B-CANF: Adaptive B-frame Coding with Conditional Augmented Normalizing Flows

Mu-Jung Chen, Yi-Hsin Chen, and Wen-Hsiao Peng, *Senior Member, IEEE*

Abstract—Over the past few years, learning-based video compression has become an active research area. However, most works focus on P-frame coding. Learned B-frame coding is under-explored and more challenging. This work introduces a novel B-frame coding framework, termed B-CANF, that exploits conditional augmented normalizing flows for B-frame coding. B-CANF additionally features two novel elements: frame-type adaptive coding and B*-frames. Our frame-type adaptive coding learns better bit allocation for hierarchical B-frame coding by dynamically adapting the feature distributions according to the B-frame type. Our B*-frames allow greater flexibility in specifying the group-of-pictures (GOP) structure by reusing the B-frame codec to mimic P-frame coding, without the need for an additional, separate P-frame codec. On commonly used datasets, B-CANF achieves the state-of-the-art compression performance as compared to the other learned B-frame codecs and shows comparable BD-rate results to HM-16.23 under the random access configuration in terms of PSNR. When evaluated on different GOP structures, our B*-frames achieve similar performance to the additional use of a separate P-frame codec.

Index Terms—Neural video coding, conditional coding, B-frame coding.

I. INTRODUCTION

THE great success of learned image compression has spurred a new wave of research and development for learned video compression. Most existing methods [1]–[7] address low-delay P-frame coding based on hybrid-based coding architecture, which comprises uni-directional temporal prediction followed by deep residual coding. Ladune *et al.* [8] recently show that residual coding is sub-optimal from the information-theoretic perspective, proposing to condition the variational autoencoder (VAE)-based deep compression on motion-compensated frames [8], [9] with the aim of reaching the lower conditional entropy.

In contrast to the rapid progress in learned P-frame coding, learned B-frame coding, which allows bidirectional referencing for higher coding efficiency, is an under-explored topic. The coding of B-frames usually involves frame interpolation with explicit motion modeling and coding. Residual or conditional coding may then follow to achieve inter-frame prediction. Table I dissects some notable B-frame coding schemes in terms of their strategies for motion and inter-frame coding, to give an overview of recent developments in

this area. As shown, most of the existing approaches [10]–[13] adopt residual-based inter-frame coding, while many still use intra motion coding (e.g. encoding optical flow maps as individual images). Notably, the first B-frame coding scheme [14] already adopts conditional inter-frame coding, although its motion coding is still an intra-based approach. Built on top of [8], the more recent VAE-based conditional coding framework [9] features both conditional motion and inter-frame coding. However, its compression performance is inferior to the two state-of-the-art methods [12], [13], both adopt residual-based inter-frame coding. We argue that the full potential of conditional coding is yet to be seen. How it is implemented can make a big difference in coding performance.

Inspired by our previous work [15], we propose a novel learned B-frame coding framework, termed B-CANF, that exploits conditional augmented normalizing flows (CANF) for B-frame coding. We choose CANF as the coding backbone because of its *generality* and *reversibility*. It is shown in [15] that CANF is able to achieve greater expressiveness by stacking multiple conditional VAE encoders and decoders.

B-CANF extends [15] and differs from most prior works (e.g. [1]–[7], [16]–[18]) in several significant ways:

- B-CANF addresses primarily B-frame coding rather than P-frame coding. Similar to [15], it adopts conditional motion and inter-frame coding. However, the way in which the motion and inter-frame predictors are formulated has been adapted to B-frame coding.
- B-CANF newly introduces frame-type adaptive coding that learns better bit allocation for hierarchical B-frame coding.
- B-CANF newly proposes a special type of B-frame, called B*-frame, to mimic P-frame coding. This tool feature allows greater flexibility in specifying the group-of-pictures (GOP) structure, without the need for an additional, separate P-frame codec.

The non-trivial application of these novel elements—namely, CANF, frame-adaptive coding, and B*-frame—to B-frame coding achieves higher compression performance. To our best knowledge, B-CANF is also the first learned B-frame coding scheme that shows comparable PSNR results to HM-16.23 under the random access configuration with intra-period 32 and GOP size 16. Furthermore, this work presents a comprehensive study on B-frame coding, the results of which are not reported in our previous work [15] for P-frame coding. These facts distinct this work from [15] and highlight its contributions to advancing B-frame coding.

TABLE I
DISSECTION OF SOME NOTABLE B-FRAME CODING SCHEMES.

	Publication	Motion Coding for B-frames	Inter-frame Coding for B-frames	Modules with Shared Param.
[14]	ECCV'18	VAE-based Intra	Cond. VAE-based	No
[10]	ICCV'19	VAE-based Intra	VAE-based Residual	I, B
HLVC [11]	CVPR'20	VAE-based Intra & DIRECT	VAE-based Residual	No
[9]	ICLRW'21	One-stage, Cond. VAE-based	Cond. VAE-based	I, P, B
B-EPIC [12]	ICCV'21	VAE-based Intra	VAE-based Residual	B, P
LHBDC [13]	TIP'22	VAE-based Residual	VAE-based Residual	No
B-CANF (Ours)	-	Cond. ANF-based	Cond. ANF-based	B, B*

It is to be noted that this work is an expanded version of our conference publication [19], which adopts CANF for B-frame coding but focuses specifically on YUV 4:2:0 content, as opposed to RGB content in this work. Moreover, this work includes extensive ablation studies and complexity analyses, which are not covered in [19].

The remainder of this paper is organized as follows: Section II reviews learned video compression and the basics of ANF-based image compression (ANFIC [20]). Section III elaborates the design of B-CANF. Section IV compares B-CANF with the state-of-the-art methods and presents ablation experimental results. Finally, we provide concluding remarks in Section V.

II. RELATED WORK

A. Learned P-frame Coding

End-to-end learned video compression has recently attracted lots of attention. Most prior works [1], [3]–[7], [11], [16], [21]–[25] focus on low-delay, P-frame coding, where a target frame is coded based on information propagated from the past decoded frames. In common, they share a similar temporal predictive coding architecture to the conventional codecs [26], [27]. As such, improving temporal prediction is one of the central research themes. In this aspect, there have been schemes such as motion-compensation networks [1], scale-space warping [3], one-stage motion estimation [7], feature-domain warping [21], multi-hypothesis prediction [4], compound spatiotemporal representation [17], and recurrent autoencoding [22], [23]. Because temporal prediction often relies on explicit motion modeling, some research efforts are dedicated to reducing motion overhead by predictive motion coding [4], incremental flow map coding [6], resolution-adaptive flow coding [5], [25] and coarse-to-fine motion coding [25]. Other notable P-frame coding techniques include the temporal prior [7], [22] that leverages a recurrent neural network to propagate causal, temporal information for entropy coding, adaptive residual skip coding [24], [25], and transformer-based autoencoders [28].

More recently, conditional coding [8], [15], [29], [30] achieves a breakthrough in inter-frame coding. It conditions the inter-frame autoencoder on the motion-compensated frame in forming a non-linear prediction of the target frame, as opposed to subtracting the motion-compensated frame from the target frame for residual coding, which is shown to be sub-optimal from the information-theoretic perspective [8]. Lately,

Mentzer *et al.* [31] take an interesting approach to conditional coding, utilizing transformers [32] to model the dependencies between the latents of video frames for entropy coding. It has the striking feature of not requiring motion coding and warping operations. Although its initial idea dates back to [14], research on conditional coding remains very active.

B. Learned B-frame Coding

In comparison with learned P-frame coding, learned B-frame coding [9]–[14], which is the main focus of this paper and targets higher coding efficiency by allowing temporal prediction from both the future and past decoded frames, is relatively under-explored. Its process normally involves frame interpolation with explicit motion modeling and coding, followed by inter-frame residual coding [10]–[13] or conditional coding [9], [14]. Due to the needs for bi-directional temporal prediction, B-frame coding usually incurs more motion overhead than P-frame coding. VAE-based intra motion coding [10], [11], [14], which encodes (bi-directional) optical flow maps as individual images, is very common. To minimize motion overhead, several strategies are proposed. Yang *et al.* [11] reuse the motion of a future or a past reference frame in a way similar to the DIRECT mode in [33]. Yılmaz *et al.* [13] adopt motion residual coding, with the predicted flow maps derived from the two reference frames. Ladune *et al.* [9] introduce one-stage, conditional motion coding, where the flow maps are estimated directly from and coded conditionally on the two reference frames without the use of optical flow estimation networks. Pourreza *et al.* [12] reuse the P-frame codec for B-frame coding, sending for each B-frame only one flow map that characterizes the motion between the target B-frame and its predicted frame interpolated by a pre-trained network. Efforts have also been made to share networks for coding different types of frames, e.g. I-frame, P-frame and B-frame [9], [12].

Table I contrasts the differences between our B-CANF and some notable learned B-frame coding schemes. Extending from our previous work [15] for P-frame coding, this paper introduces a conditional coding framework that exploits CANF for B-frame coding. As opposed to the other conditional coding schemes [9], [14], our method applies conditional coding to both motion and inter-frame coding. Moreover, it features frame-type adaptive coding, which adapts the coding process to the B-frame type. We also introduce a special type of B-frame, called B*-frame, which mimics P-frame coding by reusing the B-frame codec.
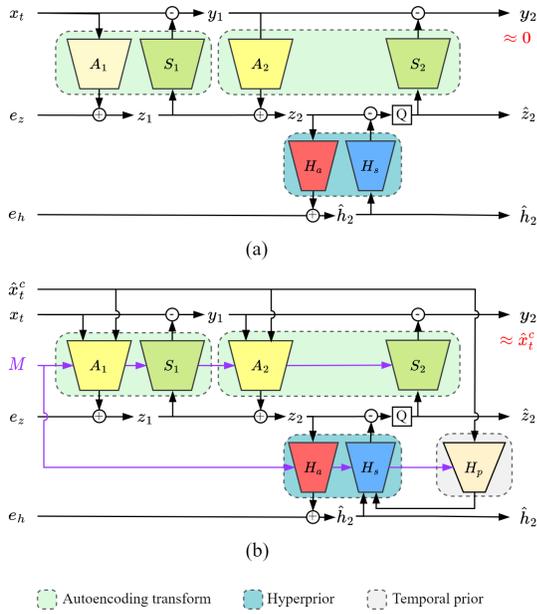
Fig. 1. Illustration of (a) the *unconditional* ANF for image coding in ANFIC [20] and (b) the proposed *conditional* ANF for B-frame coding, where M signals the B-frame type for frame-type adaptive coding.

