

ELF-VC: Efficient Learned Flexible-Rate Video Coding

Oren Rippel*, Alexander G. Anderson*, Kedar Tatwawadi, Sanjay Nair, Craig Lytle, Lubomir Bourdev
WaveOne, Inc.

{oren, alex, kedar, sanjay, craig, lubomir}@wave.one

Abstract

While learned video codecs have demonstrated great promise, they have yet to achieve sufficient efficiency for practical deployment. In this work, we propose several novel ideas for learned video compression which allow for improved performance for the low-latency mode (I- and P-frames only) along with a considerable increase in computational efficiency. In this setting, for natural videos our approach compares favorably across the entire R-D curve under metrics PSNR, MS-SSIM and VMAF against all mainstream video standards (H.264, H.265, AV1) and all ML codecs. At the same time, our approach runs at least 5x faster and has fewer parameters than all ML codecs which report these figures.

Our contributions include a flexible-rate framework allowing a single model to cover a large and dense range of bitrates, at a negligible increase in computation and parameter count; an efficient backbone optimized for ML-based codecs; and a novel in-loop flow prediction scheme which leverages prior information towards more efficient compression.

We benchmark our method, which we call ELF-VC (Efficient, Learned and Flexible Video Coding) on popular video test sets UVG and MCL-JCV under metrics PSNR, MS-SSIM and VMAF. For example, on UVG under PSNR, it reduces the BD-rate by 44% against H.264, 26% against H.265, 15% against AV1, and 35% against the current best ML codec. At the same time, on an NVIDIA Titan V GPU our approach encodes/decodes VGA at 49/91 FPS, HD 720 at 19/35 FPS, and HD 1080 at 10/18 FPS.

1. Introduction

The trends of growth of video capture and consumption are staggering. Every day, 1.5 billion hours of videos are watched across YouTube, Netflix and Facebook, and 23 million new cameras are added into circulation [44, 20, 36, 21].

In the last few years, ML-based compression algorithms have shown promise in their potential to mitigate some of this global video congestion. The ML subfield of end-to-

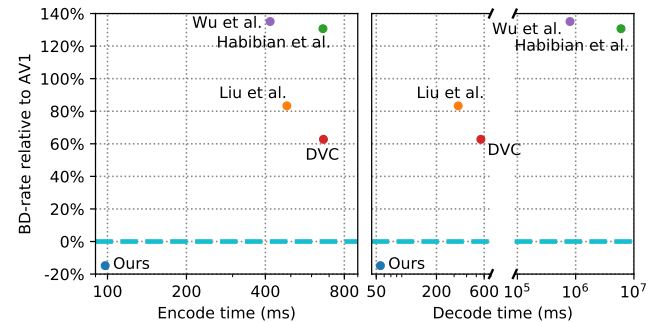


Figure 1: BD-Rate for ML-based codecs relative to AV1 as a function of encode/decode time on HD 1080 videos [41, 13, 28, 26] (UVG dataset, PSNR metric). Our approach reduces the BD-rate by 54% relative to the current fastest ML codec which reports speed [26], while running 5x faster.

end methods for image compression has grown rapidly with hundreds of papers (for instance, [4, 32, 5]) which demonstrate unequivocally that learned approaches can achieve improved coding efficiency relative to their hard-coded counterparts.

These approaches have, in turn, planted the seeds for ML-based video coding algorithms. Even though end-to-end video coding research has only taken its first few steps, it is clear that learned approaches have the potential to yield significant bitrate savings over the existing standards [33, 28, 13, 12, 27]. However, there still exists an elephant in the room: is it possible for ML-based approaches to achieve sufficient flexibility and efficiency to become practical in the real world?

We propose a new ML video codec, ELF-VC (Efficient, Learned and Flexible-Rate Video Coding) for the low-latency mode, which aims to improve three key weaknesses of ML-based video compression: bitrate flexibility, compression efficiency, and speed.

Bitrate flexibility Traditional codecs can dynamically adjust the bitrate to achieve a target bandwidth or target compression quality as a function of the complexity of the video and changing network conditions. Most existing ML codecs, however, represent each point on the R-D curve with a separate model, which is impractical due to param-

*Equal contribution

eter explosion and model loading inefficiency. In contrast, ELF is able to support a large and dense range of bitrates on a per-frame basis with a single set of parameters.

Compression Efficiency While ML-based codecs have shown improved compression efficiency over H.265, no ML codec has yet to outperform the standards across the entire PSNR curve. Moreover, no benchmarks have been presented on the VMAF metric, nor against AV1, as ML codecs have not compared favorably in these settings. We benchmark ELF on popular video test sets UVG and MCL-JCV under metrics PSNR, MS-SSIM and VMAF. In the low-latency mode, for natural videos ELF compares favorably across the entire R-D curve against the standards, and all other ML codecs. For example, on UVG under PSNR, ELF reduces the BD-rate by 44% against H.264, 26% against H.265, 15% against AV1, 35% against the next-best ML codec [2], and 54% against the next-fastest ML-codec [26] (while running 5x faster).

Speed For any practical application encoding must run at a reasonable frame rate, and decoding must run in real-time. Research on learned video compression has focused on improving the R-D curve, often at the expense of speed. For instance, autoregressive approaches inherently cannot be parallelized, resulting in methods that take many seconds to decode a single frame. In contrast, ELF runs at least 5x faster than all other ML codecs which report timings, and with fewer parameters. On an NVIDIA Titan V GPU, ELF encodes/decodes VGA at 49/91 FPS, HD 720 at 19/35 FPS, and HD 1080 at 10/18 FPS.

Our primary contributions are:

1. A novel framework for efficient rate control for learned video coding. This allows a single model to encode each frame with a wide range of bitrates, at a negligible increase to computation and parameter count.
2. A backbone specifically optimized to achieve strong performance on compression tasks while remaining computationally efficient.
3. An *in-loop flow predictor*, a novel module that utilizes previously transmitted information to get a strong initial estimate of the motion for the current frame.

