of bits for latent representation compression. Two popular image quality assessment metrics, namely, PSNR and MS-SSIM (Wang, Simoncelli, and Bovik 2003), are used to evaluate the distortion between the reconstructed and original images. We also report the Bjontegaard Delta bitrate (BD-rate) (Bjontegaard 2001) results to show average bitrate savings at equivalent distortion levels. In addition, both BD-PSNR and BD-MSSSIM are utilized to indicate average image quality improvements at constant bitrate.

**Implementation Details.** Leveraging CompressAI (Bégaint et al. 2020), we train our models with 6 different  $\lambda$  values (256, 512, 1024, 2048, 4096, 8192 for the MSE metric; 8, 16, 32, 64, 128, 256 for the MS-SSIM metric). For MSE-optimized models, they are trained for 400 epochs with the Adam optimizer (Kingma and Ba 2015). The batch size is set as 4. The initial learning rate is  $1e^{-4}$  and decayed by a factor of 2 every 100 epochs. To accelerate the experiments, both the image encoder and decoder utilize the pretrained single image compression model ELIC (He et al. 2022a). For MS-SSIM evaluation, the MSE-optimized models are fine-tuned for 300 epochs using the MS-SSIM distortion loss with the initial learning rate as  $5e^{-5}$ . During inference, we set the number of decoding steps as 8. Our experiments are conducted with NVIDIA V100 GPUs using PyTorch.

**Benchmarks.** We compare our CAMSIC with a variety of traditional and learning-based codecs. These competitive baselines are roughly split into four categories: single image

Methods	BD-PSNR $\uparrow$	BD-rate $\downarrow$	BD-MSSSIM $\uparrow$	BD-rate $\downarrow$
HEVC	0.499dB	-13.533%	0.465dB	-14.453%
VVC	1.621dB	<u>-37.213%</u>	1.538dB	-38.075%
MV-HEVC	-1.693dB	61.915%	-0.736dB	24.305%
ELIC	1.128dB	-25.196%	1.198dB	-28.346%
DispSIC	0.402dB	11.865%	2.948dB	-56.670%
SASIC	-0.021dB	0.661%	2.208dB	-43.415%
BCSIC	1.605dB	-35.891%	3.355dB	-60.292%
LDMIC	1.600dB	-36.264%	<u>3.573dB</u>	<u>-64.398%</u>
ECSIC	<u>1.700dB</u>	-36.665%	3.274dB	-59.543%
Proposed	<b>2.019dB</b>	<b>-41.961%</b>	<b>3.715dB</b>	<b>-64.519%</b>

(a) Cityscapes dataset

Methods	BD-PSNR $\uparrow$	BD-rate $\downarrow$	BD-MSSSIM $\uparrow$	BD-rate $\downarrow$
HEVC	0.103dB	-4.784%	0.092dB	-4.260%
VVC	0.802dB	-29.125%	0.669dB	-26.483%
MV-HEVC	-1.274dB	77.154%	-0.578dB	28.493%
ELIC	0.823dB	-29.748%	0.807dB	-30.453%
DispSIC	0.617dB	-23.084%	1.740dB	-46.916%
SASIC	0.231dB	-7.825%	1.499dB	-39.693%
BCSIC	1.249dB	-39.933%	2.228dB	-55.022%
LDMIC	1.282dB	-43.522%	<u>2.333dB</u>	<b>-59.167%</b>
ECSIC	<u>1.428dB</u>	<u>-43.551%</u>	2.149dB	-52.587%
Proposed	<b>1.440dB</b>	<b>-43.608%</b>	<b>2.426dB</b>	<u>-56.773%</u>

(b) InStereo2K dataset

Table 1: Comparison of various codecs in BD-PSNR, BD-MSSSIM, BD-rate results relative to BPG, with the best results in **bold** and second-best ones in underlined.

compression (BPG (Bellard 2014), ELIC (He et al. 2022a)), video compression (HEVC (Sullivan et al. 2012), VVC (Bross et al. 2021)), multi-view compression (MV-HEVC (Tech et al. 2015), LDMIC (Zhang, Shao, and Zhang 2022)) and stereo image compression (DispSIC (Zhai et al. 2022), SASIC (Wödlinger et al. 2022), BCSIC (Lei et al. 2022), ECSIC (Wödlinger et al. 2024)). For BPG, we adopt the default x265 encoder without chroma subsampling to independently compress each image. By treating a stereo image pair as a two-frame video sequence, we run HM-18.0 and VTM-23.0 software with *lowdelay P* configuration and YUV444 format to evaluate the coding efficiency of HEVC and VVC. For MV-HEVC, HTM-16.3 software is used to compress stereo images by setting a two-view intra mode. Unfortunately, it only supports YUV420 format, resulting in inferior compression performance. As for learning-based multi-view and stereo codecs except ECSIC, we implement them using the same training procedure. Since the current SOTA method ECSIC requires training on large-width stereo images to achieve the best results, we follow their open source library (Wödlinger et al. 2024) to train the ECSIC models.