C. Augmented Normalizing Flows for Image Compression

B-CANF and our previous work [15] are a sequel to the augmented normalized flow (ANF)-based image compression, also known as ANFIC [20], which adopts ANF [34] as the compression backbone. To ease the understanding of our B-CANF, this section reviews the basics of ANFIC [20]. Fig. 1(a) illustrates its coding architecture. Shown on the left are the three inputs of ANFIC [20]: the image to be encoded x and the two augmented inputs e_z, e_h; accordingly, the compressed outputs are a nearly zero image $y_2 \approx 0$, the quantized latent code $\hat{z}_2$, and the hyperprior $\hat{h}_2$.

Encoding: The encoding process proceeds from left to right. The input pair (x, e_z) with $e_z = 0$[1] is transformed into the output pair $(y_2, \hat{z}_2)$, where y_2 is regularized to approximate a zero image, by two autoencoding transforms (A_1, S_1) and (A_2, S_2) arranged in the form of additive coupling layers. A_1, A_2 are VAE-based analysis transforms, while S_1, S_2 are VAE-based synthesis transforms. In symbols, the intermediate states y_1, z_1 after the first autoencoding transform (A_1, S_1) are given by

$$A_1(x_t, e_z) = (x_t, e_z + \mu_{A_1}(x_t)) = (x_t, z_1), \quad (1)$$
$$S_1(x_t, z_1) = (x_t - \mu_{S_1}(z_1), z_1) = (y_1, z_1), \quad (2)$$

where μ_{A_1}, μ_{S_1} are two neural networks that output element-wise additive transform parameters. The second autoencoding transform (A_2, S_2) that converts y_1, z_1 into $y_2, \hat{z}_2$ follows a similar definition, except that z_2 is additionally predicted

---

[1]ANF [34] is initially designed to be a generative model, where the augmented noises are meant to induce a complex marginal on the input x_t and the entropy rate of the latents is not a concern. In ANFIC [20], the latents z_2 must be quantized and compressed to achieve low entropy. Injecting noise at e_z will increase the entropy of z_2. ANFIC [20] shows that having $e_z = 0$ can still achieve good compression performance.

TABLE II
COMPARISON OF B-CANF, ANFIC [20], AND CANF-VC [15].

	Codec Type	Frame-type Support	Frame-type Adaptive Coding
ANFIC [20]	ANF-based	I	N
CANF-VC [15]	CANF-based	I, P	N
B-CANF (Ours)	CANF-based	I, B, B*	Y

from the hyperprior and quantized as $\hat{z}_2$. Here H_a, H_s serve as the hyperprior autoencoder [35], which encodes the latent code z_2 into $\hat{h}_2$ for entropy coding. The reader is referred to [20] for network details. By doing so, ANFIC transforms the information from x_t into the augmented space $\hat{z}_2, \hat{h}_2$, which have much lower resolution than the original image and allow for more efficient coding.

Decoding: The decoding proceeds in reverse order by decoding first the hyperprior $\hat{h}_2$, followed by the decoding of $\hat{z}_2$ to update y_2 (initialized as 0) successively from right to left. During training, $e_h \sim \mathcal{U}(-0.5, 0.5)$ simulates the additive quantization noise of the hyperprior. Remarkably, the conventional VAE-based compression is a special case of ANFIC by retaining only the autoencoding transform (A_2, S_2). In this sense, ANFIC is able to achieve superior expressiveness to VAE-based compression systems by cascading more autoencoding transforms.

Table II summarizes the differences of ANFIC [20], our previous work CANF-VC [15], and B-CANF. Both CANF-VC [15] and B-CANF adopt CANF as the coding backbone, whereas ANFIC uses ANF. CANF differs from ANF in that it learns *conditional* distributions of video frames rather than *unconditional* distributions of images (Section III-B). Moreover, B-CANF distinguishes from CANF-VC [15] by the support of B-frames and B*-frames (Section III-C); additionally, it incorporates frame-type adaptive coding (Section III-D). Extensive experimental results and ablation studies on B-frame coding reported in this paper are not seen in [15].

III. PROPOSED METHOD

A. System Overview

We begin this section with an overview of our B-frame coding framework. Specifically, we target the two common hierarchical Group-of-Pictures (GOP) prediction structures shown in Fig. 2, where Fig. 2(a) depicts the case with an intra-period of the same size as the GOP and Fig. 2(b) the other case having multiple GOPs in one intra-period. In the latter case, we introduce a special type of B-frame, known as B*-frame, the coding of which reuses the same framework as for regular B-frames (including both reference and non-reference ones). Our B*-frames are conceptually similar to *generalized* B-frames [36], which allow the two reference frames to be both from the past (and even the same) decoded frames. The use of B*-frame is to mimic P-frame coding with our B-frame coding framework, an approach that is in direct contrast to [12], which uses P-frame coding for B-frames. We remark that B-frame coding is more flexible and general than P-frame coding.
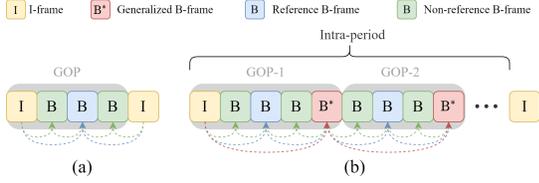
Fig. 2. Illustration of the GOP structures supported in our scheme. The GOP size is set to 4 for illustrative purposes. (a) Hierarchical B-frame coding. (b) Hierarchical B-frame coding with multiple GOPs in an intra-period. In both cases, reference and non-reference B-frames are coded recursively.

Fig. 3(a) illustrates our B-frame coding framework. Consider the coding of a B-frame x_t, where the subscript indicates the time index, given its two reference frames $\hat{x}_{t-k}, \hat{x}_{t+k}$, i.e., the previous and subsequent decoded reference frames at higher hierarchical levels, respectively. Our first step is to interpolate a predicted frame $\hat{x}_t^c$ through bidirectional motion compensation in pixel domain. To this end, we estimate and signal two optical flow maps $m_{t\to t-k}^e, m_{t\to t+k}^e$ for bidirectional backward warping. Our main focus here is on reducing motion overhead through conditional coding. To obtain the conditioning signals, we introduce a motion prediction network that derives the (predicted) optical flow maps $m_{t\to t-k}^p, m_{t\to t+k}^p$ from observing the two reference frames $\hat{x}_{t-k}, \hat{x}_{t+k}$ without having access to x_t. Our second step is to code the B-frame x_t efficiently by conditioning the inter-frame coding on the predicted frame $\hat{x}_t^c$. Our framework is comprised of the following major components:

- **Motion estimation network (ME-Net)** estimates the backward flow maps $m_{t\to t-k}^e, m_{t\to t+k}^e$ between x_t and the two reference frames $\hat{x}_{t-k}, \hat{x}_{t+k}$, respectively. In this work, we adopt SPyNet [37] for motion estimation.
- **Motion prediction network** predicts two optical flow maps $m_{t\to t-k}^p, m_{t\to t+k}^p$, which serve jointly as the condition for coding $m_{t\to t-k}^e, m_{t\to t+k}^e$, from observing the two reference frames $\hat{x}_{t-k}, \hat{x}_{t+k}$ without having access to x_t. We follow [38] in using a convolutional neural network to generate $m_{t\to t-k}^p, m_{t\to t+k}^p$ in a coarse-to-fine fashion.
- **CANF-based motion codec**, conditioned on the predicted flow maps $m_{t\to t-k}^p, m_{t\to t+k}^p$, encodes the flow maps $m_{t\to t-k}^e, m_{t\to t+k}^e$ jointly to output the reconstructed flow maps $\hat{m}_{t\to t-k}, \hat{m}_{t\to t+k}$ (Section III-B).
- **Frame synthesis network** synthesizes the predicted frame $\hat{x}_t^c$ by GridNet [39], which takes as inputs the motion-compensated reference frames and their features generated based on the reconstructed flows $\hat{m}_{t\to t-k}, \hat{m}_{t\to t+k}$.
- **CANF-based inter-frame codec** encodes the target frame x_t, conditioned on the predicted frame $\hat{x}_t^c$ (Section III-B).

B. CANF-based Inter-frame and Motion Codecs

Fig. 1(b) presents the architecture of our CANF-based codec. Both our inter-frame and motion codecs adopt a similar design. To keep the notation uncluttered, we take the inter-frame codec as an illustrative example.

CANF-based inter-frame codec seeks to learn the conditional distribution $p(x_t|\hat{x}_t^c)$ of the B-frame x_t, given the bidirectionally predicted frame $\hat{x}_t^c$. As compared to ANFIC [20], which employs an unconditional ANF to learn the image distribution $p(x_t)$ for image compression, our CANF-based inter-frame codec introduces two novel elements: (1) conditional autoencoding transforms and (2) conditional entropy coding with the combined hyperprior and temporal prior.

From Fig. 1(b), the conditional autoencoding transform is designed to transform the B-frame x_t into the latent y_2, which in the present case is regularized to approximate the predicted frame $\hat{x}_t^c$ instead of a zero image as with ANFIC [20] (cp. Fig. 1(a)). In the course of transformation, the latent $\hat{z}_2$ encodes the information necessary for instructing the transformation from x_t to $\hat{x}_t^c$ and vice versa, while the hyperprior $\hat{h}_2$ is utilized to model the distribution of $\hat{z}_2$ for entropy coding. Taking $\{A_1, S_1\}$ in Fig. 1(b) as an example, we have

$$A_1(x_t, e_z|\hat{x}_t^c) = (x_t, e_z + \mu_{A_1}(x_t, \hat{x}_t^c)) = (x_t, z_1), \quad (3)$$
$$S_1(x_t, z_1) = (x_t - \mu_{S_1}(z_1), z_1) = (y_1, z_1), \quad (4)$$

where μ_{A_1}, μ_{S_1} are two neural networks that output elementwise additive transform parameters. It is worth noting that the encoding process μ_{A_1} is conditioned on $\hat{x}_t^c$ by concatenating $\hat{x}_t^c, x_t$ to form the encoder input. The rationale behind the design is to ease the transformation from x_t to $\hat{x}_t^c$ by supplying the target $\hat{x}_t^c$ as an auxiliary signal. This process is repeated by taking the resulting y_1, z_1 as inputs to the next autoencoding transform $\{A_2, S_2\}$.

