



HAL
open science

ModeNet: Mode Selection Network For Learned Video Coding

Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, Olivier Déforges

► **To cite this version:**

Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, Olivier Déforges. ModeNet: Mode Selection Network For Learned Video Coding. Machine Learning for Signal Processing (MLSP) 2020, Sep 2020, Espoo, Finland. hal-02888453v2

HAL Id: hal-02888453

<https://hal.science/hal-02888453v2>

Submitted on 24 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MODENET: MODE SELECTION NETWORK FOR LEARNED VIDEO CODINGThéo LADUNE^{*†}, Pierrick PHILIPPE^{*}, Wassim HAMIDOUCHE[†], Lu ZHANG[†], Olivier DÉFORGES[†]^{*}Orange Labs, France[†]Univ. Rennes, INSA Rennes, CNRS, IETR – UMR 6164, Rennes, France

theo.ladune@orange.com

ABSTRACT

In this paper, a mode selection network (ModeNet) is proposed to enhance deep learning-based video compression. Inspired by traditional video coding, ModeNet purpose is to enable competition among several coding modes.

The proposed ModeNet learns and conveys a pixel-wise partitioning of the frame, used to assign each pixel to the most suited coding mode. ModeNet is trained alongside the different coding modes to minimize a rate-distortion cost. It is a flexible component which can be generalized to other systems to allow competition between different coding tools. ModeNet interest is studied on a P-frame coding task, where it is used to design a method for coding a frame given its prediction. ModeNet-based systems achieve compelling performance when evaluated under the *Challenge on Learned Image Compression 2020 (CLIC20)* P-frame coding track conditions.

Index Terms— Video Coding, Autoencoder, Mode Selection

1. INTRODUCTION AND RELATED WORKS

Modern video compression systems widely adopt coding mode competition to select the best performing tool given the signal. Coding performance improvements of MPEG/ITU video codecs (AVC, HEVC and VVC) [1, 2, 3] are mainly brought by increasing the number of coding modes. These modes include prediction mode (Intra/Inter), transform type and block shape. This concept allows to perform signal adapted processing.

In recent years, image coding standards such as BPG (HEVC-based image coding method) have been outperformed by neural networks-based systems [4, 5, 6]. Most neural networks-based systems are inspired by Ballé and Minnen’s works [4, 7, 8]. They rely on an Auto-Encoder (AE) architecture that maps the input signal to latent variables. Latent variables are then transmitted with entropy coding, based on a probability model conveyed as an Hyper-Prior (HP). Such systems are denoted as AE-HP systems in the remaining of the paper. AE-HP systems are learned in an

end-to-end fashion: all components being trained according to a global objective function, minimizing a trade off between distortion and rate. Training of AE-HP systems is often performed following Ballé’s method [7] to circumvent the presence of non-differentiable elements in the auto-encoder.

As learned image compression already exhibits state-of-the-art performance, learned video compression has started to attract the research community’s attention. Authors in [9, 10] proposed a method to compress Groups Of Pictures (GOPs) inspired by standard video coding methods *i.e.* by decomposing GOPs into intra frames, without dependency, and inter frames which are coded based on previously decoded frames. Intra frames are coded with AE-HP systems while Inter frames processing is widely inspired by classical codecs approaches, replacing traditional coding tools by neural networks. First, motion vectors (representing the motion between the current frame and the reference frames), are estimated by an optical flow network [11, 12]. Motion vectors are encoded using a AE-HP system and used to perform a prediction of the current frame. Finally, the residue (prediction error) is computed either in image or latent domain and coded using an other AE-HP system. Liu *et al.* [13] tackle a similar problem and show that using a single network for both flow estimation and coding achieves performance similar to HEVC.

Although learned video coding already demonstrates appealing performance, it does not exploit all usual video coding tools. Particularly, inter frames are fully transmitted through motion compensation and residual coding even though it may not be the best option. This is different from classical encoders, where inter frame coding relies on a combination of *Skip Mode* (direct copy of the motion compensated reference), intra coding and residual inter coding.

