

Neural Video Compression with Feature Modulation

Jiahao Li, Bin Li, Yan Lu
Microsoft Research Asia

{li.jiahao, libin, yanlu}@microsoft.com

Abstract

The emerging conditional coding-based neural video codec (NVC) shows superiority over commonly-used residual coding-based codec and the latest NVC already claims to outperform the best traditional codec. However, there still exist critical problems blocking the practicality of NVC. In this paper, we propose a powerful conditional coding-based NVC that solves two critical problems via feature modulation. The first is how to support a wide quality range in a single model. Previous NVC with this capability only supports about 3.8 dB PSNR range on average. To tackle this limitation, we modulate the latent feature of the current frame via the learnable quantization scaler. During the training, we specially design the uniform quantization parameter sampling mechanism to improve the harmonization of encoding and quantization. This results in a better learning of the quantization scaler and helps our NVC support about 11.4 dB PSNR range. The second is how to make NVC still work under a long prediction chain. We expose that the previous SOTA NVC has an obvious quality degradation problem when using a large intra-period setting. To this end, we propose modulating the temporal feature with a periodically refreshing mechanism to boost the quality. Notably, under single intra-frame setting, our codec can achieve 29.7% bitrate saving over previous SOTA NVC with 16% MACs reduction. Our codec serves as a notable landmark in the journey of NVC evolution. The codes are at <https://github.com/microsoft/DCVC>.

1. Introduction

Traditional standard codec, which relies on a hybrid residual coding-based framework, has been in development for over 30 years and is still being refined. However, the improvement in compression ratio has diminished, while the increase in complexity has grown significantly [43]. This makes further advancements within the traditional framework more and more challenging. Recently, neural video codec (NVC) has gained considerable attention, as it holds the potential to break this development bottleneck.

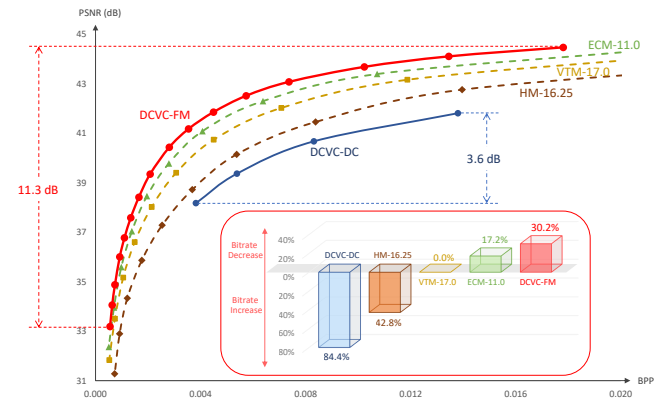


Figure 1. Rate-Distortion curve and BD-Rate comparisons with H.265/HM, H.266/VTM, ECM, and the previous SOTA NVC DCVC-DC [26]. The test dataset is HEVC E (600 frames) with single intra-frame setting (i.e. intra-period = -1) and YUV420 color space. DCVC-DC has a large performance drop under this setting while our new codec DCVC-FM still can significantly surpass ECM. Meanwhile, the quality range of our DCVC-FM is much larger than that of DCVC-DC.

The early NVC model DVC [32] still follows traditional codec and uses the residual coding-based framework. Later, many works [5, 18, 27, 28] are also based on this paradigm and propose stronger sub-modules to improve the performance. By contrast, the emerging conditional coding [22, 24] shows a lower entropy bound than residual coding and has larger potential. The condition can be freely defined and learned rather than being limited to the predicted frame in pixel domain. Meanwhile, the condition can be flexibly used to help encoding, decoding, and entropy modelling. The recent DCVC-DC model [26] already achieves a better compression ratio than both H.266/VTM [10] and the under-development ECM (the prototype of next generation traditional standard) by mining diverse spatial and temporal contexts as the condition.

Despite the progress made by DCVC-DC, we find it is still far from practical usage. The first blocking issue is lacking the wide quality range support. Although DCVC-DC can support multiple quality levels in a single model, its quality range is quite limited and only has an average

of 3.8 dB on various datasets, which definitely cannot meet the needs of a practical product. To solve this problem, we modulate the latent feature of the current frame with the learnable quantization scaler. During the training phase, we not only increase the lambda range, but also specially design a uniform quantization parameter sampling mechanism to let NVC experience various trade-offs between rate and distortion. This can enhance the harmonization between encoding and quantization processes. Consequently, the finely-controllable quantization scaler is obtained and our NVC can seamlessly adjust the quality level under a wide quality range, i.e. about 11.4 dB. Furthermore, with the support of a wide quality range, this paper also showcases the capability of rate control in a single model, when given the specified target bitrate. This effectively demonstrates the practicality of our NVC in real-world scenarios.

The second issue is how to enable NVC to cope with the long prediction chain effectively. Most existing NVC models struggle to address the temporal error accumulation problem therein. To alleviate this issue, many of them [16, 17, 32–34] rely on using a small intra-period setting (e.g. 10 or 12) to insert high-quality intra-frame more frequently. However, a smaller intra-period setting harms the whole compression efficiency. For example, [25, 44] show that intra-period 32 has an average of 23.8% bitrate saving over intra-period 12 for H.265/HM. Thus, traditional standard committee [9] strictly defines the intra-period as -1 , namely only single intra-frame is encoded for the whole video. We think NVC should also follow this setting. This also makes it more fair when comparing NVC and traditional codec. However, we find previous SOTA DCVC-DC has large quality degradation under the intra-period -1 setting, as shown in Fig. 1. To solve this problem, two countermeasures are proposed. One is increasing the video frame number to better learn the long-distance temporal correlation during the training. Another is that we propose modulating the temporal feature by periodically refreshing it, which can significantly alleviate the error propagation.