The conditional entropy coding is achieved by the hierarchical autoencoding transform $\{H_a, H_s\}$, which estimates the conditional distribution of z_2 given the hyperprior $\hat{h}_2$ and the interpolated frame $\hat{x}_t^c$ for entropy coding. Its operation is governed by

$$H_a(z_2, e_h) = (z_2, e_h + \mu_{H_a}(z_2)) = (z_2, \hat{h}_2), \quad (5)$$
$$H_s(z_2, \hat{h}_2|\hat{x}_t^c) = (\lfloor z_2 - \mu_{H_s}(\hat{h}_2, H_p(\hat{x}_t^c)) \rceil, \hat{h}_2) = (\hat{z}_2, \hat{h}_2), \quad (6)$$

where $\lfloor \cdot \rceil$ (depicted as Q in Fig. 1(b)) denotes the nearest-integer rounding at inference time for coding z_2 in a lossy manner. In particular, $\hat{h}_2 = e_h + \mu_{H_a}(z_2)$ denotes the quantized hyperprior and is entropy coded by a learned factorized distribution [40], with $e_h \sim \mathcal{U}(-0.5, 0.5)$ simulating additive quantization noise applied to the hyperprior latent $\mu_{H_a}(z_2)$. $\mu_{H_s}(\hat{h}_2, H_p(\hat{x}_t^c))$ models the mean of the latent z_2, which is assumed to follow a Gaussian distribution with the standard deviation determined by another output $\sigma(\hat{h}_2, H_p(\hat{x}_t^c))$ tied to the same backbone network as μ_{H_s}. We remark that the combination of $\hat{h}_2, \hat{x}_t^c$ in deriving the coding probabilities for z_2 exerts a combined effect of the hyperprior and temporal prior.

In Fig. 1(b), another conditioning factor M is utilized to signal the frame type of x_t–namely, reference B-frames, non-reference B-frames, or B*-frames–and to adapt feature distributions for frame-type adaptive coding (Section III-D).

CANF-based motion codec follows a similar design to our CANF-based inter-frame codec. The changes include (1) replacing the coding B-frame x_t with the concatenation of the
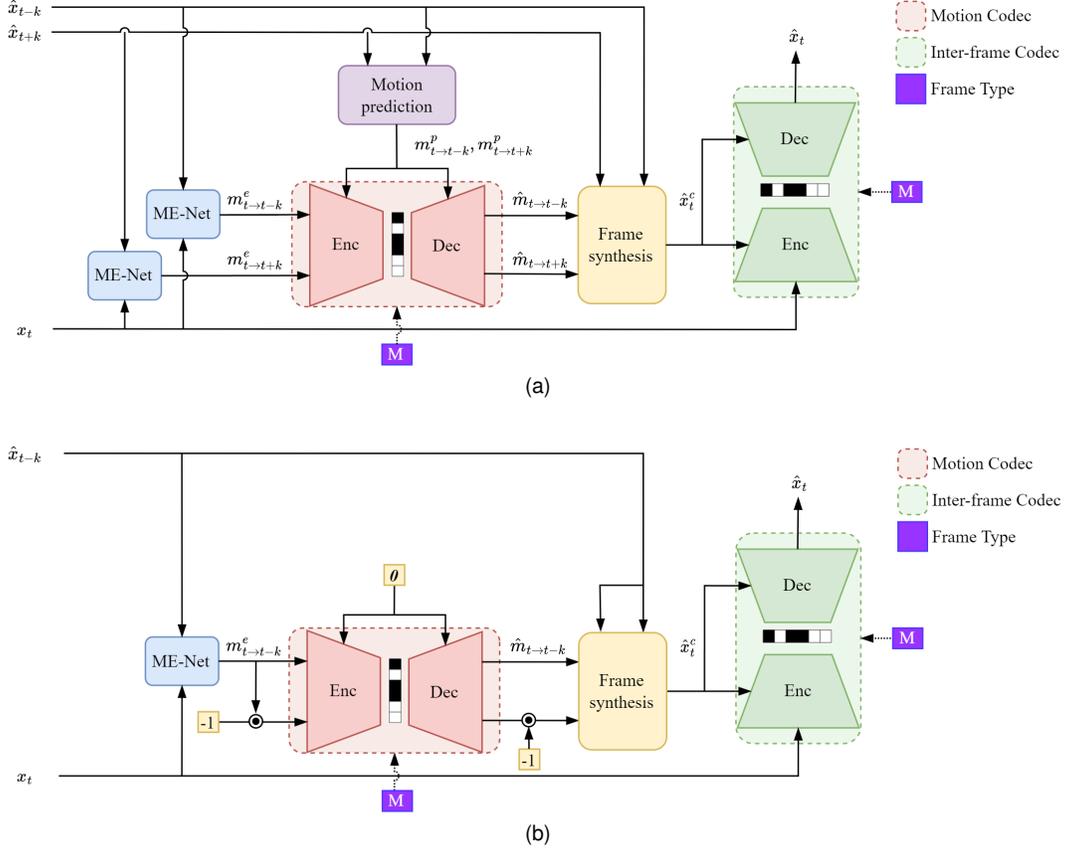
Fig. 3. Illustration of the proposed B-frame codec configured for coding (a) B-frames (reference and non-reference) and (b) B*-frames. x_t represents the current coding frame and $\hat{x}_{t-k}, \hat{x}_{t+k}$ are the previously reconstructed reference frames. ME-Net estimates the optical flow maps $m^e_{t \to t-k}, m^e_{t \to t+k}$ between x_t and its reference frames $\hat{x}_{t-k}, \hat{x}_{t+k}$, respectively. The motion prediction network outputs the predicted optical flow maps $m^p_{t \to t-k}, m^p_{t \to t+k}$, which serve as the conditioning signals for the motion codec. The frame synthesis network fuses the reference frames using the reconstructed optical flow maps $\hat{m}_{t \to t-k}, \hat{m}_{t \to t+k}$ to generate the predicted frame $\hat{x}^c_t$, which acts as the conditioning signal for the inter-frame codec. M indicates the frame type (reference B, non-reference B, B*-frame).

two optical flow maps $\{m^e_{t \to t-k}, m^e_{t \to t+k}\}$ to be coded, and (2) replacing the interpolated frame $\hat{x}^c_t$ with the concatenation of the two predicted optical flow maps $\{m^p_{t \to t-k}, m^p_{t \to t+k}\}$. It is important to note that $\{m^e_{t \to t-k}, m^e_{t \to t+k}\}$ are coded jointly based on the predicted flows $\{m^p_{t \to t-k}, m^p_{t \to t+k}\}$, allowing the motion codec to exploit correlations between these flow maps for reducing motion overhead.

The major differences between the CANF-based codecs of our B-CANF and [15] are summarized as follows:

- **Bidirectional joint motion coding:** B-CANF addresses B-frame coding. There are two optical flow maps (rather than one in [15]) to be coded. Instead of encoding these flow maps independently, they are concatenated as the input to our CANF-based motion codec for joint coding. Likewise, the motion prediction module outputs two flow map predictors. They are also concatenated as the joint conditioning signal. This allows our CANF-based motion codec to best use the contextual information from both flow map predictors for joint coding. As a result, the channels of the input, conditional input, and output are twice as many as those of their counterparts in [15].

- **Frame-type adaptive coding:** B-CANF additionally incorporates a frame-type adaption module in every con-

volutional layer of the CANF-based motion and inter-frame codecs, in order to adapt the feature distributions for frame-type adaptive coding. We will discuss this in Section III-D.

- **Network configuration:** The network settings, such as the number of input/output channels, kernel size, and stride, for the CANF-based motion and inter-frame codecs have been changed in this work, in order to strike a better balance among coding efficiency, model size, and computational complexity. More details will be provided in the released code.

C. B*-frame Extension

B*-frames allow to support multiple GOPs in an intra-period as depicted in Fig. 2(b), to strike a balance between coding performance and latency. A straightforward approach is to encode P-frames periodically using an additional P-frame codec. In doing so, however, the total model size is increased. To address this issue, we propose B*-frames, which reuse our CANF-based B-frame codec to mimic P-frame coding.

To encode a video frame x_t in B* mode, we create a predicted frame from the past decoded frame x_{t-k} just like coding a vanilla P-frame. To this end, the flow map $m^e_{t \to t-k}$

between x_t and x_{t-k} is estimated and coded. As shown in Fig. 3(b), the main inputs to our conditional motion codec should ideally comprise two optical flow maps. We generate the other "virtual" input by flipping the sign of $m_{t \to t-k}^e$ (i.e. $-1 \times m_{t \to t-k}^e$), in order to match the characteristics of $m_{t \to t+k}^e$ (for ordinary B-frame coding). Because only one reference frame x_{t-k} is used, we disable motion prediction network and set the predicted flow maps to constant 0. This reduces our motion coding scheme for B*-frames to intra-based motion coding. Although it is possible to use a motion extrapolation network to generate a flow map predictor from previously coded B*-frames, this calls for long-term motion extrapolation. It is to be noted that B*-frames are usually temporally distant from each other or may even be coded only once in an intra-period as with the HM randomaccess configuration [41]. The predicted frame $\hat{x}_t^c$ is then synthesized from the reference frames based on the coded $m_{t \to t-k}^e$ and its virtual counterpart (after sign reversal), which are found empirically to be similar but not exactly the same (see ablation experiments in Section IV-D). As such, the B*-frame is by definition a bi-prediction frame.

Our ablation experiments in Section IV-D show that B*-frames have similar compression performance to P-frames, which require an additional, separate P-frame codec.

D. Frame-type Adaptive Coding

Frame-type adaptive coding aims to achieve adaptive coding according to the reference types of B-frames. In traditional codecs, the *reference* B-frames are usually coded at higher quality than the *non-reference* B-frames by operating the same B-frame codec in different modes[2]. Following a similar strategy, we weight more heavily the distortions of the reference B-frames and B*-frames during training (Section III-E). In addition, we observe that using a fixed, shared B-frame codec for coding all the types of B-frames is unable to achieve frame-type adaptive coding. On the other hand, training separate models for different B-frames as with [14] is prohibitively expensive.

