

Coarse-to-fine Deep Video Coding with Hyperprior-guided Mode Prediction

Zhihao Hu¹, Guo Lu², Jinyang Guo³, Shan Liu⁴, Wei Jiang⁴, Dong Xu^{3†}
¹Beihang University, China ²Beijing Institute of Technology, China
³The University of Sydney, Australia ⁴Tencent, America

Abstract

The previous deep video compression approaches only use the single scale motion compensation strategy and rarely adopt the mode prediction technique from the traditional standards like H.264/H.265 for both motion and residual compression. In this work, we first propose a coarse-to-fine (C2F) deep video compression framework for better motion compensation, in which we perform motion estimation, compression and compensation twice in a coarse to fine manner. Our C2F framework can achieve better motion compensation results without significantly increasing bit costs. Observing hyperprior information (i.e., the mean and variance values) from the hyperprior networks contains discriminant statistical information of different patches, we also propose two efficient hyperprior-guided mode prediction methods. Specifically, using hyperprior information as the input, we propose two mode prediction networks to respectively predict the optimal block resolutions for better motion coding and decide whether to skip residual information from each block for better residual coding without introducing additional bit cost while bringing negligible extra computation cost. Comprehensive experimental results demonstrate our proposed C2F video compression framework equipped with the new hyperprior-guided mode prediction methods achieves the state-of-the-art performance on HEVC, UVG and MCL-JCV datasets.

1. Introduction

Video compression systems are becoming more and more important for various practical applications due to the rapidly increasing demand for transmitting and storing huge amount of videos. While the conventional methods like H.264 [43], H.265 [33] and the recent standard H.266 [32] have achieved promising results based on different hand-crafted techniques, they cannot be end-to-end optimized by using large-scale video datasets.

Recently, a large number of deep video compression works [3, 12, 15, 23] have been proposed (see Section 2 for

more details), and most of them follow the hybrid coding framework [32, 33, 43], in which both motion compensation and residual compression modules are used to reduce the spatio-temporal redundancy. Therefore, two aspects are critical when designing new deep video codecs: 1) how to generate more accurate motion information for better motion compensation and 2) how to design more effective motion compression and residual compression approaches.

The state-of-the-art learning based video compression methods [3, 15, 23] only use the single scale motion estimation and compensation strategy. Considering that the motion patterns in videos may be complex, these single-scale deep video codecs may not work well for compressing videos from complex scenarios with significant motion patterns. Motivated by the successful applications of the coarse-to-fine strategy for various tasks (e.g., optical flow estimation [29] and video super-resolution [41]), in this work, we first propose a new coarse-to-fine deep video compression framework by adopting a two-stage motion compensation strategy to better generate the predicted feature. At the coarse level, given the low-resolution features from the reference frame and the current frame, we perform motion estimation to produce the low-resolution offset features, which are then compressed after using the motion compression module. After upsampling the reconstructed offset features, we further perform coarse-level motion compensation to wrap the high-resolution reference feature as the intermediate predicted feature. Based on this intermediate predicted feature and the high-resolution feature from the current frame, we perform these major operations (i.e., motion estimation, compression, and compensation) again at the fine level to additionally warp this intermediate predicted feature for better motion compensation. Our two-stage coarse-to-fine motion compensation strategy can generate better predicted feature for the subsequent residual compression module without significantly increasing the bit cost, which leads to better video compression performance.

To further improve video compression performance, we also propose two efficient mode prediction methods for both motion compression and residual compression, which are motivated by the success of the rate-distortion (RD) opti-

† Dong Xu is the corresponding author.

mization based mode prediction methods in the traditional codecs [32,33,43] and the recent work [14] (see Section 2.1 for more details). Instead of using the computationally expensive RD optimization technology as in [14], in this work, we propose to train two prediction networks for coding mode prediction, which bring negligible extra computational cost and can thus support more types of coding modes. Specifically, we use discriminant hyperprior information (*i.e.*, the mean and variance values from the hyperprior network [27]) as the the input of the mode prediction networks as it represents the statistical characteristics of different patches and it does not introduce any additional bit cost. Our proposed mode prediction network can be readily used to adaptively select the optimal resolution of each block in motion compensation or decide whether to skip residual information from each block in residual compression.

Our contributions are summarized as follows: (1) We propose a simple and strong C2F deep video compression framework by performing two-stage motion compensation in a coarse-to-fine fashion. (2) We propose two hyperprior-guided mode prediction methods, in which we learn two mode prediction networks by using discriminant hyperprior information as the input. Our hyperprior-guided mode prediction methods do not introduce any additional bit cost, bring negligible computational cost, and can be readily used to predict the optimal coding modes (*i.e.*, the optimal block resolution for motion coding and the “skip”/“non-skip” mode for residual compression). (3) Comprehensive experiments on the HEVC, UVG and MCL-JCV datasets demonstrate our C2F framework equipped with the newly proposed hyperprior-guided mode prediction methods achieves comparable video compression performance with H265(HM) [1] in terms of PSNR and generally outperforms the latest standard VTM [2] in terms of MS-SSIM.

2. Related Work

2.1. Image and Video Compression

To reduce the transmission and storage requirement from huge volumes of image/video data, different image and video compression standards such as JPEG [39], JPEG2000 [35], H.264 [43], H.265 [33] and H.266 [32] were proposed. As these methods rely on hand-crafted techniques like DCT, they cannot be end-to-end optimized with other networks designed for various machine vision tasks (*e.g.*, object detection).

Recently, a few learning-based image compression [4, 5, 7, 9, 10, 27, 28, 36–38] and video compression [3, 8, 12, 13, 20–25, 30, 31, 44, 46] methods have achieved promising compression performance. For example, Ballé *et al.* [5] proposed to use the hyperprior network to reduce the entropy and save bits for better image compression. For video compression, the work FVC [15] performed all major op-

erations in the feature space for better motion compensation and residual compression. However, the state-of-the-art video compression approaches [3, 15, 23] only use the single scale motion compensation strategy. In contrast to these works [3, 15, 23], we propose a new deep video compression framework, in which we perform motion compensation twice in a coarse-to-fine fashion.