With these effective Feature Modulation techniques, we build a new codec DCVC-FM, based on DCVC-DC. In addition, our DCVC-FM involves other improvements toward a versatile NVC. Most existing NVCs are only optimized for RGB colorspace. Actually, traditional codecs and practical applications mainly adopt YUV colorspace. To this end, we design an NVC that can support both RGB and YUV without any fine-tuned training. In addition, via improved implementation, this paper also showcases the low-precision inference which can significantly reduce the running time and memory cost with a negligible compression ratio degradation. Experiments show that our DCVC-FM can outperform VTM by 25.5% under intra-period -1 setting, and also achieve non-trivial advantage over ECM. When compared with previous SOTA NVC DCVC-DC,

29.7% bitrate saving is achieved while the MACs (multiply-accumulate operations) are reduced by 16%.

In summary, our contributions are:

- We modulate the latent feature via learnable quantization scaler, where a uniform quantization parameter sampling mechanism is proposed to help its learning. It enables our DCVC-FM to support a wide quality range in a single model, and the rate control capability is demonstrated.
- We not only exploit the training with longer video but also module the temporal feature with a periodically refreshing mechanism to boost the quality. These help our DCVC-FM to tackle the long prediction chain.
- To further improve the practicality, we enable DCVC-FM to support both RGB and YUV colorspaces within a single model. Moreover, we demonstrate the low-precision inference with negligible bitrate increase.
- Our DCVC-FM can outperform all traditional codecs under intra-period -1 setting. When compared with the previous SOTA NVC, 29.7% bitrate reduction is achieved with 16% MAC reduction. Our codec is an important milestone in the development of NVC.

2. Related Work

2.1. Neural Image Compression

Most recent neural image codec (NIC) models follow hyperprior [7] and adopt a hierarchical framework design. Some works [21, 31, 41, 54] use transformer to strengthen the autoencoder or entropy model. Better optimization algorithms [52, 53] are also investigated. Recently, the diffusion model [14, 39, 46, 50] is explored to improve the generation ability. In addition, the light-weight models [47, 51] are also proposed. Now, NIC is quite powerful and its standardization process is already under consideration [6].

2.2. Neural Video Compression

The success of NIC also pushes the development of NVC. The early DVC [32] follows a traditional residual coding-based framework, and uses NIC to code the motion vector and residual, separately. Many NVCs [5, 13, 18, 27–29, 34, 36, 42] also adopt this paradigm and design stronger sub-modules to boost the compression ratio. For example, the optical flow estimation in scale space [5] is proposed to handle the complex motion area. Multiple reference frames are utilized to improve the temporal prediction [27]. The block-based prediction mode selection is proposed in [28].

When compared with residual coding, the conditional coding [15, 23–26, 30, 37, 40, 44] shows larger potential because its temporal context is not limited to the predicted frame in pixel-domain and it does not rely on the sub-optimal subtraction to reduce redundancy. The feature-domain temporal context can be flexibly designed and its correlation with the current frame can be automatically

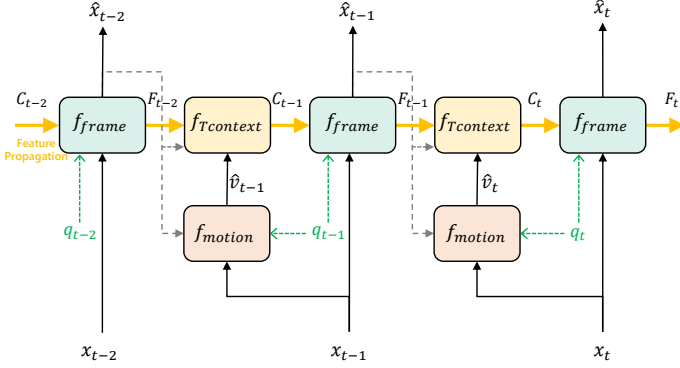


Figure 2. The framework of our DCVC-FM, built on DCVC-DC.

learned. In [30], temporally conditional entropy models is designed. The DCVC series [24–26, 40, 44] propose using high-dimension context to improve encoding, decoding, as well as the entropy modelling. The DCVC-TCM [44] introduces temporal feature propagation [11, 20]. The DCVC-HEM [25] designs a powerful entropy model utilizing both spatial and temporal contexts. The latest DCVC-DC [26] already outperforms the under-developing ECM by continually boosting the context diversity.

However, there still exist several critical problems that block the practicality of NVC. The first is the quality range problem. Although DCVC-HEM and DCVC-DC support variable bitrates in a single mode, their quality range is quite limited and cannot meet the various quality requirements. In addition, most existing NVCs including previous SOTA DCVC-DC still use a small intra-period setting (e.g., 10, 12, and 32), which is far from a practical scenario. By contrast, our proposed DCVC-FM solve both of these two critical problems via feature modulation.

3. Proposed Method

3.1. Overview

Our DCVC-FM adopts the conditional coding-based framework and is built on DCVC-DC [26]. The overall framework is shown in Fig. 2. For coding each input frame x_t (t is the frame index), there are three primary functions: f_{motion} , $f_{Tcontext}$, and f_{frame} . At first, f_{motion} employs an optical flow network to estimate the motion vector v_t between x_t and the previous reconstructed frame \hat{x}_{t-1} . v_t needs to be encoded, transmitted, and decoded as \hat{v}_t . Subsequently, $f_{Tcontext}$ uses \hat{v}_t to extract the temporal context C_t from the propagated feature F_{t-1} coming from the previous frame. Finally, conditioned on the motion-aligned temporal context C_t , x_t is encoded, transmitted, and decoded as \hat{x}_t via the function f_{frame} . Meanwhile, f_{frame} generates the F_t for the need of the next frame. When compared with DCVC-DC, we make the supporting quality range increase

