

Hyomin Choi and Ivan V. Bajić  
School of Engineering Science, Simon Fraser University  
Burnaby, BC, Canada

**Abstract**—Video content is watched not only by humans, but increasingly also by machines. For example, machine learning models analyze surveillance video for security and traffic monitoring, search through YouTube videos for inappropriate content, and so on. In this paper, we propose a scalable video coding framework that supports machine vision (specifically, object detection) through its base layer bitstream and human vision via its enhancement layer bitstream. The proposed framework includes components from both conventional and Deep Neural Network (DNN)-based video coding. The results show that on object detection, the proposed framework achieves 13-19% bit savings compared to state-of-the-art video codecs, while remaining competitive in terms of MS-SSIM on the human vision task.

**Index Terms**—video compression, video coding for machines

## I. INTRODUCTION

Video analytics is an essential technology for various applications such as traffic monitoring, visual surveillance, and autonomous navigation. Automated machine vision pipelines increasingly analyze video streams uploaded to the cloud. If one is only interested in (machine-based) visual analysis, pre-computed features can be compressed and transmitted [1], instead of the full video. There are standard codecs to compress computed features, either handcrafted or neural network-based [2], [3], although these were developed prior to the current wave of interest on the topic. A problem with these approaches, however, is that when human viewing is needed, input video must also be coded and transmitted, and the overall system becomes less efficient.

A recent trend in image/video coding is to utilize deep neural networks (DNNs) to replace either specific units within conventional codecs, or even the whole codec. Over the past few years, DNNs have made inroads in this area, often demonstrating promising coding results compared to fully-engineered conventional approaches [4]–[8]. However, most DNN-based codecs have focused on compression for human vision, just like traditional codecs. At the same time, there are many DNN-based vision analysis methods [9] but they are usually developed without regard for compression. To establish a consolidated framework that supports both human and machine vision, a standardization activity, MPEG-VCN (Video Coding for Machines) [10], has recently been initiated. Several recent proposals for image coding [11]–[13] examined scalable compression for multiple tasks. In these methods, the base layer features are used to perform machine vision. With

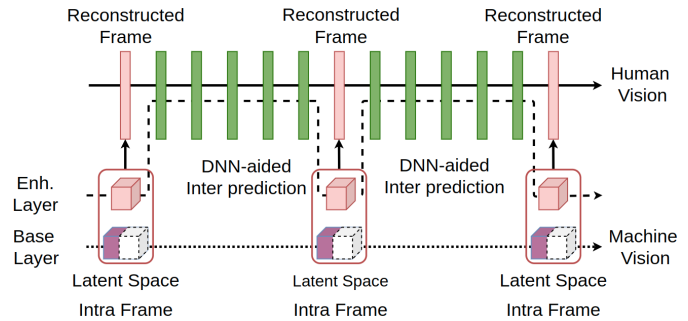


Fig. 1. Proposed video coding for humans and machines.

additional information in the enhancement layer, these methods also support high-quality input reconstruction for human vision. Meanwhile, [14], [15] tried to recover the input image directly from features, without an additional bitstream: [14] from intermediate-layer activations of YOLOv2 [16], and [15] from cuboidal features targeted at YOLOv2. If  $\mathbf{X}$  denotes the input image,  $\mathbf{Y}$  the latent space features,  $\hat{\mathbf{X}}$  the reconstructed image, and  $T$  the machine task output, then a typical machine vision pipeline can be described by a Markov chain  $\mathbf{X} \rightarrow \mathbf{Y} \rightarrow \hat{\mathbf{X}} \rightarrow T$ . Applying the data processing inequality [17] to this chain, we obtain  $I(\mathbf{Y}; \hat{\mathbf{X}}) \geq I(\mathbf{Y}; T)$ , where  $I(\cdot; \cdot)$  is the mutual information, suggesting that less information (fewer bits) is needed for machine vision than for input reconstruction.<sup>1</sup> This agrees with the scalable approaches, where both base and enhancement layers are used for input reconstruction, but only the base layer (i.e., fewer bits) for machine vision.

In this paper, we develop the first (to our knowledge) video codec for human and machine vision based on the concept of *latent-space scalability* [18]. The codec utilizes multi-task DNN-based compression for intra coding and a combination of DNN and conventional techniques for inter-frame coding. In Section II, we present the proposed video coding system, along with the explanation of its building blocks for intra and inter coding. Experimental results are presented in Section III, followed by conclusions in Section IV.

## II. PROPOSED METHODS

### A. Group-of-Pictures structure

The Group-of-Pictures (GoP) structure of our proposed video coding system is shown in Fig. 1. The GoP consists

This work was partially supported by the Natural Sciences and Engineering Council (NSERC) of Canada.

<sup>1</sup>Intuitively, we don't need the details of every pixel in order to detect objects.