We tackle this issue by introducing a frame-type adaptation (FA) module to every convolutional layer, aiming to share the convolutional layers between B-frames while adapting their features to the B-frame type. Specifically, according to the B-frame type (namely, reference B-frames, non-reference B-frames, or B*-frames), our FA module applies a channel-wise affine transformation to the output features of every convolutional layer. Take as an example the analysis transform in Fig. 4. The FA module is inserted between the convolutional layer and the generalized divisive normalization (GDN) layer. It takes B-frame type M (in one-hot vector) as input and generates the affine transform parameters γ and β for feature F adaptation by

$$\mathrm{FA}(F|M) = \gamma(M) \odot F \oplus \beta(M) \qquad (7)$$

[2]The reference B-frames refer to those B-frames which serve as reference frames for the subsequent frames in coding order, whereas the non-reference B-frames are not used for reference.
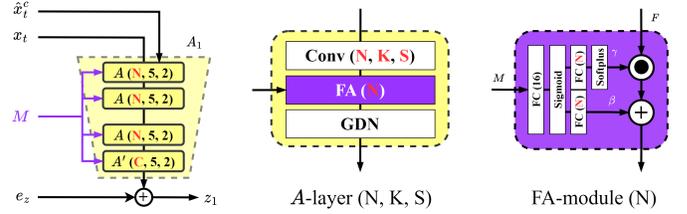
Fig. 4. Illustration of the analysis transform (A_1) with the frame-type adaptation (FA) module. The symbol $'$ indicates that the activation function (i.e. GDN) is omitted. Conv(N, K, S) stands for a convolutional layer with the number of output channels N, kernel size K, and stride S.

where $\odot$, $\oplus$ are channel-wise multiplication and addition. The FA modules appear in the autoencoding transforms and the hyperprior models. They are trained end-to-end together with the other networks. Note that similar feature adaptation techniques are widely used in computer vision tasks [42], [43] and are also adopted in end-to-end learned image compression [44]. Unlike these prior works, which take a set of feature maps as inputs to perform element-wise affine transformation, our FA module performs channel-wise affine transformation to adapt feature distributions with respect to a scalar indicating the input frame's type.

Section IV-D shows that frame-type adaptive coding improves bit allocation because our B-frame codec is able to adapt feature distributions on a frame-by-frame basis. In passing, this tool feature is to be distinguished from training separate models for different rate points, which is still the case in all the experiments.

E. Training

The model training is done on 5-frame training sequences (with their frames denoted by $x_0, x_1, \ldots, x_4$). In coding every training sequence, we follow the GOP-1 structure in Fig. 2(b), which encodes x_0 as an I-frame by a pre-trained image codec [20] (not included for training), x_4 as a B*-frame, and x_2 as a reference B-frame. Moreover, we randomly choose between x_1 and x_3 for coding, in order to have a balanced distribution over B-frame types. Assuming that x_3 is chosen, we formulate our training objective as follows:

$$\mathcal{L} = \lambda_1 \times D + R + \lambda_2 \times F, \qquad (8)$$

where

$$D = d(x_4, \hat{x}_4) + \frac{1}{\alpha_r} d(x_2, \hat{x}_2) + \frac{1}{\alpha_{nr}} d(x_3, \hat{x}_3), \qquad (9)$$

$$R = \sum_{t=2}^{4} r(\hat{x}_t), \qquad (10)$$

$$F = f_4 + \frac{1}{\alpha_r} f_2 + \frac{1}{\alpha_{nr}} f_3, \qquad (11)$$

with D denoting the weighted sum of frame distortions $\{d(x_t, \hat{x}_t)\}_{t=2}^{4}$, R the accumulated bits $\{r(\hat{x}_t)\}_{t=2}^{4}$ consumed by both the conditional motion and inter-frame coding, and F the sum of regularization losses $\{f_t = ||y_2 - \hat{x}_t^c||_2^2\}_{t=2}^{4}$ requiring that y_2 in Fig. 1 (b) should approximate their respective conditions, i.e. $\{\hat{x}_t^c\}_{t=2}^{4}$. This regularization loss is

TABLE III
SETTINGS OF THE COMPETING METHODS.

	I-frame codec (PSNR-RGB/MS-SSIM-RGB)	Coding structure	Intra-period	GOP
x265 (veryslow)	-	IPPP...	32	32
HM (LDP)	-	Hierarchical P	32	8
HM (randomaccess)	-	Hierarchical B	32	16
VTM (randomaccess)	-	Hierarchical B	32	32
DCVC [29]	cheng2020 [47]/hyperprior [35]	IPPP...	32	32
DCVC (ANFIC)	ANFIC [20]/ANFIC	IPPP...	32	32
CANF-VC [15]	ANFIC/ANFIC	IPPP...	32	32
Sheng'22 [30]	cheng2020 [47]/cheng2020	IPPP...	32	32
Sheng'22 (ANFIC)	ANFIC [20]/ANFIC	IPPP...	32	32
Li'22 [48]	Li'22 [48]/Li'22	IPPP...	32	32
Li'22 (ANFIC)	ANFIC [20]/ANFIC	IPPP...	32	32
LHBDC [13]	mbt2018-mean [47]/-	Hierarchical B	32	16
LHBDC (ANFIC)	ANFIC/-	Hierarchical B	32	16
B-CANF (Ours)	ANFIC/ANFIC	Hierarchical B	32	16

TABLE IV
BD-RATE COMPARISON WITH INTRA-PERIOD 32. THE ANCHOR IS X265 IN VERYSLOW MODE. THE BEST PERFORMER IS MARKED IN RED AND THE SECOND BEST IN BLUE.

	BD-rate (%) PSNR-RGB				BD-rate (%) MS-SSIM-RGB			
	UVG	MCL-JCV	HEVC-B	CLIC'22	UVG	MCL-JCV	HEVC-B	CLIC'22
HM (LDP)	-30.6	-26.4	-23.0	-16.8	-21.6	-17.3	-18.0	-8.8
HM (randomaccess)	-53.6	-48.2	-51.4	-41.3	-46.3	-40.8	-53.3	-39.2
VTM (randomaccess)	-68.2	-63.6	-66.4	-54.6	-69.6	-65.6	-71.1	-60.9
DCVC [29]	-2.3	-7.6	-3.0	28.7	-45.2	-49.5	-53.6	-25.9
DCVC (ANFIC)	-3.6	-9.0	-2.3	25.5	-47.7	-51.6	-55.8	-29.0
CANF-VC [15]	-28.7	-21.3	-21.4	3.3	-49.8	-50.4	-55.0	-30.3
Sheng'22 [30]	-45.7	-38.0	-43.1	-9.5	-63.6	-63.8	-73.9	-58.2
Sheng'22 (ANFIC)	-41.8	-38.0	-40.9	-9.9	-63.0	-64.7	-73.3	-57.1
Li'22 [48]	-62.4	-54.0	-57.6	-26.9	-73.0	-72.8	-80.4	-67.9
Li'22 (ANFIC)	-59.2	-51.2	-54.7	-26.3	-72.1	-71.9	-79.6	-66.6
LHBDC [13]	-18.6	-5.3	4.4	38.8	NA	NA	NA	NA
LHBDC (ANFIC)	-18.6	-6.7	3.4	38.2	NA	NA	NA	NA
B-CANF (Ours)	-47.5	-41.1	-46.9	-33.8	-60.2	-61.2	-67.1	-59.6

to have the latent representations $\hat{z}_2, \hat{h}_2$ (Fig. 1 (b)) capture the information needed to signal the transformation between the coding frame x_t and its predicted frame $\hat{x}_t^c$. The hyper-parameters λ_1, λ_2 are rate-point dependent, while α_r, α_{nr} are for bit allocation among B-frames. These hyper-parameters are detailed in Section IV.

IV. EXPERIMENTS

A. Settings

a) Training Details: For training, we use Vimeo-90k sep-tuplet dataset [45], which includes 91,701 7-frame sequences of size 448×256. The training sequences are randomly cropped to 256×256, and five consecutive frames are selected. We adopt Adam [46] optimizer with learning rate 10^{-4} and batch size 16. To optimize our models for PSNR-RGB, we choose $\lambda_2 = 0.01\lambda_1$ and $\lambda_1 = 128, 512, 1024, 2048$ to train separate models for different rate points; likewise, to optimize them for MS-SSIM-RGB, we choose $\lambda_2 = 0.01\lambda_1$ and $\lambda_1 = 2, 8, 16, 32$. For all the experiments, $\alpha_r = 1, \alpha_{nr} = 2$. These hyper-parameters are chosen to weight more heavily the quality of reference B-frames (including B*-frames).

b) Evaluation Methodologies: We evaluate our method on UVG [49], MCL-JCV [50], and HEVC Class B [51].

In addition, we also evaluate on CLIC'22 test dataset [52], which contains some challenging sequences with characteristics rarely seen in Vimeo-90k training dataset [45]. To compare fairly with HM (HEVC Test Model Version 16.23) [41] under the common test conditions, we follow the coding structure of the *encoder_randomaccess_main* configuration. That is, the intra-period is 32 and GOP size is 16 for all the test sequences. For every test sequence, we encode most of the video frames to the extent that the coded sequence contains the maximum number of intra-periods (32). As an example, for a 600-frame test sequence, we will encode 577 frames (1 I-frame + 18 intra-periods). Note that the source videos are in YUV420 format. The traditional codecs operate internally in YUV420, whereas the learned codecs perform YUV420-to-RGB444 conversion prior to coding. The reconstructed quality is measured in PSNR-RGB and MS-SSIM-RGB, and the bit-rate in bits-per-pixel (bpp). Table III summarizes the settings of these competing methods.

c) Baseline Methods: The baseline methods include traditional codecs and learning-based methods. The traditional codecs are x265 [53] in *veryslow* mode (zero-latency), HM [41] with the *encoder_lowdelay_P_main* and *encoder_randomaccess_main* configurations, and VTM [54] with the *encoder_randomaccess_main* configuration. The
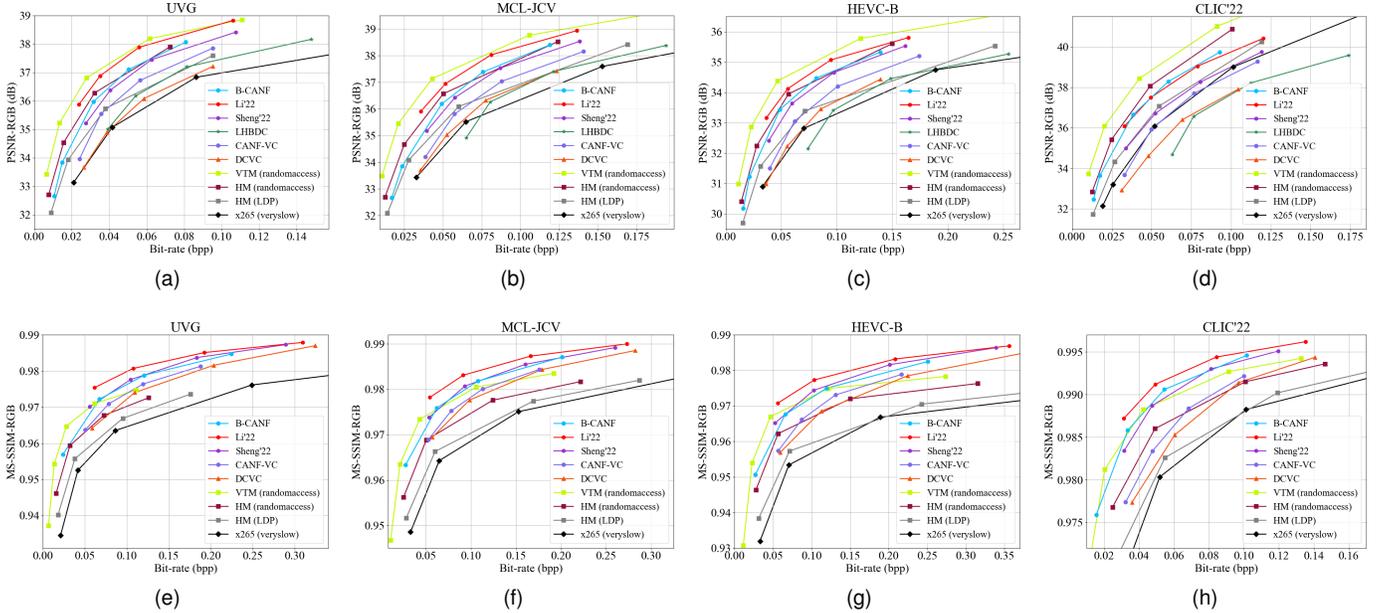
Fig. 5. Rate-distortion plots on UVG, MCL-JCV, HEVC Class B and CLIC'22 test dataset in terms of PSNR-RGB and MS-SSIM-RGB

TABLE V
BD-RATES OF B-CANF AS COMPARED AGAINST DIFFERENT ANCHORS UNDER THEIR RESPECTIVE SETTINGS.