While the coding mode prediction technology is commonly employed in the traditional codes H.264/H.265, it is rarely used in deep video compression approaches except the recent work RaFC [14], in which the rate-distortion (RD) optimization technology is used for the resolution-adaptive motion coding. However, there are two drawbacks in [14]. First, it is computationally expensive to go through the whole network to calculate the RD value for each coding mode and thus it cannot support a large number of coding modes (*e.g.*, only three coding modes are supported in [14]). Second, additional bits are required for encoding the predicted coding modes (*i.e.*, the learnt binary masks representing the coding modes at different blocks). In contrast to [14], we propose to learn two networks for coding mode prediction by using discriminant hyperprior information as the input. Our approach does not introduce any additional bit costs, and brings negligible computational costs, which can support more types of coding modes and thus achieves better compression performance.

2.2. Coarse-to-fine Strategy in Computer Vision

The coarse-to-fine strategy is commonly used for various computer vision tasks including optical flow estimation [16, 29, 34] and super-resolution [41]. Using the pyramid structure, the optical flow estimation method [29] adopted the coarse-to-fine strategy to extract more accurate motion information. The video super-resolution method [41] also adopted the coarse-to-fine alignment strategy to generate better high-resolution frames. However, how to use the coarse-to-fine strategy for video compression is not explored in the existing video compression methods. In contrast to [3, 15, 23], in this work, we propose a new coarse-to-fine video compression framework, in which our approach aims to address a different major challenge, namely, how to improve the motion compensation results without significantly increasing the bit cost.

3. Method

3.1. overview

Given the input video sequence $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_{t-1}, \mathbf{X}_t, \dots\}$, video compression systems aim to reconstruct each video frame with high quality at any bit-rate. In this work, we directly use the existing image compression method [6] to reconstruct the I frame and then employ the reconstructed previous frame as the reference frame for compressing the current frame. The

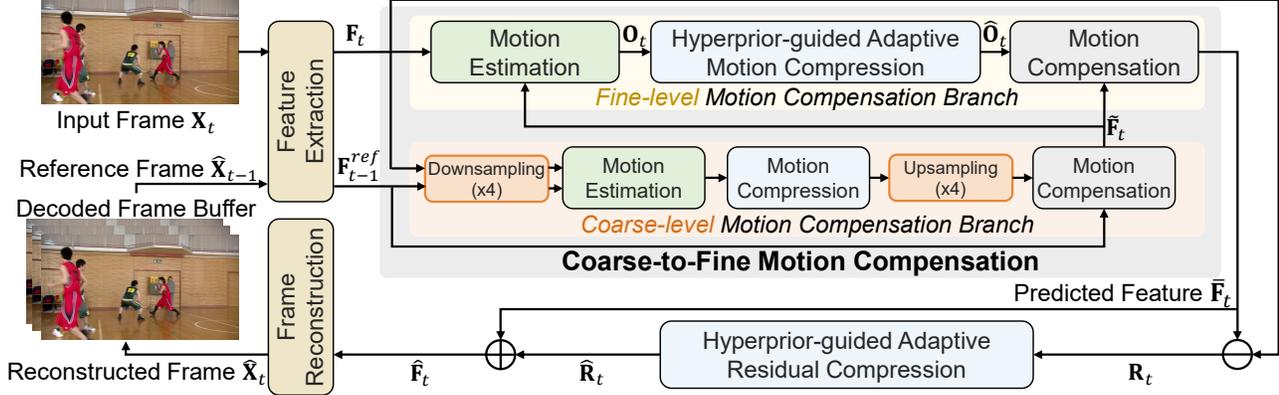


Figure 1. Overview of our proposed video compression framework. Taking the input frame X_t to be compressed and the reference frame \hat{X}_{t-1} from the decoded frame buffer, we first perform the Feature Extraction operation to generate the input feature F_t and the reference feature F_{t-1}^{ref} . Then the two-stage **Coarse-to-Fine Motion Compensation** module is used to compensate the reference feature in a coarse-to-fine fashion and generate the predicted feature \bar{F}_t , in which a hyperprior-guided adaptive motion compression method is proposed for better motion compression. Last, the hyperprior-guided adaptive residual compression module is proposed to compress the residual feature R_t between the input feature F_t and the predicted feature \bar{F}_t . After the frame reconstruction module, we produce the reconstructed frame \hat{X}_t , which is then stored in the decoded frame buffer.

overview of our proposed framework is shown in Fig. 1 and summarized as follows:

Feature Extraction. Following the previous work FVC [15], we first transform the input frame X_t and the reference frame \hat{X}_{t-1} into the input feature F_t and the reference feature F_{t-1}^{ref} , respectively. The network structure of the feature extraction module is the same as that in FVC [15], which consists of a convolution layer with stride 2 and a few residual blocks.

Coarse-to-Fine Motion Compensation. In order to produce more accurate motion compensation results, we propose the two-stage coarse-to-fine motion compensation module. At the coarse level, we first generate two low resolution features by downsampling the input feature F_t and the reference feature F_{t-1}^{ref} and then we perform motion estimation, motion compression and the upsampling operation to generate the reconstructed coarse-level offset map, based on which we perform motion compensation to warp the reference feature F_{t-1}^{ref} and eventually we generate the intermediate predicted feature \tilde{F}_t . Based on \tilde{F}_t and the input feature F_t , in the fine-level motion compensation module, we perform three major operations including motion estimation, motion compression and motion compensation again at the fine level to generate the final predicted feature \bar{F}_t . The network structure of the fine-level modules are the same as those in FVC [15], except that we adopt the newly proposed hyperprior-guided adaptive motion compression module (see Section 3.3 for more details), in which we learn a prediction network based on hyperprior information to decide the optimal block resolution for better motion coding.