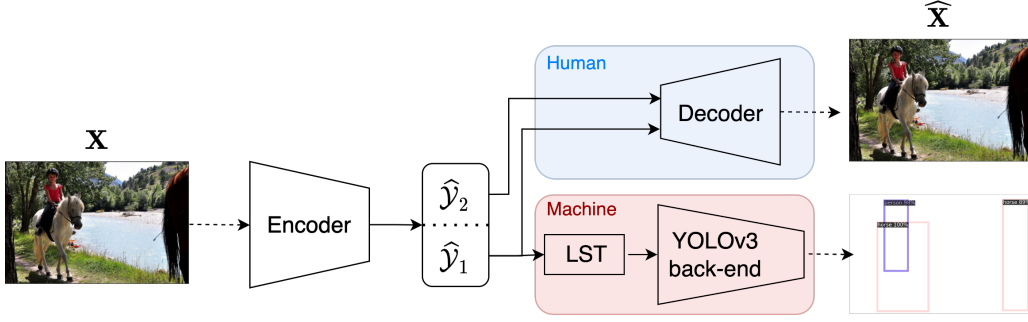


Fig. 2. Scalable multi-task intra coding.

of an intra-coded frame (shown as pink in Fig. 1) and a number of subsequent inter-coded frames (shown as green in Fig. 1). Intra frames are coded in a multi-task scalable manner using the concept of latent-space scalability, as detailed below in Section II-B. The base layer of an intra-frame bitstream supports machine vision, in our case object detection by YOLOv3 [19]. When only machine vision is needed, only the base layer bitstream needs to be decoded. The enhancement layer provides additional information for high-quality input reconstruction for human viewing. Inter frames are exclusively used for human viewing. Hence, inter-frame bitstreams are considered a part of the enhancement layer.

The intended application of such a system is as follows. Under normal operation, only the machine vision task is active. For example, a cloud-based automated analysis pipeline is monitoring an airport lounge through a surveillance camera. The camera sends the base-layer bitstream to the cloud, where the machine vision pipeline performs object detection. When a situation of interest is detected – for example, unattended luggage<sup>2</sup> – enhancement layer is requested by the cloud. This means that both the base and the enhancement layer bitstream of the next intra frame are sent to the cloud, together with subsequent inter frames. Security personnel can then view the scene and decide when to switch back to machine-only mode.

This example was given as a specific illustration of the intended use of the proposed system. The use is not limited to public area surveillance – similar situations arise in traffic monitoring, smart home security, and so on.

### B. Multi-task scalable intra coding

The structure of the multi-task scalable intra frame encoder follows that of [13], and is illustrated in Fig. 2. Intra frame  $\mathbf{X}$  is encoded into a latent representation  $\hat{\mathbf{Y}}$  as:

$$\mathbf{Y} = g_a(\mathbf{X}), \quad \hat{\mathbf{Y}} = \mathcal{Q}(\mathbf{Y}), \quad (1)$$

where  $g_a(\cdot)$  is a trainable analysis transform [4]–[6] and  $\mathcal{Q}(\cdot)$  is the quantizer, implemented as rounding to the nearest integer. The architecture of the encoder can be taken from any end-to-end trainable image codec; we have used the architecture from [6]. The main difference is that we retrain the encoder

to produce a structured latent space  $\hat{\mathbf{Y}} = \{\hat{\mathbf{Y}}_1, \hat{\mathbf{Y}}_2\}$ , such that  $\hat{\mathbf{Y}}_1$  contains the information needed for object detection (base layer), and  $\hat{\mathbf{Y}}_2$  contains enhancement information which, together with  $\hat{\mathbf{Y}}_1$ , enables reconstruction of the input frame. Out of the 192 channels in  $\hat{\mathbf{Y}}$ , 128 are assigned to  $\hat{\mathbf{Y}}_1$  and the remaining 64 to  $\hat{\mathbf{Y}}_2$ .

From such latent representation, two decoders – one for machine vision, the other for human viewing – can efficiently reconstruct the required information. The object detection decoder decodes only  $\hat{\mathbf{Y}}_1$ :

$$\tilde{\mathcal{F}} = \text{LST}(\hat{\mathbf{Y}}_1), \quad T = \mathcal{O}(\tilde{\mathcal{F}}), \quad (2)$$

where LST stands for the Latent Space Transform [18] that maps the encoder’s latent space into a latent space of the object detection network, and  $\mathcal{O}(\cdot)$  is the back-end of the object detection network producing detection output  $T$ . In our case, we have chosen layer 13 of YOLOv3 as the target latent space, so  $\mathcal{O}(\cdot)$  consists of all YOLOv3 layers after layer 13. Several possible architectures for the LST were presented in [13], [18], here we have used the one from [13].

Meanwhile, the decoder for human viewing decodes the entire latent space  $\hat{\mathbf{Y}} = \{\hat{\mathbf{Y}}_1, \hat{\mathbf{Y}}_2\}$ ,

$$\hat{\mathbf{X}} = \mathcal{D}(\hat{\mathbf{Y}}), \quad (3)$$

where  $\mathcal{D}$  is the decoder and  $\hat{\mathbf{X}}$  is the approximation to the original input frame  $\mathbf{X}$ . The architecture of the decoder  $\mathcal{D}$  can be taken from the same end-to-end codec where the encoder came from; in our case, this is [6].

