



HAL
open science

A Hybrid Deep Animation Codec for Low-Bitrate Video Conferencing

Goluck Konuko, Stéphane Lathuilière, Giuseppe Valenzise

► **To cite this version:**

Goluck Konuko, Stéphane Lathuilière, Giuseppe Valenzise. A Hybrid Deep Animation Codec for Low-Bitrate Video Conferencing. IEEE International Conference on Image Processing, Oct 2022, Bordeaux, France. hal-03720713v1

HAL Id: hal-03720713

<https://hal.science/hal-03720713v1>

Submitted on 27 Jul 2022 (v1), last revised 29 Dec 2022 (v2)

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.

A HYBRID DEEP ANIMATION CODEC FOR LOW-BITRATE VIDEO CONFERENCING

Goluck Konuko[†], Stéphane Lathuilière[‡] Giuseppe Valenzise[†]

[†] Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes
91190, Gif-sur-Yvette, France

[‡] LTCI, Télécom Paris, Institut Polytechnique de Paris, France

ABSTRACT

Deep generative models, and particularly facial animation schemes, can be used in video conferencing applications to efficiently compress a video through a sparse set of keypoints, without the need to transmit dense motion vectors. While these schemes bring significant coding gains over conventional video codecs at low bitrates, their performance saturates quickly when the available bandwidth increases. In this paper, we propose a layered, hybrid coding scheme to overcome this limitation. Specifically, we extend a codec based on facial animation by adding an auxiliary stream consisting of a very low bitrate version of the video, obtained through a conventional video codec (e.g., HEVC). The animated and auxiliary videos are combined through a novel fusion module. Our results show consistent average BD-Rate gains in excess of -30% on a large dataset of video conferencing sequences, extending the operational range of bitrates of a facial animation codec alone.

Index Terms— Video compression, video animation, fusion module, video conferencing

1. INTRODUCTION

In the pursuit of higher video compression performance at low bitrates, deep generative models have been recently employed in image and video compression to overcome the intrinsic limitations of traditional video coding schemes in modeling complex pixel dependencies [1, 2]. In particular, for video conferencing applications, we have shown in our previously proposed Deep Animation Codec (DAC) [3] that it is possible to code videos with talking heads at bitrates as low as 5 kbps, and reconstruct them with good visual quality, using a face animation scheme [4].

Face animation employs a set of sparse keypoints to encode the motion of faces, instead of dense motion vectors used in traditional pixel-based codecs. At the decoder side, the keypoints are used to synthesize a dense motion field, which is then employed to warp a source frame (coded as an Intra frame with a conventional codec), producing a realistic, high-quality synthesized picture. A similar scheme has also been applied in concurrent or later works [5, 6].

A main shortcoming of face animation codecs is that their performance tends to saturate quickly when the available bandwidth increases, limiting the achievable video reconstruction quality. This is partly due to the open loop coding structure, which makes it difficult to model long-term pixel dependencies as well as significant displacements in the background, introducing a quality drift. A possible solution consists in adding Intra-refresh frames to reset this drift periodically [3]. However, this entails a higher bitrate, which makes this option noncompetitive compared with conventional codecs.

In this work, we solve the problem of long-term dependencies and background motion by an alternative, hybrid approach. We augment the DAC bitstream with an auxiliary stream, obtained by compressing at very low quality the original video, using a state-of-the-art conventional codec (in our work, we use HEVC). The DAC output and auxiliary decoded video are fed into a novel fusion module, which combines them to reconstruct the final result. Despite its very low bitrate (and the consequent poor visual quality), the auxiliary video provides enough information to regularize the animation and compensate for the motion estimation errors in the DAC bitstream. Using this Hybrid Deep Animation Codec (H-DAC), we obtain BDRate reductions over HEVC in excess of 30%, and similar performance to VVC, on two test datasets composed by video conferencing sequences.

2. RELATED WORK

The application of deep learning to image and video compression has received a great deal of attention in the past few years [7, 8, 9], thanks to its ability to represent complex pixel dependencies and obtain good visual quality at low bitrate [10]. Deep neural networks can be applied either to replace/improve specific coding tools (e.g., spatial prediction [11]), or to optimize the whole coding pipeline in an end-to-end fashion, typically using variational auto-encoders [12, 13]. This work falls in the former category, as we employ a deep neural network to optimize the motion prediction and compensation through face animation, but we rely on conventional entropy coding, intra-frame prediction

and an auxiliary standard video stream to reconstruct the decoded video.