In this paper, a mode selection network (ModeNet) is proposed. Its role is to select the most suited coding mode for each pixel. ModeNet is based on a lightweight AE-HP system, which is trained end-to-end alongside the networks performing the different coding modes. It learns to assign each pixel to the coding mode that provides the best rate-distortion tradeoff. Consequently, the proposed ModeNet can be integrated seamlessly into any neural-based coding scheme to se-




Fig. 1: Architecture of the complete system.

lect the most appropriate coding mode.

ModeNet behavior and benefits are highlighted through an architecture composed of two competing modes: each pixel is either copied from the prediction (*Skip Mode* in classical codecs) or conveyed through an AE-HP coder. We show that using ModeNet achieves compelling performance when evaluated under the *Challenge on Learned Image Compression 2020* (CLIC20) P-frame coding track conditions [14].

2. PRELIMINARY

This work focuses on *P-frame coding* with two *AE-HP systems*. These two concepts are briefly summarized below.

AE-HP system This coding scheme is composed of a convolutional encoder which maps the input signal to latents and a convolutional decoder which reconstructs the input signal from quantized latents. Latents are transmitted with entropy coding based on a latents probability model. To improve performance, the probability model is conditioned on side-information [8] and/or on previously received latents [4].

P-frame coding Let $(\mathbf{x}_{t-1}, \mathbf{x}_t) \in \mathbb{R}^{2 \times C \times H \times W}$ be the previous frame and the frame to be coded, respectively. C , H and W denote the number of color channels, height and width of the image, respectively. The previous frame \mathbf{x}_{t-1} has already been transmitted and it thus available at the decoder side to be used as a reference frame $\tilde{\mathbf{x}}_{t-1}$. Since this work follows the CLIC20 P-frame coding test conditions, the coding of \mathbf{x}_{t-1} frame is considered lossless *i.e.* $\tilde{\mathbf{x}}_{t-1} = \mathbf{x}_{t-1}$. P-frame coding is the process to encode \mathbf{x}_t knowing $\tilde{\mathbf{x}}_{t-1}$. A prediction $\tilde{\mathbf{x}}_t$ of \mathbf{x}_t is made available, based on $\tilde{\mathbf{x}}_{t-1}$ and side-information (such as motion). The conditional entropy of \mathbf{x}_t and $\tilde{\mathbf{x}}_t$ verifies:

$$H(\mathbf{x}_t | \tilde{\mathbf{x}}_t) = H(\mathbf{x}_t) - I(\mathbf{x}_t, \tilde{\mathbf{x}}_t) \leq H(\mathbf{x}_t), \quad (1)$$

where H is the Shannon entropy and I is the mutual information. Thus using information from $\tilde{\mathbf{x}}_t$ allows to lower the uncertainty about \mathbf{x}_t , resulting in better coding performance. This work aims at minimizing a rate-distortion trade-off under a lossy P-frame coding objective:

$$\mathcal{L}(\lambda) = D(\hat{\mathbf{x}}_t, \mathbf{x}_t) + \lambda R, \text{ with } \hat{\mathbf{x}}_t = f(\tilde{\mathbf{x}}_t, \mathbf{x}_t), \quad (2)$$

where D is a distortion measure, $\hat{\mathbf{x}}_t$ is the reconstruction from an encoding/decoding process f with an associated rate R weighted by a Lagrange multiplier λ .