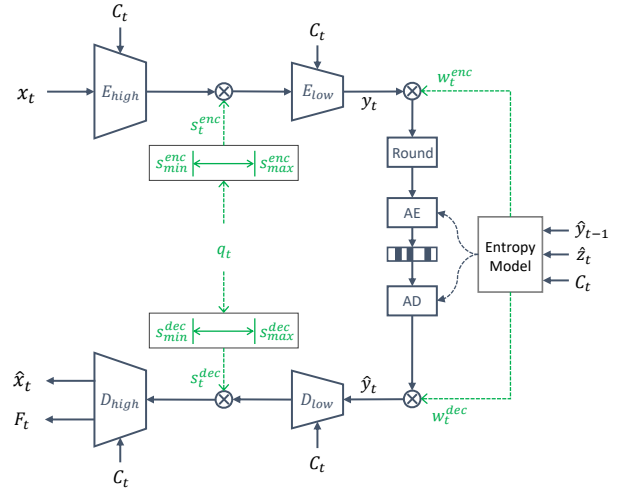


Figure 3. The framework of our frame coding function f_{frame} . E_{high} and E_{low} are encoder at high and low resolution, respectively. D_{high} and D_{low} are corresponding decoder. AE and AD are arithmetic encoder and decoder. The quantization and inverse quantization processes are also applied to f_{motion} in a similar way.

from 3.8 dB to 11.4 dB (Section 3.2). In addition, to effectively cope with the intra-period -1 setting, we improve the feature propagation mechanism (Section 3.3). Section 3.4 shows the details of new capabilities: the single model for both RGB and YUV colorspace, and the low-precision inference via our improved implementation.

3.2. Wide Quality Range in a Single Model

A basic quantization and inverse quantization processes can be formulated as:

$$\hat{I} = QS \cdot \lfloor \frac{I}{QS} \rfloor \quad (1)$$

I and QS are input value and quantization step, respectively. $\lfloor \cdot \rfloor$ is the rounding operation. The QS controls the reconstruction quality of output \hat{I} . To enable variable quality for NVC, we also incorporate similar mechanism into the encoding and decoding processes to modulate the latent feature. The key challenge is how to decide the corresponding QS and also let it support a wide range.

Fig. 3 shows our design for the frame coding function f_{frame} . As shown in this figure, during the encoding process, there are two values related to the quantization step. One is s_t^{enc} and the other is w_t^{enc} . Both of them are used to modulate the latent feature, but they are generated in different ways and applied in different granularities. s_t^{enc} is the global quantization scaler and is generated according to the quantization parameter q_t from the user input, which is similar with the concept of QP in the traditional codec. The q_t is an integer scalar ranging from $[0, q_num-1]$, where q_num is the adjustable number of q_t value and is set as

64 in the implementation. Similar to the traditional video codec where the quantization step size exponentially grows with the linear increase in QP, we use the following equation to interpolate s_t^{enc} based on q_t and the quantization scale range $[s_{min}^{enc}, s_{max}^{enc}]$ via:

$$s_t^{enc} = s_{min}^{enc} \cdot \left(\frac{s_{max}^{enc}}{s_{min}^{enc}} \right)^{\frac{q_t}{q_{num}-1}} \quad (2)$$

To improve the numerical stability, we change it to:

$$s_t^{enc} = e^{\ln s_{min}^{enc} + \frac{q_t}{q_{num}-1} \cdot (\ln s_{max}^{enc} - \ln s_{min}^{enc})} \quad (3)$$

In our design, s_{min}^{enc} and s_{max}^{enc} are learned during training.

We know the optimization object of a codec is usually $Loss_{RD} = R + \lambda \cdot D$ ¹, where R and D represent the bitrate and distortion, respectively. The λ is used to control the trade-off between R and D . At present, most existing NVCs need to train a separate model for each quality level with the corresponding predefined and constant λ value. As for our model, to support variable quality levels, the λ is also variable during the training. We predefine a λ range as $[\lambda_{min}, \lambda_{max}]$. In our implementation, $[\lambda_{min}, \lambda_{max}]$ is set as $[1, 768]$. For each training step, we randomly select q_t and interpolate the λ value from this range. The interpolation is similar with Eq. 3 and calculated as:

$$\lambda = e^{\ln \lambda_{min} + \frac{q_t}{q_{num}-1} \cdot (\ln \lambda_{max} - \ln \lambda_{min})} \quad (4)$$

For each step, we uniformly sample the integral q_t value from the range $[0, q_{num}-1]$, and then obtain the λ value via Eq. 4 to calculate the $Loss_{RD}$. If we compare Eq. 4 and Eq. 3, we can know that the $[\lambda_{min}, \lambda_{max}]$ will guide the learning of $[s_{min}^{enc}, s_{max}^{enc}]$ via the back-propagation of $Loss_{RD}$. By controlling the λ value range, we can easily adjust the value range of quantization scaler. By employing the uniformly sampling mechanism for q_t , the codec is able to experience different s_t^{enc} values and explore various trade-offs between the R and D during training. This enhances the harmonization between the encoding and quantization processes, resulting in that our codec is able to learn finely grained and controllable quantization scaler s_t^{enc} to modulate the latent feature.

It should be noted that, not like Eq. 1, our codec does not have the restriction that the same quantization step value must be used during encoding and decoding because our quantization and inverse quantization are performed at latent feature domain. Removing this restriction enables a larger flexibility to smoothly adjust the quality. Thus, the corresponding s_t^{dec} are separately learned during the decoding. In addition, just because of removing this restriction, we can let s_t^{enc} be multiplied to modulate the latent feature rather than being a divisor during the encoding. This helps

¹Sometimes λ is applied to R .