Structuring of the latent space is achieved by training all components of the encoder in Fig. 2, except  $\mathcal{O}$ , the object detection back-end, jointly using the loss function

$$\mathcal{L} = R + \lambda \cdot \text{MSE}(\mathbf{X}, \hat{\mathbf{X}}) + \lambda \cdot \gamma \cdot \text{MSE}(\mathcal{F}, \tilde{\mathcal{F}}), \quad (4)$$

where  $R$  is a rate estimate (obtained using the entropy model from [6]),  $\lambda$  controls the trade-off between rate and distortion, and  $\gamma$  controls the trade-off between latent-space distortion for object detection,  $\text{MSE}(\mathcal{F}, \tilde{\mathcal{F}})$ , and pixel-domain distortion for input reconstruction,  $\text{MSE}(\mathbf{X}, \hat{\mathbf{X}})$ . Since  $\tilde{\mathcal{F}}$  is computed only from  $\hat{\mathbf{Y}}_1$ , as shown in (2), by gradient-based training, information related to object detection will be steered only into  $\hat{\mathbf{Y}}_1$ . Meanwhile, since  $\hat{\mathbf{X}}$  is obtained from the entire latent space  $\hat{\mathbf{Y}}$ , as shown in (3), information related to input

<sup>2</sup>A piece of luggage without a human detected near it.

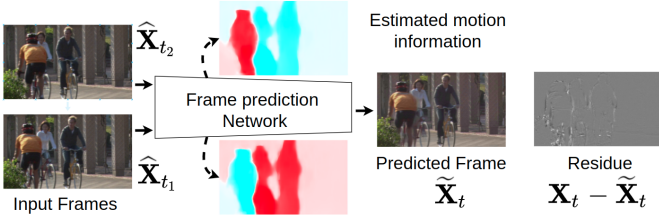


Fig. 3. Bi-directional affine transformation-based deep frame prediction from [21].

reconstruction will be spread throughout the entire latent space, including  $\hat{\mathbf{y}}_2$ .

### C. Inter-frame coding

Inter-frame coding is based on High Efficiency Video Coding (HEVC) [20], specifically HEVC test model HM-16.20.<sup>3</sup> Besides conventional HEVC coding tools, we also integrate DNN-based affine frame prediction from [21] into the coding pipeline. The operation of this network is illustrated in Fig. 3. The network takes two previously coded frames  $\hat{\mathbf{X}}_{t_1}$  and  $\hat{\mathbf{X}}_{t_2}$  as input, and predicts a half-way frame  $\hat{\mathbf{X}}_t$  between the input frames, where  $t_1 < t < t_2$ . Internally, the network estimates motion between the two input frames and the desired output frame, as visualized in the figure, along with adaptive filter kernels. With computed motion and filter kernels, the network produces an estimated frame such that the residual signal to code is minimized. The reader is referred to [21] for details about the prediction network.

The structure of the overall video decoder is shown in Fig. 4. To identify intra frames, we still use the network abstraction layer (NAL) header with 2 bytes as in conventional HEVC. Within the header data, reserved 6-bit word to support HEVC scalability is also re-used to distinguish layer IDs for our task scalability. When only the base layer of an intra frame is received,  $\hat{\mathbf{y}}_1$  is reconstructed and fed to the object detection pipeline (LST + YOLOv3 back-end) to detect objects. When the enhancement layer of an intra frame is also received,  $\hat{\mathbf{y}}_2$  can be reconstructed. Then, input frame is reconstructed from the full latent representation  $\hat{\mathbf{y}} = \{\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2\}$ . The reconstructed frame is registered in decoded picture buffer (DPB) so that it can be used as a reference for inter frames.

For inter-frame coding, the coded bitstream is parsed through the entropy decoder followed by inverse quantization and transformation. Reconstructed residual signal is added to the predictor to reconstruct the input frame. For block-level prediction in the inter frame, there are conventional HEVC intra and inter prediction tools. The DNN-predicted frame can be used in the HEVC coding pipeline in several ways. One is the “direct” mode, in which the frame is used as an additional prediction mode, indicated by a flag bit. In this case, no additional motion information is needed. As a result, this mode ends up being selected up to 57% of the time in the HEVC rate-distortion (RD) optimization, according to [21].

<sup>3</sup>[http://hevc.hhi.fraunhofer.de/svn/svn\\_HEVCSoftware/tags/HM-16.20+SCM-8.8](http://hevc.hhi.fraunhofer.de/svn/svn_HEVCSoftware/tags/HM-16.20+SCM-8.8)

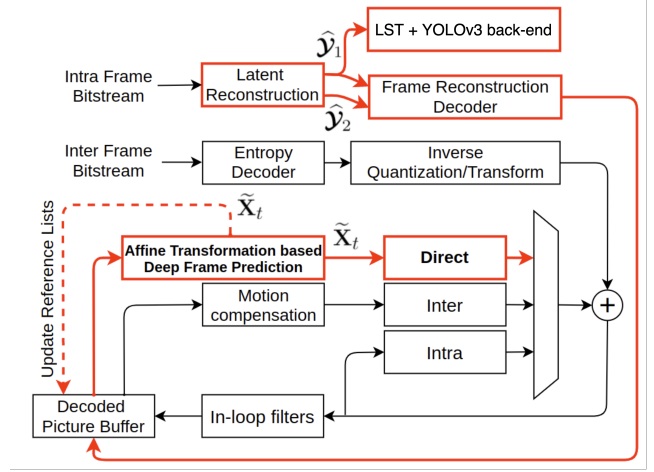


Fig. 4. Block diagram of the overall video decoder. DNN-based components are shown in red.