## Experimental Results

**Compression performance.** Fig. 4 shows the RD curves of these methods on Cityscapes and InStereo2K datasets. We report BD-rate, BD-PSNR and BD-MSSSIM scores of each codec relative to BPG in Table 1. It is observed that the single image codec ELIC achieves 25.196% and 29.748% bitrate saving in terms of PSNR on Cityscapes and InStereo2K datasets. By contrast, our CAMSIC has larger improvements, i.e., 41.961% and 43.608%. This demonstrates that introducing a stereo conditional entropy model to capture




Figure 5: Complexity of learning-based codecs on the InStereo2K dataset. Both encoding and decoding time are evaluated by an NVIDIA V100 GPU.

the correlations between stereo images effectively enables a more accurate estimation of the probability distribution.

Without manually designed and complex linear or non-linear transformations, our method outperforms most of these compression baselines, highlighting the importance of a powerful spatial-disparity entropy model. Specifically, a series of unidirectional codecs (HEVC, VVC, MV-HEVC, DispSIC, SASIC) necessitate explicit warping to remove disparity redundancies. Notably, our CAMSIC surpasses VVC, the best unidirectional codec, across all bpp ranges on the InStereo2K and Cityscapes dataset. In addition, bidirectional codecs (BCSIC, ECSIC) adopt bidirectional attention to reduce inter-view redundancies. Remarkably, even with a straightforward image encoder-decoder architecture, our CAMSIC still saves more bits over the SOTA method ECSIC in terms of PSNR and MS-SSIM. Moreover, compared with the latest distributed codec LDMIC that adopts the joint decoding method based on global cross-attention mechanism, our proposed method improves by about 0.419dB and 0.158dB PSNR at the same bpp level on two datasets (Cityscapes and InStereo2K). These results affirm the superiority of our proposed entropy model in exploiting data correlations for enhancing compression performance.

**Coding Complexity.** We compare the computation complexity of six learning-based multi-view and stereo codecs on InStereo2K dataset. The running time on GPU includes the execution latency of entropy coder. As illustrated in Fig. 5, our CAMSIC method is orders of magnitude faster than other auto-regressive approaches (DispSIC, BCSIC and LDMIC) by leveraging the parallel coding advantages of the MIM technique. Furthermore, by simplifying the stereo image encoding and decoding processes, our approach runs faster than SASIC and ECSIC. These results showcase a better trade-off between performance and speed among contemporary learning-based codecs.

### Ablation Study

**Different Entropy Models.** Table 2 presents the BDBR results of different entropy models on the Cityscapes and InStereo2K datasets. Specifically, we replace our entropy model with the CNN-based ECSIC, SASIC and BCSIC entropy models as well as the content-irrelevant encoder-decoder Transformer entropy model. Note that the vanilla MIM-based Transformer entropy model outperforms the ECSIC entropy model, which is attributed to the Trans-

Variant	Cityscapes	InStereo2K
Ours	0%	0%
(V1) w/ ECSIC entropy model	48.89%	22.38%
(V2) w/ SASIC entropy model	36.16%	19.44%
(V3) w/ BCSIC entropy model	5.36%	5.76%
(V4) w/ Vanilla MIM entropy model	36.65%	19.48%
(V5) w/o RPE + w/o ID	3.700	1.316
(V6) w/o ID	2.739	0.501

Table 2: Ablation studies for different components. The first row is set as the anchor to measure BD-rate. RPE means relative position embedding. ID indicates identity embedding.

former’s inherent capability to capture long-range dependencies more effectively than CNN architectures. However, simply migrating the vanilla encoder-decoder Transformer architecture and the naive MIM technique to the entropy model leads to inefficient prior information propagation, which makes the vanilla MIM entropy model inferior to the BCSIC entropy model with bidirectional view interaction. To address this challenge, we introduce a novel content-aware MIM technique to design an advanced decoder-free Transformer entropy model. The improvements demonstrate that our content-aware MIM significantly surpasses the content-irrelevant MIM, which underscores the importance of bidirectional interactions between prior information and estimated tokens in improving compression performance.