3. MODE SELECTION FOR P-FRAME CODING

3.1. Problem formulation

Let us define \mathcal{S} as a set of pixels of frame \mathbf{x}_t verifying the following inequality:

$$d(\tilde{\mathbf{x}}_t, \mathbf{x}_t; i) \leq d(\hat{\mathbf{x}}_t, \mathbf{x}_t; i) + \lambda r(\mathbf{x}_t | \tilde{\mathbf{x}}_t; i), \quad (3)$$

where $d(\cdot, \cdot; i)$ is the i -th pixel distortion and $r(\mathbf{x}_t | \tilde{\mathbf{x}}_t; i)$ the rate of the i -th pixel of \mathbf{x}_t knowing $\tilde{\mathbf{x}}_t$. The set \mathcal{S} gives the zones of \mathbf{x}_t preferably conveyed by using $\tilde{\mathbf{x}}_t$ copy (*Skip Mode*) rather than by an encoder-decoder system. \mathcal{S} is rewritten as:

$$\mathcal{S} = \{x_{t,i} \mid x_{t,i} \in \mathbf{x}_t, \ell(\tilde{\mathbf{x}}_t, \mathbf{x}_t; i) \leq \lambda\},$$

$$\text{with } \ell(\tilde{\mathbf{x}}_t, \mathbf{x}_t; i) = \frac{d(\tilde{\mathbf{x}}_t, \mathbf{x}_t; i) - d(\hat{\mathbf{x}}_t, \mathbf{x}_t; i)}{r(\mathbf{x}_t | \tilde{\mathbf{x}}_t; i)}. \quad (4)$$

The partitioning function ℓ is a rate-distortion comparator, which assigns a coding mode (either copy or transmission) to each pixel. It is similar to the RD-cost used to arbitrate different coding modes in traditional video coding. In the remaining of the article, $\bar{\mathcal{S}}$ is the complement set of \mathcal{S} , used to denote all pixels not in \mathcal{S} *i.e.* pixels for which transmission results in a better rate-distortion trade-off than copy from $\tilde{\mathbf{x}}_t$.

Hand-crafting the partitioning function ℓ is not trivial. Indeed, both the rate and the distortion of the i -th pixel depends on choices made for previous and future pixels. Classical codecs circumvent this issue by computing ℓ on blocks of pixels assumed independent from each others.

The purpose of this work is to introduce a convolutional mode selection network (ModeNet), whose role is both to indicate which pixels belong to \mathcal{S} and to convey this partitioning. This performs a pixel-wise partitioning of \mathbf{x}_t , allowing both causal and anti-causal dependencies, learned by minimizing a global rate-distortion objective function.

3.2. Proposed system

The proposed system is built around ModeNet, which learns a pixel-wise weighting α allowing to choose among two different coding methods for each pixel. Here, the two methods in competition are copying prediction pixel from $\tilde{\mathbf{x}}_t$ or coding pixels of \mathbf{x}_t using an AE-HP system (CodecNet).

An overview of the system architecture is shown in Fig. 1. ModeNet and CodecNet architecture are described in details in section 3.3. ModeNet is defined as a function m :

$$R_m, \alpha = m(\mathbf{x}_{t-1}, \mathbf{x}_t), \quad (5)$$

where $\alpha \in [0; 1]^{H \times W}$ is the pixel-wise weighting and R_m the rate needed to convey α . The pixel-wise weighting α is continuously valued in $[0; 1]^{H \times W}$ performing smooth transitions between coding modes to avoid blocking artifacts.