However, this mode is constrained to use square blocks. For this reason, we supplement the pipeline by another approach indicated by the dashed red line in Fig. 4. Here, the frame with the largest picture order count (POC) difference from the current frame in the DPB is replaced by the DNN-generated frame  $\hat{\mathbf{X}}_t$ . As such, the DNN-generated frame becomes a reference frame in the DPB. All inter-coding modes can be used in this case, but additional motion information may be required, as is the case with conventional inter-prediction. For block-level coding in inter frames, all coding modes, including the modes utilizing the DNN-generated frame  $\hat{\mathbf{X}}_t$ , compete in the HEVC RD optimization process.

## III. EXPERIMENTS

### A. Implementation and training

All DNNs are implemented in Pytorch. Python embedding library<sup>4</sup> is used to embed the DNNs into HM-16.20 implemented in C++. While compressing input video, forward operations of the DNNs to perform the multi-task compression and the frame prediction are executed across CPU and GPU.

The intra-frame codec (Fig. 2) was trained on the CLIC,<sup>5</sup> JPEG AI [22], and VIMEO-90K [23] datasets, on randomly cropped patches of size  $256 \times 256$ , following the procedure in [13]. Adam optimizer with a learning rate of  $10^{-4}$  was used. The codec was trained with  $\gamma = 0.006$  in (4), for six values of  $\lambda$  shown in Table I. The DNN-based frame prediction network (Fig. 3) was trained on randomly selected triplets of  $152 \times 152$  patches cropped from the videos from Xiph<sup>6</sup> and VIMEO-90K [23], following the procedure in [21]. AdaMax optimizer with a learning rate of  $10^{-3}$  was used.

To use the pre-trained YOLOv3 [19] back-end in the evaluation of object-detection performance, the input resolution was resized to  $512 \times 512$  using bilinear interpolation without letter-boxing. Compression benchmarks are the two latest video

<sup>4</sup><https://docs.python.org/3/extending/>

<sup>5</sup><http://www.compression.cc/>

<sup>6</sup><https://media.xiph.org/video/derf/>

TABLE I  
λ VALUES FOR TRAINING INTRA-FRAME CODING MODELS

Quality Index	1	2	3	4	5	6
λ	0.0018	0.0035	0.0067	0.013	0.025	0.0483

TABLE II  
COMBINATIONS OF MODEL INDEX (λ) AND QP FOR INTRA AND INTER FRAME, RESPECTIVELY, IN RANDOM ACCESS CODING

Intra	Model index ( $\lambda$ )	6 (0.0483)		5 (0.025)		4 (0.013)	
Inter	QP	18	22	26	30	34	38

coding standards: HEVC [20] (specifically, HM-16.20) and Versatile Video Coding (VVC) [24] (specifically, VTM-10.0). We encode the test sequences using the all-intra and random access configurations with intra period of 8. For benchmarks, we encode the sequences with  $QP \in \{18, 22, 26, 30, 34, 38\}$  and  $QP \in \{20, 24, 28, 32, 36, 40\}$  for HEVC and VVC, respectively. For our proposed system in the all intra configuration, we encode the sequences using the six models whose λ values are shown in Table I. For the random access case, to achieve the range of bit rates comparable to the benchmarks, we used several combinations of λ values for the intra codec and QP for inter coding, as shown in Table II. Inter frames are coded with  $QP \in \{18, 22, 26, 30, 34, 38\}$  plus the QP offsets related to the hierarchical reference structure defined in the HEVC Common Test Conditions (CTC) [25]. Bjøntegaard Delta (BD) metrics [26], [27] are used to evaluate the performance against the benchmarks in terms of rate-distortion and rate-accuracy.

### B. Simultaneous evaluation for human and machine vision

First, we evaluate the performance of our system against the benchmarks simultaneously on human and machine vision. We do this on the SFU-HW-Objects-v1 dataset [28], which contains COCO<sup>7</sup>-style object labels for a set of HEVC raw video test sequences. This dataset is also being used in MPEG-VCM [29]. Table III summarizes the performance of our coding system versus the benchmarks, with best results indicated in bold. Since this experiment involves object detection, for which our system uses only the base layer of the intra-coded frames, the test is carried out in the all-intra configuration. Benchmark codecs code intra frames, and decoded frames are fed to YOLOv3. In our system, only the base layer of intra frames is decoded and fed via LST to the YOLOv3 back-end, as shown in Fig. 2.