1.1. Related work

Traditional video compression There has been a long history of hand-designed video codecs, such as H.263 and H.264 [40], which form the basis of video standards widely used today. More recently, H.265 [35], VP9 [30] and AV1 [8] have made significant improvements over the legacy video standards, and continue to be an active research

area. The traditional codecs have been exceptionally well-engineered and tuned, and have been difficult for the ML community to match both in terms of compression fidelity but also computational efficiency.

ML-based flexible-rate compression Compared to previous work, this paper is the first efficient flexible-rate video compression paper. Beyond the loss formulation, the fundamental challenge of learned flexible-rate compression is supporting a large dynamic range of bitrates in a single model. As the range gets larger and denser, performance often degrades. The topic of flexible-rate modeling has been addressed in a number of previous learned image coding papers which use conditional convolution [9], variable quantization width [10, 1], and recurrent networks [37]. While the former methods are slow due to an autoregressive probability model for the coderlayer or a recurrent network, our work doesn't use such methods.

In the learned video compression literature, [33] use a spatial multiplexer to achieve flexible rates with a single model. However, this approach requires a slow and complex search at encode-time in order to produce the spatial multiplexer map.

ML-based video compression Due to the inherent complexity of designing video compression algorithms, various works formulate end-to-end solutions to conquer subsets of the grand problem. One class of approaches such as [41, 11, 31] focuses directly on interpolation and omit P-frames. Another popular direction (which this work follows) is to design a low-latency ML-based codec, which only features keyframe compression and forward frame extrapolation (i.e I/P-frames only) [28, 23, 26]. Promising recent directions involve modeling motion using scale-space flow [2] and resolution-adaptive flow [24, 15], propagating a latent state [33, 12], and explicitly mitigating error propagation [27]. Yet another promising approach [13] revolves around using spatiotemporal autoencoders to encode chunks of frames.

Efficient ML-based compression Research on computationally efficient ML codecs is still in its nascence. In the image compression world, several approaches explore efficient codec modeling, relying on architecture optimizations [32, 18, 25]. To our knowledge, the only other ML video codec work with focus on the topic is [26], whose approach is to remove inter-dependency between frames and instead rely on entropy conditioning to capture redundancy.

2. Problem setup and baseline model

We aim to encode a video with frames $\mathbf{x}_1, \dots, \mathbf{x}_T \in [0, 1]^{3 \times H \times W}$ using the low-latency mode including I- and P-frames only. In this section, we describe a baseline model

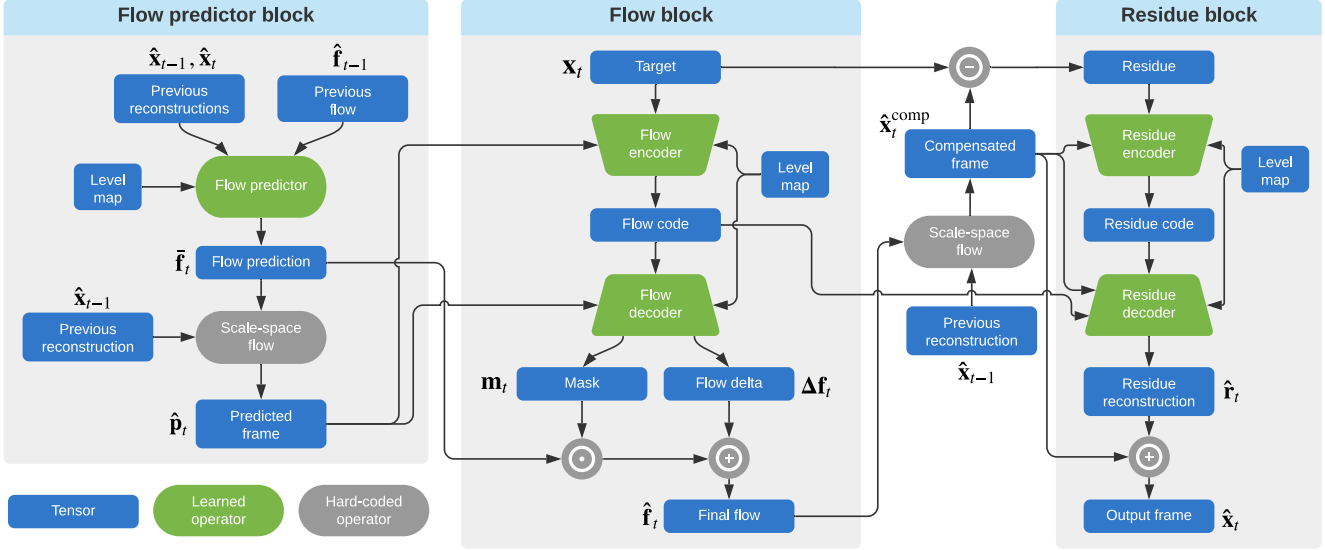


Figure 2: The overall architecture of ELF-VC. The **predictor block** (Section 3.3) uses previously transmitted information to get a strong initial estimate of the motion for the current frame, without sending any bits. The **flow block** refines this motion estimate and the **residue block** reconstructs the remaining residual. The **level map** is provided as input to different parts of the model to facilitate conditional modeling (Section 3.1). The learned operators are powered by the neural network backbone described in Section 3.2. The details of the state propagation and layer specifications are omitted for clarity and can be found Appendix B.

which is a combination of ideas from previous state-of-the-art ML codecs.

Baseline I-frame model The keyframe compressor (I-frame) is an image codec which encodes individual frames. As discussed in Section 1.1, there is a well-established body of work on learned techniques for image compression. For the baseline I-frame model, we use the backbone presented in Section 3.2, combined with a simplified variant of the hyperprior coding scheme presented in [5] (see Appendix A for details) and avoid any context modeling.

Baseline P-frame model The starting point for our baseline P-frame model is the common flow-residue model [28, 3], where an autoencoder (“flow block”) is used to reconstruct an optical flow field, and the *residue* — the leftover difference with the target — is compressed with a second autoencoder (“residue block”). Similar to [33], we propagate a state across blocks and across frames, which results in a significant improvement in the model (Table 3).