Hyperprior-guided Adaptive Residual Compression. The residual between the input feature F_t and the final predicted feature \bar{F}_t is denoted by the residual feature R_t and it will be compressed by the hyperprior-guided adaptive resid-

ual compression module (see Section 3.3 for more details), in which based on hyperprior information we also learn a prediction network to predict the “skip”/“non-skip” mode for better encoding residual features. Adding back the reconstructed residual feature \hat{R}_t to the final predicted feature \bar{F}_t , we produce the reconstructed feature \hat{F}_t .

Frame Reconstruction. Feeding the reconstructed feature \hat{F}_t into the frame reconstruction module that consists of a few residual blocks and a deconvolution layer [15], we generate the reconstructed frame \hat{X}_t , which is then stored in the decoded frame buffer for processing the next frame.

Entropy Coding. The encoded features from the coarse-level motion compensation, the fine-level motion compensation and the residual compression modules will be transformed into the bit-streams. During the training process, we use the bit-rate estimation network to predict the bit-rate. More details will be discussed in Section 3.4.

3.2. Coarse-to-Fine Motion Compensation

To produce high quality features after motion compensation, we propose the C2F video compression framework by performing motion compensation in a coarse-to-fine fashion with only little computation and bit cost at the coarse level. As shown in Fig. 1, we first use the coarse-level motion compensation branch to generate the intermediate predicted feature \tilde{F}_t , and then take \tilde{F}_t as the new reference feature to perform fine-level motion compensation.

Specifically, as shown in Fig. 2, at the coarse-level motion compensation branch, the downsampling operation, which contains two convolution layers with stride 2, will transform the two features F_{t-1}^{ref} and F_t with the resolution of $H \times W$ into the coarse features with the resolution of $\frac{1}{4}H \times \frac{1}{4}W$. Based on the coarse features, we will go through the motion estimation module consisting of two convolu-

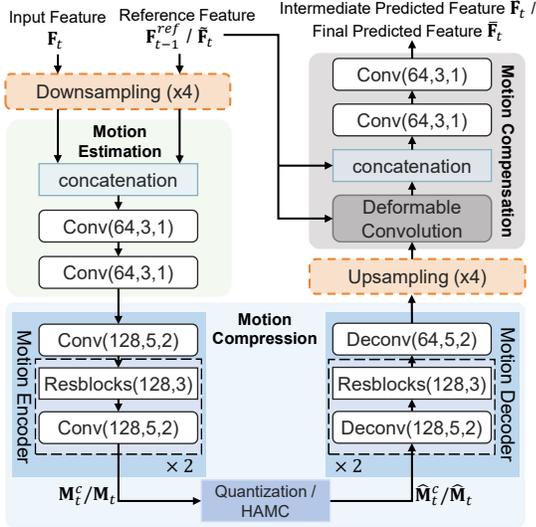


Figure 2. Network structure of our proposed coarse-to-fine framework including three major modules (*i.e.*, the motion estimation/compression/compensation modules). Coarse-level motion compensation takes the reference feature \mathbf{F}_{t-1}^{ref} and the input feature \mathbf{F}_t as the input to generate the intermediate predicted feature $\tilde{\mathbf{F}}_t$, while fine-level motion compensation takes $\tilde{\mathbf{F}}_t$ and \mathbf{F}_t as the input to generate the final predicted feature $\hat{\mathbf{F}}_t$. Note we use the dashed boxes for both “upsampling” and “downsampling” operations because the two modules are only used in the coarse-level motion compensation branch. “HAMC” means Hyperprior-guided Adaptive Motion Compression.

tional layers, the motion compression module and the upsampling operation to generate the reconstructed coarse-level offset map with the resolution of $H \times W$. Here the coarse-level encoded motion feature \mathbf{M}_t^c will be quantized and used for entropy coding. Finally, we follow FVC [15] and use the deformable convolution [11] operation for feature space motion compensation, which takes the reconstructed offset map as the input to control the sampling location in the reference feature map. As in [15], we additionally concatenate the output of the deformable convolution layer with the reference feature and use two convolution layers to generate the intermediate predicted feature $\tilde{\mathbf{F}}_t$.

The fine-level motion compensation branch is similar to the coarse-level motion compensation branch except that the input reference feature \mathbf{F}_{t-1}^{ref} is replaced by the intermediate predicted feature $\tilde{\mathbf{F}}_t$ and the downsampling and upsampling modules are removed. In order to more effectively compress the motion information at the fine level, we propose a new hyperprior-guided adaptive motion compression method to adaptively compress the encoded motion feature \mathbf{M}_t^c into bit-stream, which will be discussed in Section 3.3.

By using the coarse-level motion compensation branch to first roughly compensate the reference feature \mathbf{F}_{t-1}^{ref} with little computation and bit costs, we can more accurately produce the predicted feature $\tilde{\mathbf{F}}_t$ at the fine-level motion compensation branch, which eventually leads to better video

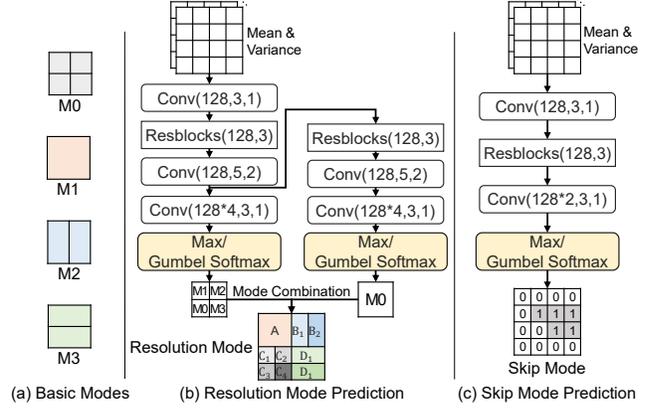


Figure 3. (a) Four basic modes used in the resolution mode prediction network. (b) The resolution mode prediction network that predicts the optimal resolution for each block in the hyperprior-guided adaptive motion compression module and (c) the skip mode prediction network that predicts whether to skip residual information from each block in the hyperprior-guided adaptive residual compression module. For better illustration, we assume the size of the encoded motion/residual feature is 4×4 and only visualize the predicted mode from one channel.

compression performance.

