

# Artificial Intelligence based Video Codec (AIVC) for CLIC 2022

Théo Ladune, Gordon Clare, Pierrick Philippe, Félix Henry  
Orange Innovation, France

firstname.lastname@orange.com

## Abstract

*This paper presents the AIVC submission to the CLIC 2022 video track. AIVC is a fully-learned video codec based on conditional autoencoders. The flexibility of the AIVC models is leveraged to implement rate allocation and frame structure competition to select the optimal coding configuration per-sequence. This competition yields compelling compression performance, offering a rate reduction of -26 % compared with the absence of competition.*

## 1. Introduction

Over the last few years, deep neural networks (DNN) have been proposed to replace the handcrafted operations performed by conventional video coding algorithms (e.g. HEVC [1] and VVC [2]). Thanks to their expressiveness and their ability to be globally optimized, DNN-based coding schemes have been showing fast-paced performance improvement. For instance, the learned image coder proposed by Cheng *et al.* [3] is competitive with the image coding configuration of VVC.

Yet, the introduction of an additional temporal dimension makes video coding substantially more difficult than image coding. Even though numerous previous works address learned video coding [4–9], conventional algorithms remain the best performing approaches. This assessment is supported by the CLIC 2021 video track results, which was won by a perceptually optimized VVC submission [10].

The processing of a video frame by a conventional and a learned video codec share many similarities. Both approaches compute a prediction of the frame to be coded, based on previously transmitted frames. Then only the unpredicted part is sent. While conventional and learned codecs rely on the same basic principles, a significant difference arises when one takes a closer look. For each step of the processing, conventional approaches feature a large number of coding modes (*i.e.* different operations available), enabling the selection of the most suited operation for the signal to compress. This selection is performed at a

pixel level, allowing for fine-grain content adaptation. On the contrary, learned approaches tend to perform the same processing on all input data, which limits content adaptation.

The goal of the CLIC 2022 video track is to compress several 10-second video sequences into an allocated bits budget, while obtaining the best reconstruction quality. These videos exhibit a great variety of spatial and temporal complexity. As such, it is possible to save bits for simple signals (which require fewer bits to be properly reconstructed) and invest these bits into more complex content. While this competition is easily achieved for conventional codecs, it is more challenging for learned ones.

In this paper, we propose to use AIVC [9], an open-source learned video codec, and to adapt it for the CLIC 2022 video track. In order to cope with the wide variety of the video sequences, we carry out large-scale competition among the different possibilities offered by AIVC: multiple quality levels and flexible frame structures (intra-frame period, GOP size). Moreover, a simple filtering is added as a preprocessing option to target the low bitrate track. In short, this paper illustrates how to obtain the best results from a fully learned video codec in a practical context.

## 2. The AIVC codec

This section describes AIVC, an end-to-end learned video codec. More technical details are provided in [9], open-source models and several illustrations are available at <https://orange-opensource.github.io/AIVC/>.

### 2.1. Architecture overview

AIVC performs the coding of one video frame  $\mathbf{x}_t$ , while exploiting information from up to two already transmitted (*i.e.* available at the decoder) reference frames  $\hat{\mathbf{x}}_p$  and  $\hat{\mathbf{x}}_f$ . AIVC architecture is presented in Fig. 1. It is based on two *conditional* autoencoders: MNet and CNet. Unlike the usual analysis-synthesis autoencoder, a conditional autoencoder [11] features a decoder-side conditioning transform. This conditioning transform aims to compute relevant information from decoder-side quantities, which reduces the amount of bits sent from the encoder to the decoder.