The flow block is an autoencoder which takes in the target \mathbf{x}_t , previous reconstruction $\hat{\mathbf{x}}_{t-1}$ and previous state \mathbf{s}_{t-1} as inputs. Through a bottleneck, it updates the state and then produces a flow output $\hat{\mathbf{f}}_t$ to compensate from $\hat{\mathbf{x}}_{t-1}$ to \mathbf{x}_t using the warping transform $\mathbf{F}(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{f}}_t)$. We build on [2] and for $\mathbf{F}(\cdot, \cdot)$ adopt the scale-space flow operator which elegantly handles uncertainty in motion estimation. Such a flow $\hat{\mathbf{f}}_t$ is made of 3 channels: horizontal and vertical displacements along with a blurring param-

eter (scale-space sigma). The compensated frame is then $\hat{\mathbf{x}}_t^{\text{comp}} = \mathbf{F}(\hat{\mathbf{x}}_{t-1}, \hat{\mathbf{f}}_t)$.

The residue block then takes in the leftover signal $\mathbf{x}_t - \hat{\mathbf{x}}_t^{\text{comp}}$ and the state returned by the flow block to produce a residue $\hat{\mathbf{r}}_t$ which is added to the warping to produce final reconstruction $\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^{\text{comp}} + \hat{\mathbf{r}}_t$. It also outputs a final state \mathbf{s}_t that is passed to the next frame.

The architecture for this baseline is essentially the flow and residual blocks in Figure 2 without level maps. The backbones for both blocks are presented in 3.2. Each block of the P-frame model uses the same hyperprior coding scheme as the I-frame model. Similarly to the I-frame model, autoregressive context models are avoided because they are prohibitively slow for practical use in their current formulation.

3. Novel Contributions

In this section, we describe new ideas building on top of the baseline model (Section 2). Section 4.4 presents ablation studies of their individual contributions.

3.1. Flexible-rate framework for ML video codecs

Rate control, or the optimization of visual quality under real-world constraints (e.g. bandwidth and latency), is essential for the practical deployment of a video codec. For example, in order to minimize bandwidth-induced latency, it is important to constrain the average bitrate of the video *and* the maximum bitrate for any frame.

While there exist methods for flexible-rate ML-based image compression (Sec. 1.1), it is challenging to extend these to video without suffering a loss of performance. Our work introduces two main innovations: a novel loss modulation scheme which improves BD-rate by 10% (see 4.4) by better training the model for high rates, and a novel embedding scheme which allows the rate to be smoothly varied across the rate-distortion curve.

For each frame we initially aim to support L different points on the R-D curve (we refer to them as levels) using a single model. In order for the network to optimize a level-dependent loss, the discrete level is converted into a one-hot vector of dimension L . This vector is tiled spatially and concatenated to the input of each of the neural network encoder and decoder (*level map*, Figure 2), and as an input to the in-loop flow predictor (Sec. 3.3). The I-frame model also uses a learned level and channel-dependent quantization width for the quantized coder, similar to [10]. Variable quantization width was not sufficient to achieve competitive performance — we found that giving the rate as input to all of the listed places to the network was necessary for optimal performance.

Multi-level loss setup The multi-level loss consists of the typical rate and distortion terms, with added level conditioning:

$$\mathcal{L}_{\text{rec}} = \frac{1}{T} \sum_t \mathbb{E}_l \left[\mu_t^{(l)} D(\mathbf{x}_t, \hat{\mathbf{x}}_t^{(l)}) \right] \quad (1)$$

$$\mathcal{L}_{\text{ent}} = \frac{1}{T} \sum_t \mathbb{E}_l \lambda_{\text{reg}}^{(l)} R_t^{(l)}. \quad (2)$$

The compression level l for frame t is sampled during training (see 4.1). $D(\cdot, \cdot)$ is the distortion metric (e.g. MSE loss), and $R_t^{(l)}$ is the total codelength of the encoded frame (see Appendix A) for the flow and residual blocks. $\mu_t^{(l)}$ is a dynamically-chosen weight that substantially improves multi-level training as explained below (see Table 3). The loss encourages the model to achieve different points on the R-D curve by using different regularization weights $\lambda_{\text{reg}}^{(l)}$ for different levels.

Dynamic loss modulation During training, the reconstruction loss for each frame and level is multiplied by a dynamically-changing weight $\mu_t^{(l)} \geq 1$ in order to encourage better performance for the levels and frames where the PSNR of the P frame is much lower than the PSNR for the I frame. The motivation for this idea is that when training the baseline unmodulated models (with $\mu_t^{(l)} = 1$), we observed that the higher bitrate levels of the variable-rate model trained much more slowly than lower ones. We further noticed that frames later in the GOP train more slowly than the earlier frames.

One natural idea for the loss modulation was to increase the weight associated with the parts of the model that were

under-performing. We developed the following method in order to avoid needing to set the weights manually. We fix the I-frame weights $\mu_0^{(l)}$ for each level to be 1. Then for a particular level, we increase the P-frame weight if the PSNR of the P-frame is much lower than the PSNR of the I-frame for that level. Specifically, the moving averages $\mathbb{E}D_t^{(l)}$ for the t -th frame and l -th level are computed. If $\text{PSNR}_t^{(l)} < \text{PSNR}_0^{(l)} - \delta$, then $\mu_t^{(l)}$ is increased by a small value. If the opposite is true, $\mu_t^{(l)}$ is decreased. See the ablation studies 4.4 and Appendix D for additional details and discussion.