Mean Average Precision (mAP) [19] is used as the object detection accuracy metric. Unlike the Peak Signal-to-Noise-Ratio (PSNR), mAP vs. bit rate curves are not always concave, or even monotonic [29], which makes it impossible to compute a valid BD-rate-mAP value. One example is shown in Fig. 5(a) for the sequence FourPeople, where we see that HEVC and VVC curves are non-concave and non-monotonic. For this

TABLE III  
BD PERFORMANCE OF THE PROPOSED VIDEO CODING SYSTEM AGAINST HEVC AND VVC IN THE ALL-INTRA CONFIGURATION

Benchmark		HEVC (HM-16.20)			VVC (VTM-10.0)		
		Machine Vision	Human Vision		Machine Vision	Human Vision	
Class	Sequence	BD-rate-			BD-rate-		
		mAP	PSNR	MS-SSIM	mAP	PSNR	MS-SSIM
A	PeopleOnStreet Traffic	<b>-37.17%</b> 33.82%	8.55% 16.80%	<b>-22.93%</b> <b>-20.72%</b>	<b>-29.52%</b> 61.09%	36.47% 44.38%	<b>-6.34%</b> <b>-4.09%</b>
	Average	<b>-1.68%</b>	12.67%	<b>-21.83%</b>	15.78%	40.42%	<b>-5.21%</b>
B	BQTerrace	16.37%	29.84%	<b>-18.33%</b>	<b>-2.26%</b>	73.32%	7.84%
	BasketballDrive	<b>-49.91%</b>	24.57%	<b>-13.63%</b>	<b>-47.16%</b>	64.10%	9.47%
	Cactus	<b>-30.68%</b>	20.79%	<b>-19.18%</b>	<b>-46.64%</b>	55.70%	2.28%
	Kimono	<b>-75.00%</b>	1.37%	<b>-15.72%</b>	<b>-70.98%</b>	24.91%	0.74%
	ParkScene	<b>-35.81%</b>	14.63%	<b>-16.45%</b>	<b>-20.30%</b>	40.05%	<b>-0.63%</b>
	Average	<b>-35.01%</b>	18.24%	<b>-16.66%</b>	<b>-37.47%</b>	51.62%	3.94%
C	BQMall	<b>-51.04%</b>	1.07%	<b>-20.80%</b>	<b>-51.96%</b>	31.80%	0.95%
	BasketballDrill	<b>-37.45%</b>	0.62%	<b>-22.76%</b>	<b>-46.88%</b>	46.70%	5.09%
	PartyScene	<b>-8.01%</b>	15.60%	<b>-12.54%</b>	<b>-12.25%</b>	43.87%	5.33%
	RaceHorses	27.07%	8.49%	<b>-11.43%</b>	<b>-36.60%</b>	38.90%	8.37%
	Average	<b>-17.36%</b>	6.44%	<b>-16.88%</b>	<b>-36.92%</b>	40.32%	4.94%
D	BQSquare	<b>-6.51%</b>	7.39%	<b>-25.10%</b>	<b>-15.38%</b>	32.52%	<b>-10.52%</b>
	BasketballPass	<b>-57.82%</b>	<b>-2.33%</b>	<b>-16.14%</b>	<b>-55.58%</b>	29.18%	6.82%
	BlowingBubbles	<b>-15.49%</b>	1.08%	<b>-15.26%</b>	<b>-2.86%</b>	30.57%	5.72%
	RaceHorses	21.69%	<b>-4.15%</b>	<b>-11.10%</b>	<b>-22.45%</b>	27.46%	11.82%
	Average	<b>-14.53%</b>	0.50%	<b>-16.90%</b>	<b>-24.07%</b>	29.93%	3.46%
E	Johnny	116.35%	7.87%	<b>-19.50%</b>	86.62%	47.54%	7.45%
	KristenAndSara	<b>-39.08%</b>	7.48%	<b>-29.17%</b>	<b>-8.03%</b>	42.40%	<b>-8.88%</b>
	Average	38.64%	6.21%	<b>-24.90%</b>	39.29%	41.19%	<b>-2.60%</b>
Avg. (A - D)		<b>-20.40%</b>	9.62%	<b>-17.47%</b>	<b>-26.65%</b>	41.33%	2.86%
Avg. (A - E)		<b>-13.45%</b>	9.05%	<b>-18.71%</b>	<b>-18.89%</b>	41.31%	1.95%

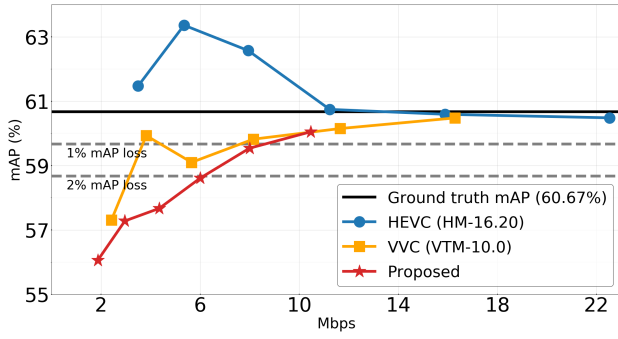
reason, the sequence FourPeople has been excluded from the results. Other sequences had well-behaved mAP vs. bit rate curves, like the one shown in Fig. 5(b) for BasketballPass.