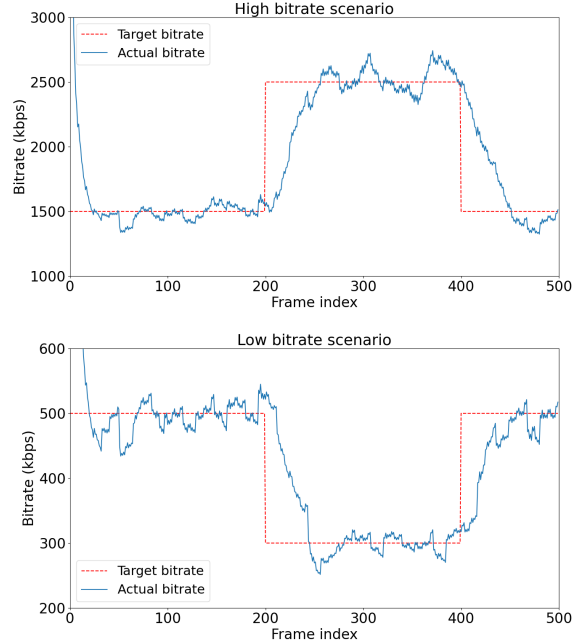


Figure 4. Rate control examples using *BasketballDrive* video sequence (1080p, 50fps, 500 frames). The above (below) example is with a relatively high (low) target bitrate scenario.

stabilize the training process by circumventing the potential division by zero issues.

However, s_t^{enc} is the same for all spatial positions, which may ignore the different spatial characteristics for the video content. So we follow [19, 25] and use entropy model to learn spatial-channel-wise quantization scaler w_t^{enc} . w_t^{enc} not only helps achieve precise modulation at each position but also is adaptive to the video content of each frame. This kind of content-adaptive and dynamic feature modulation can also improve the final compression efficiency.

Benefited from these advanced designs, our DCVC-FM finally can support quality adjustment in a wide range. This capability is also the prerequisite of rate control, which is the core functionality of a practical codec. Fig. 4 shows two rate control examples with different target bitrate scenarios. As we can adjust the q_t for each frame in our codec, the fluctuated target bitrate can be supported, as shown in Fig. 4. The actual bitrate gets close to the target bitrate. It is noted that we just showcase the feasibility of rate control using DCVC-FM rather than focusing on designing new rate control algorithm for NVC. Currently, we just simply adjust the q_t based on the buffer fullness (more details are in supplementary materials). In the future, more advanced rate control algorithms can be proposed based on DCVC-FM.

3.3. Long Prediction Chain

The temporal quality degradation is a fundamental problem for all video codecs, but is especially serious for NVC. Pre-

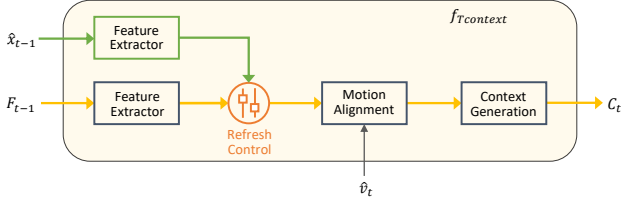


Figure 5. Temporal feature modulation via periodical refresh.

vious NVCs already began to address this problem. For example, DCVC-DC follows traditional codec and introduces the widely-used hierarchical quality structure to periodically improve the quality. This can help alleviate the error propagation but is not enough. As shown in Fig. 6, DCVC-DC has serious quality degradation when only single intra frame is used (i.e. intra-period = -1). To solve this problem, two countermeasures are proposed. It should be noted that our NVC is conditional coding-based framework, whose major advantage is using temporal feature as the context and condition. In particular, the temporal feature can be propagated across many frames. So the key is how to design a more effective temporal feature propagation mechanisms. Our first improvement is increasing the video frame number during the training. Although it is a simple modification, it is quite helpful. A longer video enables recognizing similar patterns over long distances of time and then better exploring the temporal correlation.

Feature propagation is also double-edged sword, as the propagated feature may be contaminated by accumulated errors or contain some uncorrelated information. To this end, we propose modulating the the propagated feature, and force NVC refresh it periodically (the refresh period is set as 32 in the implementation). As shown in Fig. 5, for the input frame x_t , we will not let $f_{Tcontext}$ extract temporal context from the propagated feature F_{t-1} if we expect to perform the feature refresh. Instead, we use a separate feature extractor module to extract the temporal context from \hat{x}_{t-1} . As \hat{x}_{t-1} only has pixel information with 3 dimensions, which contains much less information than F_{t-1} . So, to facilitate the coding of x_t , it will force this separate module to extract correlated temporal context from \hat{x}_{t-1} as much as possible. The new extracted temporal context will be propagated to future frames. Such refresh-based modulation mechanism effectively alleviates the error propagation problem. As shown in Fig. 6, our DCVC-FM can maintain the quality across frames with lower bitrate cost when compared with DCVC-DC.

3.4. Implementation

Single model for both RGB and YUV colorspaces. Although DCVC-DC supports a single network structure for both RGB and YUV colorspaces, the separate model weights with separate training are still required. To improve

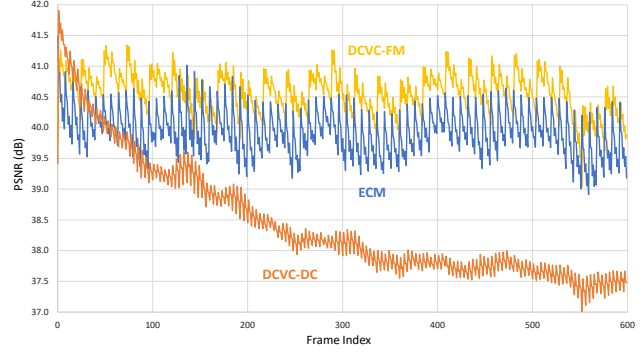


Figure 6. Quality comparison across frames. The test video is *KristenAndSara* from HEVC E dataset (video conferencing scenario). The average bpp (bits per pixel) results of DCVC-DC, ECM, and proposed DCVC-FM are 0.0037, 0.0029, and 0.0026, respectively. Intra-period = -1.

the versatility of NVC, we directly train a single model for both RGB and YUV without additional fine-tuning. To support it, our training loss covers both colorspaces as: $Loss_{RD} = R + \lambda \cdot (k \cdot D_{YUV} + (1-k) \cdot D_{RGB})$. D_{YUV} and D_{RGB} are the distortion in YUV and RGB, respectively. k is the hyper-parameter for weighting, and set as 0.8 in the implementation. To use the same input interface for RGB and YUV, the UV contents will be up-sampled if the input is YUV420 content. Correspondingly, UV contents are down-sampled after obtaining the reconstructed frame.