Anchor	I-frame codec (PSNR-RGB/MS-SSIM-RGB)	Intra-period	GOP	BD-rate (%) PSNR-RGB				BD-rate (%) MS-SSIM-RGB			
				UVG	MCL-JCV	HEVC-B	CLIC'22	UVG	MCL-JCV	HEVC-B	CLIC'22
HLVC [11]	BPG [55]/ICLR'19 [56]	10	10	-53.7	-55.4	-50.3	-56.9	-44.2	-48.5	-38.7	-44.8
HLVC (ANFIC)	ANFIC [20]/ANFIC	10	10	-53.9	-54.9	-50.3	-56.9	-44.7	-49.0	-39.0	-44.1
RLVC [22]	BPG/ICLR'19	13	13	-50.6	-49.4	-46.5	-47.8	-36.6	-41.9	-33.6	-48.9
RLVC (ANFIC)	ANFIC/ANFIC	13	13	-47.8	-46.9	-44.1	-46.9	-34.0	-39.3	-35.3	-46.1
LHBDC [13]	mbt2018-mean [47]/mbt2018-mean	8	8	-33.6	-36.1	-43.1	-41.2	-31.4	-	-	-
LHBDC (ANFIC)	ANFIC/ANFIC	8	8	-30.9	-32.9	-40.7	-39.9	-	-	-	-
B-EPIC [12]	NA/NA	infinity	12	-32.3	-33.8	-32.5	-	-15.5	-21.2	-26.5	-

learning-based methods include CANF-VC [15], DCVC [29] and two concurrent works, Sheng'22 [30] and Li'22 [48], which are the state-of-the-art learned P-frame coding schemes. We also compare B-CANF with LHBDC [13], the state-of-the-art learned B-frame coding scheme. To align the I-frame codec, we additionally evaluate their performance using ANFIC [20] as the I-frame codec. We refer to these variants as DCVC (ANFIC), Sheng'22 (ANFIC), Li'22 (ANFIC) and LHBDC (ANFIC), respectively. Results for the baselines and their variants are produced with the code released by the authors.

B. Rate-Distortion Performance

When comparing our B-CANF with CANF-VC [15] and LHBDC [13], we observe that B-CANF surpasses them in terms of both PSNR-RGB and MS-SSIM-RGB across all the test datasets. However, B-CANF falls short of achieving the same level of performance as Li'22 [48], which is considered the state-of-the-art learned P-frame codec. This discrepancy in performance can be attributed to the more advanced entropy coding model employed by Li'22 [48] and the domain shift issue that arises between the training and testing phases. For a comprehensive analysis of this domain shift, please refer to Section IV-C.

Table V further compares our method with HLVC [11], RLVC [22], LHBDC [13], and B-EPIC [12] following their suggested settings. All the baseline methods adopt or support B-frame coding. In this experiment, these baseline methods are used as anchors for BD-rate evaluation. In addition, we note that the codes for B-EPIC [12] are unavailable. We thus copy the results from their paper in evaluating BD-rates.

In HLVC [11], a three-level hierarchical coding structure with P- and B-frames is used for encoding GOP's of size 10. We follow exactly the same coding structure as HLVC [11] to perform encoding with B-CANF, except that we use B*-frames to replace P-frames. Recall that B*-frames are used to mimic P-frames. In RLVC [22], a bidirectional P-frame coding structure is proposed to take advantage of the past and future frames for better coding efficiency. For the same purpose, our B-CANF proposes B-frame coding. To demonstrate the best achievable performance of B-CANF, we adopt hierarchical B-frame coding for encoding GOP's of size 13 in comparing with RLVC [22]. In LHBDC [13], a hierarchical B prediction structure with GOP 8 and intra-period 8 is proposed. In comparison with B-EPIC [12], which adopts a GOP size of 12 with an infinite intra-period (only one I-frame in the beginning), our B-CANF uses the same GOP size. However, we set the intra-period to 12 (instead of infinity) to mitigate the temporal error propagation. Despite that this leads to

| | BD-rate (%) PSNR-RGB | | | | | |
| | UVG | | MCL-JCV | | HEVC-B | |
	I32	I8	I32	I8	I32	I8
B-CANF (ANFIC)	0.0	0.0	0.0	0.0	0.0	0.0
Sheng'22 (ANFIC)	14.6	20.4	9.2	13.6	14.0	22.1
Li'22 (ANFIC)	-14.3	4.0	-13.7	-4.3	-12.6	7.5

Fig. 6. Profiles of the per-frame PSNR's and the per-frame bit rates averaged over GOPs on UVG dataset, with and without frame-type adaptive (FA) coding at high rates ($\lambda_1 = 2048$). Frame 0, 32 are coded as I-frame, Frame 16 is coded as B*-frame, and the odd-numbered frames are non-reference B-frames while the even-numbered frames are reference B-frames.

more frequent coding of I-frames, its advantages outweigh the temporal error propagation. Under these various settings, our B-CANF (which adopts ANFIC [20] as the I-frame codec) still achieves considerable BD-rate reductions over HLVC [11], RLVC [22], LHBDC [13], and B-EPIC [12]. This is evidenced by the negative BD-rate numbers in the upper half of Table V.

C. Domain Shift in Training B-frame Codecs

The performance gap between our B-CANF and the works of Sheng'22 [30] and Li'22 [48] deserves further investigation. A crucial aspect that is often overlooked when comparing the compression performance of learned P-frame and B-frame codecs is the training datasets. Most learned codecs (P-frame and B-frame) are trained on Vimeo dataset [45], which consists of short video sequences, each containing only 7 frames. Consequently, we use 5-frame GOPs with 3-level B-frame coding for training. Notably, these training B-frames are temporally adjacent to each other.

However, during testing, we employ a large GOP of size 32 (i.e. the intra-period is 32) with 5 hierarchical temporal levels in most cases. This results in a substantial number of B-frames being predicted from distant reference frames, leading to a significant domain shift (i.e. a significant change in statistics) between the training and test scenarios. This discrepancy can negatively affect the generalization of our ME-Net, motion prediction network, and frame synthesis network, especially when dealing with videos featuring substantial motion. In contrast, learned P-frame codecs like Li'22 [48] and Sheng'22 [30] have more consistent training and test scenarios in terms of temporal prediction distance. In P-frame coding, the domain shift can occur after coding several frames because the quality of the reference frame often degrades over time. In B-frame coding, it happens right from the first coding B-frame, affecting all the subsequent frames. This early domain shift in B-frame coding has a more immediate and pronounced impact on compression performance.

To mitigate the domain shift between the training and test scenarios, we present additional results in Table VI for a smaller GOP size of 8. In order to ensure a fair comparison, we utilize ANFIC as the I-frame codec for all the competing methods. The results demonstrate that our B-CANF performs comparably to or better than Li'22 [48] under GOP=8, even without utilizing advanced entropy coding. Note that these results are not intended to advocate the use of small GOPs or short intra-periods. Rather, they serve to highlight the limitations of Vimeo dataset [45] in training B-frame codecs and demonstrate the potential of B-CANF.

D. Ablation Experiments

We conduct ablation experiments with GOP 16 and intra-period 32. Unless otherwise specified, the BD-rates are reported against x265 (veryslow) in terms of PSNR-RGB.

a) Conditional vs. Residual Coding: Table VII analyzes the impact of our conditional coding scheme on compression performance. We test four variants, switching between residual and conditional coding for motion and inter-frame coding. When residual coding is selected, we use ANFIC [20] to code flow map or frame prediction residuals. Comparing variants A and D, turning both motion and inter-frame codecs into residual-type ones results in a considerable performance drop. In particular, our conditional coding scheme is seen to benefit inter-frame coding more than motion coding (from D to B vs. from D to C). A breakdown analysis of the bitstream composition explains that the motion part represents only 5-15% of the entire bitstream. Interestingly, the gains of our conditional coding in improving motion coding and inter-frame coding are nearly additive.

b) Frame-type Adaptive Coding: As mentioned in Section III-D, there are two mechanisms that contribute to the better bit allocation of our B-CANF. The first is to weight differently the distortions of different types of B-frames during training. The second is the FA module, which allows both the conditional motion and inter-frame codecs to behave differently according to the input frame's type. Table VIII presents an ablation study to shed lights on how they contribute to the rate-distortion performance. Comparing setting A (our B-CANF) with setting B, we see that turning off the FA module leads to a significant rate-distortion performance loss. In this case, the same (amortized/average) motion and inter-frame codecs are used to process different types of B-frames. Likewise, in comparison with the full model (setting A), having only the FA module without weighting the distortions according to the frame type (setting C) during training is sub-optimal too. These results suggest that they both contribute significantly to the resulting rate-distortion performance.

TABLE VII
ABLATION STUDY OF CONDITIONAL CODING. SETTING A IS OUR PROPOSED B-CANF MODEL.

Setting	Motion coding		Inter-frame coding		BD-rate (%) PSNR-RGB			
	Residual	Conditional	Residual	Conditional	UVG	MCL-JCV	HEVC-B	CLIC'22
A		✓		✓	-47.5	-41.1	-46.9	-33.8
B	✓			✓	-41.8	-36.3	-42.2	-30.0
C		✓	✓		-33.4	-26.8	-34.6	-23.0
D	✓		✓		-28.2	-22.8	-30.3	-20.0

TABLE VIII
ABLATION STUDY OF THE DISTORTION WEIGHTING AND THE FRAME-TYPE ADAPTATION (FA) MODULE. SETTING A IS OUR PROPOSED B-CANF MODEL.