Level embedding While a discrete set of L levels is useful for accumulating moving averages to modulate the loss during training, for practical purposes typical values of $L \approx 8$ are too coarse for precise rate control. We introduce a technique which allows smoothly varying the bitrate at encode time, without any additional training. We achieve this by linearly interpolating the one-hot level vectors to achieve intermediate levels arbitrarily fine-grained. Having trained the model with L_e levels, during inference we achieve a denser sampling of the bitrate range by embedding a more fine-grained level $l \in \{0, 1, \dots, L - 1\}$ within the original L_e -dimensional space. Specifically, we implement the linear interpolation scheme

$$l_e = (1 - \{s_l\}) \cdot \text{onehot}(u_l | L_e) + \{s_l\} \cdot \text{onehot}(v_l | L_e) \quad (3)$$

where $s_l = l \frac{L_e - 1}{L - 1}$, $u_l = \lfloor s_l \rfloor$, $v_l = u_l + 1$, $\{s_l\}$ is the fractional part of s_l , and $\text{onehot}(n | N)$ is an N -dimensional one-hot vector where the n -th component is one. See Appendix E for details.

Previous work in the image compression world [9] that used two knobs (discrete rate and quantization width) resulted in non-monotonic behavior. In contrast, our method has one continuous knob that results in the quality increasing monotonically as the bitrate increases. Furthermore, our method does not require any additional training in order to interpolate the rate with competitive performance.




Figure 3: The flexible-rate nature of ELF facilitates deployment in the presence of different constraints. We present examples of 3 different rate controllers: maximum BPP (caps of 0.081 on I, 0.037 on P), minimum quality (PSNR of 37), and constant level. A complex event around frame 300 (Jockey video, UVG dataset) causes the BPP to spike under guaranteed minimum quality (orange curve). If we constrain the bandwidth, the quality drops (green curve).




Figure 4: Visual comparison for an HD 1080 video encoded using the same model at two different levels. The **top row** shows the final flow $\hat{\mathbf{f}}_t$, the **middle row** shows difference between the target and the final reconstruction, and the **bottom row** shows the spatial bit allocation (Appendix A). The error is boosted by a factor of 4 and clipped for visualization purposes. The spatial bits plots use the same color mapping where yellow corresponds more bits spent.

3.2. Compression-centric backbone

We propose an efficient backbone that achieves competitive results for ML-based coding, which lowers the BD-rate by 30% and improves speed by 75% (Table 4.4) relative to popular video backbones. This is based on a block that we refer to as the Delayed Merge (DM) block (Figure 6). We experimented with various common backbones such as DenseNet [17], ResNet [14], multiscale dual path [7, 16, 33], and inverted residual modules [34], and found that a DenseNet-like (and optionally multiscale) representation achieves a balanced tradeoff between expressivity and computational performance. However, in its original formulation, DenseNet features many concatenation operations and convolutions with a small number of filters. As such, it does not lend itself to efficient computation. This issue also affects multiscale approaches presented in [16, 33].

Instead, after experimentation (see Appendix) we converge on a mix between residual and dense blocks where several convolutions, each with C channels, are ran with additive accumulation, at which point the intermediate activations are concatenated only a single time and the dimensionality is reduced back down. We refer to such a block as a DM- C block.

In addition to the backbone structure, we also explored various forms of attention popularized for image compression — but did not find these to enhance coding efficiency.

For increased efficiency, all blocks are exclusively ran in

scales at least 4x downsampled relative to pixel space. Detailed descriptions of the end-to-end architectures are found in Appendix B.

Optimizing for a light decoder There exists an inherent asymmetry between encoding and decoding, since in many use-cases the decoder must run in computationally-constrained environments (such as phones). In order to optimize the decoding speed given a fixed encoding speed, we shift computation from the decoder branches (inputs \rightarrow coderlayer) to the encoder branches (coderlayer \rightarrow outputs). In our ablations, we find that an asymmetric encoder-decoder pair has better R-D performance in addition to being faster during decoding as compared to a symmetric encoder-decoder pair (Table 3). For both flow and residue encoders we use three DM-256 blocks sequentially in scales 4x-8x-16x. For the flow and residue decoders we use three DM-64 or DM-128 respectively in scale sequence 16x-8x-4x. Complete layer specification is found in Appendix B.

3.3. In-loop flow prediction

We propose a novel *in-loop flow predictor* which predicts the current flow from previously transmitted frames and flows. It does not transmit any additional bits itself, and is ran prior to flow/residual encoding or decoding. It allows for BD-rate savings of 13% (Table 3).

Intuitively, designing video codecs revolves around exploiting redundancy across frames. Clearly, such similarities are captured well using optical flow; however, we observe that there also exists further redundancy among the optical flow fields *themselves*. Since consecutive flow fields are similar to one another due to linearity for motion, the current flow field can be predicted reasonably well from information already available on the decoder side — without transmitting any additional information. This predicted flow can then be refined using the flow autoencoder.

The predictor structure can be found in Figure 2. It takes in as inputs the previous flow $\hat{\mathbf{f}}_{t-1}$ and two previous frame reconstructions $\hat{\mathbf{x}}_{t-2}$ and $\hat{\mathbf{x}}_{t-1}$, and produces a base flow prediction $\bar{\mathbf{f}}_t$. The flow autoencoder produces a sigmoid map \mathbf{m}_t masking the predicted flow element-wise, and sparse flow delta $\Delta\hat{\mathbf{f}}_t$, which is added to the masked predicted flow. The final flow is then $\hat{\mathbf{f}}_t = \mathbf{m}_t \odot \bar{\mathbf{f}}_t + \Delta\hat{\mathbf{f}}_t$. The original baseline model in Section 2, then, is a special case with a predicted flow of zero and no mask. The predictor computation is shared by both the encoder and decoder.

In Figure 5, it is seen that the zero-bit predicted flow (top left) is already similar to the final flow (top right). Hence, the flow block only needs to output sparse flow touchups, thereby spending less bits. The neural network backbone used within the predictor is based upon the DM block described in Section 3.2 but it additionally computes features at two scales for each block (see Appendix B for details).