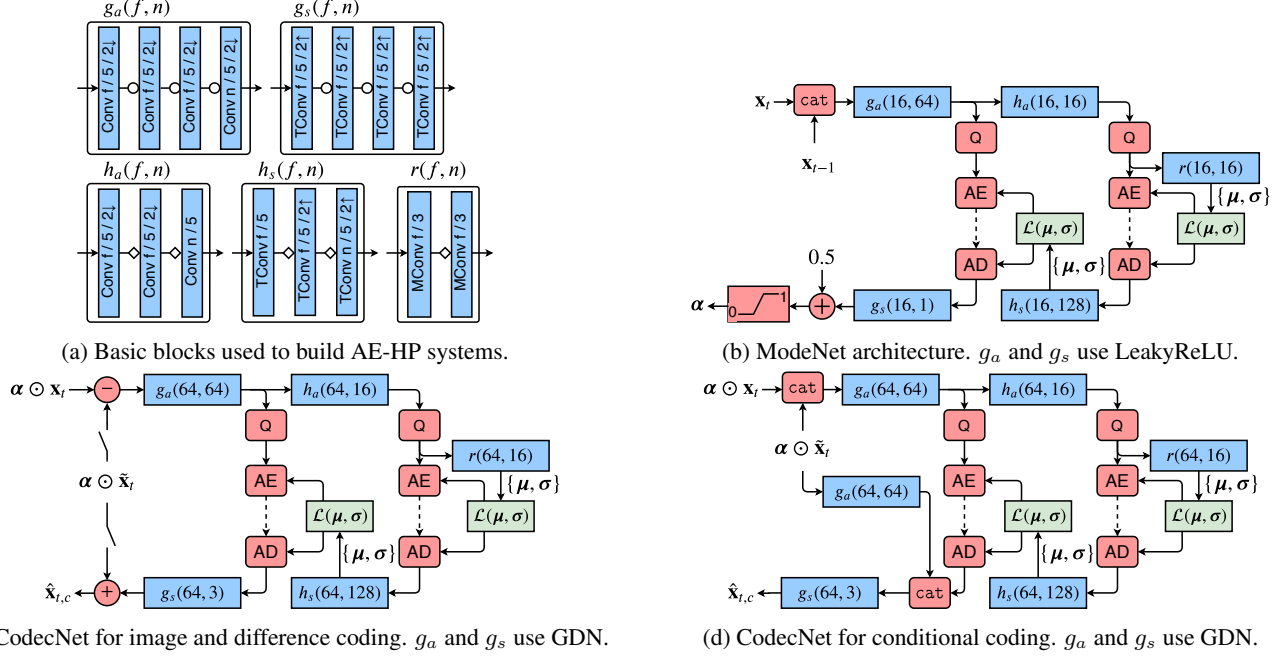


Fig. 2: Detailed architecture of all used networks. **Top left figure:** Building blocks of all subsystems. g_a and g_s are the main encoder/decoder. h_a and h_s are the hyperprior encoder/decoder. r is an auto-regressive module as in [4]. Each block is set-up by f (number of internal features) and n (number of output features). Squared arrows denote LeakyReLU, rounded arrows refer to either LeakyReLU or GDN [7]. Convolutions parameters: filters number \times kernel size / stride. TConv and MConv stand respectively for Transposed convolution and Masked convolution. cat stands for concatenation along feature axis, Q for quantization, AE and AD for arithmetic encoding/decoding with a Laplace distribution \mathcal{L} .

CodecNet is similarly defined as a function c , which codes areas \bar{S} of \mathbf{x}_t (selected through α) using information from $\tilde{\mathbf{x}}_t$:

$$R_c, \hat{\mathbf{x}}_{t,c} = c(\alpha \odot \tilde{\mathbf{x}}_t, \alpha \odot \mathbf{x}_t). \quad (6)$$

Element-wise matrix multiplication is denoted by \odot , $\hat{\mathbf{x}}_{t,c} \in \mathbb{R}^{C \times H \times W}$ is the reconstruction of $\alpha \odot \mathbf{x}_t$ and R_c the associated rate. The same α is used to multiply all C color channels. ModeNet is used to split \mathbf{x}_t between what goes through CodecNet and what is directly copied from $\tilde{\mathbf{x}}_t$, without transmission. Thus the complete system output is:

$$\hat{\mathbf{x}}_t = (1 - \alpha) \odot \tilde{\mathbf{x}}_t + c(\alpha \odot \tilde{\mathbf{x}}_t, \alpha \odot \mathbf{x}_t). \quad (7)$$

This equation highlights that the role of α is to zero areas from \mathbf{x}_t before transmission to spare their associated rate. The whole system is trained in an end-to-end fashion to minimize the rate-distortion trade-off:

$$\mathcal{L}(\lambda) = D(\hat{\mathbf{x}}_t, \mathbf{x}_t) + \lambda(R_m + R_c), \quad (8)$$

where D denotes a distortion metric. Following the CLIC20 P-frame test conditions, the Multi Scale Structural Similarity Metric (MS-SSIM)[15] is used:

$$D(\hat{\mathbf{x}}_t, \mathbf{x}_t) = 1 - \text{MS-SSIM}(\hat{\mathbf{x}}_t, \mathbf{x}_t).$$

As this work focuses on mode selection, a naive prediction $\tilde{\mathbf{x}}_t = \hat{\mathbf{x}}_{t-1} = \mathbf{x}_{t-1}$ is used. This allows to not add the burden of motion estimation to the system. Results shown in this paper would still hold when working with a more relevant prediction issued from motion compensation.

