Bidirectional Stereo Image Compression with Cross-Dimensional Entropy Model

Zhening Liu[®], Xinjie Zhang[®], Jiawei Shao[®], Zehong Lin^{*}[®], and Jun Zhang[®]

Abstract. With the rapid advancement of stereo vision technologies, stereo image compression has emerged as a crucial field that continues to draw significant attention. Previous approaches have primarily employed a unidirectional paradigm, where the compression of one view is dependent on the other, resulting in imbalanced compression. To address this issue, we introduce a symmetric bidirectional stereo image compression architecture, named BiSIC. Specifically, we propose a 3D convolution based codec backbone to capture local features and incorporate bidirectional attention blocks to exploit global features. Moreover, we design a novel cross-dimensional entropy model that integrates various conditioning factors, including the spatial context, channel context, and stereo dependency, to effectively estimate the distribution of latent representations for entropy coding. Extensive experiments demonstrate that our proposed BiSIC outperforms conventional image/video compression standards, as well as state-of-the-art learning-based methods, in terms of both PSNR and MS-SSIM.

Keywords: Stereo Image Compression · Bidirectional Architecture · Cross-Dimensional Entropy Model

1 Introduction

Stereo vision, which mimics human binocular vision, has emerged as an important area in computer vision. It enables a wide range of applications, such as 3D movies, remote sensing [29], autonomous driving [38], and augmented/virtual reality (AR/VR) [1]. With the rapid development of stereo cameras, stereo images have become essential for providing immersive and convincing visual presentations that enhance user experience and benefit downstream applications. Meanwhile, the proliferation of stereo cameras leads to a significant surge in the volume of stereo images, resulting in substantial storage costs and posing challenges to transmission efficiency. Consequently, stereo image compression plays a pivotal role in ensuring the practicality and effectiveness of stereo vision.

^{*} Corresponding author

In recent years, numerous methods have been developed for single image compression, including traditional codecs [6, 34] and learning-based methods [2–4, 10, 24]. These methods, however, focus on extracting individual features for compression and fail to capture the cross-view correlation. Consequently, they are less effective for stereo image compression that requires the extraction of both intra-view information and inter-view information. To exploit the inherent shared information in stereo images, traditional methods for stereo image coding follow a predictive compression procedure of video standards [9,31]. In this process, one view is compressed first and then used to predict the other view. Subsequently, the disparity between the predicted and actual views is compressed. Nonetheless, these methods typically rely on hand-crafted modules and lack end-to-end optimization, resulting in suboptimal performance.

Inspired by the success of learning-based single image compression, researchers have proposed novel stereo image compression methods using deep neural networks (DNNs) to enhance compression performance [12, 13, 20, 21, 36, 37, 40, 41]. Initially, extensive research focuses on unidirectional learned stereo image compression methods [12, 13, 21, 36, 37, 40], where one view is compressed and used as a reference to compress the other view. Despite achieving high overall ratedistortion (RD) performance, these unidirectional learned compression methods yield a noticeable imbalance in compression quality between stereo views, similar to traditional predictive compression methods. This imbalance is unfavorable for human vision and undermines downstream machine vision tasks [23]. Moreover, some of these methods [12, 13, 37, 40] heavily rely on sequential coding order and accurate estimation of the explicit warping relationship between views. To avoid these issues, a recent study [20] proposes a bidirectional compression method that leverages 2D convolution to extract individual features and facilitate information sharing between views. Nonetheless, the compression performance is limited due to the separate processing of convolutions and the plain spatial context based entropy model.

Based on the above discussion, we propose a novel bidirectional method for stereo image compression, named BiSIC, to address the aforementioned issues. BiSIC achieves improvements in four key aspects: encoder/decoder structure, mutual stereo learning, entropy model, and coding speed. Firstly, regarding the codec structure, although 2D convolutions have proven effective in various computer vision tasks, they fail to capture the aligned morphs and features between stereo views. To overcome this limitation, we propose using 3D convolution as the backbone for the encoder/decoder, which takes both views as input to capture the inter-view dependencies that are inaccessible in previous works based on 2D convolution. Secondly, for mutual stereo learning, we propose a mutual attention block to facilitate the transfer of features between views. This module enhances the integration and cooperation of the stereo views, leading to improved compression performance. Thirdly, for the entropy model, prior methods primarily use hyperprior and spatial context as conditions, overlooking the rich inter-view dependencies across other dimensions. To exploit these dependencies,



Fig. 1: The proposed bidirectional stereo image compression architecture. AE and AD are arithmetic coder for entropy coding. The backbones of encoder \mathbf{E} , decoder \mathbf{D} , hyper encoder h_a , and hyper decoder h_s are constructed with 3D convolution to model local features. 3DConvTr denotes transposed 3D convolution. To enhance global feature extraction, bidirectional mutual attention blocks are inserted between 3D convolutional layers. Moreover, a novel cross-dimensional entropy model is utilized to capture complex inter-view dependencies.

we construct a cross-dimensional entropy model that aggregates hyperprior, spatial context, channel context, and stereo dependency.