A similar idea was proposed by [22]. We note that this approach does not include learned mask \mathbf{m}_t to allow turning

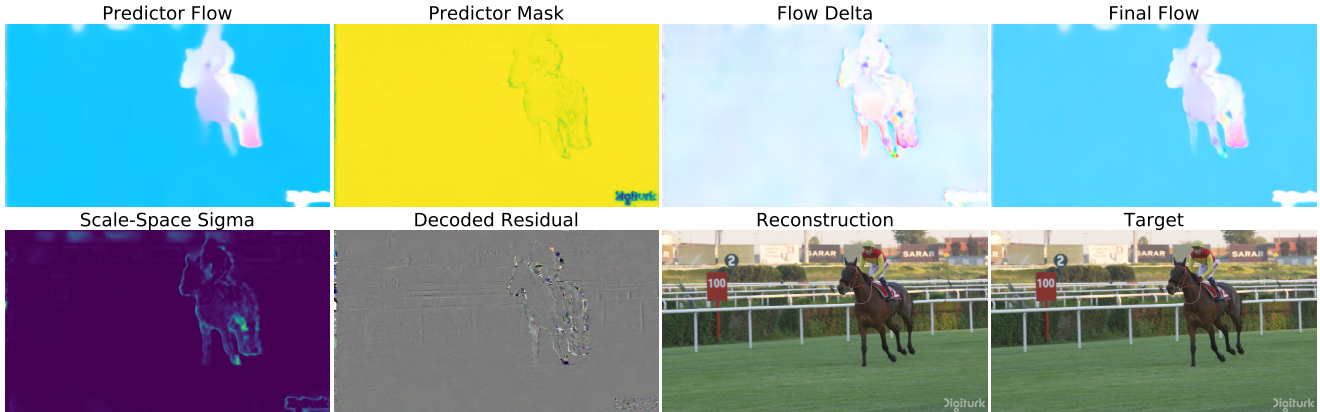


Figure 5: Intermediate tensors in the P-frame decoder for a typical HD 1080 video. The in-loop flow predictor uses previously transmitted information to generate an initial estimate of the motion (**predictor flow**), which is then masked by the **predictor mask**. A **flow delta** is decoded from the state in the first block and added to the predicted flow to generate the **final flow**. It can be seen that the predicted flow, which was computed without transmitting any additional bits, captures much of the motion — allowing for a sparser flow delta. The flow and the **scale-space sigma** (also decoded from the state) are applied to the previous frame to generate a compensated frame. The last block generates an estimate of the remaining **residue** and adds it to the compensated frame to generate the **final reconstruction**. The residue is multiplied by a factor of 4 and clipped for visualization purposes. All flows are normalized so that a fully saturated color is a flow with a magnitude of 15 pixels.

the predictor off in areas that are difficult to predict. Also, in this approach *only* the predictor has a runtime of 17 FPS for resolution 320×256 . In contrast, our entire decoder, predictor included, runs in 139 FPS for resolution 352×288 (which has 20% more pixels).

4. Results

4.1. Experimental setup

We train our models on the Vimeo-90k dataset [42] and on MSE using the Adam optimizer [19] for a total of 800,000 iterations with a batch size of 8, crop sizes of 256×256 , and GOP of IPPP. During inference we unroll over P-frames to complete the 16-frame GOP. We train in RGB for PSNR and MS-SSIM, and in YUV for VMAF. We start with a momentum of 0.9, and learning rate of 7×10^{-5} which we lower by 5x at 80% and 95% into the training.




Figure 6: The Delayed Merge (DM- C) block with C base channels. Our CNN backbones use this as the primary building block with processing across different scales. The DM block allows for the benefits of DenseNet while being faster due to the reduced number of concatenation operations (Section 3.2).

Codec	PSNR		MS-SSIM	
	UVG	MCL-JCV	UVG	MCL-JCV
H.264	-44.3%	-33.7%	-52.9%	-51.2%
H.265	-26.1%	-17.0%	-46.0%	-46.1%
AV1	-14.8%	-1.0%	-50.0%	-50.9%
Agustsson et al. (2020) [2]	-35.3%	-29.4%	-32.6%	-34.8%
RLVC (2020) [43]	-35.0%	-35.3%	-28.6%	-38.4%
NVC (2020)[24]	-37.9%	-	-31.1%	-
Lu et al. (2020) [27]	-	-42.3%	-40.4%	-
Liu et al. (2020) [26]	-54.2%	-46.8%	-60.0%	-60.7%
DVC (2019) [28]	-47.6%	-	-59.9%	-
Habibian et al. (2019) [13]	-57.3%	-	-53.6%	-
Wu et al. (2018) [41]	-62.9%	-	-69.5%	-

Table 1: BD-rate savings of ELF relative to common video standards, and state-of-the-art ML codecs (full R-D curves available in Figure 7). Numbers are reported on the PSNR and MS-SSIM metrics for the UVG and MCL-JCV datasets.

We train for another 40,000 iterations on a crop size of 320×320 on a Vimeo90k-like dataset we generated with a larger crop size (Appendix F). Finetuning on a larger crop size improved results a small amount. For models reporting on MS-SSIM or VMAF, we fine-tune for 80,000 iterations on the respective metrics.

We train a total of 2 models which together cover the entire bitrate range. During training, each model aims to optimize $L = 8$ different points on the R-D curve where the regularization weights λ_{reg} are chosen linearly in logspace in the range $[10^{-1.7}, 10^{-3.7}]$ for the lower BPP model and $[10^{-3}, 10^{-5}]$ for the higher. The level for the initial I-frame, l_0 , is chosen randomly. The following levels are then chosen as $l_t = l_{t-1} + v_t$ rounded to the nearest integer and clipped to the range $[0, L - 1]$, where $v_t \sim \mathcal{N}(\mu = 0, \sigma = 0.5)$ is