On object detection, our coding system shows significant bit savings of 13.45% and 18.85% on average against HEVC and VVC, respectively, when averaged over all sequence classes. Without Class E sequences, average bit reduction is even higher – 20.40% and 26.65%, respectively, against HEVC and VVC. Surprisingly, we save more bits against VVC compared to HEVC, which implies that advanced coding tools adopted in VVC are less machine vision-friendly.

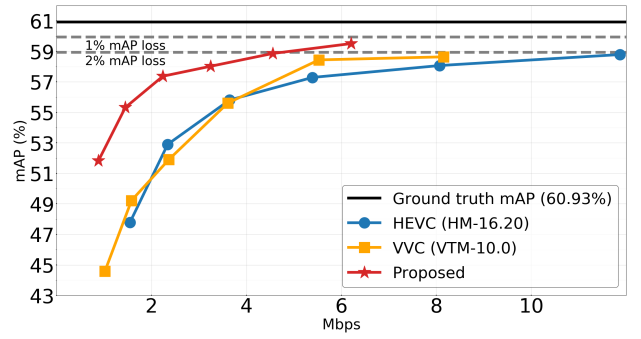
In terms of input reconstruction for human viewing, standard codecs perform better, as expected, because that is what they are optimized for. In terms of BD-rate-PSNR, our system increases bits by about 9% and 41% against HEVC and VVC, respectively. In other words, the compression efficiency of VVC is far superior to other methods in terms of rate-PSNR. Meanwhile, our system performs reasonably well against HEVC, achieving some gains on two sequences in Class D. Overall, our method shows better performance in classes C and D compared to other classes. We suspect that this is due to input scaling with bilinear interpolation, which may cause some artifacts to the sequences in Class C ( $832 \times 480$ ) and D ( $416 \times 240$ ) compared to sequences with higher resolution.

In terms of BD-rate-MS-SSIM, our system outperforms HEVC by 18.71% and has marginally worse performance (by 1.95%) compared to VVC. If we consider MS-SSIM a more relevant metric for human viewing experience, then one could argue that our system provides comparable or better performance for human viewing while achieving gains on machine vision. Indeed, it has been known for a while that DNN-based codecs do well on MS-SSIM, and our system

<sup>7</sup><https://cocodataset.org>



(a) FourPeople



(b) BasketballPass

Fig. 5. Examples of rate-mAP curves: (a) shows the case when these curves are non-concave and non-convex. Moreover, there are no overlaps on the mAP axis between the blue curve and the other two curves. Hence, BD-rate-mAP cannot be reliably computed. (b) shows the case where the curves rate-mAP curves have similar characteristics as the rate-PSNR curves, so BD-rate-mAP can be reliably computed.

TABLE IV  
INPUT RECONSTRUCTION PERFORMANCE OF THE PROPOSED VIDEO CODING SYSTEM AGAINST HEVC AND VVC IN THE RANDOM ACCESS CONFIGURATION WITH THE INTRA PERIOD OF 8

Benchmark		HEVC (HM-16.20)		VVC (VTM-10.0)	
Class	Sequence	BD-rate (PSNR)	BD-rate (MS-SSIM)	BD-rate (PSNR)	BD-rate (MS-SSIM)
A	PeopleOnStreet	-1.27%	-12.15%	20.82%	9.41%
	Traffic	21.88%	8.90%	48.65%	33.31%
	Average	10.30%	-1.63%	34.74%	21.36%
B	BQTerrace	21.70%	3.32%	55.15%	32.94%
	BasketballDrive	5.85%	-2.02%	42.65%	31.89%
	Cactus	16.54%	-1.89%	49.58%	27.42%
	Kimono	0.50%	-9.96%	29.06%	14.88%
	ParkScene	14.13%	0.86%	39.48%	23.98%
	Average	11.74%	-1.94%	43.18%	26.22%
C	BQMall	3.14%	-9.64%	40.89%	22.20%
	BasketballDrill	10.91%	-4.05%	56.60%	54.33%
	PartyScene	12.99%	-0.45%	43.24%	24.76%
	RaceHorses	4.23%	-1.58%	37.94%	31.42%
	Average	7.82%	-3.93%	44.67%	33.18%
D	BQSquare	7.38%	-9.49%	50.49%	19.02%
	BasketballPass	-2.86%	-9.68%	36.77%	23.01%
	BlowingBubbles	4.18%	-6.94%	39.37%	21.03%
	RaceHorses	-2.71%	-4.75%	38.38%	31.18%
	Average	1.50%	-7.71%	41.25%	23.56%
E	FourPeople	11.52%	-11.51%	45.47%	13.16%
	Johnny	17.84%	-2.49%	62.58%	32.28%
	KristenAndSara	14.26%	-16.50%	53.67%	11.36%
	Average	14.54%	-10.17%	53.90%	18.94%
Avg. (A - D)		7.77%	-3.97%	41.94%	26.72%
Avg. (A - E)		8.90%	-5.00%	43.93%	25.42%

benefits from DNN-based intra coding in this experiment.