Specifically, we introduce a masked 3D convolution to capture the stereo spatial context and utilize an attention block to process the stereo channel context. Finally, to improve coding speed and reduce auto-regressive steps, we design a fast variant based on a stereo-checkerboard structure. This design converts autoregressive inference into a two-fold operation, improving the coding efficiency. The contributions of this paper are summarized as follows:

- We propose a novel 3D convolution based codec backbone that processes stereo images as a whole, enabling the learning of local stereo features and effective capture of shared patterns across views.
- We design a bidirectional mutual attention block that facilitates the exchange of global features between stereo views. This block produces compact attention maps and is robust to input images of different sizes.
- We develop a powerful symmetric cross-dimensional entropy model, which incorporates hyperpriors, spatial context, channel context, and stereo dependency as conditioning factors, to enhance the probability estimation.
- To accelerate coding, we propose a stereo-checkerboard design that transforms the spatial auto-regressive inference into a two-fold process. This design strikes a balance between time consumption and compression quality.

2 Related Works

2.1 Single Image Compression

Image compression is an important topic that continuously attracts significant research interest. With the increasing demand for vision products and applica-

tions, steady progress has been made in traditional image compression standards, evolving from JPEG [28,34] and BPG [6] to video compression standards such as HEVC [30] and VVC [9]. These traditional methods establish a standard compression pipeline comprising transform, quantization, and entropy coding. Over the past five years, deep learning has revolutionized various vision applications. As demonstrated in the seminal work [3], DNNs have the potential to replace components in traditional compression methods and enable end-to-end compression. This work introduces uniform noise to mimic the distortion caused by quantization, thereby paving the way for end-to-end optimization in image compression.

In learned image compression models [3,4,10,14–17,22,24–26,42], neural networks are employed to parameterize the probability distribution of transformed latents. Accurate estimation of this distribution is crucial for achieving compact compression via entropy coding. To enhance probability estimation, various conditions have been integrated into entropy models [2], including hyperpriors [4], spatial context [24], and channel-wise context [25], to capture a broader range of dependencies. To be specific, hyperprior is a lightweight subsidiarity information transmitted alongside the encoded latents. Both spatial context and channel-wise context leverage auto-regressive inference, which is conditioned on the preceding segments of data by iteratively interleaving the inference. Checkerboard methods [15], on the other hand, avoid the spatial auto-regressive process and accelerate the coding speed by constraining the spatial context to the near neighbors of each pixel instead of spatial causal components. These dependencies operate on diverse scales and can be progressively combined to further improve effectiveness, as demonstrated in [14, 17].

2.2 Stereo Image Compression

Stereo image compression is an extension of single image compression that considers capturing both intra-view features and inter-view correlations. Traditional approaches employ key-frame based predictive coding, akin to video coding [18,33]. However, these methods rely on hand-crafted feature extractors and classical optimization techniques. As these modules are optimized individually, they often fail to achieve satisfactory RD performance. Moreover, the predictive compression pipeline also leads to imbalanced compression.

In recent years, learned stereo image compression methods have gained increasing attention [12,13,20,21,36,37,40,41]. These methods can be broadly categorized into unidirectional and bidirectional approaches. Unidirectional methods first compress one view image and then use it as a reference to compress the other view, significantly reducing the required bits. The pioneering work [21] introduces a skip module to warp the features from one view to encode the other. Subsequent works have utilized various techniques to represent the warping relationship between stereo views, such as homography matrices [12,13], disparity maps [40], and horizontal shifts [37].

These unidirectional methods heavily rely on the precise estimation of the spatial transformations. A very recent work, ECSIC [36], avoids explicit para-

metric transformation in favor of constructing a unidirectional entropy model that conditions the right image on the known left image. However, the decompressed results of these unidirectional methods often have imbalanced quality due to the one-way dependency between stereo views. On the other hand, bidirectional methods, while remaining less explored, produce more balanced quality by extracting shared information from both views and compressing stereo images simultaneously. One such approach [20] introduces a contextual transform module for feature transfer between views. However, it employs separate downsampling networks and only takes spatial context into account, which limits the compression performance. In contrast, to fully exploit the shared information and establish stronger dependencies, we employ a joint downsampling codec and a cross-dimensional entropy model.

3 Proposed Methods

The proposed architecture for stereo image compression, as depicted in Fig. 1, consists of two main components: a joint codec and an entropy model. The codec employs 3D convolutions to transform the original stereo images into compact latent representations, effectively exploiting the correlation between the two views. Subsequently, the entropy model estimates the probability distribution of these latent representations, enabling the generation of a compact bitstream through entropy coding. In the following subsections, we elaborate on the codec and the entropy model.

3.1 Joint Codec

3D Convolution Backbone. The codec acts as a non-linear transformation that downsamples the raw images into compact latent representations and generates hyperpriors for the entropy model. Different from previous approaches to stereo image compression that employ 2D convolutional layers, we adopt 3D convolutions as the backbone of our codec. Unlike 2D convolution, which operates separately on each view, 3D convolution concurrently processes both stereo images. Consequently, it inherently possesses the capability to extract inter-view correlations.