3.3. Hyperprior-guided Adaptive Motion Compression and Residual Compression

In order to more effectively compress motion information at the fine-level motion compensation branch, we propose the hyperprior-guided adaptive motion compression (HAMC) method to automatically predict the optimal block resolution for different spatial locations and different channels based on hyperprior information (*i.e.*, the mean and variance values decoded from the hyperprior network), which can better compress motion information with no extra bits and negligible extra computation cost.

Our method aims to decide the optimal partition mode for each target block, which can be predicted by using our resolution mode prediction network. For better illustration, below we assume the size of the encoded motion feature is 4×4 . Specifically, our resolution mode prediction network consists of two branches, which predict the optimal mode for each 2×2 subblock and the overall 4×4 block, respectively. As shown in Fig. 3(a), there are four basic modes for each 2×2 subblock (note similarly, we also have four basic modes for the target 4×4 block). Then the predicted modes from the two branches will be combined to generate a large number of possible coding modes for the target block. When the predicted mode for the target block is M0, we additionally use the four modes predicted at each of the 2×2 subblocks to generate the updated coding mode for this target block. Otherwise, we directly use the mode predicted for this target block. An example is shown in Fig. 3(b), as the mode predicted for the 4×4 target block is M0, we

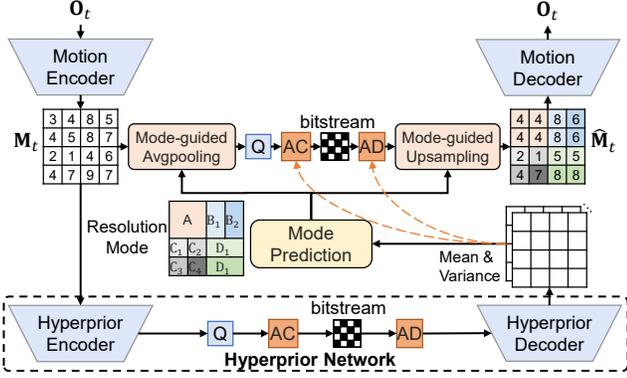


Figure 4. The network structure of the hyperprior-guided adaptive motion compression method. AC and AD denote arithmetic coding and arithmetic decoding, respectively. Integers are used in \mathbf{M}_t and we only visualize one channel for better illustration.

use the four predicted modes [M1, M2; M0, M3] from the four corresponding 2×2 subblocks to generate the final predicted mode of this target block.

To predict the optimal mode, we propose a mode prediction network to automatically decide the resolution for each block based on hyperprior information, which represents the statistical information of each block. As shown in Fig. 3(b), we take the mean and variance values from the hyperprior decoder as the input of the mode prediction network to generate the confidence score of each mode. The channel number of the final convolution layer (*i.e.*, before the gumbel softmax layer) is 128×4 , which represents the confidence scores for 4 modes (shown in Fig. 3(a)) over 128 channels for each 2×2 or 4×4 block. During the inference stage, we directly decide the optimal mode based on the maximum confidence score. However, the max operation is undifferentiable, so the whole network cannot be end-to-end optimized through back-propagation. To address this issue, during the training process, we adopt the Gumbel softmax strategy [17] to decide the optimal mode as this Gumbel softmax module is differentiable, which thus enables end-to-end optimization for the whole network. Finally, after using the max/Gumbel softmax operation to generate the optimal mode for each 4×4 or 2×2 block and for each channel, we combine the predicted basic modes from each 4×4 or 2×2 block to generate the optimal resolution mode for this target block (see Fig. 3(b)).

The whole network structure of our proposed HAMC method is shown in Fig. 4. We take the offset map \mathbf{O}_t (*i.e.*, the output of the fine-level motion estimation module) as the input of the motion encoder to generate the encoded motion feature \mathbf{M}_t . In our proposed HAMC, we predict the optimal resolution mode based on hyperprior information from the hyperprior network. According to the predicted resolution mode, we can more effectively transform the encoded motion feature into the bitstream. Taking subblock A (*i.e.*, the top-left 2×2 subblock) in Fig. 4 as an example, we first

perform the mode-guided avgpooling operation to average pool the four values “3, 4, 4, 5” in the left-top 2×2 subblock into only one value “4”, which is then quantized and transmitted as the bitstream by using the arithmetic coding (AC) operation. After the arithmetic decoding (AD) operation, we perform the mode-guided upsampling operation to generate the four values of “4” in the left-top 2×2 subblock. Considering that we only transmit one value “4” instead of four values “3, 4, 4, 5” to the decoder side (see the top-left 2×2 subblocks in \mathbf{M}_t and $\hat{\mathbf{M}}_t$), our method uses much less bits. Similar operations are also performed for other three subblocks. Consequently, our proposed HAMC method can automatically select large block sizes in smooth areas with less significant motion patterns for bit-rate saving and use small block sizes for areas around moving object boundaries for achieving more accurate motion compensation results. In this way, we can effectively reduce the number of bits for transmitting the encoded motion feature \mathbf{M}_t without substantially degrading the quality of the reconstructed features, which leads to better compression results.

Adaptive Residual Compression. We also propose the hyperprior-guided adaptive residual compression (HARC) method to more effectively compress the sparse residual information, in which we use the “skip”/“non-skip” mode prediction network to predict whether to skip the residual information in each block. As shown in Fig. 3(c), the skip mode prediction network only use one-branch network instead of two branches as in HAMC. It also takes hyperprior information as the input to predict the “skip”/“non-skip” mode for each entry at each channel of the encoded residual feature. As a result, the areas consisting of insignificant residual information will be predicted as the “skip” mode for bit-rate saving, while the significant residual information will still be transmitted to the decoder side for better reconstruction result.