### C. Input reconstruction with random access coding

Table IV summarizes input reconstruction performance in terms of BD-rate metrics for the random access configuration with the intra period of 8. Here, benchmark codecs perform better than they did on the object detection task, because they

TABLE V  
THE EFFECT OF DNN-AIDED FRAME PREDICTION IN THE RANDOM ACCESS CONFIGURATION

Class	BD-rate-PSNR	BD-rate-MS-SSIM
A	-2.19%	-3.61%
B	0.35%	0.60%
C	-1.02%	-1.33%
D	-0.79%	-0.38%
E	-1.33%	-1.46%
Average	-0.77%	-0.86%

were optimized for this kind of use. In terms of BD-rate-PSNR, our system increases the rate by about 8.9% on average against HEVC. Recall that our inter-coding pipeline is built upon HEVC. Considering the fact that conventional scalable extensions of HEVC increase the bit rate by 15%–25% per layer [30], our scalable system for human and machine vision performs well within this margin. The performance against VVC in terms of BD-rate-PSNR is correspondingly lower, as expected, with about 44% rate increase. Our codec performs much better in terms of MS-SSIM. In fact, in this case, it provides BD-rate savings of 5%, on average, against HEVC, and the loss against VVC is now reduced to about 25%.

### D. Ablation study

Here we examine the effect of DNN-aided frame prediction within our system, by comparing the full version of the system against a stripped-down version, which does not include DNN-based frame prediction. The results are shown in Table V for the random access configuration with the intra period of 8. DNN-aided frame prediction brings 0.8%–0.9% bit savings on average, both in terms of PSNR and MS-SSIM.

### E. Break-even points

In earlier sections we saw that, compared with conventional HEVC or VVC coding, our system achieves compression gains when only the machine vision task is needed, but suffers

TABLE VI  
BREAK-EVEN POINTS AGAINST HEVC AND VVC

Benchmark	HEVC		VVC	
Metric	PSNR	MS-SSIM	PSNR	MS-SSIM
Break-even point	0.5978	1.0	0.3138	0.9064

coding loss in certain cases when input reconstruction is needed for human viewing. Thus, in practice, the question of whether or not our system will provide bit savings depends on how frequently input reconstruction is needed compared to machine vision. In this section, we quantify this trade-off in terms of the maximum fraction of time that input reconstruction is needed, on average, while still allowing compression gains for our system. We call this fraction of time the *break-even point*.

According to Table III, our system is 13.45% more efficient (i.e., uses 0.8655 the amount of bits), on average, compared to HEVC when only object detection is required. At the same time, it is 9.05% less efficient (i.e., uses 1.0905 the amount of bits), on average, when input reconstruction is needed, if reconstruction quality is measured by PSNR. Let  $t_h \in [0, 1]$  be the fraction of time that input reconstruction is needed. The amount of bits used by our system will be less than or equal to that used by HEVC if

$$(1 - t_h) \cdot 0.8655 + t_h \cdot 1.0905 \leq 1. \quad (5)$$

Solving for  $t_h$  that achieves equality in (5), we obtain the break-even point of  $t_h = 0.5978$ . That is to say, if input reconstruction is needed less than 59.78% of the time, our system will provide overall bit savings over HEVC. In fact, if input reconstruction quality is measured by MS-SSIM instead of PSNR, our system would always provide gains over HEVC, since the corresponding BD-rate-MS-SSIM is negative in Table III. Repeating the same calculation for other cases, we obtain the break-even points shown in Table VI. The smallest break-even point across all cases is 0.3138. Hence, if input reconstruction is needed less than 30% of the time, our system will provide savings even against VVC.

#### IV. CONCLUSION

We developed a new scalable video coding system supporting machine vision (object detection) in the base layer and input reconstruction for human viewing in the enhancement layer. The system was benchmarked against the two most recent video coding standards - HEVC and VVC. The results show that the proposed system provides savings of 13-19% on the object detection task. At the same time, it provides comparable or better performance in terms of MS-SSIM in the all-intra coding configuration, as well as the random access configuration against HEVC. Against VVC in the random access configuration, there is a loss of about 44% in terms of BD-rate-PSNR and about 25% in terms of BD-rate-MS-SSIM. But even against VVC, our system provides rate savings, so long as human viewing is needed less than 30% of the time.