Our encoder, denoted by **E**, comprises four 3D convolutional layers, with bidirectional mutual attention blocks inserted after the second and the fourth layers to facilitate global feature communication between the two views. The encoder transforms a pair of stereo images, i.e., the left and right images denoted by $\boldsymbol{x}_l, \boldsymbol{x}_r \in \mathbb{R}^{B \times 3 \times H \times W}$, into compact latent representations $\boldsymbol{y}_l, \boldsymbol{y}_r \in \mathbb{R}^{B \times N \times \frac{H}{16} \times \frac{W}{16}}$ and quantizes them into $\hat{\boldsymbol{y}} := \{\hat{\boldsymbol{y}}_l, \hat{\boldsymbol{y}}_r\}$ for efficient transmission. Conversely, the decoder **D** transforms the quantized latents $\hat{\boldsymbol{y}}$ back into images $\hat{\boldsymbol{x}}_l$ and $\hat{\boldsymbol{x}}_r$. The encoding and decoding processes are expressed as:

$$\boldsymbol{y}_l, \boldsymbol{y}_r = \mathbf{E}(\boldsymbol{x}_l, \boldsymbol{x}_r), \tag{1}$$

$$\hat{\boldsymbol{x}}_l, \hat{\boldsymbol{x}}_r = \mathbf{D}(\hat{\boldsymbol{y}}_l, \hat{\boldsymbol{y}}_r).$$
(2)

To capture the dependencies among elements in the latents $\hat{\boldsymbol{y}}$ for efficient entropy coding [4,7], hyperpriors are generated through another 3D convolution based hyper encoder h_a . The generated hyperpriors $\boldsymbol{z}_l, \boldsymbol{z}_r \in \mathbb{R}^{B \times M \times \frac{H}{64} \times \frac{W}{64}}$ are quantized into $\hat{\boldsymbol{z}}_l, \hat{\boldsymbol{z}}_r$ and then transmitted alongside the quantized latents $\hat{\boldsymbol{y}}$. At the decoder side, a 3D convolution based hyper decoder h_s upscales the quantized hyperpriors $\hat{\boldsymbol{z}}_l, \hat{\boldsymbol{z}}_r$ into $\tilde{\boldsymbol{z}}_l, \tilde{\boldsymbol{z}}_r$, which are subsequently utilized in the entropy model, as elaborated in Sec. 3.2. The processes of the hyper encoder and hyper decoder are expressed as:

$$\boldsymbol{z}_l, \boldsymbol{z}_r = h_a(\boldsymbol{y}_l, \boldsymbol{y}_r), \tag{3}$$

$$\tilde{\boldsymbol{z}}_l, \tilde{\boldsymbol{z}}_r = h_s(\hat{\boldsymbol{z}}_l, \hat{\boldsymbol{z}}_r).$$
(4)

Bidirectional Mutual Attention Block. Convolutional layers have an excellent ability for local modeling, but their ability for long-range feature modeling is relatively limited. On the other hand, attention modules [32] excel at extracting long-range features. Therefore, we propose combining the attention mechanism with convolutional layers to compensate for the shortcomings of the latter and leverage the advantages of both convolution and attention. To this end, we design a novel bidirectional mutual attention block, as depicted in Fig. 2, to be inserted between convolutional layers as the pink blocks shown in Fig. 1.

As illustrated in Fig. 2, the bidirectional mutual attention block contains two cross-attention stages, each followed by a self-attention layer. The two cross-attention stages differ in how they concentrate on the feature map and are referred to as cross-key attention and cross-query attention. Before each cross-attention stage, the input stereo data y_l, y_r first pass through a residual block, which are then embedded into low-dimensional query, key, and value tensors $Q, K, V \in \mathbb{R}^{B \times C \times (H \times W)}$. To avoid the prohibitive computation of classical attention, which yields an $(H \times W) \times (H \times W)$ sized attention map, we adopt the efficient attention approach proposed in [27]. This approach computes the product of the key and the transposed value, thereby reducing the attention map size to $C_K \times C_V$.

In the cross-key attention stage, for each view, we generate an attention map using its key and the value from the other view, which is then queried by the current view. By doing so, we can extract the cross-key feature Φ from the stereo views as:

$$\boldsymbol{\Phi}_{r \to l} = (\sigma(\boldsymbol{K}_r) \times \boldsymbol{V}_l^T)^T \times \sigma(\boldsymbol{Q}_r), \ \boldsymbol{\Phi}_{l \to r} = (\sigma(\boldsymbol{K}_l) \times \boldsymbol{V}_r^T)^T \times \sigma(\boldsymbol{Q}_l), \quad (5)$$

where σ is the softmax function. This approach calculates the attention map using the key and value from different views, thereby facilitating inter-view alignment and the identification of common patterns. When a feature in one view is correlated with the corresponding feature in the other view, it is effectively captured through the attention map.