3.4. Loss Function and Entropy Coding

The whole network is end-to-end optimized by minimizing the following rate-distortion loss,

$$\mathcal{L} = \mathbb{H}(\hat{\mathbf{M}}_t^c) + \mathbb{H}(\hat{\mathbf{M}}_t) + \mathbb{H}(\hat{\mathbf{Y}}_t) + \lambda d(\mathbf{X}_t, \hat{\mathbf{X}}_t), \quad (1)$$

where $\mathbb{H}(\cdot)$ denotes the number of bits for encoding the features including the quantized encoded coarse-level motion feature $\hat{\mathbf{M}}_t^c$, the quantized encoded fine-level motion feature $\hat{\mathbf{M}}_t$ and the quantized encoded residual feature $\hat{\mathbf{Y}}_t$. $d(\mathbf{X}_t, \hat{\mathbf{X}}_t)$ denotes the distortion between the input frame \mathbf{X}_t and the reconstructed frame $\hat{\mathbf{X}}_t$. λ is the hyperparameter that controls the trade-off between the bit-rate and distortion. During the training process, we adopt the bit-rate estimation network from [27] without using the time-consuming auto-regressive model to estimate the bits for compressing $\hat{\mathbf{M}}_t$ and $\hat{\mathbf{Y}}_t$. Considering that the resolution of coarse-level motion feature is relative small, we di-

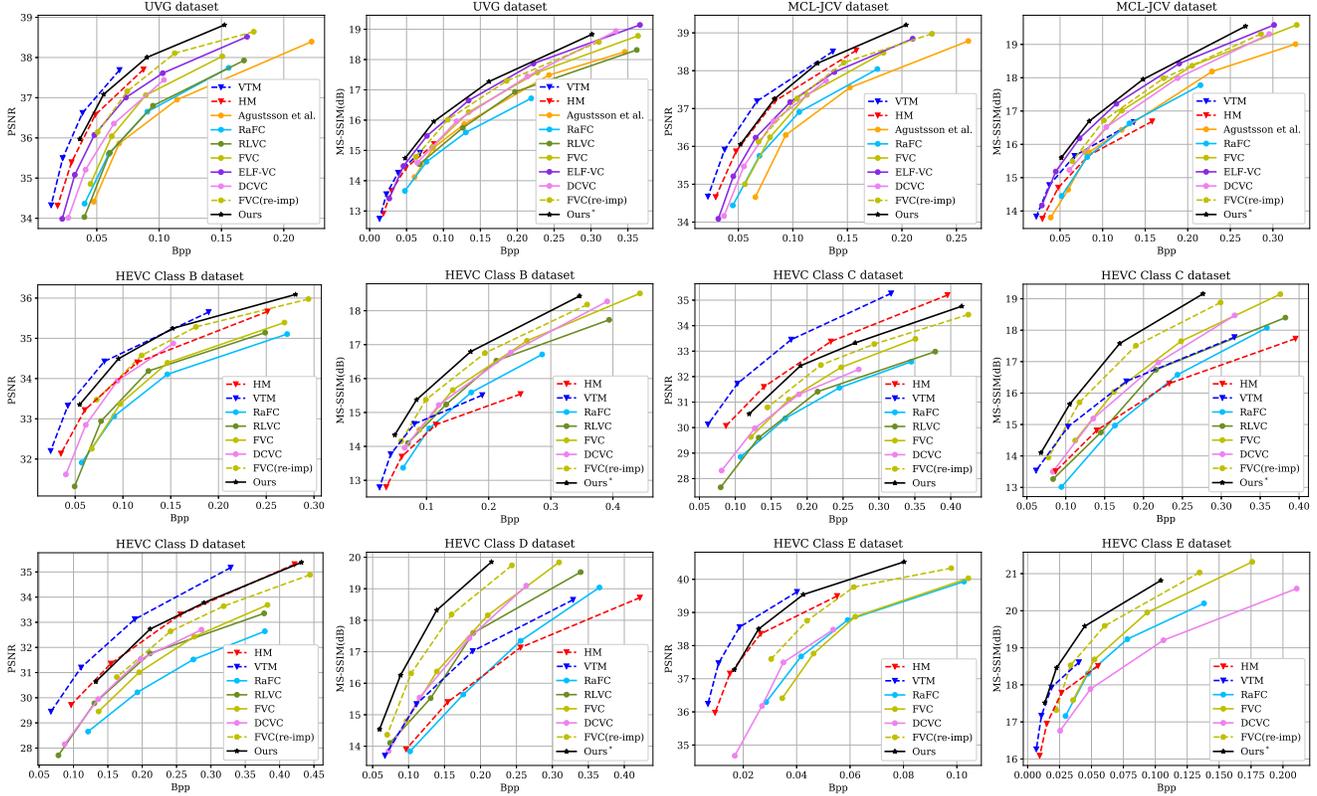


Figure 5. The experimental results on the UVG, MCL-JCV, HEVC Class B, Class C, Class D and Class E datasets.

rectly use the simple bit-rate estimation network in [4] for estimating the bits for M_t^c .

4. Experiments

4.1. Experimental Setup

Training Dataset. Following the previous works [15, 23], we use the Vimeo-90K dataset [45] during the training stage. This dataset contains 89,800 video sequences with each video sequence consisting of 7 consecutive frames with the resolution of 488×256 . For data augmentation, the video sequences are random flipped and random cropped into 256×256 patches before feeding into the network.

Testing Datasets. We evaluate our performance on multiple datasets including the HEVC [33] Class B, C, D, E, UVG [26] and MCL-JCV [40] datasets. The HEVC standard datasets [33] contain different types of video sequences with various resolutions including 1920×1080 (Class B), 832×480 (Class C), 416×240 (Class D) and 1280×720 (Class E). The UVG dataset [26] contains seven 1080p video sequences with high frame rate and the MCL-JCV dataset [40] contains thirty 1080p video sequences, which are widely used for learning-based video codec evaluation.

Evaluation Metric. PSNR and MS-SSIM [42] are used to evaluate the video quality. PSNR is the most popular metric for evaluating the video sequence distortion and MS-SSIM is commonly adopted for subjective visual quality

evaluation. Bit per pixel (bpp) is used to evaluate the number of bits for compressing motion information and residual information.