#### REFERENCES

- [1] L. Duan, J. Liu, W. Yang, T. Huang, and W. Gao, "Video coding for machines: A paradigm of collaborative compression and intelligent analytics," *IEEE Trans. Image Process.*, vol. 29, pp. 8680–8695, 2020.
- [2] MPEG-CDVS, "Compact descriptors for visual search," 2015, ISO/IEC JTC 1 15938-13.
- [3] MPEG-CDVA, "Compact descriptors for video analysis," 2019, ISO/IEC JTC 1 15938-15.
- [4] J. Ballé, V. Laparra, and E. Simoncelli, "End-to-end optimized image compression," in *Proc. ICLR'17*, 2017.
- [5] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, "Variational image compression with a scale hyperprior," in *Proc. ICLR'18*, 2018.
- [6] Z. Cheng, H. Sun, M. Takeuchi, and J. Katto, "Learned image compression with discretized gaussian mixture likelihoods and attention modules," in *Proc. IEEE CVPR'20*, 2020.
- [7] D. Dandan, M. Zhan, C. Di, C. Qingshuang, L. Zoe, and Z. Fengqing, "Advances in video compression system using deep neural network: A review and case studies," *arXiv preprint arXiv:2101.06341*, Jan. 2021.
- [8] S. Ma, X. Zhang, C. Jia, Z. Zhao, S. Wang, and S. Wang, "Image and video compression with neural networks: A review," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 6, pp. 1683–1698, June 2020.
- [9] W. Zhang, T. Yao, S. Zhu, and A. E. Saddik, "Deep learning-based multimedia analytics: a review," *ACM TOMCCAP*, vol. 15, 2019.
- [10] M. Raffie, L. Yu, Y. Zhang, and S. Liu, "[VCM] Call for Evidence for Video Coding for Machines," in *ISO/IEC JTC 1/SC 29/WG 2, m55605*, Oct. 2020.
- [11] Y. Hu, S. Yang, W. Yang, L.-Y. Duan, and J. Liu, "Towards coding for human and machine vision: A scalable image coding approach," in *2020 IEEE International Conference on Multimedia and Expo (ICME)*, 2020.
- [12] K. Liu, D. Liu, L. Li, N. Yan, and H. Li, "Semantics-to-signal scalable image compression with learned reversible representations," *Int. J. Comput. Vis.*, pp. 1–17, Jun. 2021.
- [13] H. Choi and I. V. Bajić, "Scalable image coding for humans and machines," *IEEE Trans. Image Process.*, vol. 31, pp. 2739–2754, 2022.
- [14] H. Choi and I. V. Bajić, "Near-lossless deep feature compression for collaborative intelligence," in *Proc. IEEE MMSP*, Aug. 2018.
- [15] A. Ahmmed, M. Paul, M. Murshed, and D. Taubman, "Human-machine collaborative video coding through cuboidal partitioning," in *Proc. IEEE ICIP'21*, Sep. 2021, pp. 2074–2078.
- [16] J. Redmon and A. Farhadi, "YOLO9000: better, faster, stronger," in *Proc. IEEE CVPR'17*, Jul. 2017.
- [17] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley, 2006.
- [18] H. Choi and I. V. Bajić, "Latent-space scalability for multi-task collaborative intelligence," in *Proc. IEEE ICIP'21*, Sep. 2021, pp. 3562–3566.
- [19] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, Apr. 2018.
- [20] Int. Telecommun. Union-Telecommun. (ITU-T) and Int. Standards Org./Int/Electrotech. Commun. (ISO/IEC JTC 1), "High efficiency video coding," 2019, rec. ITU-T H.265 and ISO/IEC 23008-2, 2019.
- [21] H. Choi and I. V. Bajić, "Affine transformation-based deep frame prediction," *IEEE Trans. Image Process.*, vol. 30, pp. 3321–3334, 2021.
- [22] "JPEG AI dataset," [Online]: <https://jpeg.org/jpegai/dataset.html>.
- [23] T. Xue, B. Chen, J. Wu, D. Wei, and W. T. Freeman, "Video enhancement with task-oriented flow," *IJCV*, vol. 127, pp. 1106–1125, 2019.
- [24] Int. Telecommun. Union-Telecommun. (ITU-T) and Int. Standards Org./Int/Electrotech. Commun. (ISO/IEC JTC 1), "Versatile video coding," 2020, rec. ITU-T H.266 and ISO/IEC 23090-3, 2020.
- [25] F. Bossen, "Common HM test conditions and software reference configurations," in *ISO/IEC JTC1/SC29 WG11, JCTVC-L1100*, Jan. 2013.
- [26] G. Bjøntegaard, "VCEG-M33: Calculation of average PSNR differences between RD-curves," in *ITU-T VCEG-M33*, Apr. 2001.
- [27] C. Hollmann, S. Liu, W. Gao, and X. Xu, "[VCM] On VCM reporting template," in *ISO/IEC JTC 1/SC 29/WG 2, m56185*, Jan. 2021.
- [28] H. Choi, E. Hosseini, S. R. Alvar, R. A. Cohen, and I. V. Bajić, "A dataset of labelled objects on raw video sequences," *Data in Brief*, vol. 34, p. 106701, 2021.
- [29] C. Hollmann, J. S. amd M. Damghanian, P. Wennersten, and R. Sjöberg, "[VCM] Adjustments for the SFU-HW dataset," in *ISO/IEC JTC 1/SC 29/WG 2, m57974*, Oct. 2021.
- [30] J. M. Boyce, Y. Ye, J. Chen, and A. K. Ramasubramonian, "Overview of SHVC: scalable extensions of the high efficiency video coding standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 26, pp. 20–34, 2016.