In the cross-query attention stage, the attention map is obtained using the key and value from a single view, and then queried by the other view, yielding the cross-query feature Ψ as:

$$\boldsymbol{\Psi}_{r \to l} = (\boldsymbol{\sigma}(\boldsymbol{K}_l) \times \boldsymbol{V}_l^T)^T \times \boldsymbol{\sigma}(\boldsymbol{Q}_r), \ \boldsymbol{\Psi}_{l \to r} = (\boldsymbol{\sigma}(\boldsymbol{K}_r) \times \boldsymbol{V}_r^T)^T \times \boldsymbol{\sigma}(\boldsymbol{Q}_l).$$
(6)



Fig. 2: Overview of the proposed bidirectional mutual attention block. The blue and green lines represent features extracted from the left view and right view, respectively. The network structures of the residual block, the basic efficient attention unit, and the combine block are illustrated on the right, from top to bottom.

The attention map in this stage is generated within one view, thus maintaining an inner-focused relationship. This process preserves the individual self-features of each view, while the other view serves as a reference for determining the allocation of attention weights.

After obtaining the cross-key attention and cross-query attention features Φ, Ψ , we employ self-attention to augment these features as $\tilde{\Phi}, \tilde{\Psi}$ and use a shared-parameter combine block to get the desired mutual attention output, as shown in the bottom right of Fig. 2. Specifically, the combine block concatenates the augmented features $\tilde{\Phi}, \tilde{\Psi}$ and the residual component, and shrinks the channel number to be consistent with the block input.

This bidirectional mutual attention block captures diverse features and provides a global receptive field for the codec. It is suitable for any stereo information transfer scenarios, provided that the inputs are two features of equal size, making it a generalized plug-and-play block. Moreover, the size of the attention maps is determined only by the embedding channel C and does not vary with the size of input images.

3.2 Cross-Dimensional Entropy Model

The proposed codec effectively transforms stereo images into small-sized latent representations. To achieve efficient compression, the next crucial step is to encode these latents into compact bitstreams using accurate probability distribution and entropy coding [4,24]. Thus, an entropy model is critical for parameterizing and estimating probabilities. In learned image compression, the probability distribution is commonly modeled as Gaussian distributions [24], with the mean μ and variance σ^2 being estimated. To achieve accurate estimation, it is critical to provide appropriate references. Apart from hyperpriors, previous stereo image compression methods primarily use spatial context [13, 20, 41] and unidirectional stereo dependency as references [21,36], while abundant features on the



Aggregation

 $\mu_l \neq \sigma_l^2$

Fig. 3: Illustration of the proposed symmetric cross-dimensional entropy model. It jointly aggregates the hyperprior (blue), stereo spatial context (red), and stereo channel context (yellow) as conditions for an effective probability distribution estimation. The hyper decoder h_s for hyperprior is shown in Fig. 1. The masked 3D convolution for spatial context is provided in Fig. 4a. Moreover, the mutual attention block for channel context is detailed in Fig. 2.

Hyperprior

Spatial Context

Channel Context

Aggregation

 $\mu_r \mathbf{v} \sigma_r^2$

channel dimension are overlooked. In this work, we propose a symmetric crossdimensional entropy model that integrates hyperprior, spatial context, channel context, and stereo dependency in a fully symmetric bidirectional way.

The architecture of the proposed cross-dimensional entropy model is depicted in Fig. 3. We use $\hat{y}_{l,i}^k$ to denote the *i*-th (i = 1, 2, ..., n) entry of the *k*-th (k = 1, 2, ..., K) channel slice from the received left view latent \hat{y}_l , and likewise for the others. The cross-dimensional condition formulation is expressed as:

$$p_{\hat{y}_{l}}(\hat{y}_{l,i}^{k}|\hat{z}_{l}, \hat{y}_{l}^{(7)$$

$$p_{\hat{\boldsymbol{y}}_r}(\hat{\boldsymbol{y}}_{r,i}^k | \hat{\boldsymbol{z}}_r, \hat{\boldsymbol{y}}_r^{< k}, \hat{\boldsymbol{y}}_{r,
(8)$$

For the hyperprior component, we generate the dependency $\tilde{\boldsymbol{z}}_l, \tilde{\boldsymbol{z}}_r$ from the hyper decoder using Eq. (4). To obtain the channel context, we evenly split the stereo latents on the channel dimension into K slices. These latent slices share a great similarity, providing abundant conditions [25]. To exploit this similarity, the probability estimation is performed slice by slice, using the previous slices as references for the following one. In addition, the decoded slices of the other view also serve as references. Therefore, we propose to use the mutual attention block introduced in the previous subsection to aggregate the channel-wise features Θ_l, Θ_r from both views, conditioning on $\hat{\boldsymbol{y}}_l^{< k}$ and $\hat{\boldsymbol{y}}_r^{< k}$, as depicted by the yellow lines in Fig. 3.

To capture the spatial context within each channel slice, we consistently employ 3D convolution to jointly learn spatial features from the causal areas $\hat{y}_{< i}^k$. Specifically, during the coding process, the current compression/decompression only accesses the areas that have already been encoded/decoded. To maintain the causality of the compression, we mask out the uncoded part, which is referred to as masked 3D convolution, as shown in Fig. 4a. In this approach, for the *i*-th entry \hat{y}_i^k , the weights of uncoded part are set to zero, and we obtain spatial

8

Z. Liu et al.



(a) Masked 3. convolution.





(c) Stereo-checkerboard pattern in BiSIC-Fast.