**Operations for Token Interaction.** In our framework, we enable relative position embedding and introduce identity embeddings to facilitate token interaction. To show the effects of these operations, we undertake a series of ablation experiments. As shown in Table 2, compared with removing relative position embeddings (variant V5), adding relative position embeddings (variant V6) saves more bitrate, which indicates that providing the position cues for token interaction enables a more accurate estimation of probability distribution. Based on variant V6, we attach two identity embeddings to the input token sequence to form our final version, which further reduces the bitrate consumption. It implies that explicitly distinguishing content-aware tokens and estimated tokens boosts token interaction.

### Conclusion

In this paper, we propose a stereo image compression framework, namely CAMSIC, with a succinct yet potent Transformer entropy model as the core to effectively capture the spatial-disparity correlations. Our CAMSIC consists of a straightforward encoder-decoder architecture to independently transform each image. A novel content-aware masked image modeling (MIM) technique is presented to fully exploit the power of Transformer entropy model. Our content-aware MIM unleashes bidirectional interactions between the prior information and estimated tokens, which allows us to eliminate an extra Transformer decoder. Thus, a decoder-free Transformer entropy model is proposed to improve the efficiency. Experimental results show that our approach achieves the state-of-the-art compression performance with fast encoding and decoding.

## Acknowledgments

This work was supported by the General Research Fund (Project No. 16209622) from the Hong Kong Research Grants Council.

## References

- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2015. Density modeling of images using a generalized normalization transformation. *arXiv preprint arXiv:1511.06281*.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2017. End-to-end optimized image compression. 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*.
- Bao, W.; Wang, W.; Xu, Y.; Guo, Y.; Hong, S.; and Zhang, X. 2020. Instereo2k: a large real dataset for stereo matching in indoor scenes. *Science China Information Sciences*, 63: 1–11.
- Bégaint, J.; Racapé, F.; Feltman, S.; and Pushparaja, A. 2020. CompressAI: a PyTorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029*.
- Bellard, F. 2014. BPG image format. <https://bellard.org/bpg/>. Accessed: 2024-05-10.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Bjontegaard, G. 2001. Calculation of average PSNR differences between RD-curves. *ITU-T VCEG-M33, April, 2001*.
- Bross, B.; Wang, Y.-K.; Ye, Y.; Liu, S.; Chen, J.; Sullivan, G. J.; and Ohm, J.-R. 2021. Overview of the versatile video coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10): 3736–3764.
- Chang, H.; Zhang, H.; Jiang, L.; Liu, C.; and Freeman, W. T. 2022. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11315–11325.
- Chen, F.; Xu, Y.; and Wang, L. 2022. Two-stage octave residual network for end-to-end image compression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 3922–3929.
- Chen, T.; Liu, H.; Ma, Z.; Shen, Q.; Cao, X.; and Wang, Y. 2021. End-to-end learnt image compression via non-local attention optimization and improved context modeling. *IEEE Transactions on Image Processing*, 30: 3179–3191.
- Cheng, Z.; Sun, H.; Takeuchi, M.; and Katto, J. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7939–7948.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- Deng, X.; Deng, Y.; Yang, R.; Yang, W.; Timofte, R.; and Xu, M. 2023. MASIC: Deep Mask Stereo Image Compression. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Deng, X.; Yang, W.; Yang, R.; Xu, M.; Liu, E.; Feng, Q.; and Timofte, R. 2021. Deep homography for efficient stereo image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1492–1501.
- Fehn, C. 2004. Depth-image-based rendering (DIBR), compression, and transmission for a new approach on 3D-TV. In *Stereoscopic displays and virtual reality systems XI*, volume 5291, 93–104. SPIE.
- Feichtenhofer, C.; Fan, H.; Li, Y.; and He, K. 2022. Masked Autoencoders As Spatiotemporal Learners. In *Advances in Neural Information Processing Systems*.
- Gupta, A.; Tian, S.; Zhang, Y.; Wu, J.; Martín-Martín, R.; and Fei-Fei, L. 2023. Maskvit: Masked visual pre-training for video prediction. In *International Conference on Learning Representations*.
- He, D.; Yang, Z.; Peng, W.; Ma, R.; Qin, H.; and Wang, Y. 2022a. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5718–5727.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022b. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*, 4171–4186.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Lei, J.; Liu, X.; Peng, B.; Jin, D.; Li, W.; and Gu, J. 2022. Deep stereo image compression via bi-directional coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19669–19678.
- Liang, J.; Wu, C.; Hu, X.; Gan, Z.; Wang, J.; Wang, L.; Liu, Z.; Fang, Y.; and Duan, N. 2022. NUWA-Infinity: Autoregressive over Autoregressive Generation for Infinite Visual Synthesis. In *Advances in Neural Information Processing Systems*.
- Liu, J.; Huang, X.; Liu, Y.; and Li, H. 2022. Mixmim: Mixed and masked image modeling for efficient visual representation learning. *arXiv preprint arXiv:2205.13137*.
- Liu, J.; Sun, H.; and Katto, J. 2023. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14388–14397.