Deep generative models are a family of methods that aim at learning (or sampling from) the data distribution [14]. They have been used in video compression to code a picture/video at very low bitrate, e.g., by hallucinating parts of the video outside the region of interest [15]. In this work, we focus instead on synthesizing the foreground pixels of a talking head in video conferencing. To this end, we employ image animation models, which are a specific kind of generative models able to produce realistic, high-quality videos of moving faces [16]. A typical application consists of transferring the movements of a driving sequence to a source frame [17, 18] in order to swap faces in videos and produce deep fakes [19].

More recently, image animation models have been employed in video coding to enable ultra-low bitrate video conferencing [3, 5, 20, 21]. In these works, the face motion is represented by a set of sparse keypoints, which are encoded and transmitted as bitstream. At the decoder side, the received keypoints are used to reconstruct a dense optical flow, which is then used to warp a source frame (intra coded) and produce an estimate of the reconstructed frame. Among these proposed methods, only the Deep Animation Codec (DAC) in [3] offers the possibility to vary the bitrate and quality to a certain extent, by modulating the frequency of intra refresh. In this work, we build on the basic architecture of DAC, but we extend it with an additional auxiliary stream to handle motion in the background and long-range temporal dependencies, and to increase its operational range of bitrates and qualities.

3. PROPOSED CODING METHOD (H-DAC)

Our goal at test time is to compress an input video sequence F_0, \dots, F_t corresponding to a video conferencing scenario. To this aim, we consider that at training time, we have at our disposal many videos containing talking faces.

Our codec is divided in three main modules illustrated in Figure 1. First a *conventional video codec* (light red module) is used on both encoder and decoder sides to transmit the input video with a very low bitrate. In this module, any conventional video codec can be employed (in our experiments, we adopt HEVC with a low-delay configuration).

In the second module, an image animation model (green module) is employed. This module is based on the *Deep Animation Codec* (DAC) proposed in our previous work [3]. It provides to the decoder the initial frame of the video and keypoints that are learned in an unsupervised manner in order to describe the motion between the initial frame and the current frame. Our image animation model is described in Section 3.1. Note that, at every time step t , the current frame is encoded via the current HEVC P-frame and the keypoint estimated in the current image F_t .

Finally, a *fusion module* combines the low-quality video provided by the conventional video codec with the output of

the image animation module. The fusion and image animation modules are jointly learned by minimizing a reconstruction loss on the decoder side. Practically we employ a combination of a perceptual and adversarial losses as in [18]. These three modules are detailed in the following.

3.1. Animation Module

Image animation is based on the idea that a source image depicting a single object such as human faces can be employed to create a video sequence by transferring motion encoded as the displacement of keypoints. In our video compression context, the image animation framework is used as follows: the initial frame, interchangeably referred to as the source frame or intra frame in a group of pictures is transmitted to the decoder and all the other frames are encoded via the displacement of keypoints estimated independently in every frame.

Intra frame compression. The intra frame is compressed by any off-the-shelf codec. We apply the BPG codec here, which is equivalent to HEVC Intra. The intra frame used in the animation process is the decoded frame received from the BPG codec.

Keypoint extraction, quantization and coding. As in [3], a keypoint detection network is used to predict the location of the keypoints in every frame. Each keypoint is associated with a 2×2 Jacobian matrix as described by [18], which describes the orientation of the local motion vectors. Dense optical flow maps are generated through a first-order Taylor expansion around the sparse keypoint locations. The estimated optical flow map is used to deform the features of the source frame before further refinement is applied. As in [3], we use 10 facial keypoints. The intra frame keypoints are used as reference since they are available both at the encoder and decoder side without the need for transmission. This enables an efficient 8-bit quantization process that greatly reduces the bitrate contribution of the motion keypoints. The quantized keypoint vectors are entropy-coded using a standard arithmetic codec.

3.2. Reconstruction and Fusion Module

The animation-based codec alone cannot handle well background motion and disocclusions. Conversely, the HEVC base layer provides poor texture details but can encode changes in the background or object disocclusions while being robust to complex non-rigid motions (e.g., hair movements) that are not captured by keypoints. In this section we describe our reconstruction and fusion module that leverages both streams to reconstruct a good quality frame at the decoder side.