Low precision inference. Most NVCs only report the results based on 32-bit floating point implementation. We hope to use 16-bit to accelerate the NVC. However, to support 16-bit, we need to improve the implementation of *grid_sample* function, which is widely used for motion alignment. Our DCVC-FM is implemented with PyTorch. In the PyTorch², the “grid” parameter in *grid_sample* represents the absolute position in the frame. When using 16-bit precision, only 10-bit significand is not enough to represent³. To solve this problem, we use 16-bit precision to represent the relative offset and reimplement the *grid_sample*. It enables the NVC inference at 16-bit precision, and can significantly save the memory and complexity with negligible compression ratio change.

Structure optimization. To reduce computation, we further make two additional improvements based on DCVC-DC. One is reducing the convolution kernel size at the high-resolution feature for the motion estimation module. Another is that we use more depthwise separable convolutions which can reduce the computation cost and alleviate the over-fitting simultaneously [12]. The section 4.3 verifies that these two improvements bring non-trivial MAC reduction with quite small bitrate increase.

²Currently up to PyTorch-2.1 version.

³https://en.wikipedia.org/wiki/Half-precision_floating-point_format

Table 1. BD-Rate (%) comparison in RGB colorspace. 96 frames with intra-period=32.

	UVG	MCL-JCV	HEVC B	HEVC C	HEVC D	HEVC E	Average
VTM-17.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HM-16.25	38.4	45.6	40.3	37.9	32.4	41.0	39.3
ECM-5.0	-10.6	-13.2	-11.5	-12.6	-11.2	-9.8	-11.5
CANF-VC	73.0	70.8	64.3	76.2	63.1	120.5	78.0
DCVC	166.1	121.5	123.0	143.0	98.2	272.9	154.1
DCVC-TCM	44.1	51.0	40.2	66.3	37.0	82.7	53.6
DCVC-HEM	1.1	8.6	5.1	22.2	2.4	20.5	10.0
DCVC-DC	-19.1	-11.3	-12.0	-10.3	-26.1	-18.0	-16.1
DCVC-FM	-17.0	-5.6	-14.3	-23.7	-36.7	-24.5	-20.3

Note: Some numbers are slightly different with those in [26] as we are using a wider quality range to calculate BD-Rate.

Table 2. BD-Rate (%) comparison in YUV420 colorspace. 96 frames with intra-period=32.

	UVG	MCL-JCV	HEVC B	HEVC C	HEVC D	HEVC E	Average
VTM-17.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HM-16.25	38.0	45.9	39.3	34.6	29.0	37.5	37.4
ECM-5.0	-11.5	-15.0	-12.7	-13.7	-12.2	-11.0	-12.7
DCVC-DC	-17.8	-12.0	-10.8	-12.4	-28.5	-20.4	-17.0
DCVC-FM	-21.6	-11.4	-16.3	-25.8	-39.3	-30.1	-24.1

4. Experimental Results

4.1. Experimental Settings

Datasets. In line with most existing NVCs, we utilize the Vimeo-90k dataset [49] for the training. Besides the ready-made 7-frame videos in Vimeo-90k, we also process the raw Vimeo videos [4] to generate additional 6,658 videos, and each consists of 32 frames. The 7-frame video is used for training first, followed by the use of the 32-frame for fine-tuning, which further boosts the quality. For testing, we use the common HEVC B~E [9], UVG [38], and MCL-JCV [48] datasets.

Test Conditions. All tests are conducted under low-delay coding settings. We employ the BD-Rate metric [8] to measure the compression ratio change, where positive values are bitrate increase, and negative values indicate bitrate savings. The video quality is evaluated using PSNR. Our benchmarks include traditional codecs H.265/HM [2], H.266/VTM [3], and ECM [1].

For the RGB colorspace, to facilitate the comparison with existing methods, we follow the test condition in [26] and test 96 frames for each video with intra-period 32. We also compare our DCVC-FM with previous SOTA NVC models, including CANF-VC [15], DCVC [24], DCVC-TCM [44], DCVC-HEM [25], and DCVC-DC [26].

For the YUV420 colorspace, we initially follow [26] and test 96 frames with intra-period 32. As aforementioned,

such a small intra-period is far from practical applications, and the traditional standard committee [9] strictly sets the intra-period as -1 . Thus, we advocate for testing NVCs under this intra-period -1 setting. Unlike [26], which only tests 96 frames, we also test all frames with intra-period -1 . It is quite challenging for existing NVCs but is the most fair setting for the comparison with traditional codec. It is noted that all coding tools and reference structure of traditional codecs use the best settings to represent their best compression ratio. Both ECM-5.0 and the recent ECM-11.0 are tested. For NVC, most existing models are only optimized for the RGB colorspace, and DCVC-DC [26] is the only NVC that has released a model for the YUV420 colorspace. The BD-Rate calculations are based on the weighted PSNR for the three color components, with weights of $(6, 1, 1) / 8$, consistent with the standard committee [45].