Fig. 4: (Left) Illustration of masked 3D convolution. The blue region is used as the condition for the probability estimation of the current target depicted in red. The weights of the convolution kernel in blue are valid, while the white region is masked to maintain the causal spatial context. (Middle) Demonstration of the auto-regressive process in the cross-dimensional entropy model. The yellow and red arrows represent the channel context and spatial context, respectively. The spatial context is processed entry by entry. (Right) Demonstration of the stereo-checkerboard in the fast variant, where the blue and red parts are each processed only once.

context features $\Upsilon_{l,i}$, $\Upsilon_{r,i}$ by considering the causal context $\hat{y}_{l,<i}^k$ and $\hat{y}_{r,<i}^k$ from both views simultaneously, as depicted by the red lines in Fig. 3.

Upon extracting various dependency features, we employ an aggregation network G_{ag} to fuse the conditions and obtain the probability estimations as:

$$\boldsymbol{\mu}_l, \boldsymbol{\sigma}_l^2 = \mathbf{G}_{\mathrm{ag}}(\tilde{\boldsymbol{z}}_l, \boldsymbol{\Theta}_l, \boldsymbol{\Upsilon}_l), \ \boldsymbol{\mu}_r, \boldsymbol{\sigma}_r^2 = \mathbf{G}_{\mathrm{ag}}(\tilde{\boldsymbol{z}}_r, \boldsymbol{\Theta}_r, \boldsymbol{\Upsilon}_r).$$
(9)

The entire procedure in the entropy model is designed to be fully symmetric and bidirectional, which aligns with our motivation for a bidirectional design.

3.3 Loss Function

The proposed stereo image compression network is trained using the rate-distortion loss. The rate is defined by bit-per-pixel (BPP), which includes bits for transmitting latents $\hat{\boldsymbol{y}}_l$, $\hat{\boldsymbol{y}}_r$ and hyperpriors $\hat{\boldsymbol{z}}_l$, $\hat{\boldsymbol{z}}_r$. The distortion is measured using the mean squared error (MSE) and the multi-scale structural similarity index measure (MS-SSIM) [35]. The overall loss function is given by:

$$\mathcal{L} := \lambda D + R = \lambda \sum_{l,r} d(\boldsymbol{x}, \hat{\boldsymbol{x}}) + \sum_{l,r} (r(\hat{\boldsymbol{y}}) + r(\hat{\boldsymbol{z}})),$$
(10)

where $d(\cdot)$ denotes the distortion function with MSE or MS-SSIM, $r(\cdot)$ calculates the bitrate using the estimated probability of \hat{y} and \hat{z} [4], and λ is a hyperparameter that balances the rate and distortion.

3.4 Fast Variant

The proposed entropy model effectively captures multiple references and estimates the probability distribution. However, the auto-regressive iterations in



Fig. 5: Illustration of the entropy model for the proposed fast variant based on stereocheckerboard. The blue and red parts represent the stereo anchor and stereo nonanchor, respectively. By concurrently processing two views, the 3D convolution based anchor context network effectively captures the dependency from the stereo anchor.

the spatial context are time-consuming. To strike a better balance between compression runtime and performance, we need to reduce the auto-regressive executions while maintaining a strong condition for the entropy model. To this end, we propose a fast variant of our model, named BiSIC-Fast, by extending the checkerboard structure [15] into stereo-checkerboard and utilizing 3D convolutions to jointly model the spatial context of two views with only two iterations.

This fast variant maintains the codec as BiSIC, while the entropy model network is modified as illustrated in Fig. 5. In this approach, we split the latents \boldsymbol{y} into two parts, stereo anchor part and stereo non-anchor part, as represented in blue and red in Fig. 5, respectively. The entropy parameter estimation for the stereo anchor part relies on the hyperprior condition and channel context. The stereo anchor is decompressed first, and then the stereo non-anchor utilizes the decompressed stereo anchor as a dependency. Consequently, the entropy parameter estimation for the stereo non-anchor part is based on the hyperprior, channel context, and stereo anchor context.

By adopting this approach, we avoid the spatial auto-regressive process in favor of a two-fold operation, as demonstrated in Fig. 4b and Fig. 4c. Although the spatial auto-regressive process is omitted to save runtime, the neighbors of each pixel in the latents are still maintained, preserving a good perceptive field.

4 Experiments

4.1 Experimental Implementation

Datasets. We evaluate the performance of our methods on two benchmark datasets: InStereo2K [5] and Cityscapes [11]. InStereo2K contains 2,060 indoor scene stereo image pairs of size $1,080 \times 860$. It is split into 2,010 image pairs for training and 50 for testing. Cityscapes contains 5,000 outdoor scene stereo image pairs of size $2,048 \times 1,024$, where 2,975 pairs are allocated for training and 1,525 for testing.

11

Benchmarks. We compare our methods with traditional codecs, such as BPG [6], H.265/HEVC [30], MV-HEVC [31], and H.266/VVC [9]. BPG encodes each view of the stereo pair independently, while H.265/HEVC and H.266/VVC compress the left and right view images jointly. In addition, we show the results of learning-based stereo image compression schemes, including HESIC+ [13], SASIC [37], BCSIC [20], LDMIC [41], and ECSIC [36].