- Liu, J.; Wang, S.; and Urtasun, R. 2019. Dsic: Deep stereo image compression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3136–3145.
- 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. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Zhang, X.; Shao, J.; Lin, Z.; and Zhang, J. 2025. Bidirectional stereo image compression with cross-dimensional entropy model. In *European Conference on Computer Vision*, 480–496. Springer.
- Mentzer, F.; Toderici, G.; Minnen, D.; Caelles, S.; Hwang, S. J.; Lucic, M.; and Agustsson, E. 2022. VCT: A Video Compression Transformer. In *Advances in Neural Information Processing Systems*.
- 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*.
- Minnen, D.; and Singh, S. 2020. Channel-wise autoregressive entropy models for learned image compression. In *2020 IEEE International Conference on Image Processing (ICIP)*, 3339–3343. IEEE.
- Mital, N.; Özyilkan, E.; Garjani, A.; and Gündüz, D. 2023. Neural distributed image compression with cross-attention feature alignment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2498–2507.
- Skodras, A.; Christopoulos, C.; and Ebrahimi, T. 2001. The JPEG 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5): 36–58.
- Stepanov, D.; and Tishchenko, I. 2016. The concept of video surveillance system based on the principles of stereo vision. In *2016 18th Conference of Open Innovations Association and Seminar on Information Security and Protection of Information Technology (FRUCT-ISPIT)*, 328–334. IEEE.
- Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12): 1649–1668.
- Tech, G.; Chen, Y.; Müller, K.; Ohm, J.-R.; Vetro, A.; and Wang, Y.-K. 2015. Overview of the multiview and 3D extensions of high efficiency video coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(1): 35–49.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Video-MAE: Masked Autoencoders are Data-Efficient Learners for Self-Supervised Video Pre-Training. In *Advances in Neural Information Processing Systems*.
- Vetro, A.; Wiegand, T.; and Sullivan, G. J. 2011. Overview of the stereo and multiview video coding extensions of the H. 264/MPEG-4 AVC standard. *Proceedings of the IEEE*, 99(4): 626–642.
- Wallace, G. K. 1991. The JPEG still picture compression standard. *Communications of the ACM*, 34(4): 30–44.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Wödlinger, M.; Kotera, J.; Keglevic, M.; Xu, J.; and Sablatnig, R. 2024. ECSIC: Epipolar Cross Attention for Stereo Image Compression. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3436–3445.
- Wödlinger, M.; Kotera, J.; Xu, J.; and Sablatnig, R. 2022. Sasic: Stereo image compression with latent shifts and stereo attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 661–670.
- Xiang, J.; Tian, K.; and Zhang, J. 2023. Mimt: Masked image modeling transformer for video compression. In *International Conference on Learning Representations*.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Bao, J.; Yao, Z.; Dai, Q.; and Hu, H. 2022. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9653–9663.
- Yin, H.; Wang, Y.; Tang, L.; Ding, X.; Huang, S.; and Xiong, R. 2020. 3d lidar map compression for efficient localization on resource constrained vehicles. *IEEE transactions on intelligent transportation systems*, 22(2): 837–852.
- Zhai, Y.; Tang, L.; Ma, Y.; Peng, R.; and Wang, R. 2022. Disparity-based Stereo Image Compression with Aligned Cross-View Priors. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2351–2360.
- Zhang, X.; Shao, J.; and Zhang, J. 2022. LDMIC: Learning-based Distributed Multi-view Image Coding. In *The Eleventh International Conference on Learning Representations*.
- Zhu, Y.; Yang, Y.; and Cohen, T. 2021. Transformer-based transform coding. In *International Conference on Learning Representations*.
- Zou, R.; Song, C.; and Zhang, Z. 2022. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17492–17501.