4.2. Comparisons with Previous SOTA Methods

RGB colorspace. Table 1 shows the performance comparison for RGB colorspace under 96 frames with intra-period 32. From this table, we can see that our DCVC-FM obtains the best compression ratio. Our DCVC-FM achieves an average of 20.3% bitrate saving over VTM, and also has significant advantage over ECM. In addition, our DCVC-FM is also better DCVC-DC as the bitrate saving of DCVC-DC is 16.1% over VTM. Under this setting, still outperforming DCVC-DC is not an trivial thing considering

Table 3. BD-Rate (%) comparison in YUV420 colorspace. 96 frames with intra-period = -1.

	UVG	MCL-JCV	HEVC B	HEVC C	HEVC D	HEVC E	Average
VTM-17.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HM-16.25	39.5	48.3	41.5	40.3	32.6	41.5	40.6
ECM-5.0	-13.3	-16.4	-14.6	-15.6	-14.1	-12.6	-14.4
DCVC-DC	-13.6	-7.9	-7.8	-5.6	-27.6	-10.3	-12.2
DCVC-FM	-25.4	-11.6	-17.1	-24.4	-41.5	-31.6	-25.3

Table 4. **BD-Rate (%) comparison in YUV420 colorspace. All frames with intra-period = -1.**

	UVG	MCL-JCV	HEVC B	HEVC C	HEVC D	HEVC E	Average
VTM-17.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
HM-16.25	40.1	48.6	47.6	41.0	34.5	42.8	42.4
ECM-5.0	-14.9	-17.0	-17.3	-16.6	-16.1	-14.1	-16.0
ECM-11.0	-20.0	-22.1	-22.2	-21.2	-20.4	-17.2	-20.5
DCVC-DC	5.7	-5.0	12.2	-4.2	-16.5	84.4	12.8
DCVC-FM	-20.7	-10.3	-18.2	-32.2	-41.2	-30.2	-25.5

that our codec has much wider quality range and support both RGB and YUV in single model when compared with DCVC-DC.

YUV420 colorspace. Table 2 shows the comparison for YUV420 colorspace under 96 frames with intra-period 32. As shown in this table, our DCVC-FM can surpass all traditional codecs and DCVC-DC. Our bitrate saving over VTM is 24.1% on average while that of DCVC-DC is 17.0%. Table 1 and Table 2 show the consistent improvement over DCVC-DC for both RGB and YUV420 colorspace.

However, we focus more on the evaluation using intra-period -1 setting. Table 3 shows the comparison under 96 frames with intra-period -1. From this table, we can find that our DCVC-FM can maintain the bitrate saving, i.e. 25.3%, over VTM, when compared with the number 24.1% under intra-period 32. By contrast, the bitrate saving of DCVC-DC over VTM decreases from 17.0% to 12.2% when the intra-period changes from 32 to -1.

But only testing 96 frames is not enough to evaluate the behavior of NVC when handling the long prediction chain. So we also test the all frames for each test video. Table 4 shows the corresponding comparison. From this table, we can see that our DCVC-FM still can outperform VTM by 25.5%. By contrast, DCVC-DC has large performance drop, i.e. 12.8% bitrate increase over VTM, when testing all frames. In particular, DCVC-DC has 84.8% bitrate increase for HEVC E dataset. By contrast, our codec can cope with this dataset well. In addition, if we use DCVC-DC as anchor in Table 4, our DCVC-FM can achieve an average of 29.7% bitrate saving for all datasets. As far as we know, our DCVC-FM is the first NVC that can achieve such high

compression ratio under intra-period -1 setting.

Fig. 7 shows the rate-distortion curves. From these curves, we can see that our DCVC-FM achieves much wider quality range than DCVC-DC. For example, the quality range of DCVC-DC on MCL-JCV dataset is 3.7 dB ([38.66, 42.32]), which cannot meet the requirement of higher compression ratio. By contrast, DCVC-FM is 11.4 dB ([31.40, 42.80]), which is much wider. Table 5 shows more comparisons and we can see the obvious quality range expansion on each test dataset. The quality range of the proposed model is similar to that encoded with VTM using QP 25~49.

4.3. Ablation Study

Table 6 shows the ablation study on each improvement. In this table, the baseline model M_a is DCVC-DC. We first test the structure optimization on reducing motion estimation module and using more depthwise convolution. This optimization (M_b) has 0.8% bitrate increase but can save the MACs by 16% reduction over DCVC-DC (M_a), as shown in Table 8. Table 6 also shows that, when further enabling wider quality range support (M_c), there is 3.4% bitrate increase. If also supporting single model for both RGB and YUV colorspace (M_d), the bitrate increase gets to 4.8% because these functionalities are not cost-free. Training with longer video can bring a large performance improvement, and M_e can achieve 15.5% bitrate saving, which demonstrates the benefits of utilizing longer temporal correlation. Based on M_e , refreshing the temporal feature (M_f) makes the bitrate saving increase to 29.7% as it can effectively alleviate the quality degradation problem.