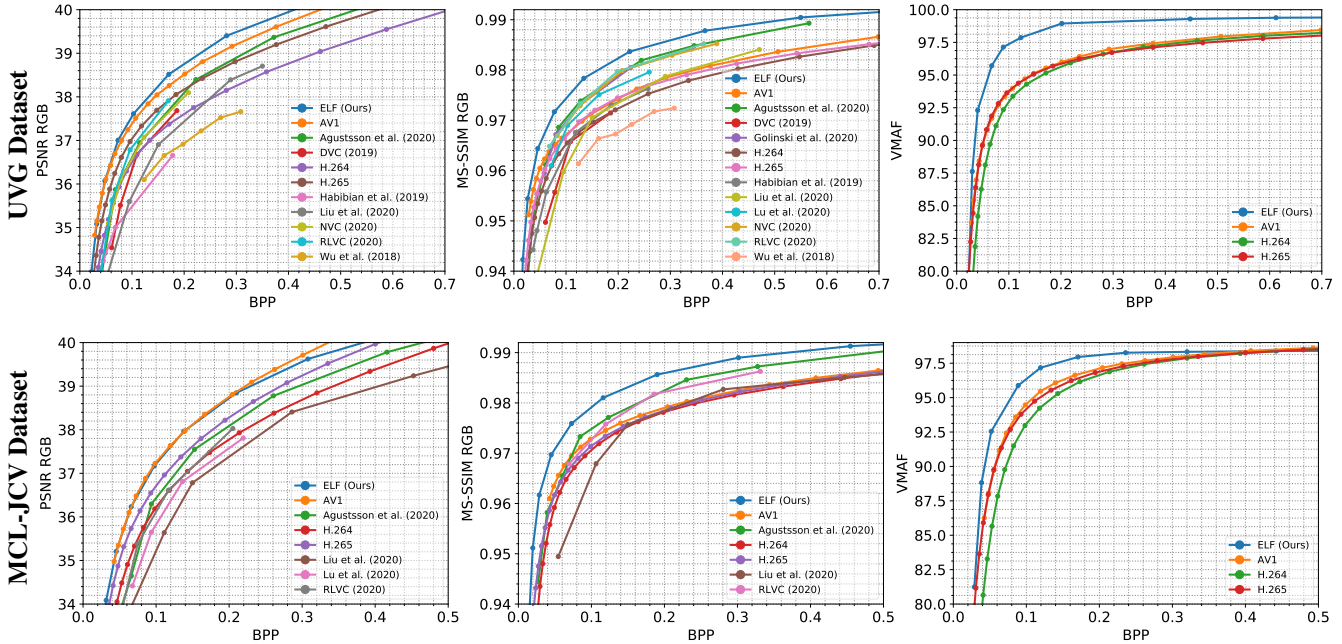


Figure 7: Rate-distortion curves of traditional codecs and state-of-the-art ML codecs [2, 28, 12, 13, 26, 27, 24, 43, 41] on the UVG and MCL-JCV video datasets.

sampled from a normal distribution.

To generate all R-D curves in this paper (apart from the maximum bitrate and minimum quality rate controllers in Figure 7), we sweep over levels and keep the level constant across all frames of each video.

The model is trained in FP32. The graph is converted to FP16, apart from the subnet decoder where the output probabilities are sensitive to minor perturbations. This does not lead to any reduction in R-D performance. At inference time, we run all models using TensorRT. Entropy encoding and decoding are implemented on CPU and are parallelized over channels and space (see Appendix A). The entropy coding portion takes about 10% of the total runtime.

4.2. Benchmarking procedure

Baseline codecs We benchmark against modern commercial codecs H.264/AVC, H.265/HEVC, and AV1, as well as the most competitive ML codecs to our knowledge. We use FFMpeg to encode H.264 and H.265⁴. We emphasize that while some existing works restrict the baselines to the `veryfast` preset, we use the default preset, which is a much more competitive and realistic baseline. We do not constrain the codecs in any way apart from disabling B-frames. We use the SVT-AV1 encoder for AV1. See Appendix I for the exact commands used. For the ML-based

¹[43] only reports runtimes for 240p; to enable comparison on CIF the numbers were scaled proportionally to the number of pixels.

²For [26], we only count their fast C++ implementation; the authors further report that the Python interface leads to an overhead of 1,190ms for encoding and 650ms for decoding, but we ignore these counts.

Method	FPS		Method	FPS	
	Encode	Decode		Encode	Decode
CIF 352x288			HD 720 1280x720		
RLVC ¹ [43]	12	25	Rippel et al. [33]	0.5	3
Rippel et al. [33]	6	30	ELF no-opt (Ours)	10	21
DVC [28]	25	41	ELF (Ours)	19	35
ELF no-opt (Ours)	45	106	HD 1080 1920x1080		
ELF (Ours)	71	139	Habibian et al. [13]	1.5	10 ^{-3.7}
VGA 640x480			Wu et al. [41]	2.4	10 ⁻³
Rippel et al. [33]	2	10	Rippel et al. [33]	0.2	1.0
ELF no-opt (Ours)	24	59	DVC [28]	1.5	1.8
ELF (Ours)	47	91	Liu et al. ² [26]	2.0	3.0
			ELF no-opt (Ours)	5.0	12
			ELF (Ours)	10	18

Table 2: Comparison of runtimes for different resolutions of all ML codecs which report timings, for BPP 0.2. For HD 1080, ELF runs at least 5x faster than other ML codecs. “ELF no-opt” is the fully unoptimized variant: FP32, no TensorRT, and non-parallelized entropy coding.

codecs, we compare against all recent approaches including [41, 28, 13, 12, 27, 2, 24, 26, 43]³.

Metrics We evaluate all reconstructions on popular video metrics PSNR, MS-SSIM [39], and VMAF⁵. We run VMAF using its hook into FFMpeg. In order to compare against existing ML-based approaches, PSNR and MS-SSIM are evaluated in the RGB colorspace. This is not ideal

³[13]’s PSNR results on the UVG dataset were taken from [26].

⁴We use the latest FFMpeg to benchmark the standards. Our H.264 and H.265 curves are slightly better than the ones in [2], but we validated we do match them with an older FFMpeg version.

by any means and YUV 4:2:0 is preferable for perceptual quality optimization.

We found that it is easy for ML codecs to perform well on MS-SSIM and VMAF by the numbers, but that perceptual quality is not commensurate with these gains. We took extra measures to avoid overfitting on VMAF by training it jointly with PSNR.