Implementation Details. We train our model in three stages. At the first stage, we use two consecutive frames including one I frame and one P frame to train our model for 2,000,000 steps without adopting both HAMC and HARC schemes. Then we extend the length of the training video sequence to 7 frames at the second stage for another 300,000 steps. Finally, we add our newly proposed HAMC and HARC schemes and train our complete model for 200,000 steps. The initial learning rate is set as $5e-5$, which is decreased by 80% at the 1,900,000th step and the 2,400,000th step. We set the batch size as 4 for the first stage and 2 for other stages. We use the Adam optimizer [18] based on PyTorch with CUDA support. We use the mean square error as the distortion loss for the PSNR results and additionally fine-tune the PSNR models by using MS-SSIM as the distortion loss for 100,000 steps to produce the MS-SSIM results. When training our model on the machine with a single 2080TI GPU, it takes about 4.5 days, 2 days and 1.3 days for the first stage, the second stage and the final stage, respectively, and it costs 15 hours for fine-tuning the MS-SSIM results. In order to minimize the influence of the video length mismatch between the training sequence and the testing sequence, we additionally adopt the random

shift and different distortion weight strategies as suggested in [25].

4.2. Experimental Results

To evaluate the effectiveness of our proposed method, we compare our proposed method with the state-of-the-art learning-based methods including Agustsson *et al.* [3], RaFC [14], RLVC [47], FVC [15], DCVC [19] and ELFVC [30]. Based on the same setting as our proposed method, we further provide the re-implementation results of FVC [15] without adopting the multi-frame feature fusion module, which is denoted by “FVC(re-imp)” and used as our baseline method. For the conventional methods, we directly use the standard H.265(HM) [1] and VTM [2] with the low delay P configuration for comparison. Different from the previous methods that use the commercial software FFmpeg to generate the results of x265, HM and VTM are the standard versions that can achieve much better performance but are extremely slow.

In order to fairly compare with the standard HM and VTM, we set the GoP size as 100 for all datasets and use BPG [6] for I frame compression. To minimize the cumulative error, we also follow HM and VTM to use a better P-frame compression model for compressing the fourth frame of every 4 frames. When using MS-SSIM for performance evaluation, our model is further fine-tuned by using the MS-SSIM loss as the distortion loss, which is denoted by “Ours*”.

As shown in Fig. 5, our method outperforms all other learning-based methods by a large margin in terms of PSNR. When compared with the recently proposed ELFVC on the UVG dataset, our proposed method achieves 0.5dB improvement at 0.1bpp. Compared with the conventional method H.265(HM) in terms of PSNR, our method achieves better results on most datasets. VTM [2] is the latest video compression standard, which executes various hand-designed modules to achieve the current best performance and thus runs extremely slow (less than 0.001fps). Although the performance of VTM is better than our method in terms of PSNR, we observe that our results are close to VTM at high bit-rate on all high-resolution datasets (*i.e.*, UVG, MCL-JCV, HEVC Class B and Class E). Besides, our method runs at 3.41fps, which is 3000x faster than VTM. Our method generally outperforms all baseline methods in terms of MS-SSIM. Additionally, when using H.265(HM) as the anchor method, our average bit-rate saving over the HEVC Class B,C,D,E datasets is 4.58%.

4.3. Ablation Study

The ablation study of our proposed method is shown in Fig. 6. We take FVC(re-imp) as our baseline method. When compared with FVC(re-imp), our coarse-to-fine framework C2F achieves 0.3dB improvement at 0.08bpp, which indicates the effectiveness of our proposed coarse-

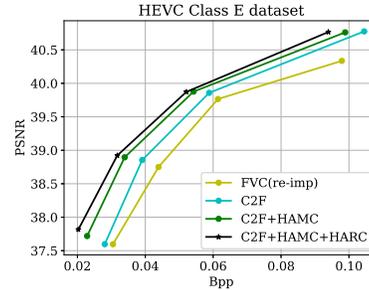


Figure 6. Ablation study on the HEVC Class E dataset. (1) **FVC(re-imp)**: Our re-implemented baseline method FVC [15]. (2) **C2F**: Our proposed coarse-to-fine video compression framework. (3) **C2F+HAMC**: Our proposed C2F framework equipped with the hyperprior-guided adaptive motion compression (HAMC) method. (4) **C2F+HAMC+HARC**: Our proposed C2F framework equipped with both HAMC and the hyperprior-guided adaptive residual compression (HARC) method.

to-fine strategy for motion compensation. When comparing C2F+HAMC with C2F, the newly proposed hyperprior-guided adaptive motion compression (HAMC) method brings 0.6dB improvement at 0.03bpp. In addition, we observe that the coarse-to-fine strategy improves more at higher bit-rate, while HAMC achieves more improvement at lower bit-rate. One possible explanation is that the coarse-to-fine strategy aims to improve the quality of the predicted feature and it is thus more beneficial to use the C2F framework for generating high quality reference frame, while our HAMC scheme focuses on bit-rate saving and it can thus reduce more redundancy at lower bit-rate. Finally, our proposed coarse-to-fine framework equipped with both HAMC and the hyperprior-guided adaptive residual compression (HARC) methods achieves the best results and outperforms the baseline FVC(re-imp) by 1.2dB at 0.03bpp, which demonstrates the effectiveness of our proposed C2F framework and two new methods HAMC and HARC.

4.4. Model Analysis

Visualization of Coarse-to-fine Motion Compensation. To verify the effectiveness of our proposed coarse-to-fine strategy for motion compensation, we take the 5th reconstructed P frame from the first video of the HEVC Class B dataset to visualize both motion information and the corresponding motion compensation results. For fair comparison, here we use the same fine-level motion compression network as FVC(re-imp) without adopting the HAMC method in our C2F framework. It is observed that our coarse-level offset map (see Fig. 7(c)) consists of coarse patch-level motion information, which is easy to be compressed and only costs 0.003bpp. We also observe that the fine-level offset map (see Fig. 7(d)) can capture more detailed pixel-level motion information than the coarse-level offset map, which can better handle more complex motion patterns. As a result, our proposed two-stage coarse-to-