The reconstruction and fusion module is depicted in the right yellow box in Figure 1. It processes the outputs of the baseline video codec and the animation module. It is composed by 5 blocks. The optical flow encoder-decoder network




Fig. 1: Basic scheme of our proposed hybrid codec. The bitstream is composed by the low-bitrate HEVC video, the Intra frame (IF) and the entropy-coded (EC) keypoints. The animation module employs IF, encoded by a traditional codec (e.g. HEVC Intra), to animate the subsequent inter-frames, using a set of sparse keypoints to encode motion. The fusion module concatenates the warped (\mathcal{W}) intra frame features and the low-quality HEVC frame features, and decodes a predicted output frame.

produces a dense optical flow between the intra frame and the current one based on the decoded keypoints. The convolutional encoders E_A and E_B extract features from the intra frame and the base layer decoded frame, respectively. The warping operator \mathcal{W} [4] applies motion compensation in the feature domain. The warped features from the animated frame and those obtained by E_B are concatenated, and the decoder D_F produces the final reconstructed frame after feature fusion. The source code for our coding framework is available at [goluck-konuko.github.io](https://github.com/goluck-konuko).

4. EXPERIMENTS AND RESULTS

Datasets. Following [3], we use two datasets in our experiments. The **VoxCeleb2** dataset is a large audio-visual dataset of talking heads with about 22k videos extracted from Youtube at a resolution of 256x256 pixels [18]. 90 sequences with complex motion patterns from the VoxCeleb2 test set are used for testing the compression framework. For more technical details, please refer to [3]. In addition, the **Xiph.org**¹ dataset is also used as test dataset. In this case, we employ the model trained on VoxCeleb2. We downloaded the videos from Xiph.org talking humans [3] and selected 16 sequences from which we crop the region of interest around the human face at a resolution of 256x256.

Implementation details. Regarding the network architecture, we use the motion transfer network architecture described by [4] and design a fusion module with spatial attention layers from [22]. The HEVC base layer is coded with a fixed QP value of 50 for all sequences. The network is trained in an end-to-end fashion for 100 epochs.

¹<https://media.xiph.org/video/derf/>

Table 1: Bjontegaard-Delta Performance of H-DAC over HEVC

	VoxCeleb	Xiph.org
	BD quality / BD rate	BD quality / BD rate
PSNR ↑	1.07 / -33.36	0.97 / -30.7
SSIM ↑	0.02 / -33.41	0.02 / -28.33
msVGG ↓	-19.16 / -48.84	-20.04 / -41.64

Metrics. In addition to the widely used PSNR and SSIM metrics, we adopt the msVGG loss that is the multi-scale LPIPS loss used in [4].

Comparison to state-of-the-art video codecs. We evaluate the coding framework performance with respect to the HEVC codec under a low-delay configuration. Qualitative comparison are reported as Bjontegaard-Delta rate over HEVC in Table 1. We observe that we obtain better performance in terms of BD rate for the three metrics and the two datasets. The gain in BD rate is especially clear in terms of msVGG since this metric is employed as training loss.

We show the RD curves using PSNR in Figure 3, where we also report the VVC performance for reference. We see that DAC [3] achieves very low bitrate, but its quality saturates quickly and cannot reach PSNR values above 34dB even when the bitrate increases. Conversely, the use of a conventional video codec stream shifts the range of bitrates over which our framework can operate before reaching a saturation point. Regarding the comparison with HEVC, H-DAC consistently obtains better scores at all the considered bitrates. Interestingly, our approach achieves close performance to VVC (which has a significantly higher complexity than HEVC). Notice that we did employ only HEVC in H-DAC; however, the conventional video coding module and the intra frames in



Fig. 2: Qualitative comparison of reconstructed images with our codec and state-of-the-art codecs at a similar bitrate. H-DAC significantly improves DAC in reproducing face expressions with good fidelity (e.g., the open mouth in the second row and the teeth in the third row), and is robust to non-face objects such as the hand in the first row. In addition, the synthesized images display lower distortion than HEVC and VVC.




Fig. 3: RD curve for 105 test videos sampled from VoxCeleb and Xiph.org test sets. The DAC is our coding framework proposed in [3] which uses adaptive refresh for quality scalability.

the animation module might be replaced with VVC, leading to further coding gains.

To conclude this section, Figure 2 shows some qualitative compression results. At comparable bitrates, we observe a considerable difference in reconstruction quality, especially in scenes where we may have other elements than just faces, such as the speaker’s hand within the target frame in the first row. While the conventional codecs display significant blur

(VVC) and blocking (HEVC) artifacts at this bitrate, H-DAC does not produce these distortions and synthesizes images with better perceptual quality. In particular, the resulting visual quality is better than VVC despite the RD curves computed in Figure 3. However, a more rigorous subjective study needs to be conducted to confirm this observation, and is left for future work.