Setting	Distortion weighting	FA module	BD-rate (%) PSNR-RGB			
			UVG	MCL-JCV	HEVC-B	CLIC'22
A	✓	✓	-47.5	-41.1	-46.9	-33.8
B	✓		-39.4	-32.1	-41.0	-28.7
C		✓	-39.3	-35.4	-38.1	-25.0
D			-32.0	-30.4	-36.3	-22.1

TABLE IX
ABLATION STUDY OF FRAME-TYPE ADAPTIVE CODING. SETTING A IS OUR PROPOSED B-CANF MODEL.

Setting	Frame-type adaptive coding		BD-rate (%) PSNR-RGB			
	Motion codec	Inter-frame codec	UVG	MCL-JCV	HEVC-B	CLIC'22
A	✓	✓	-47.5	-41.1	-46.9	-33.8
B	✓		-41.6	-32.7	-40.6	-28.4
C		✓	-42.3	-35.9	-43.2	-30.9
D			-39.4	-32.1	-41.0	-28.7

Fig. 6 further visualizes how the proposed FA module may impact the bit allocation among B-frames in a GOP and their decoded quality. In this and the following experiments, the distortion weighting is enabled during training. We see that disabling the FA module (i.e., setting B in Table VIII) will learn an "average" codec that tends to allocate more bits to the non-reference B-frames. As a result, the PSNR distribution becomes more smooth and less hierarchical.

Table IX further presents BD-rate results by turning on and off frame-type adaptive coding in the motion and/or inter-frame codecs. Interestingly, when applied to either the inter-frame codec (from D to C) or the motion codec (from D to B) alone, it is not as effective as the case when it is enabled for both codecs (from D to A).

To gain insights into this observation, Fig. 7 analyzes the bit allocation between motion coding and inter-frame coding for B-frames. Comparing setting C with setting D, we observe that enabling frame-type adaptive coding for the inter-frame codec only is able to reduce effectively the bit rate of inter-frame coding (green) while keeping the motion bit rate (red) relatively untouched. Likewise, the comparison of settings B and D shows that applying it to the motion codec only reduces effectively the motion bit rate (red). In this case, the slightly increased bit rate of inter-frame coding (green) may be attributed to the use of a shared inter-frame codec, which is unable to adapt to the changes in the inter-frame correlation that result from frame-type adaptive motion coding. Turing on frame-type adaptive coding for both codecs allows the greatest degree of flexibility in allocating bits to different parts, achieving a synergy effect.

c) Coded Flow Maps for B*-frames: Figs. 8(b) and (c) visualize the two coded flow maps for B*-frames. As shown
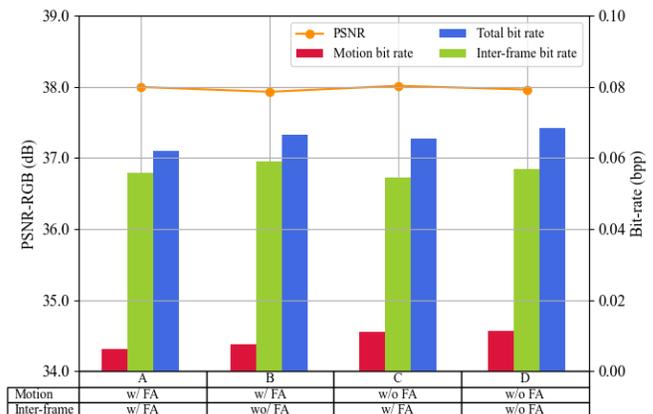
Fig. 7. Comparison of PSNR-RGB and bit allocation by turning on and off frame-type adaptive coding in the motion and/or inter-frame codecs at high rates ($\lambda_1 = 2048$). The PSNR-RGB and bit-allocation results are averaged over B-frames of the test sequences in UVG dataset. FA refers to frame-type adaptive coding. Setting A is our B-CANF model.

in Fig. 8(d), they are seen to disagree slightly with each other around object boundaries. Table X presents BD-rate results for three variants of motion compensation. *Both* refers to using two hypotheses, i.e. the coded flow map $\hat{m}^e_{t \to t-k}$ and its virtual counterpart after sign reversal, for motion compensation, whereas *First* or *Second* refers to using solely one of them. We see that using solely $\hat{m}^e_{t \to t-k}$ (*First*) or its virtual counterpart (*Second*) for motion compensation leads to less rate saving than using both simultaneously. This may be attributed to the fact that the input flow map is usually less reliable in boundary regions. Performing multi-hypothesis prediction with slightly different motion estimates in these

TABLE X
ABLATION STUDY OF THE NUMBER OF HYPOTHESES USED FOR MOTION
COMPENSATING B*-FRAME UNDER INTRA-PERIOD 32 AND GOP SIZE 4.

Hypothesis	BD-rate (%) PSNR-RGB			
	UVG	MCL-JCV	HEVC-B	CLIC'22
Both	-40.6	-34.6	-39.4	-18.7
First	-35.0	-31.0	-35.1	-14.3
Second	-35.2	-30.9	-35.1	-14.7

TABLE XI
THE EFFECT OF GOP SIZE ON B-CANF WITH AND WITHOUT A SEPARATE
P-FRAME CODEC UNDER INTRA-PERIOD 32.

GOP	BD-rate (%) PSNR-RGB							
	UVG		MCL-JCV		HEVC-B		CLIC'22	
	w/o	w/	w/o	w/	w/o	w/	w/o	w/
16	**-47.5**	-46.4	**-41.1**	**-41.1**	-46.9	**-47.7**	-33.8	**-36.5**
8	**-45.7**	-44.8	**-39.5**	-39.3	-45.1	**-46.3**	-29.7	**-30.8**
4	**-40.6**	-34.7	**-34.6**	-31.7	**-39.4**	-37.2	**-18.7**	-12.5
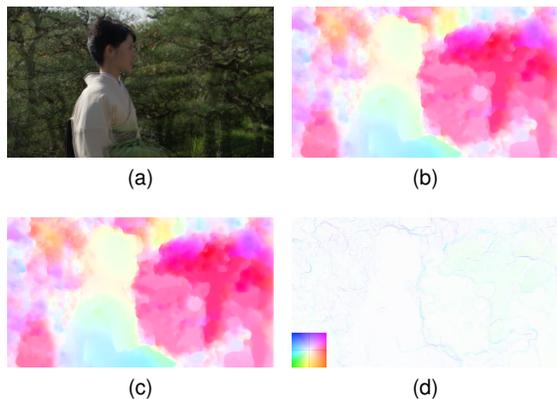
Fig. 8. Visualization of (a) a coding frame overlaid with its reference frame, (b) the coded flow map $\hat{m}_{t\to t-k}^e$ (cp. Fig. 3(b)), (c) its virtual counterpart after sign reversal, and (d) the differences between (b) and (c). White color represents zero motion.

Fig. 9. Rate-distortion comparison of GOP sizes 4, 8, 16 on UVG dataset under intra-period 32.

regions is an effective means of mitigating the reliability issue.

d) B*-frame vs. Separate P-frame Coding: Table XI investigates the effectiveness of B*-frame coding as a substitute for P-frame coding under various GOP sizes. The former reuses the B-frame codec while the latter needs to train end-to-end a separate, dedicated P-frame codec together with the other components of B-CANF. In most cases, our B*-frame mechanism is seen to achieve similar BD-rate results to the use of a separate P-frame codec. In particular, under GOP size 4, it performs even better than P-frame coding. Recall that B*-frames are essentially a multi-hypothesis prediction technique.

e) GOP Size vs. Rate-distortion Performance: We evaluate the effect of GOP size on the performance of B-CANF. A number of GOP sizes, including 4, 8, 16, are tested with intra-period 32. The BD-rates are summarized in Table XI (see the results w/o a separate P-frame codec). The corresponding rate-distortion curves on UVG dataset are presented in Fig. 9. From Fig. 9, the rate-distortion performance of B-CANF is seen to improve with the increased GOP size. The improvement is most obvious at low rates. Like P-frames, our B*-frames suffer more from temporal error propagation with smaller GOP sizes (in which cases, B*-frames are sent more frequently), especially at low rates where poor reconstruction and motion quality is expected. Increasing GOP size decreases the frequency of B*-frames, thereby reducing temporal error propagation.

f) The Number of Autoencoding Transforms: Table XII compares the BD-rates between 2-step and 1-step B-CANF to explore the effect of the number of autoencoding transforms on compression performance. The 1-step B-CANF is obtained by skipping the first autoencoding transform $\{A_1, S_1\}$ in

Fig. 1(b). In order to have a similar model size to the 2-step B-CANF, it has more channels in each autoencoding transform. Remarkably, the 1-step B-CANF reduces to a specific implementation of conditional VAE (CVAE). From Table XII, we observe that the gain of the 2-step B-CANF over the 1-step B-CANF is obvious across the datasets. The result justifies our use of CANF rather than CVAE.

E. Subjective Quality

Fig. 10 presents the subjective quality comparison between HM (randomaccess), LHBDC [13] and our B-CANF. LHBDC (MSE) and B-CANF (MSE) are trained to optimize PSNR-RGB while B-CANF (SSIM) is our model trained with MS-SSIM-RGB. All the schemes are evaluated with GOP 16 and intra-period 32, and all the learning-based methods use ANFIC as the I-frame codec. It is seen that our B-CANF (MSE) achieves comparable or even better subjective quality than LHBDC (MSE), with its bit rate being nearly one order of magnitude smaller than that of LHBDC (MSE). Compared with HM (randomaccess), our B-CANF (SSIM) preserves more texture details (cf. patterns on fingers in the first row, pillars in the second row, and textures in the last row) at a lower bit rate.