Metrics. The quality of the reconstructed images is evaluated by the peak signal-to-noise ratio (PSNR) and MS-SSIM [35]. The rate is measured by BPP. Besides, we use the widely employed Bjøntegaard Delta bitrate (BDBR) and Bjøntegaard Delta PSNR (BD-PSNR) [8] to compare the rate-distortion (RD) performance among the methods. Specifically, BDBR measures the average bit saving at the same distortion level, while BD-PSNR reflects the relative PSNR gap at an equivalent bitrate. A lower value of BDBR and a higher value of BD-PSNR indicate better performance.

Implementation Details. During the training process, we optimize the model using the MSE loss over 1.5 million steps. We employ the Adam optimizer [19] and set the batch size to 8. The learning rate is initialized at 10^{-4} and is reduced by half after every 300,000 steps. The hyperparameter λ is chosen from the set {256, 512, 1024, 2048, 3072, 4096} for different bitrates. To achieve better MS-SSIM performance, we fine-tune the pre-trained models for an additional 500,000 steps. More implementation details are provided in the supplementary material.

4.2 Results

RD Performance. Fig. 6 shows the RD curves of different compression methods on the InStereo2K and Cityscapes datasets. In terms of the MS-SSIM metric, our proposed method BiSIC outperforms all the other compression methods on both datasets. Besides, BiSIC achieves higher PSNR in the high BPP range compared to the baselines, while maintaining competitive PSNR values at low bitrates. The BDBR results are presented in Tab. 1, where BPG is set as the baseline. Our BiSIC achieves more than 40% and 50% BDBR reduction on InStereo2K and Cityscapes, respectively. Notably, compared to the state-of-the-art video codec VVC, BiSIC achieves an additional 12.76% BDBR reduction for PSNR and an additional 30.08% BDBR reduction for MS-SSIM on InStereo2K. Compared with unidirectional methods, BiSIC achieves significant bit savings by observing a holistic view and mutually sharing features, which helps remove redundancies in each view. Moreover, compared with the state-of-the-art bidirectional coding scheme BCSIC [20], our method achieves an additional 6.5% to 15% BDBR reduction. This is because the employed 3D convolution captures more correlations than the separately deployed 2D convolution used in BCSIC [20], and our cross-dimensional entropy model effectively aggregates various abundant dependencies from different dimensions, providing accurate probability estimation. In addition, the proposed fast variant also demonstrates satisfactory performance, with a marginal degradation in compression performance compared to our BiSIC. Qualitative Results. Fig. 7 visualizes the reconstructed images of different methods on the Cityscapes dataset. For a fair comparison, all images are com-



Fig. 6: Rate-distortion curves in terms of PSNR and MS-SSIM on the InStereo2K and Cityscapes datasets.

pressed to similar bitrates. We observe that the proposed method not only achieves higher PSNR for the reconstructed images at reduced bitrates but also maintains consistent PSNR across stereo views, thanks to its bidirectional architecture. In comparison, the SASIC method displays a PSNR discrepancy of 0.6487 dB between views, as it compresses the right image based on the left one. The VVC codec compresses images in a predictive manner, resulting in an even larger PSNR gap of 2.5769 dB between views. We leave additional qualitative analysis to the supplementary material.

4.3 Runtime Comparison

In this part, we evaluate the computational complexity of various compression algorithms. To ensure a fair comparison, each method is executed on an Intel Xeon Platinum 8370C CPU. Tab. 2 reports the average runtime on the InStereo2K test set, which includes the encoding and decoding latency. We observe that our BiSIC-Fast significantly reduces the runtime compared with the proposed BiSIC. This is because we replace the spatial auto-regressive interweaving in the entropy model with a two-fold stereo-checkerboard, which simplifies the dependency structure within the latent representations. Note that BiSIC-Fast still maintains competitive performance, as demonstrated in Tab. 1. Therefore, it achieves a better trade-off between the compression runtime and performance.

	InStereo2K		Cityscapes	
Method	PSNR	MS-SSIM	PSNR	MS-SSIM
HEVC [30]	-18.80%	-12.59%	-27.46%	-24.36%
MV-HEVC [31]	-7.69%	2.14%	-18.02%	-17.13%
VVC [9]	-35.31%	-31.05%	-56.25%	-44.04%
HESIC+[13]	-14.96%	-43.22%	-23.83%	-50.79%
SASIC [37]	-18.40%	-23.87%	-21.47%	-29.78%
BCSIC [20]	-41.22%	-54.67%	-42.62%	-60.72%
LDMIC [41]	-41.95%	-58.98%	-41.92%	-61.90%
ECSIC [36]	-43.71%	-55.65%	-52.06%	-64.96%
BiSIC (ours)	-48.07%	-61.13%	-57.49%	-67.98%
BiSIC-Fast (ours)	<u>-45.35%</u>	<u>-59.36%</u>	-51.96%	<u>-65.56%</u>

Table 1: BDBR values of different compression methods. **Bold** indicates best results, and <u>underlined</u> values are the second-best ones.



Fig. 7: Qualitative comparison of image reconstruction quality across various codecs.