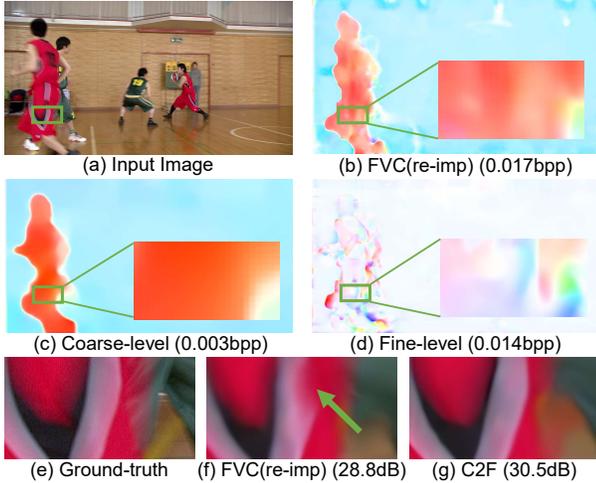


Figure 7. Visualization of (a) the input frame, (b) the offset map from the baseline method FVC(re-imp) using single scale motion compensation and (c) the coarse-level offset map and (d) the fine-level offset map from our coarse-to-fine framework on the HEVC Class B dataset. We also provide (e) one ground-truth patch and its corresponding motion compensation results by using (f) the baseline method FVC(re-imp) and (g) our C2F framework. Bpps for compressing the corresponding offset maps and PSNRs of the corresponding motion compensation results are also reported.

fine motion compensation framework can achieve more accurate motion compensation result (see Fig. 7(g)). When compared with FVC(re-imp) (see Fig. 7(f)), the PSNR of our method is improved by 1.7dB with similar bit-rate cost, which demonstrates that our coarse-to-fine strategy can significantly improve the motion compensation performance. Additionally, when compared with the baseline method FVC(re-imp) on the HEVC Class E dataset, on average the motion compensation result after using our C2F framework is improved by 1.14dB, while costing similar bpps.

Visualization of the Predicted Modes. In Fig. 8, we take the first P frame of the 3rd video from the HEVC Class E dataset as an example and visualize the reconstructed offset and residual maps, as well as their corresponding predicted resolution and “non-skip” modes. We also visualize their corresponding mean and variance values from the hyperprior networks in HAMC and HARC, respectively. From Fig. 8(e), we observe that the areas around moving object boundaries contain large mean and/or variance values. Based on such discriminant hyperprior information, our method HAMC will use the small block sizes when encoding motion information in such areas (see Fig. 8(c)). In contrast, larger block sizes will be preferred for the background areas with small mean and/or variance values. In Fig. 8(b), a few areas contain significant residual information (see the large mean or variance values in Fig. 8(f)). Based on hyperprior information, our HARC method prefers the “non-skip” mode for these blocks when encoding residual information (see Fig. 8(d)).

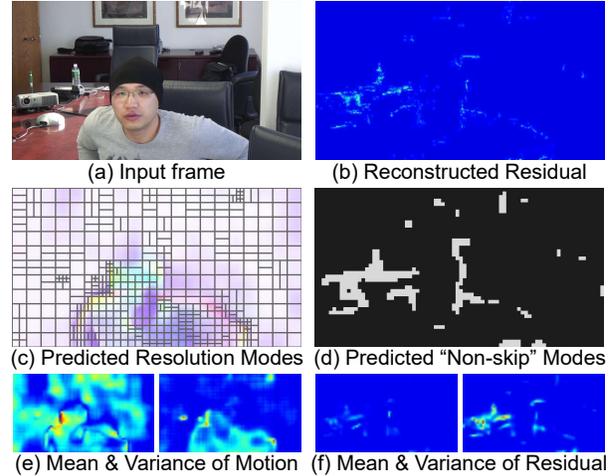


Figure 8. Visualization of the predicted modes for the first P frame (the input frame and the reconstructed residual are shown in (a) and (b)) of the 3rd video from the HEVC Class E dataset by using our proposed methods HAMC (c) and HARC (d). The corresponding mean and variance values (e,f) from the hyperprior networks are also shown. In (b,e,f), the red color (*resp.*, the blue color) denotes large (*resp.*, small) residual/mean/variance value at the corresponding location. In (d), the white color denotes the predicted “non-skip” mode.

Running Speed. We report the inference speed of our proposed framework based on the videos with the resolution of 1920×1080 on the machine with a single 2080TI GPU. Our proposed C2F+HAMC+HARC method runs at 3.41fps while our basic C2F framework without adopting both HAMC and HARC methods runs at 3.43fps, which demonstrate that our proposed hyperprior-guided mode prediction methods bring negligible extra computation cost.

5. Conclusion

In this work, we have proposed a new coarse-to-fine (C2F) video compression framework equipped with two newly proposed hyperprior-guided mode prediction schemes HAMC and HARC, which are respectively used for more accurate motion compensation and for compressing motion and residual information with less bit cost. Comprehensive experiments demonstrate our method achieves comparable performance with H.265(HM) in terms of MS-SSIM. Considering that the mode selection strategy is widely used in the conventional codecs like H.265, our work opens a new door for the subsequent researchers to use/extend our hyperprior-guided method to decide other types of optimal modes for better video compression performance.

Acknowledgement This work was supported by the National Key Research and Development Project of China (No. 2018AAA0101900) and in part by the National Natural Science Foundation of China under Grant 62102024.