3.3. Networks architecture

Both ModeNet and CodecNet networks are built from standard AE-HP blocks described in Fig. 2a. The ModeNet role is to process the previous and current frames to transmit the pixel-wise weighting α . It is implemented as a lightweight AE-HP system (cf. Fig. 2b), with \mathbf{x}_{t-1} and \mathbf{x}_t as inputs. A bias of 0.5 is added to the output as it makes training easier. To assure that $\alpha \in [0, 1]^{H \times W}$ a clipping function is used. There are 200 000 parameters in ModeNet, which represents around 10 % of CodecNet number of parameters.

In order to transmit pixels in \bar{S} , three different configurations of CodecNet are investigated. Two of them are based on the architecture depicted in Fig. 2c. They consist in either plain image coding of \mathbf{x}_t or in difference coding of $(\mathbf{x}_t - \tilde{\mathbf{x}}_t)$ (prediction error coding). Last configuration is conditional coding denoted as $(\mathbf{x}_t | \tilde{\mathbf{x}}_t)$, shown Fig. 2d. This configuration theoretically results in better performance. Indeed, from a source coding perspective:




Fig. 3: Details on the subdivision performed by ModeNet. The pair of frames $(\mathbf{x}_{t-1}, \mathbf{x}_t)$ represents a singer moving in front of a static background. The microphone in the foreground is also motionless. Frame $\tilde{\mathbf{x}}_{t-1} = \mathbf{x}_{t-1}$ is used as prediction $\tilde{\mathbf{x}}_t$.

$$H(\mathbf{x}_t | \tilde{\mathbf{x}}_t) \leq \min(H(\mathbf{x}_t), H(\mathbf{x}_t - \tilde{\mathbf{x}}_t)). \quad (9)$$

Therefore, coding \mathbf{x}_t while retrieving all information from $\tilde{\mathbf{x}}_t$ results in less information to transmit than difference or image coding.

4. NETWORK TRAINING

All networks are trained in an end-to-end fashion to minimize the global loss function stated in eq. (8). Non-differentiable parts are approximated as in Ballé’s work [7, 8] to make training possible. End-to-end training allows ModeNet to learn to partition \mathbf{x}_t , without the need of an auxiliary loss or a hand-crafted criterion. Due to the competition between signal paths, some care is taken when training. The training process is composed of two stages:

Warm-up. Training of CodecNet only (*i.e.* ModeNet weights are frozen). Unlike copy, CodecNet is not immediately ready to process its input. Thus CodecNet has to be trained for a few epochs so the competition between copy and CodecNet is relevant.

Alternate training. Alternate training of ModeNet and CodecNet, one epoch for each (*i.e.* the other network weights are frozen).

The training set is constructed from the CLIC20 P-frame training dataset [14]. Half a million 256×256 pairs of crops are randomly extracted from consecutive frames. The batch size is set to 8 and an initial learning rate of 10^{-4} is used. The learning rate is divided by 5 at 50 % and 75 % of the training.

5. MODE VISUALISATION

This section details the processing of a pair of frames $(\mathbf{x}_{t-1}, \mathbf{x}_t)$ by the proposed system. Frames are from the sequence *CoverSong_720P-3261* extracted from the CLIC20 P-frame dataset. The system used for generating the visuals is implemented as in Fig. 1.