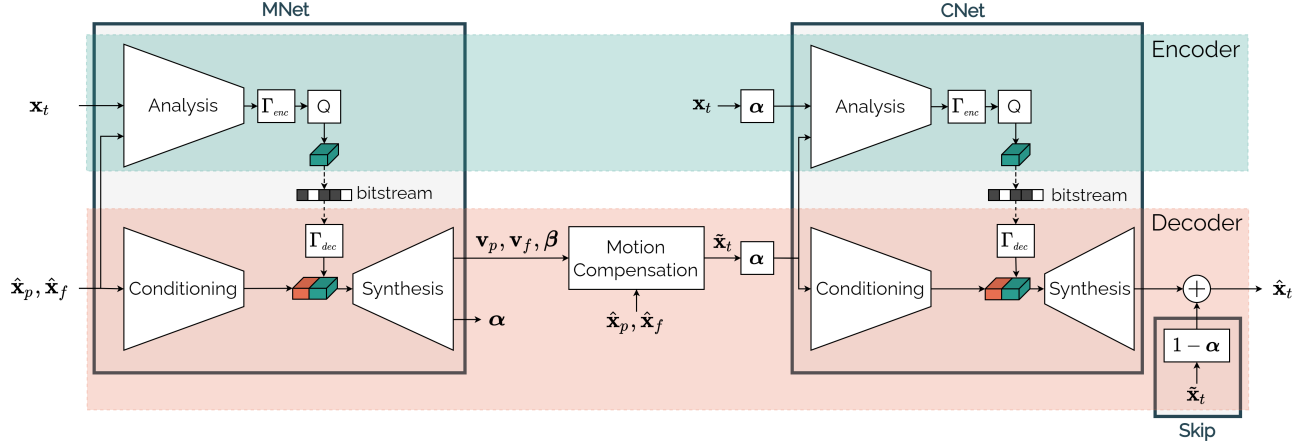


Figure 1. Block diagram of AIVC.  $\Gamma_{enc}$  and  $\Gamma_{dec}$  are feature-wise quantization gains as proposed in [9].

MNet stands for motion network and is responsible for estimating and conveying motion information. This information is then used by a motion compensation algorithm to obtain a temporal prediction  $\tilde{x}_t$ . Motion information is composed of two pixel-wise motion maps  $\mathbf{v}_p$  and  $\mathbf{v}_f$  (one for each reference frame) as well as a bi-directional prediction weighting  $\beta$ . Furthermore, MNet also computes and conveys a coding mode selection  $\alpha$ , which arbitrates between two possible coding modes.

The first coding mode is called Skip mode. It is a direct copy of the temporal prediction, with no transmission associated. Consequently, Skip mode is particularly dedicated to the stationary areas, well predicted enough to be directly copied. The second coding mode relies on CNet to transmit the areas of  $\mathbf{x}_t$  not present in the prediction  $\tilde{x}_t$  *i.e.* the areas for which the prediction is not accurate enough. In the end, both coding mode contributions are summed up to obtain the decoded frame  $\hat{x}_t$ .

The analysis and synthesis architecture proposed in [3] is used for MNet and CNet. The conditioning transform replicates the analysis architecture. These transforms are based on convolutional layers with attention modules and residual blocks. Moreover, AIVC relies on hyperpriors to perform the entropy coding. The AIVC encoder has around 17 million parameters. The AIVC decoder has 34 million parameters which requires 136 MBytes of storage.

## 2.2. Training

The purpose of the training stage is to prepare the model to compress frames with zero, one or two reference frames available. To this end, a training example consists of the successive coding of 3 frames (see Fig. 2). This coding configuration features an I-frame (no reference), a P-frame (one reference) and B-frame (two references). The model

parameters are learned through a stochastic gradient descent which aims to minimize the following loss function over the 3 frames:

$$\mathcal{L}_\lambda = \sum_{t=0}^2 D(\mathbf{x}_t, \hat{\mathbf{x}}_t) + \lambda(R_m + R_c). \quad (1)$$

In the above equation,  $D$  denotes the distortion function between the original and compressed frame. It is implemented based on the MS-SSIM. The hyperparameter  $\lambda$  balances the distortion with the sum of MNet rate  $R_m$  and CNet rate  $R_c$ . The whole model is learned end-to-end from scratch and 8 different hyperparameters  $\lambda$  are used to obtain 8 encoder-decoder pairs, addressing various qualities.




Figure 2. Training configuration.