F. Complexity Analysis

Table XIII characterizes the complexity of the proposed B-CANF in terms of runtimes, MACs and model size. We also measure the memory footprint following [28]. This includes the size of the decoded picture buffer, which stores decoded frames, flow maps, and/or context features. The unit of measurement is "full-res," where one reconstructed frame occupies
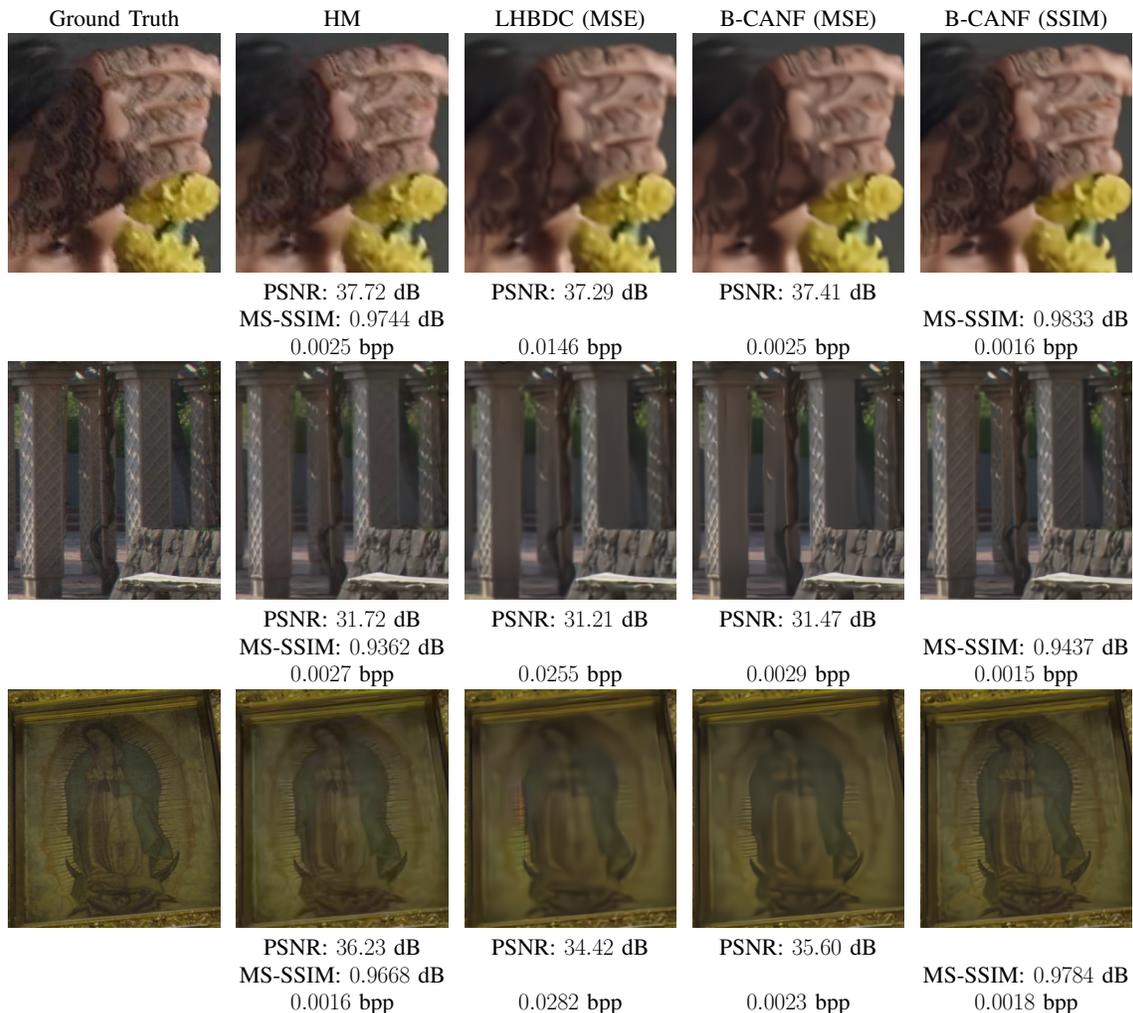
Fig. 10. Subjective quality comparison of HM (randomaccess), LHBDC [13] and our B-CANF. The parentheses "(MSE)" and "(SSIM)" indicate that the model is trained with MSE and MS-SSIM, respectively. Zoom in for better visualization.

TABLE XII
ABLATION STUDY OF THE NUMBER OF AUTOENCODING TRASFORMS

Motion and Inter-frame codecs	BD-rate (%) PSNR-RGB			
	UVG	MCL-JCV	HEVC-B	CLIC'22
1-step B-CANF (CVAE)	-36.7	-30.1	-37.3	-24.9
2-step B-CANF (Ours)	-47.5	-41.1	-46.9	-33.8

the equivalent of 3 full-res. The peak memory represents the maximum buffering requirement for generated features during compression. The reported numbers in Table XIII were obtained using an NVIDIA GeForce RTX 2080Ti GPU for the learned codecs, and an Intel(R) Core(TM) i7-9700K CPU @ 3.60GHz for HM [41] and VTM [54]. The experiments were conducted using 1080p input videos.

For frame-type P/B*, we use the prediction structure IPPP... under intra-period 32 and GOP size 32. The encoding/decoding times are then averaged over the first 100 P/B*-frames of Beauty sequence in UVG dataset. For frame-type B, the encoding/decoding times are averaged over the first 100 B-frames of the same sequence with intra-period 32, GOP size 16, and hierarchical B prediction.

The following observations can be made from Table XIII.

(1) The prolonged runtimes of DCVC [29] are attributed to the use of an autoregressive model for entropy coding. In contrast, CANF-VC [15], LHBDC [13] and our B-CANF do not use any auto-regressive model for motion coding or inter-frame coding. (2) In comparison with VAE-based coding, DCVC [29] and LHBDC [13], the larger MAC of CANF-VC [15] and B-CANF comes from stacking multiple autoencoding transforms. Nevertheless, the comparable encoding/decoding runtimes suggest that these autoencoding transforms are parallel-friendly. (3) Our B-CANF has lower buffering requirements and peak memory when compared to other approaches. (4) With our B-CANF, the encoding time and MAC of B-frames are seen to be higher than those of B*-frames (the last two rows in Table XIII). This is because motion estimation is performed twice for B-frames, whereas it is carried out only once for B*-frames (just like P-frames). In contrast, the decoding times and MAC of both types of frame are quite close, even though B-frames incur extra computations in the motion prediction network (Fig. 3(a) vs. Fig. 3(b)). This indicates that the motion prediction network is not the major computation bottleneck during decoding. (5) As compared with CANF-VC [15], which supports only P-frame coding, the B-frame coding in B-

TABLE XIII
COMPARISON OF THE MODEL COMPLEXITY IN TERMS OF ENCODING/DECODING MACs, RUNTIMES, MODEL SIZE, DECODED PICTURE BUFFER, AND PEAK MEMORY. "FULL-RES" STANDS FOR THE SIZE OF INPUT RESOLUTION.

Method	Frame-type	Encode		Decode		Model Size	Decoded Picture Buffer	Peak Memory
		Time	MACs	Time	MACs			
HM [41]	B	33.28s	-	0.04s	-	-	-	-
VTM [54]	B	731.33s	-	0.07s	-	-	-	-
DCVC [29]	P	7.70s	1.16M/pixel	28.97s	0.77M/pixel	8M	3 full-res	128 full-res
CANF-VC [15]	P	1.45s	2.45M/pixel	1.07s	1.77M/pixel	31M	13 full-res	64 full-res
Sheng'22 [30]	P	0.82s	1.42M/pixel	0.58s	0.92M/pixel	10.7M	67 full-res	128 full-res
Li'22 [48]	P	0.83s	1.68M/pixel	0.61s	1.25M/pixel	17.5M	67 full-res	128 full-res
LHBDC [13]	B	1.19s	1.70M/pixel	0.73s	1.12M/pixel	23.5M	15 full-res	96 full-res
B-CANF (Ours)	B*	1.44s	2.68M/pixel	1.06s	2.01M/pixel	24M	3 full-res	64 full-res
B-CANF (Ours)	B	1.69s	3.08M/pixel	1.09s	2.10M/pixel	24M	15 full-res	64 full-res

CANF has higher encoding MAC because motion estimation is performed twice. In contrast, our B*-frames, similar to P-frames, have comparable encoding MAC to CANF-VC [15]. The decoding MAC's of both B-frames and B*-frames are seen to be comparable to that of CANF-VC [15]. (6) In terms of model size, B-CANF is similar to LHBDC [13].

V. CONCLUSION

In this paper, we propose a CANF-based B-frame coding framework, known as B-CANF, that exploits the notion of conditional coding for both motion and inter-frame coding. It features B*-frames and frame-type adaptive coding. We show that (1) B*-frames allow greater flexibility in supporting various GOP sizes without the need for an extra P-frame codec, and (2) frame-type adaptive coding improves the bit allocation among B/B*-frames. Extensive experimental results confirm the superiority of B-CANF to the other state-of-the-art B-frame coding schemes. How to achieve even higher compression performance by addressing the domain shift issue and reducing the model's complexity is our future work.

REFERENCES

[1] G. Lu, W. Ouyang, D. Xu, X. Zhang, C. Cai, and Z. Gao, "Dvc: An end-to-end deep video compression framework," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 11 006–11 015.

[2] G. Lu, X. Zhang, W. Ouyang, L. Chen, Z. Gao, and D. Xu, "An end-to-end learning framework for video compression," *IEEE transactions on Pattern Analysis and Machine Intelligence*, 2020.

[3] E. Agustsson, D. Minnen, N. Johnston, J. Balle, S. J. Hwang, and G. Toderici, "Scale-space flow for end-to-end optimized video compression," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8503–8512.

[4] J. Lin, D. Liu, H. Li, and F. Wu, "M-lvc: multiple frames prediction for learned video compression," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3546–3554.

[5] Z. Hu, Z. Chen, D. Xu, G. Lu, W. Ouyang, and S. Gu, "Improving deep video compression by resolution-adaptive flow coding," in *European Conference on Computer Vision*. Springer, 2020, pp. 193–209.

[6] O. Rippel, A. G. Anderson, K. Tatwawadi, S. Nair, C. Lytle, and L. Bourdev, "Elf-vc: Efficient learned flexible-rate video coding," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 14 479–14 488.

[7] H. Liu, M. Lu, Z. Ma, F. Wang, Z. Xie, X. Cao, and Y. Wang, "Neural video coding using multiscale motion compensation and spatiotemporal context model," *IEEE Transactions on Circuits and Systems for Video Technology*, 2020.

[8] T. Ladune, P. Philippe, W. Hamidouche, L. Zhang, and O. Déforges, "Optical flow and mode selection for learning-based video coding," in *2020 IEEE 22nd International Workshop on Multimedia Signal Processing (MMSP)*. IEEE, 2020, pp. 1–6.

[9] T. Ladune, P. Philippe, W. Hamidouche, L. Zhang, and O. Déforges, "Conditional coding for flexible learned video compression," in *Neural Compression: From Information Theory to Applications–Workshop@ ICLR 2021*, 2021.

[10] A. Djelouah, J. Campos, S. Schaub-Meyer, and C. Schroers, "Neural inter-frame compression for video coding," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[11] R. Yang, F. Mentzer, L. V. Gool, and R. Timofte, "Learning for video compression with hierarchical quality and recurrent enhancement," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6628–6637.

[12] R. Pourreza and T. Cohen, "Extending neural p-frame codecs for b-frame coding," *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6660–6669, 2021.