Test sets We benchmark all codecs on popular video datasets UVG [29] and MCL-JCV [38]. These datasets are commonly used for video codec evaluation, and contain respectively 7 and 30 diverse HD 1080 videos with totals of 3,900 and 4,115 frames.

4.3. Performance

Coding efficiency Figure 7 presents the rate-distortion curves for all approaches, and Table 1 provides BD-rate [6] summaries of these curves. It can be seen that ELF compares favorably against all standards and ML codecs, under all metrics — with the exception of BPP (bits per pixel) 0.25 and higher for the MCL-JCV dataset against AV1. In analyzing as to why, we observed that our model, similar to [2], performs very poorly on the four non-photorealistic/cartoon videos within MCL-JCV. Excluding these videos we outperform AV1 across the entire range (Appendix G).

Computational efficiency We benchmark the runtime of our codec across different resolutions on an NVIDIA Titan V GPU, and include all time spent on network execution, entropy encoding/decoding, and so on. We only exclude CPU↔GPU memory transfer overhead: this is not a fundamental limitation but rather an artifact of TensorFlow’s inability to pin data to GPU memory across model iterations. The benchmarks and comparisons against all other approaches which report timings⁶ can be found in Table 2. For example, for HD 1080, ELF runs at 98ms/frame for encoding and 55ms/frame for decoding; this is 5x faster than the second-fastest ML codec [26], while reducing the BD-rate by 55% on average relative to it (Table 1). The tables also provide benchmarks for “ELF no-opt”, which is the fully unoptimized variant: FP32, no TensorRT, and non-parallelized entropy coding.

4.4. Ablation studies

We study the individual contributions of the proposed ideas. For the ablation environment we follow exactly same training and inference procedures described in 4.1, apart from training each model for 250,000 iterations and only training the lower bitrate range model. The results can be found in Table 3.

⁵See repository at <https://github.com/Netflix/vmaf>.

⁶Many of the timings were taken from [26], who were able to gather these from the original authors of the respective papers. We were further able to contact [33] who generously provided detailed benchmarks. [28, 43, 26] use an NVIDIA 1080 Ti GPU, and [33] uses an NVIDIA Titan V.

Property	Option	BD-rate Increase	FPS for HD 720		# Param	
			Enc.	Dec.	Enc.	Dec.
Flow predictor	Yes	0%	19	35	38M	11M
	No	13%	21 (+10%)	37 (+6%)	37M (-3%)	10M (-9%)
Bitrate coverage	Range	0%	19	35	38M	11M
	Point	1%	19	35	220M (+500%)	70M (+500%)
Backbone	DM	0%	19	35	38M	11M
	Common 128	76%	18 (-5%)	24 (-31%)	24M (-37%)	12M (+9%)
	Common 192	43%	14 (-26%)	20 (-43%)	28M (-18%)	14M (+27%)
	DM-Sym 128	31%	20 (+5%)	28 (-20%)	25M (-34%)	13M (+18%)
	DM-Sym 192	18%	15 (-21%)	22 (-37%)	35M (-8%)	17M (+55%)
Loss modulation	Yes	0%	19	35	38M	11M
	No	10%	19	35	38M	11M
State	Yes	0%	19	35	38M	11M
	No	37%	19	36 (+3%)	37M (-3%)	11M

Table 3: Ablation studies of different modeling choices. The BD-rate values were computed under PSNR on UVG.

Using 2 flexible-rate models instead of 12 single-rate models reduces the number of required parameters by 6x without harming compression performance at all.

Swapping to the common backbone often used in the image/video coding literature [5, 3], with base channels 128 or 192 (“Common” in the ablation table) worsens BD-rate by at least 43% and slows the decoder down by 31% (we do our best to tune the model for this backbone).

Swapping to a symmetric DM-based backbone where the encoder and decoder have the same number of channels (“Sym” in the ablation table) increases the BD-rate by at least 18% and slows down the decoder by at least 20%.

Removing the predictor results in a BD-rate increase of 13%, without changing runtimes dramatically. Removing the dynamic loss modulation worsens BD-rate by 10%.

5. Conclusion

Compared to the video coding standards, ML-based video compression is still in its infancy. While ML video codecs have achieved impressive R-D performance, some of this success can be attributed to the huge investment in general in neural network research, software, and hardware. There still remains considerable work to be done towards practical deployment of ML codecs.

In this work, we propose ELF-VC, which makes significant improvements in terms of R-D performance, bitrate flexibility and speed. We believe there still exists significant room for improvement across all modeling decisions: architectural choices for the different modules, rate control, and so on. Moreover, while our approach is close to real-time on a mid-range desktop GPU, it still requires further optimization to achieve real-time performance on edge devices such as phones. Another important direction for future work in ML-based video compression is a visual perception metric for video that is suitable for backpropagation and closely aligns with human visual perception.