## 2.3. Decoder conditioning

During the encoding and decoding process, the latent variables of AIVC are mapped to a binary code using the arithmetic coder `torchac` [12]. This step is particularly sensitive to a potential drift due to floating point computations. Such a drift might arise as the encoding and decoding happen on two different devices. To circumvent this, the decoder is made resilient to cross-platform encoding/decoding through a light quantization of its internal parameters.

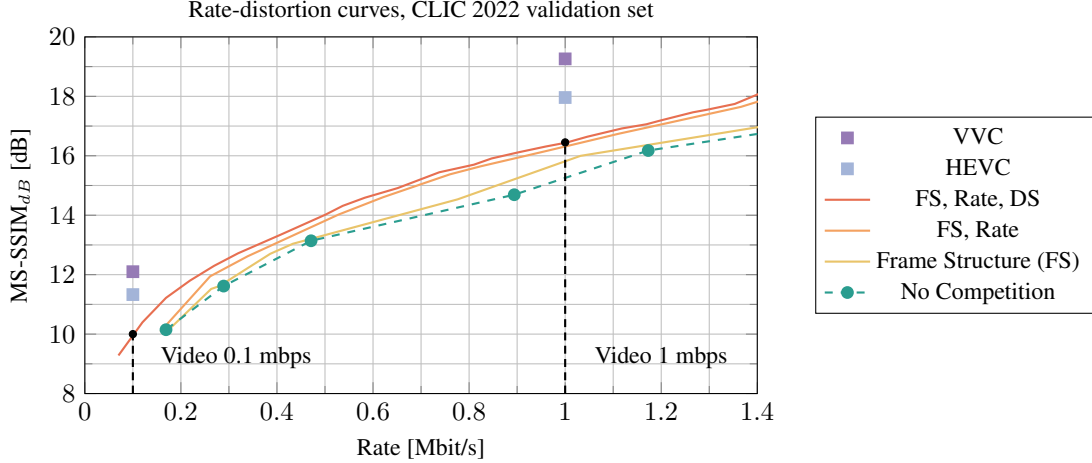


Figure 3. Benefits of per-sequence competition. FS stands for frame structure competition and DS denotes the possibility of downsampling. The black dots indicates the AIVC submissions to the challenge. Quality is measured as  $MS\text{-}SSIM_{dB} = -10 \log_{10}(1 - MS\text{-}SSIM)$ .

### 3. Coding choices competition

A per-sequence competition among the different coding choices of AIVC is implemented to obtain better quality at both challenge rate targets. Given a set of coding choices  $\mathcal{C} = \{c_1, \dots, c_N\}$ , AIVC performs the coding of the video sequence with each coding choice  $c_i$  to obtain its rate-distortion cost:

$$J_{c_i}(\lambda) = D_{c_i} + \lambda R_{c_i}, \quad (2)$$

Where  $D$  denotes the distortion (based on the MS-SSIM),  $R$  the rate and  $\lambda$  an external rate constraint. Finally, the optimal coding choice  $c^*$  is chosen in order to minimize the rate-distortion cost of the video sequence:

$$c^* = \underset{c_i \in \mathcal{C}}{\operatorname{argmin}} J_{c_i}(\lambda). \quad (3)$$

**Frames structure**—The CLIC sequences exhibit various temporal behaviors: some sequences present important motions while others have large motionless areas. In order to take into account this diversity, different frame structures are tested by the AIVC encoder. The ability of AIVC to process I, P and B frames is leveraged by evaluating several intra periods (distance between I frames) and GOP sizes (distance between references). As such, an optimal frame structure is chosen per-sequence. The green and yellow curves in Fig. 3 present the benefits of frame structure competition over the initial models with a fixed frame structure.

**Rate allocation**—Some video sequences of the CLIC 2022 dataset can be compressed into a smaller number of bits while still offering an acceptable visual quality. As such, it is relevant to save bits for these easy-to-compress sequences in order to invest those bits into more challenging sequences. In practice, each sequence is compressed

with the 8 encoder-decoder pairs and the best one is chosen sequence-wise. The orange curve in Fig. 3 presents the benefits of adding the rate allocation competition to the frame structure competition. This extended competition allows for a significant increase in quality at the 1 Mbit/s rate target. Yet, the 100 kbit/s target remains unreachable for our current systems.