4.4 Ablation Study

We conduct ablation experiments on the InStereo2K dataset to analyze the contribution of the proposed modules. Our model comprises several key components, including the 3D convolution backbone, cross-dimensional entropy model, and mutual attention block. To evaluate their impact, we replace each module with a baseline counterpart or remove it from the model, and compare the performance in terms of PSNR. The experimental results are shown in Fig. 8 and Tab. 3.

Effectiveness of 3D Convolution Backbone. To assess the capability of 3D convolutional layers in learning inter-view features, we replace 3D convolutional layers in codec with traditional 2D convolutional layers. The resultant model, referred to as Baseline (2D-Conv), experiences a PSNR degradation of 0.3218 dB at the same bitrate, as shown in Fig. 8 and Tab. 3. This indicates that the 3D convolutional layers are more effective at capturing the spatio-stereo information inherent in the data and reducing the redundancies between views.

Effectiveness of Cross-Dimensional Entropy Model. To validate the performance improvement offered by our proposed cross-dimensional entropy model, we compare it with the joint auto-regressive and hyperprior entropy model [24], which we refer to as Baseline (Minnen). As shown in Fig. 8 and Tab. 3, the

Table 2: Average ru	untime of	different	methods.
---------------------	-----------	-----------	----------

$Codec \ \big \ BPG$	HEVC VV	C SASIC BC	SIC LDMIC	BiSIC (ours)	BiSIC-Fast (ours)
Runtime 16.17s	28.16s 190.2	7s 20.24s 89.	44s 81.35s	167.25s	22.82s



Method	BD-PSNR	
BiSIC (ours)	0	
Baseline (2D-Conv)	-0.3218 dB	
Baseline (Minnen)	$-0.3475~\mathrm{dB}$	
Baseline (Ying)	-0.6554 dB	
Baseline (Lei)	$-0.4956~\mathrm{dB}$	
Baseline (w/o Atten)	$-0.7882~\mathrm{dB}$	

Fig. 8: Rate-distortion curves of different ablation baselines on InStereo2K.

Table 3: BD-PSNR [8] results rel-ative to the proposed BiSIC.

proposed cross-dimensional entropy model achieves better rate-distortion performance than Baseline (Minnen). This suggests that our model provides more accurate probability estimations, which in turn minimizes the coding overhead. **Effectiveness of Mutual Attention Block.** To evaluate the effectiveness of the proposed mutual attention block, we consider three baselines for comparisons, namely Baseline (Ying), Baseline (Lei), and Baseline (w/o Atten). Baseline (Ying) and Baseline (Lei) replace the mutual attention block with the stereo attention module [39] and the contextual transfer module [20], respectively. In particular, Baseline (w/o Atten) represents a variant without the attention module. As shown in Fig. 8 and Tab. 3, our proposed model outperforms all baselines by a large margin. Notably, removing the proposed mutual attention block results in a dramatic degradation, which demonstrates its significance.

5 Conclusion

In this work, we introduce a bidirectional stereo image compression network that employs 3D convolution to encode local stereo features and propose a mutual attention block to capture global features. Besides, we design a symmetric cross-dimensional entropy model that integrates hyperprior, spatial context, channel context, and stereo dependency. Our bidirectional design mitigates the issue of imbalanced compression quality intrinsic to unidirectional methods and eliminates the requirement for sequential coding. As demonstrated in the experimental results, our BiSIC achieves greater bit savings at the same level of compression quality compared to both traditional coding standards and existing learning-based methods. Moreover, we propose a fast variant that significantly reduces the compression runtime while maintaining competitive performance.

Acknowledgements

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

References

- Arena, F., Collotta, M., Pau, G., Termine, F.: An overview of augmented reality. Computers 11(2), 28 (2022)
- Ballé, J., Chou, P.A., Minnen, D., Singh, S., Johnston, N., Agustsson, E., Hwang, S.J., Toderici, G.: Nonlinear transform coding. IEEE Journal of Selected Topics in Signal Processing 15(2), 339–353 (2020)
- Ballé, J., Laparra, V., Simoncelli, E.P.: End-to-end optimized image compression. arXiv preprint arXiv:1611.01704 (2016)
- Ballé, J., Minnen, D., Singh, S., Hwang, S.J., Johnston, N.: Variational image compression with a scale hyperprior. arXiv preprint arXiv:1802.01436 (2018)
- Bao, W., Wang, W., Xu, Y., Guo, Y., Hong, S., Zhang, X.: Instereo2k: a large real dataset for stereo matching in indoor scenes. Science China Information Sciences 63, 1–11 (2020)
- 6. Bellard, F.: BPG image format. Website (2014), https://bellard.org/bpg/
- Bishop, C.M.: Latent variable models. In: Learning in graphical models, pp. 371– 403. Springer (1998)
- Bjontegaard, G.: Calculation of average psnr differences between rd-curves. ITU SG16 Doc. VCEG-M33 (2001)
- Bross, B., Wang, Y.K., Ye, Y., Liu, S., Chen, J., Sullivan, G.J., Ohm, J.R.: Overview of the versatile video coding (VVC) standard and its applications. IEEE Transactions on Circuits and Systems for Video Technology **31**(10), 3736–3764 (2021)
- Cheng, Z., Sun, H., Takeuchi, M., Katto, J.: Learned image compression with discretized gaussian mixture likelihoods and attention modules. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
- Deng, X., Deng, Y., Yang, R., Yang, W., Timofte, R., Xu, M.: MASIC: Deep mask stereo image compression. IEEE Transactions on Circuits and Systems for Video Technology (2023)
- Deng, X., Yang, W., Yang, R., Xu, M., Liu, E., Feng, Q., Timofte, R.: Deep homography for efficient stereo image compression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1492–1501 (June 2021)
- He, D., Yang, Z., Peng, W., Ma, R., Qin, H., Wang, Y.: 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. pp. 5718–5727 (2022)
- He, D., Zheng, Y., Sun, B., Wang, Y., Qin, H.: Checkerboard context model for efficient learned image compression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14771–14780 (2021)