Figure 3a shows the inputs of ModeNet. They are encoded and decoded as the pixel-wise weighting α . The value of α tends to be zero¹ for pixels in \mathcal{S} *i.e.* when copying $\tilde{\mathbf{x}}_t$ results in a better rate-distortion cost than transmission through CodecNet. \mathcal{S} corresponds to static areas in $(\mathbf{x}_{t-1}, \mathbf{x}_t)$ as the background and the microphone, which are well captured by α . These areas are shown in Fig. 3d.

CodecNet selected inputs are $\alpha \odot \mathbf{x}_t$ and $\alpha \odot \tilde{\mathbf{x}}_t$ depicted in Fig. 3c. Copying areas of the prediction $\tilde{\mathbf{x}}_t$ allows to zero areas in \mathbf{x}_t which prevents CodecNet to spend rate for these

¹As images are in YUV format, all-zero areas appear in green




Fig. 4: Rate-distortion performance of the systems. All systems are evaluated on CLIC20 P-frame validation dataset. Quality metric is $MS\text{-}SSIM_{dB} = -10 \log_{10}(1 - MS\text{-}SSIM)$ (the higher the better). Rate is indicated in bits per pixel (bpp). Img. denotes image, Diff. difference, Cond. conditional and HEVC LP is HEVC in low-delay P configuration.

areas. Figure (3b) shows the spatial distribution of the rate in CodecNet and clearly highlights this behavior.

In this example, the rate associated to α is 0.005 bit per pixel (bpp). This shows that ModeNet is able to convey a smooth partitioning of an arbitrary number of objects for a marginal amount of rate.

6. EXPERIMENTAL RESULTS

Performance improvements brought by ModeNet are assessed on the CLIC20 P-frame validation set, under the challenge test conditions. In order to obtain RD-curves, each system is learnt with different λ . Results are gathered in Fig. 4. For the sake of brevity, systems denoted as *NN Mode X* are complete systems (*cf.* Fig. 1) composed of both ModeNet and CodecNet in coding configuration X. Similarly, systems *NN X* denotes CodecNet only system without ModeNet (*i.e.* no copy possibility: α is an all-ones matrix).

6.1. Anchors

CodecNet performance is assessed by training and evaluating it without ModeNet, meaning that \mathbf{x}_t is completely coded through CodecNet. The three configurations of CodecNet (*cf.* section 3.3 and Fig. 2) are tested. The image configuration is compared with HEVC in All Intra configuration. Difference configuration is compared with HEVC coding the pre-computed difference image. For both comparison, HEVC encodings are performed with the HM 16.20 reference software. Results in terms of MS-SSIM versus the rate are shown in Fig.

4a. CodecNet achieves consistently better performance than HEVC for both configurations across all bitrates, proving its competitiveness.

Conditional coding achieves better performance than both difference and image coding as expected from eq. (9). This shows the relevance of performing conditional coding relative to difference coding.

6.2. Performances of ModeNet-based systems

Performances of ModeNet-based systems are shown Fig. 4b. Using ModeNet increases the performance of both image and conditional coding. Image coding of \mathbf{x}_t alone does not have any information about the previous frame. Thus, adding ModeNet and the possibility of copying areas of $\tilde{\mathbf{x}}_t$ results in an important increase of the performance.

Interestingly, NN Mode Image achieves significantly better results than NN Difference. As illustrated in Fig. 3, \mathcal{S} tends to represent the areas similar in $(\tilde{\mathbf{x}}_t, \mathbf{x}_t)$, which are well handled by difference coding. Thus, performance gap between NN Mode Image and NN Difference arises on \mathcal{S} , where image coding outperforms difference coding.

An ideal conditional coder is able to retrieve all informations about \mathbf{x}_t in $\tilde{\mathbf{x}}_t$ making $\tilde{\mathbf{x}}_t$ copy useless. However, leveraging all information in $\tilde{\mathbf{x}}_t$ is not possible for a neural network with reduced complexity. There are still areas for which $\tilde{\mathbf{x}}_t$ copy provides a smaller rate-distortion cost than transmission. Thus using ModeNet to identify them improves performance.