**Preprocessing**—In order to compress videos at lower rates, an optional preprocessing step is added to AIVC. This preprocessing consists of a bilinear spatial downsampling of the video, which removes high-frequency content known to be notably hard to compress. The video is then upsampled at the decoder-side to recover the original resolution. The usage of the downsampling preprocessing is added as a supplementary coding choice for the sequence-wise competition. The red curve in Fig. 3 presents the benefits of adding the downsampling option to the frame structure and rate allocation competitions. While this preprocessing option is added with the low-rate target in mind, it turns out to be also slightly beneficial at higher rates.

In the end, performing per-sequence competition of the available coding modes yields a significant performance increase compared to the raw AIVC models. Indeed, using coding mode competition offers a BD-rate of  $-26\%$  *i.e.* it requires  $26\%$  less rate to achieve the same quality. This results in the submissions to the challenge validation stage shown in Table 1. The details of the different coding choices selected for both submissions is presented in Fig. 4. Compared to the high-rate submission, the low-rate one more often uses the downsampling option and relies on lower quality levels for the rate allocation.

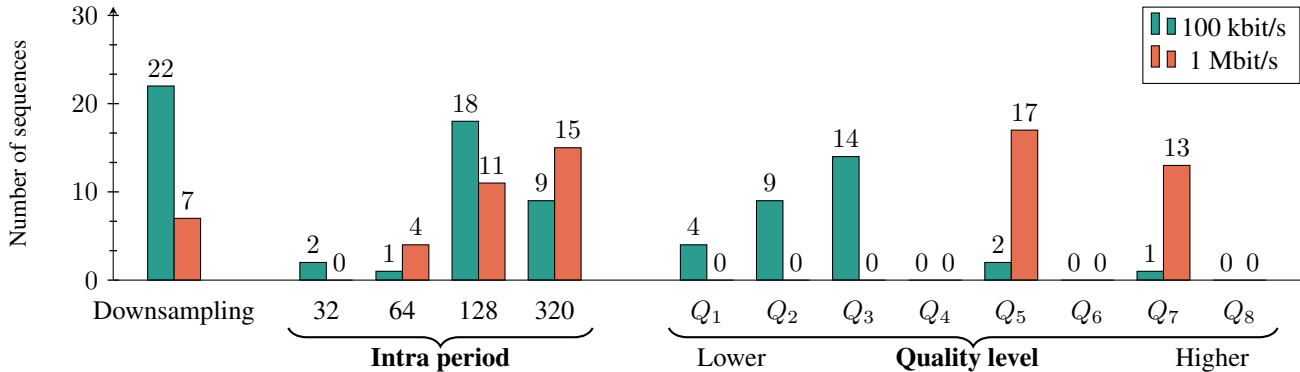


Figure 4. Coding options selected for the CLIC 2022 validation set.

Table 1. Submissions for the validation stage.

Name	Decoder size [MBytes]	Data size [MBytes]	PSNR [dB]	Decoding time [s]
AIVC	884.7	3.55	26.198	16 069
AIVC	884.9	36.63	30.994	25 364

Two modern video coding standards (HEVC and VVC) are used as anchors in Fig. 3. For fair comparison, these coding standards also benefit from rate allocation competition and also feature the downsampling option. While our learned codec AIVC offers promising performance, conventional approaches such as HEVC and VVC still obtain the best results. This may be due to their extensive usage of competition, where each pixel of a video can be processed in many different ways. This allows for fine-grain content adaptation and leads to better coding performance.

#### 4. Limitations

Even though AIVC results show that learned video coding is a promising field, a few limitations still have to be overcome. The most promising area for improvement is the introduction of more content adaptation mechanisms within the codec, as discussed in Section 3. In this section, two others limitations of AIVC are discussed.