Table 8 also shows the runtime comparison. Our MACs

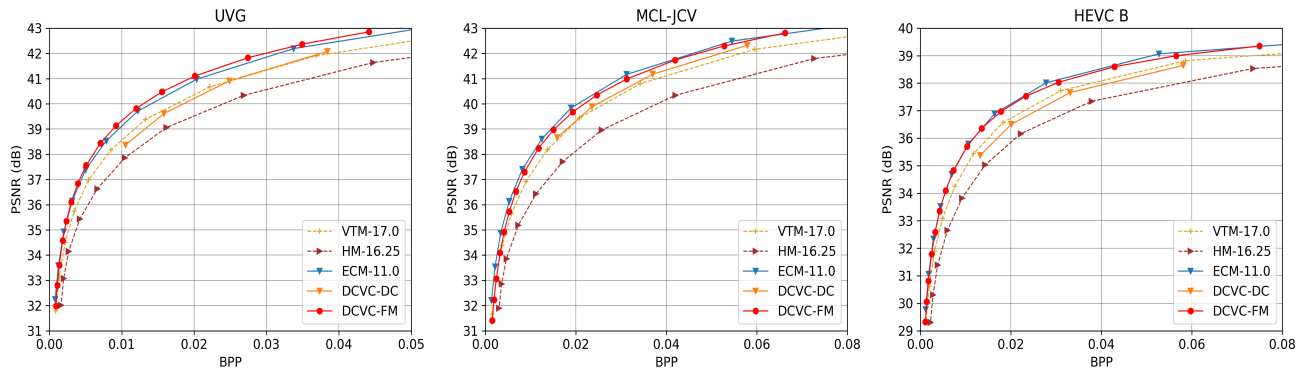


Figure 7. Rate and distortion curves for UVG , MCL-JCV, and HEVC B datasets. The comparison is in YUV420 colorspace. All frames with intra-period = -1. The curves of more datasets are in supplementary materials.

Table 5. The quality range (PSNR, dB) comparison in YUV420 colorspace. All frames with intra-period = -1.

	UVG	MCL-JCV	HEVC B	HEVC C	HEVC D	HEVC E	Average
DCVC-DC	3.7	3.7	3.3	4.2	4.6	3.6	3.8
DCVC-FM	10.9	11.4	10.0	12.4	12.5	11.3	11.4

Table 6. Ablation study using BD-Rate (%).

	M_a	M_b	M_c	M_d	M_e	M_f
Structure optimization		✓	✓	✓	✓	✓
Wider quality range support			✓	✓	✓	✓
Single model for RGB&YUV				✓	✓	✓
Training with longer video					✓	✓
Feature refresh						✓
BD-Rate(%)	0.0	0.8	3.4	4.8	-15.5	-29.7

have obvious reduction but the actual running time using 32-bit floating point inference is a little higher. This is because we use more depthwise convolution layers which have a lower computational density than the normal convolution, but can be further accelerated in the future [35]. When enabling 16-bit floating point inference with our optimized implementation for *grid_sample*, the running time has significant reduction. In particular, our NVC can save half of the memory usage if using 16-bit inference. Table 7 also shows that there is a significant 87.3% bitrate increase without our optimized implementation.

5. Conclusion and Limitation

In conclusion, this paper proposes feature modulation techniques and resolves two major challenges limiting the prac-

Table 7. BD-Rate (%) using 16-bit floating point (fp) inference.

fp32	fp16 w/o optimization	fp16 w/ optimization
0.0	87.3	0.9

Table 8. Complexity comparison.

	MACs	Encoding Time	Decoding Time
DCVC-DC w/ fp32	2642G	1005ms	765ms
DCVC-FM w/ fp32	2225G	1040ms	775ms
DCVC-FM w/ fp16	2225G	530ms	475ms

Note: Tested on NVIDIA 2080TI with using 1080p as input.

ticality of NVC. Through a uniform quantization parameter sampling mechanism to help the learning of quantization scaler, we can finely modulate the latent feature and enable wide quality range support. Meanwhile, the rate control capability is demonstrated. We also have addressed issues with long prediction chains by modulating the temporal feature with a periodically refreshing mechanism. In addition, our DCVC-FM now supports both RGB and YUV colorspace, and allows for low-precision inference. Our DCVC-FM represents a significant step in the evolution of NVC technology.

However, although DCVC-FM enables low-precision float point inference, its speed is still far from real-time. In addition, the float point inference has cross-platform issue for the entropy coding in NVC. In the future, we will investigate these topics and build a more powerful NVC which can be widely deployed in practical products.

References

- [1] ECM. <https://vcgit.hhi.fraunhofer.de/ecm/ECM/>. 6
- [2] HM. <https://vcgit.hhi.fraunhofer.de/jvet/HM/>. 6
- [3] VTm. https://vcgit.hhi.fraunhofer.de/jvet/VVCSoftware_VTM/. 6
- [4] Original vimeo links. https://github.com/anchen1011/toflow/blob/master/data/original_vimeo_links.txt. 6
- [5] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8503–8512, 2020. 1, 2
- [6] João Ascenso, Elena Alshina, and Touradj Ebrahimi. The JPEG AI Standard: Providing Efficient Human and Machine Visual Data Consumption. *IEEE Multimedia*, 30(1):100–111, 2023. 2
- [7] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *6th International Conference on Learning Representations, ICLR*, 2018. 2
- [8] Gisle Bjontegaard. Calculation of average PSNR differences between RD-curves. *VCEG-M33*, 2001. 6
- [9] Frank Bossen et al. Common test conditions and software reference configurations. In *JCTVC-L1100*, 2013. 2, 6
- [10] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. 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. 1
- [11] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4947–4956, 2021. 3
- [12] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 5
- [13] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers. Neural inter-frame compression for video coding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2
- [14] Noor Fathima Goose, Jens Petersen, Auke Wiggers, Tianlin Xu, and Guillaume Sautiere. Neural image compression with a diffusion-based decoder. *arXiv preprint arXiv:2301.05489*, 2023. 2
- [15] Yung-Han Ho, Chih-Peng Chang, Peng-Yu Chen, Alessandro Gnutti, and Wen-Hsiao Peng. Canf-vc: Conditional augmented normalizing flows for video compression. *European Conference on Computer Vision*, 2022. 2, 6
- [16] 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
- [17] 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 (CVPR)*, pages 1502–1511, 2021. 2
- [18] Zhihao Hu, Guo Lu, Jinyang Guo, Shan Liu, Wei Jiang, and Dong Xu. Coarse-to-fine deep video coding with hyperprior-guided mode prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5921–5930, 2022. 1, 2
- [19] Cong Huang, Jiahao Li, Bin Li, Dong Liu, and Yan Lu. Neural compression-based feature learning for video restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5872–5881, 2022. 4
- [20] Cong Huang, Jiahao Li, Lei Chu, Dong Liu, and Yan Lu. Disentangle propagation and restoration for efficient video recovery. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8336–8345, 2023. 3
- [21] A Burakhan Koyuncu, Han Gao, and Eckehard Steinbach. Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. *arXiv preprint arXiv:2203.02452*, 2022. 2
- [22] Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, and Olivier Déforges. Optical flow and mode selection for learning-based video coding. In *22nd IEEE International Workshop on Multimedia Signal Processing*, 2020. 1
- [23] Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, and Olivier Déforges. Conditional coding for flexible learned video compression. In *Neural Compression: From Information Theory to Applications – Workshop @ ICLR*, 2021. 2
- [24] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 3, 6
- [25] Jiahao Li, Bin Li, and Yan Lu. Hybrid spatial-temporal entropy modelling for neural video compression. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1503–1511, 2022. 2, 3, 4, 6
- [26] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with diverse contexts. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, Canada, June 18-22, 2023*, 2023. 1, 2, 3, 6
- [27] Jianping Lin, Dong Liu, Houqiang Li, and Feng Wu. M-LVC: multiple frames prediction for learned video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2
- [28] Bowen Liu, Yu Chen, Rakesh Chowdary Machineni, Shiyu Liu, and Hun-Seok Kim. Mmvc: Learned multi-mode video compression with block-based prediction mode selection and density-adaptive entropy coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18487–18496, 2023. 1, 2
- [29] Haojie Liu, Ming Lu, Zhan Ma, Fan Wang, Zhihuang Xie, Xun Cao, and Yao Wang. Neural video coding using