References

- [1] HEVC test model (HM). <https://hevc.hhi.fraunhofer.de/HM-doc/>. Accessed: 2022-03-28. **2, 7**
- [2] VVC test model (VTM). <https://jvet.hhi.fraunhofer.de/>. Accessed: 2022-03-28. **2, 7**
- [3] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Ballé, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2020. **1, 2, 7**
- [4] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *International Conference on Learning Representations (ICLR)*, 2017. **2, 6**
- [5] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *International Conference on Learning Representations (ICLR)*, 2018. **2**
- [6] Fabrice Bellard. BPG image format. URL <https://bellard.org/bpg>, 2015. **2, 7**
- [7] Zhenghao Chen, Shuhang Gu, Guo Lu, and Dong Xu. Exploiting intra-slice and inter-slice redundancy for learning-based lossless volumetric image compression. *IEEE Transactions on Image Processing*, 2022. **2**
- [8] Zhenghao Chen, Guo Lu, Zhihao Hu, Shan Liu, Wei Jiang, and Dong Xu. LSVC: A learning-based stereo video compression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 734–743, 2022. **2**
- [9] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learning image and video compression through spatial-temporal energy compaction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10071–10080, 2019. **2**
- [10] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7948, 2020. **2**
- [11] Jifeng Dai, Haozhi Qi, Y. Xiong, Y. Li, Guodong Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 764–773, 2017. **4**
- [12] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers. Neural inter-frame compression for video coding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6421–6429, 2019. **1, 2**
- [13] Runsen Feng, Yaojun Wu, Zongyu Guo, Zhizheng Zhang, and Zhibo Chen. Learned video compression with feature-level residuals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 120–121, 2020. **2**
- [14] Zhihao Hu, Zhenghao Chen, Dong Xu, Guo Lu, Wanli Ouyang, and Shuhang Gu. Improving deep video compression by resolution-adaptive flow coding. In *European Conference on Computer Vision*, pages 193–209. Springer, 2020. **2, 7**
- [15] Zhihao Hu, Guo Lu, and Dong Xu. FVC: A new framework towards deep video compression in feature space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1502–1511, 2021. **1, 2, 3, 4, 6, 7**
- [16] Tak-Wai Hui, Xiaou Tang, and Chen Change Loy. Lite-flownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8981–8989, 2018. **2**
- [17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *International Conference on Learning Representations (ICLR)*, 2017. **5**
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference for Learning Representations*, 2015. **6**
- [19] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34, 2021. **7**
- [20] Jianping Lin, Dong Liu, Houqiang Li, and Feng Wu. M-LVC: Multiple frames prediction for learned video compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3546–3554, 2020. **2**
- [21] Salvator Lombardo, Jun Han, Christopher Schroers, and Stephan Mandt. Deep generative video compression. In *Advances in Neural Information Processing Systems*, pages 9287–9298, 2019. **2**
- [22] Guo Lu, Chunlei Cai, Xiaoyun Zhang, Li Chen, Wanli Ouyang, Dong Xu, and Zhiyong Gao. Content adaptive and error propagation aware deep video compression. In *European Conference on Computer Vision*, pages 456–472. Springer, 2020. **2**
- [23] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. DVC: An end-to-end deep video compression framework. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11006–11015, 2019. **1, 2, 6**
- [24] Guo Lu, Xiaoyun Zhang, Wanli Ouyang, Li Chen, Zhiyong Gao, and Dong Xu. An end-to-end learning framework for video compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, in Press:1–1, 2020. **2**
- [25] Fabian Mentzer, Eirikur Agustsson, Johannes Ballé, David Minnen, Nick Johnston, and George Toderici. Towards generative video compression. *arXiv preprint arXiv:2107.12038*, 2021. **2, 7**
- [26] A. Mercat, Marko Viitanen, and J. Vanne. UVG dataset: 50/120fps 4k sequences for video codec analysis and development. *Proceedings of the 11th ACM Multimedia Systems Conference*, 2020. **6**
- [27] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, pages 10771–10780, 2018. **2, 5**
- [28] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *2020*

- IEEE International Conference on Image Processing (ICIP)*, pages 3339–3343. IEEE, 2020. 2
- [29] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4161–4170, 2017. 1, 2
- [30] Oren Rippel, Alexander G Anderson, Kedar Tatwawadi, Sanjay Nair, Craig Lytle, and Lubomir Bourdev. ELF-VC: Efficient learned flexible-rate video coding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3033–3042, 2021. 2, 7
- [31] Oren Rippel, Sanjay Nair, Carissa Lew, Steve Branson, Alexander G Anderson, and Lubomir Bourdev. Learned video compression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3454–3463, 2019. 2
- [32] Gary Sullivan. Versatile video coding (VVC) arrives. In *2020 IEEE International Conference on Visual Communications and Image Processing (VCIP)*, pages 1–1. IEEE, 2020. 1, 2
- [33] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012. 1, 2, 6
- [34] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. 2
- [35] David S Taubman and Michael W Marcellin. Jpeg2000: Standard for interactive imaging. *Proceedings of the IEEE*, 90(8):1336–1357, 2002. 2
- [36] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. *International Conference for Learning Representations*, 2017. 2
- [37] George Toderici, Sean M O’Malley, Sung Jin Hwang, Damien Vincent, David Minnen, Shumeet Baluja, Michele Covell, and Rahul Sukthankar. Variable rate image compression with recurrent neural networks. *International Conference for Learning Representations*, 2017. 2
- [38] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5306–5314, 2017. 2
- [39] Gregory K Wallace. The JPEG still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992. 2
- [40] Haiqiang Wang, Weihao Gan, Sudeng Hu, Joe Yuchieh Lin, Lina Jin, Longguang Song, Ping Wang, Ioannis Katsavounidis, Anne Aaron, and C-C Jay Kuo. MCL-JCV: a jnd-based H.264/AVC video quality assessment dataset. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 1509–1513. IEEE, 2016. 6
- [41] Xintao Wang, Kelvin C. K. Chan, K. Yu, C. Dong, and Chen Change Loy. EDVR: Video restoration with enhanced deformable convolutional networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1954–1963, 2019. 1, 2
- [42] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003. 6
- [43] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the H.264/AVC video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7):560–576, 2003. 1, 2
- [44] Chao-Yuan Wu, Nayan Singhal, and Philipp Krahenbuhl. Video compression through image interpolation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 416–431, 2018. 2
- [45] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019. 6
- [46] Ren Yang, Fabian Mentzer, Luc Van Gool, and Radu Timofte. Learning for video compression with hierarchical quality and recurrent enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6628–6637, 2020. 2
- [47] R. Yang, Fabian Mentzer, L. Van Gool, and R. Timofte. Learning for video compression with recurrent auto-encoder and recurrent probability model. *IEEE Journal of Selected Topics in Signal Processing*, 15:388–401, 2021. 7