To better appreciate the results, HEVC low-delay P (LP) performance is presented. HEVC LP codes \mathbf{x}_t with \mathbf{x}_{t-1} as

reference frame and is able to perform motion compensation to obtain a relevant prediction. Consequently, it outperforms all other systems which are constrained to directly use \mathbf{x}_{t-1} as their prediction, without motion compensation.

Using ModeNet with the best CodecNet configuration (conditional coding) allows to decrease the rate by 40 % compared to difference coding for the whole frame. Even though this gap would decrease when working with a motion compensated prediction, we believe that using ModeNet to arbitrate between conditional coding of $(\mathbf{x}_t | \tilde{\mathbf{x}}_t)$ and copy of $\tilde{\mathbf{x}}_t$ would improve most learned video coding methods, which still uses difference coding for the whole frame.

7. CONCLUSION AND FUTURE WORKS

In this paper, we propose a mode selection network which learns to transmit a partitioning of a frame to code, allowing to choose among different coding methods pixel-wise. ModeNet benefits are illustrated under a P-frame coding task. It is shown that coding the prediction error is not necessarily the best choice and using ModeNet to select better coding methods significantly increase performance.

This paper shows that the proposed ModeNet performs a smooth partitioning of an arbitrary number of areas in a frame, for a marginal rate and complexity overhead. It can be generalized to other coding schemes to leverage competition of complementary coding modes, which is known to be one of the most powerful tools in classical video coding.

An extension of this work is to use motion information to improve the prediction process. As the proposed method outperforms residual coding, having a competitive motion compensated prediction would result in compelling performance.

8. REFERENCES

- [1] Detlev Marpe, Thomas Wiegand, and Gary J. Sullivan, "The H.264/MPEG4 advanced video coding standard and its applications," *IEEE Communications Magazine*, vol. 44, no. 8, pp. 134–143, 2006.
- [2] Gary J. Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand, "Overview of the high efficiency video coding (hevc) standard," *IEEE Trans. Cir. and Sys. for Video Technol.*, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.
- [3] S. Kim J. Chen, Y. Ye, "Algorithm description for versatile video coding and test model 8 (vtm 8)," Jan. 2020.
- [4] David Minnen, Johannes Ballé, and George Toderici, "Joint autoregressive and hierarchical priors for learned image compression," in *Conference on Neural Information Processing Systems 2018, NeurIPS*, Montréal, Canada., pp. 10794–10803.
- [5] Jooyoung Lee, Seunghyun Cho, and Seung-Kwon Beack, "Context-adaptive entropy model for end-to-end optimized image compression," in *International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, 2019*.
- [6] Jiaheng Liu, Guo Lu, Zhihao Hu, and Dong Xu, "A unified end-to-end framework for efficient deep image compression," *CoRR*, vol. abs/2002.03370, 2020.
- [7] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli, "End-to-end optimized image compression," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 2017*.
- [8] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston, "Variational image compression with a scale hyperprior," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, 2018*.
- [9] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao, "DVC: an end-to-end deep video compression framework," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, pp. 11006–11015*.
- [10] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers, "Neural inter-frame compression for video coding," in *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), 2019, 2019*, pp. 6420–6428.
- [11] Anurag Ranjan and Michael J. Black, "Optical flow estimation using a spatial pyramid network," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, 2017*, pp. 2720–2729.
- [12] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz, "Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, 2018*, pp. 8934–8943.
- [13] Haojie Liu, Han Shen, Lichao Huang, Ming Lu, Tong Chen, and Zhan Ma, "Learned video compression via joint spatial-temporal correlation exploration," *CoRR*, vol. abs/1912.06348, 2019.
- [14] Workshop and Challenge on Learned Image Compression, "<https://www.compression.cc/>," June 2020.
- [15] Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik, "Multi-scale structural similarity for image quality assessment," in *Proc. IEEE Conf. on Signals, Systems, and Computers*, 2003, pp. 1398–1402.