- multiscale motion compensation and spatiotemporal context model. *IEEE Transactions on Circuits and Systems for Video Technology*, 2020. 2
- [30] Jerry Liu, Shenlong Wang, Wei-Chiu Ma, Meet Shah, Rui Hu, Pranaab Dhawan, and Raquel Urtasun. Conditional entropy coding for efficient video compression. In *European Conference on Computer Vision*, pages 453–468. Springer, 2020. 2, 3
- [31] Jinming Liu, Heming Sun, and Jiro Katto. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14388–14397, 2023. 2
- [32] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11015, 2019. 1, 2
- [33] 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.
- [34] 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*, 2020. 2
- [35] Gangzhao Lu, Weizhe Zhang, and Zheng Wang. Optimizing depthwise separable convolution operations on gpus. *IEEE Transactions on Parallel and Distributed Systems*, 33(1):70–87, 2021. 8
- [36] Wufei Ma, Jiahao Li, Bin Li, and Yan Lu. Uncertainty-Aware Deep Video Compression with Ensembles. *IEEE Transactions on Multimedia*, 2024. 2
- [37] Fabian Mentzer, George Toderici, David Minnen, Sung-Jin Hwang, Sergi Caelles, Mario Lucic, and Eirikur Agustsson. Vct: A video compression transformer. *arXiv preprint arXiv:2206.07307*, 2022. 2
- [38] Alexandre Mercat, Marko Viitanen, and Jarno Vanne. UVG dataset: 50/120fps 4k sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 297–302, 2020. 6
- [39] Zhihong Pan, Xin Zhou, and Hao Tian. Extreme generative image compression by learning text embedding from diffusion models. *arXiv preprint arXiv:2211.07793*, 2022. 2
- [40] Linfeng Qi, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Motion information propagation for neural video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6111–6120, 2023. 2, 3
- [41] Yichen Qian, Ming Lin, Xiuyu Sun, Zhiyu Tan, and Rong Jin. Entroformer: A transformer-based entropy model for learned image compression. *arXiv preprint arXiv:2202.05492*, 2022. 2
- [42] 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/CVF International Conference on Computer Vision (ICCV)*, pages 14479–14488, 2021. 2
- [43] Vadim Seregin, Jie Chen, Roman Chernyak, Fabrice Leanne, and Kai Zhang. JVET AHG report: ECM software development (AHG6). In *JVET-AF0006*, 2023. 1
- [44] Xihua Sheng, Jiahao Li, Bin Li, Li Li, Dong Liu, and Yan Lu. Temporal Context Mining for Learned Video Compression. *IEEE Transactions on Multimedia*, 2022. 2, 3, 6
- [45] Gary J Sullivan and Jens-Rainer Ohm. Meeting report of the fourth meeting of the joint collaborative team on video coding (jct-vc), daegu, kr, 20–28 january 2011. *Document JCTVC-D500, Daegu, KR*, 2011. 6
- [46] Lucas Theis, Tim Salimans, Matthew D Hoffman, and Fabian Mentzer. Lossy compression with gaussian diffusion. *arXiv preprint arXiv:2206.08889*, 2022. 2
- [47] Guo-Hua Wang, Jiahao Li, Bin Li, and Yan Lu. EVC: Towards Real-Time Neural Image Compression with Mask Decay. In *International Conference on Learning Representations*, 2023. 2
- [48] 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
- [49] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019. 6
- [50] Ruihan Yang and Stephan Mandt. Lossy image compression with conditional diffusion models. *arXiv preprint arXiv:2209.06950*, 2022. 2
- [51] Yibo Yang and Stephan Mandt. Computationally-efficient neural image compression with shallow decoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 530–540, 2023. 2
- [52] Jing Zhao, Bin Li, Jiahao Li, Ruiqin Xiong, and Yan Lu. A universal encoder rate distortion optimization framework for learned compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1880–1884, 2021. 2
- [53] Jing Zhao, Bin Li, Jiahao Li, Ruiqin Xiong, and Yan Lu. A universal optimization framework for learning-based image codec. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(1):1–19, 2023. 2
- [54] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17492–17501, 2022. 2