- 16 Z. Liu et al.
- Hu, Y., Yang, W., Liu, J.: Coarse-to-fine hyper-prior modeling for learned image compression. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 11013–11020 (2020)
- Jiang, W., Yang, J., Zhai, Y., Ning, P., Gao, F., Wang, R.: Mlic: Multi-reference entropy model for learned image compression. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 7618–7627 (2023)
- Kadaikar, A., Dauphin, G., Mokraoui, A.: Joint disparity and variable size-block optimization algorithm for stereoscopic image compression. Signal Processing: Image Communication 61, 1–8 (2018)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Lei, J., Liu, X., Peng, B., Jin, D., Li, W., Gu, J.: Deep stereo image compression via bi-directional coding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19669–19678 (June 2022)
- Liu, J., Wang, S., Urtasun, R.: DSIC: Deep stereo image compression. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
- Liu, J., Sun, H., Katto, J.: Learned image compression with mixed transformer-cnn architectures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14388–14397 (2023)
- Liu, Y., Ren, J., Zhang, J., Liu, J., Lin, M.: Visually imbalanced stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2029–2038 (2020)
- Minnen, D., Ballé, J., Toderici, G.D.: Joint autoregressive and hierarchical priors for learned image compression. Advances in neural information processing systems 31 (2018)
- Minnen, D., Singh, S.: Channel-wise autoregressive entropy models for learned image compression. In: 2020 IEEE International Conference on Image Processing (ICIP). pp. 3339–3343. IEEE (2020)
- Pan, G., Lu, G., Hu, Z., Xu, D.: Content adaptive latents and decoder for neural image compression. In: European Conference on Computer Vision. pp. 556–573. Springer (2022)
- Shen, Z., Zhang, M., Zhao, H., Yi, S., Li, H.: Efficient attention: Attention with linear complexities. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 3531–3539 (2021)
- Skodras, A., Christopoulos, C., Ebrahimi, T.: The jpeg 2000 still image compression standard. IEEE Signal processing magazine 18(5), 36–58 (2001)
- Song, X., Zhao, X., Fang, L., Hu, H., Yu, Y.: Edgestereo: An effective multi-task learning network for stereo matching and edge detection. International Journal of Computer Vision 128, 910–930 (2020)
- Sullivan, G.J., Ohm, J.R., Han, W.J., Wiegand, T.: Overview of the high efficiency video coding (HEVC) standard. IEEE Transactions on circuits and systems for video technology 22(12), 1649–1668 (2012)
- Tech, G., Chen, Y., Müller, K., Ohm, J.R., Vetro, A., Wang, Y.K.: 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 (2015)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)

- Vetro, A., Wiegand, T., Sullivan, G.J.: 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 (2011)
- Wallace, G.K.: The jpeg still picture compression standard. Communications of the ACM 34(4), 30–44 (1991)
- Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003. vol. 2, pp. 1398–1402. Ieee (2003)
- 36. Wödlinger, M., Kotera, J., Keglevic, M., Xu, J., Sablatnig, R.: ECSIC: Epipolar cross attention for stereo image compression. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3436–3445 (2024)
- Wödlinger, M., Kotera, J., Xu, J., Sablatnig, R.: SASIC: Stereo image compression with latent shifts and stereo attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 661–670 (June 2022)
- Yin, H., Wang, Y., Tang, L., Ding, X., Huang, S., Xiong, R.: 3d lidar map compression for efficient localization on resource constrained vehicles. IEEE transactions on intelligent transportation systems 22(2), 837–852 (2020)
- Ying, X., Wang, Y., Wang, L., Sheng, W., An, W., Guo, Y.: A stereo attention module for stereo image super-resolution. IEEE Signal Processing Letters 27, 496– 500 (2020)
- 40. Zhai, Y., Tang, L., Ma, Y., Peng, R., Wang, R.: Disparity-based stereo image compression with aligned cross-view priors. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 2351–2360 (2022)
- 41. Zhang, X., Shao, J., Zhang, J.: LDMIC: Learning-based distributed multi-view image coding. arXiv preprint arXiv:2301.09799 (2023)
- 42. Zhu, Y., Yang, Y., Cohen, T.: Transformer-based transform coding. In: International Conference on Learning Representations (2021)