[13] M. A. Yılmaz and A. M. Tekalp, "End-to-end rate-distortion optimized learned hierarchical bi-directional video compression," *IEEE Transactions on Image Processing*, vol. 31, pp. 974–983, 2022.

[14] C.-Y. Wu, N. Singhal, and P. Krahenbuhl, "Video compression through image interpolation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.

[15] Y.-H. Ho, C.-P. Chang, P.-Y. Chen, A. Gnutti, and W.-H. Peng, "Canf-vc: Conditional augmented normalizing flows for video compression," in *Proceedings of the European Conference on Computer Vision (ECCV)*, October 2022.

[16] Z. Chen, T. He, X. Jin, and F. Wu, "Learning for video compression," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 2, pp. 566–576, 2020.

[17] H. Liu, M. Lu, Z. Chen, X. Cao, Z. Ma, and Y. Wang, "End-to-end neural video coding using a compound spatiotemporal representation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 8, pp. 5650–5662, 2022.

[18] M. Afonso, F. Zhang, and D. R. Bull, "Video compression based on spatio-temporal resolution adaptation," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 29, no. 1, pp. 275–280, 2019.

[19] M.-J. Chen, H.-S. Xie, C. Chien, W.-H. Peng, and H.-M. Hang, "Learned hierarchical b-frame coding with adaptive feature modulation for yuv 4:2:0 content," *arXiv preprint arXiv:2212.14187*, 2022.

[20] Y.-H. Ho, C.-C. Chan, W.-H. Peng, H.-M. Hang, and M. Domański, "Anfic: Image compression using augmented normalizing flows," *IEEE Open Journal of Circuits and Systems*, vol. 2, pp. 613–626, 2021.

[21] Z. Hu, G. Lu, and D. 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*, 2021, pp. 1502–1511.

[22] R. Yang, F. 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*, vol. 15, no. 2, pp. 388–401, 2020.

[23] A. Golinski, R. Pourreza, Y. Yang, G. Sautiere, and T. S. Cohen, "Feedback recurrent autoencoder for video compression," in *Proceedings of the Asian Conference on Computer Vision*, 2020.

[24] Y. Shi, Y. Ge, J. Wang, and J. Mao, "Alphavc: High-performance and efficient learned video compression," 2022. [Online]. Available: https://arxiv.org/abs/2207.14678

[25] Z. Hu, G. Lu, J. Guo, S. Liu, W. Jiang, and D. 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 (CVPR)*, June 2022, pp. 5921–5930.

[26] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, "Overview of the high efficiency video coding (hevc) standard," *IEEE Transactions on circuits and systems for video technology*, vol. 22, no. 12, pp. 1649–1668, 2012.

[27] B. Bross, Y.-K. Wang, Y. Ye, S. Liu, J. Chen, G. J. Sullivan, and J.-R. Ohm, "Overview of the versatile video coding (vvc) standard and its applications," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 10, pp. 3736–3764, 2021.

[28] Y. Zhu, Y. Yang, and T. Cohen, "Transformer-based transform coding," in *International Conference on Learning Representations*, 2022. [Online]. Available: https://openreview.net/forum?id=IDwN6xjHnK8

[29] J. Li, B. Li, and Y. Lu, "Deep contextual video compression," *Advances in Neural Information Processing Systems*, 2021.

[30] X. Sheng, J. Li, B. Li, L. Li, D. Liu, and Y. Lu, "Temporal context mining for learned video compression," *IEEE Transactions on Multimedia*, pp. 1–12, 2022.

[31] F. Mentzer, G. Toderici, D. Minnen, S.-J. Hwang, S. Caelles, M. Lucic, and E. Agustsson, "Vct: A video compression transformer," 2022. [Online]. Available: https://arxiv.org/abs/2206.07307

[32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

[33] T. Wiegand, G. Sullivan, G. Bjontegaard, and A. Luthra, "Overview of the h.264/avc video coding standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 560–576, 2003.

[34] C.-W. Huang, L. Dinh, and A. Courville, "Augmented normalizing flows: Bridging the gap between generative flows and latent variable models," *arXiv preprint arXiv:2002.07101*, 2020.

[35] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, "Variational image compression with a scale hyperprior," in *International Conference on Learning Representations*, 2018.

[36] M. Flierl and B. Girod, "Generalized b pictures and the draft h.264/avc video-compression standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 587–597, 2003.

[37] A. Ranjan and M. J. Black, "Optical flow estimation using a spatial pyramid network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4161–4170.

[38] Z. Huang, T. Zhang, W. Heng, B. Shi, and S. Zhou, "Rife: Real-time intermediate flow estimation for video frame interpolation," *arXiv preprint arXiv:2011.06294*, 2020.

[39] D. Fourure, R. Emonet, E. Fromont, D. Muselet, A. Tremeau, and C. Wolf, "Residual conv-deconv grid network for semantic segmentation," *British Machine Vision Conference*, 2017.

[40] E. P. S. Johannes Ballé, Valero Laparra, "End-to-end optimized image compression," *International Conference for Learning Representations*, 2017.

[41] "Hm reference software for hevc," https://vcgit.hhi.fraunhofer.de/jvet/HM/-/tree/HM-16.23/, accessed: 2022-05-018.

[42] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. Huang, "Free-form image inpainting with gated convolution," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 4470–4479.

[43] X. Wang, K. Yu, C. Dong, and C. C. Loy, "Recovering realistic texture in image super-resolution by deep spatial feature transform," 06 2018, pp. 606–615.

[44] M. Song, J. Choi, and B. Han, "Variable-rate deep image compression through spatially-adaptive feature transform," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2021, pp. 2360–2369. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00238

[45] T. Xue, B. Chen, J. Wu, D. Wei, and W. T. Freeman, "Video enhancement with task-oriented flow," *International Journal of Computer Vision*, vol. 127, no. 8, pp. 1106–1125, 2019.

[46] J. B. Diederik P. Kingma, "Adam: A method for stochastic optimization," *International Conference for Learning Representations*, 2015.

[47] J. Bégaint, F. Racapé, S. Feltman, and A. Pushparaja, "Compressai: a pytorch library and evaluation platform for end-to-end compression research," *arXiv preprint arXiv:2011.03029*, 2020.

[48] J. Li, B. Li, and Y. Lu, "Hybrid spatial-temporal entropy modelling for neural video compression," in *Proceedings of the 30th ACM International Conference on Multimedia*. ACM, oct 2022. [Online]. Available: https://doi.org/10.1145%2F3503161.3547845

[49] A. Mercat, M. Viitanen, and J. Vanne, "Uvg dataset: 50/120fps 4k sequences for video codec analysis and development," in *Proceedings of the 11th ACM Multimedia Systems Conference*, 2020, pp. 297–302.

[50] H. Wang, W. Gan, S. Hu, J. Y. Lin, L. Jin, L. Song, P. Wang, I. Katsavounidis, A. Aaron, and C.-C. J. Kuo, "Mcl-jcv: a jnd-based h.264/avc video quality assessment dataset," in *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2016, pp. 1509–1513.

[51] F. Bossen *et al.*, "Common test conditions and software reference configurations," *JCTVC-L1100*, vol. 12, no. 7, 2013.

[52] G. Toderici, W. Shi, R. Timofte, L. Theis, J. Balle, E. Agustsson, N. Johnston, and F. Mentzer, "Workshop and Challenge on Learned Image Compression (CLIC2022)," 2022. [Online]. Available: http://compression.cc/

[53] "Ffmpeg," https://www.ffmpeg.org/, accessed: 2022-05-018.

[54] "Vtm," https://vcgit.hhi.fraunhofer.de/jvet/VVCSoftware_VTM/, accessed: 2022-03-02.

[55] F. Bellard, "BPG image format," *URL https://bellard. org/bpg*, 2015.

[56] J. Lee, S. Cho, and S.-K. Beack, "Context-adaptive entropy model for end-to-end optimized image compression," in *the 7th Int. Conf. on Learning Representations*, May 2019.

Mu-Jung Chen received his B.S. degree in electronic and computer engineering from National Taiwan University of Science and Technology (NTUST), Taiwan, in 2021. He is currently pursuing his M.S. degree in computer science and engineering, National Yang Ming Chiao Tung University (NYCU). His research interests include learning-based image/video coding, computer vision, and deep/machine learning.

Yi-Hsin Chen received her B.S. degree in applied mathematics from National Chung Hsing University (NCHU), Taiwan, and her M.S. degree in data science and engineering from National Chiao Tung University (NCTU), Taiwan, in 2018 and 2020, respectively. She is currently pursuing her Ph.D. degree in computer science and engineering, National Yang Ming Chiao Tung University (NYCU), Taiwan. Her research interests include learning-based image/video coding, image/video restoration, computer vision, and deep/machine learning.
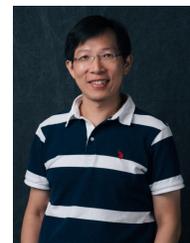
Wen-Hsiao Peng received his Ph.D. degree from National Chiao Tung University (NCTU), Taiwan, in 2005. He was with the Intel Microprocessor Research Laboratory, USA, from 2000 to 2001, where he was involved in the development of ISO/IEC MPEG-4 fine granularity scalability. Since 2003, he has actively participated in the ISO/IEC and ITU-T video coding standardization process and contributed to the development of SVC, HEVC, and SCC standards. He was a Visiting Scholar with the IBM Thomas J. Watson Research Center, USA, from 2015 to 2016. He is currently a Professor with the Computer Science Department, National Yang Ming Chiao Tung University, Taiwan. He has authored over 75+ journal/conference papers and over 60 ISO/IEC and ITU-T standards contributions. His research interests include learning-based video/image compression, deep/machine learning, multimedia analytics, and computer vision. Dr. Peng was Chair of the IEEE Circuits and Systems Society (CASS) Visual Signal Processing (VSPC) Technical Committee from 2020-2022. He was Technical Program Co-chair for 2021 IEEE VCIP, 2011 IEEE VCIP, 2017 IEEE ISPACS, and 2018 APSIPA ASC; Publication Chair for 2019 IEEE ICIP; Area Chair/Session Chair/Tutorial Speaker/Special Session Organizer for IEEE ICME, IEEE VCIP, and APSIPA ASC; and Track/Session Chair and Review Committee Member for IEEE ISCAS. He served as AEiC for Digital Communications for IEEE JETCAS and Associate Editor for IEEE TCSVT. He was Lead Guest Editor, Guest Editor and SEB Member for IEEE JETCAS, and Guest Editor for IEEE TCAS-II. He was Distinguished Lecturer of APSIPA and the IEEE CASS.