**Quality metric**—The issue of the quality metric used during the training stage is yet to be solved. For convenience, this work relies solely on the MS-SSIM to assess the quality of the decoded videos. However, recent work by Mentzer et al. [13] shows the inadequacy of traditional quality metrics (PSNR, MS-SSIM, VMAF [14], PIM [15], LPIPS [16], FID [17]) to predict the actual (subjective) quality of video content. The same work hints that an additional GAN-based distortion yields significant quality improvements. A similar conclusion is reached by the CLIC

2021 image track, where most of the top performing systems feature a GAN-based distortion term. We believe that AIVC could obtain better perceptual quality through the introduction of a GAN-based loss.

**Test-train mismatch**—A second limitation of AIVC is due to a train-test mismatch. To achieve reasonable training time, AIVC is trained using the smallest possible coding configuration which features a 2-frame intra period. For the CLIC 2022 challenge, the shortest intra period used is 32 frames and the longest is 320 frames. This causes an important change in the I frame importance between the training stage and the test stage. Consequently, AIVC learns a rate allocation strategy targeting a 2-frame intra period, which largely deviates from the test stage. This inconsistency may have an impact on the rate-distortion performance. This needs to be solved by using a bigger training configuration, at the cost of a longer training time.

#### 5. Conclusion

This paper presents the AIVC submission to the CLIC 2022 video track. AIVC is a fully learned video codec, based on two conditional autoencoders. Thanks to its flexibility in frame structure and rate allocation, AIVC allows the performing of sequence-wise competition. Consequently, the optimal coding configuration is used for each sequence. This results in a rate reduction of 26 %.

By design, learned video codecs offer some interesting features for subjective quality as they do not operate through block-based operations. Yet, learned codecs are still outperformed by conventional video coding algorithms (HEVC, VVC) for objective metrics such as MS-SSIM. We strongly believe that the importance of competition in conventional systems allows for more content adaptation, resulting in state-of-the-art performance. This hints at relevant future works for learned coding, which should foster competition and content adaptation through the introduction of additional coding modes.

## References

- [1] Gary J. Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 2012. 1
- [2] Benjamin Bross, Jianle Chen, Jens-Rainer Ohm, Gary J. Sullivan, and Ye-Kui Wang. Developments in international video coding standardization after avc, with an overview of versatile video coding (VVC). *Proceedings of the IEEE*, 2021. 1
- [3] Zhongxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1, 2
- [4] Zongyu Guo, Runsen Feng, Zhizheng Zhang, Xin Jin, and Zhibo Chen. Learning cross-scale prediction for efficient neural video compression, 2021. 1
- [5] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 1
- [6] David Alexandre, Hsueh-Ming Hang, Wen-Hsiao Peng, and Marek Domański. Deep video compression for interframe coding. In *2021 IEEE International Conference on Image Processing (ICIP)*. 1
- [7] Zhihao Hu, Zhenghao Chen, Dong Xu, Guo Lu, Wanli Ouyang, and Shuhang Gu. Improving deep video compression by resolution-adaptive flow coding. In *Computer Vision - ECCV 2020 - 16th European Conference*. 1
- [8] Zhihao Hu, Guo Lu, and Dong Xu. FVC: A new framework towards deep video compression in feature space. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021*. 1
- [9] Théo Ladune and Pierrick Philippe. AIVC: Artificial intelligence based video codec, 2022. 1, 2
- [10] P. Philippe and T. Ladune. Coding standards as anchors for the CVPR CLIC video track. In *4th Challenge on Learned Image Compression*, Jun 2021. 1
- [11] Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, and Olivier Déforges. Conditional coding for flexible learned video compression. In *Neural Compression: From Information Theory to Applications – Workshop @ ICLR 2021*. 1
- [12] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [13] Fabian Mentzer, Eirikur Agustsson, Johannes Ballé, David Minnen, Nick Johnston, and George Toderici. Neural video compression using gans for detail synthesis and propagation. 2021. 4
- [14] Netflix. VMAF - video multi-method assessment fusion. <https://github.com/Netflix/vmaf>. 4
- [15] Sangnie Bhardwaj, Ian Fischer, Johannes Ballé, and Troy T. Chinen. An unsupervised information-theoretic perceptual quality metric. In *Annual Conference on Neural Information Processing Systems, NeurIPS 2020*. 4
- [16] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 4
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017. 4