References

- [1] Variable rate image compression with content adaptive optimization. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2020-June:533–537, 2020. [2](#)
- [2] 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*, pages 8503–8512, 2020. [2](#), [3](#), [6](#), [7](#), [8](#)
- [3] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. *arXiv preprint arXiv:1804.02958*, 2018. [3](#), [8](#)
- [4] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimization of nonlinear transform codes for perceptual quality. In *Picture Coding Symposium (PCS), 2016*, pages 1–5. IEEE, 2016. [1](#)
- [5] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018. [1](#), [3](#), [8](#)
- [6] G. Bjontegaard. Calculation of average psnr differences between rd-curves. 2001. [8](#)
- [7] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. In *Advances in Neural Information Processing Systems*, pages 4467–4475, 2017. [5](#)
- [8] Yue Chen, Debargha Mukherjee, Jingning Han, Adrian Grange, Yaowu Xu, Zoe Liu, Sarah Parker, Cheng Chen, Hui Su, Urvang Joshi, et al. An overview of core coding tools in the av1 video codec. In *2018 Picture Coding Symposium (PCS)*, pages 41–45. IEEE, 2018. [2](#)
- [9] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. *Proceedings of the IEEE International Conference on Computer Vision*, 2019-Octob:3146–3154, 2019. [2](#), [4](#)
- [10] Ze Cui, Jing Wang, Bo Bai, Tiansheng Guo, and Yihui Feng. G-VAE: A Continuously Variable Rate Deep Image Compression Framework. 2020. [2](#), [4](#)
- [11] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers. Neural inter-frame compression for video coding. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6421–6429, 2019. [2](#)
- [12] Adam Golinski, Reza Pourreza, Yang Yang, Guillaume Sautiere, and Taco S Cohen. Feedback recurrent autoencoder for video compression. *arXiv preprint arXiv:2004.04342*, 2020. [1](#), [2](#), [7](#)
- [13] Amirhossein Habibi, Ties van Rozendaal, Jakub M Tomczak, and Taco S Cohen. Video compression with rate-distortion autoencoders. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7033–7042, 2019. [1](#), [2](#), [6](#), [7](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [15] Zhihao Hu, Zhenghao Chen, Dong Xu, Guo Lu, Wanli Ouyang, and Shuhang Gu. Improving deep video compression by resolution-adaptive flow coding. In *European Conference in Computer Vision (ECCV)*, 2020. [2](#)
- [16] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Weinberger. Multi-scale dense networks for resource efficient image classification. In *International Conference on Learning Representations*, 2018. [5](#)
- [17] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. [5](#)
- [18] Nick Johnston, Elad Eban, Ariel Gordon, and Johannes Ballé. Computationally efficient neural image compression. *arXiv preprint arXiv:1912.08771*, 2019. [2](#)
- [19] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [20] LDV. 2017 on netflix - a year in bingeing. 2017. [1](#)
- [21] LDV. 45 billion cameras by 2022 fuel business opportunities. 2017. [1](#)
- [22] Jianping Lin, Dong Liu, Houqiang Li, and Feng Wu. M-lvc: Multiple frames prediction for learned video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3546–3554, 2020. [5](#)
- [23] Haojie Liu, Tong Chen, Ming Lu, Qiu Shen, and Zhan Ma. Neural video compression using spatio-temporal priors. *arXiv preprint arXiv:1902.07383*, 2019. [2](#)
- [24] Haojie Liu, M. Lu, Zhan Ma, Fan Wang, Zhihuang Xie, Xun Cao, and Yao Wang. Neural video coding using multiscale motion compensation and spatiotemporal context model. *ArXiv*, abs/2007.04574, 2020. [2](#), [6](#), [7](#)
- [25] Jiaheng Liu, Guo Lu, Zhihao Hu, and Dong Xu. A unified end-to-end framework for efficient deep image compression, 2020. [2](#)
- [26] Jerry Liu, Shenlong Wang, W. Ma, Meet Shah, Rui Hu, Pranaab Dhawan, and R. Urtasun. Conditional entropy coding for efficient video compression. *European Conference on Computer Vision (ECCV)*, 2020. [1](#), [2](#), [6](#), [7](#), [8](#)
- [27] Guo Lu, Chunlei Cai, Xiaoyun Zhang, Li Chen, Wanli Ouyang, Dong Xu, and Zhiyong Gao. Content adaptive and error propagation aware deep video compression. *arXiv preprint arXiv:2003.11282*, 2020. [1](#), [2](#), [6](#), [7](#)
- [28] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11006–11015, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [29] Alexandre Mercat, Marko Viitanen, and Jarno Vanne. Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 297–302, 2020. [8](#)
- [30] Debargha Mukherjee, Jim Bankoski, Adrian Grange, Jingning Han, John Koleszar, Paul Wilkins, Yaowu Xu, and Ronald Bultje. The latest open-source video codec vp9-an overview and preliminary results. In *2013 Picture Coding Symposium (PCS)*, pages 390–393. IEEE, 2013. [2](#)
- [31] Woonsung Park and Munchurl Kim. Deep predictive video compression with bi-directional prediction. *arXiv preprint arXiv:1904.02909*, 2019. [2](#)

- [32] Oren Rippel and Lubomir Bourdev. Real-time adaptive image compression. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2922–2930, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. [1](#), [2](#)
- [33] Oren Rippel, Sanjay Nair, Carissa Lew, Steve Branson, Alexander G Anderson, and Lubomir Bourdev. Learned video compression. *arXiv preprint arXiv:1811.06981*, 2018. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)
- [34] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. [5](#)
- [35] 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*, 22(12):1649–1668, 2012. [2](#)
- [36] TechCrunch. Facebook hits 100m hours of video watched a day. 2018. [1](#)
- [37] George Toderici, Sean M O’Malley, Sung Jin Hwang, Damien Vincent, David Minnen, Shumeet Baluja, Michele Covell, and Rahul Sukthankar. Variable rate image compression with recurrent neural networks. *arXiv preprint arXiv:1511.06085*, 2015. [2](#)
- [38] Haiqiang Wang, Weihao Gan, Sudeng Hu, Joe Yuchieh Lin, Lina Jin, Longguang Song, Ping Wang, Ioannis Katsavounidis, Anne Aaron, and C-C Jay Kuo. Mcl-jcv: a jnd-based h. 264/avc video quality assessment dataset. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 1509–1513. IEEE, 2016. [8](#)
- [39] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *Signals, Systems and Computers, 2004.*, volume 2, pages 1398–1402. Ieee, 2003. [7](#)
- [40] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7):560–576, 2003. [2](#)
- [41] Chao-Yuan Wu, Nayan Singhal, and Philipp Krähenbühl. Video compression through image interpolation. In *ECCV*, 2018. [1](#), [2](#), [6](#), [7](#)
- [42] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 127(8):1106–1125, 2019. [6](#)
- [43] Ren Yang, Fabian Mentzer, Luc Van Gool, and Radu Timofte. Learning for video compression with recurrent auto-encoder and recurrent probability model. *arXiv preprint arXiv:2006.13560*, 2020. [6](#), [7](#), [8](#)
- [44] YouTube. Youtube for press. 2020. [1](#)