5. CONCLUSION

Using deep facial animation for compression in video conferencing applications leads to significant coding gains at ultra-low bitrates. However, the performance tends to saturate quickly as the available bandwidth increases, due to a poor capability of animation schemes in handling long-term temporal dependencies, disocclusions and background changes. In this paper, we propose H-DAC, a Hybrid Deep Animation Codec that combines face animation with an auxiliary, conventional bitstream at very low bitrate. Our results demonstrate that this approach can overcome the limitations of codecs using purely face animation to synthesize frames, bringing significant quality gains compared to state-of-the-art video codecs.

Acknowledgement: This work was funded by Labex DigiCosme - Université Paris-Saclay

6. REFERENCES

- [1] P. Pad and M. Unser, "On the optimality of operator-like wavelets for sparse AR (1) processes," in *IEEE ICASSP*, 2013.
- [2] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *TIP*, 2004.
- [3] G. Konuko, G. Valenzise, and S. Lathuilière, "Ultra-low bitrate video conferencing using deep image animation," in *IEEE ICASSP*, 2020.
- [4] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, and N. Sebe, "First order motion model for image animation," in *Neurips*, 2019.
- [5] T.-C. Wang and L. M.-Y. Mallya, Arun, "One-shot free-view neural talking-head synthesis for video conferencing," in *CVPR*, 2021.
- [6] M. Oquab, P. Stock, O. Gafni, D. Haziza, T. Xu, P. Zhang, and O. Celebi, "Low bandwidth video-chat compression using deep generative models," in *CVPR*, 2021.
- [7] O. Rippel and L. Bourdev, "Real-time adaptive image compression," *arXiv preprint arXiv:1705.05823*, 2017.
- [8] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, "Variational image compression with a scale hyperprior," *arXiv preprint arXiv:1802.01436*, 2018.
- [9] E. Agustsson, M. Tschannen, F. Mentzer, R. Timofte, and L. V. Gool, "Generative adversarial networks for extreme learned image compression," in *IEEE ICCV*, 2019.
- [10] G. Valenzise, A. Purica, V. Hulusic, and M. Cagnazzo, "Quality assessment of deep-learning-based image compression," in *IEEE MMSP*, 2018.
- [11] L. Wang, A. Fiandrotti, A. Purica, G. Valenzise, and M. Cagnazzo, "Enhancing hevc spatial prediction by context-based learning," in *IEEE ICASSP*, 2019.
- [12] J. Ballé, V. Laparra, and E. P. Simoncelli, "End-to-end optimized image compression," *Neurips*, 2016.
- [13] G. Lu, W. Ouyang, D. Xu, X. Zhang, C. Cai, and Z. Gao, "Dvc: An end-to-end deep video compression framework," in *IEEE CVPR*, 2019.
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, Warde-Farley, S. D., Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS*, 2014.
- [15] A. S. Kaplanyan, A. Sochenov, T. Leimkühler, M. Okunev, T. Goodall, and G. Rufo, "Deepfovea: Neural reconstruction for foveated rendering and video compression using learned statistics of natural videos," *ACM TOG*, 2019.
- [16] H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Niessner, P. Pérez, C. Richardt, M. Zollhöfer, and C. Theobalt, "Deep video portraits," *ACM TOG*, 2018.
- [17] C. Chan, S. Ginosar, T. Zhou, and A. A. Efros, "Everybody dance now," *CoRR*, vol. abs/1808.07371, 2018.
- [18] A. Siarohin, S. Lathuilière, S. Tulyakov, E. Ricci, and N. Sebe, "Animating arbitrary objects via deep motion transfer," in *IEEE/CVF CVPR*, 2019.
- [19] X. Yang, Y. Li, and S. Lyu, "Exposing deep fakes using inconsistent head poses," in *IEEE ICASSP*, 2019.
- [20] M. Oquab, P. Stock, O. Gafni, D. Haziza, T. Xu, P. Zhang, O. Celebi, Y. Hasson, P. Labatut, B. Bose-Kolanu, et al., "Low bandwidth video-chat compression using deep generative models," *arXiv preprint arXiv:2012.00328*, 2020.
- [21] B. Chen, Z. Wang, B. Li, R. Lin, S. Wang, and Y. Ye, "Beyond keypoint coding: Temporal evolution inference with compact feature representation for talking face video compression," in *DCC*, 2022.
- [22] Z. Cheng, H. Sun, M. Takeuchi, and J. Katto, "Learned image compression with discretized gaussian mixture likelihoods and attention modules," 2020.