



Conditional Entropy Coding for Efficient Video Compression

Jerry Liu^{1(✉)}, Shenlong Wang^{1,2(✉)}, Wei-Chiu Ma^{1,3(✉)}, Meet Shah^{1(✉)},
Rui Hu^{1(✉)}, Pranaab Dhawan^{1(✉)}, and Raquel Urtasun^{1,2(✉)}

¹ Uber ATG, St. Pittsburgh, USA

{jerry1,slwang,weichiu,meet.shah, rui.hu,pdhawan,urtasun}@uber.com

² University of Toronto, Toronto, Canada

³ MIT, Cambridge, USA

Abstract. We propose a very simple and efficient video compression framework that only focuses on modeling the conditional entropy between frames. Unlike prior learning-based approaches, we reduce complexity by not performing any form of explicit transformations between frames and assume each frame is encoded with an independent state-of-the-art deep image compressor. We first show that a simple architecture modeling the entropy between the image latent codes is as competitive as other neural video compression works and video codecs while being much faster and easier to implement. We then propose a novel internal learning extension on top of this architecture that brings an additional $\sim 10\%$ bitrate savings without trading off decoding speed. Importantly, we show that our approach outperforms H.265 and other deep learning baselines in MS-SSIM on higher bitrate UVG video, and against all video codecs on lower framerates, while being thousands of times faster in decoding than deep models utilizing an autoregressive entropy model.

1 Introduction

The efficient storage of video data is vitally important to a large number of settings, from online websites such as Youtube and Facebook to robotics settings such as drones and self-driving cars. This necessitates the use of good video compression algorithms. Both image and video compression are fields that have been extensively researched in the past few decades. Traditional image codecs such as JPEG2000, BPG, and WebP, and traditional video codecs such as HEVC H.265, AVC/H.264 [28, 38] are widely used and hand-engineered to work well in a variety of settings. But the lack of learning involved in the algorithm leaves room open for more end-to-end optimized solutions.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58520-4_27) contains supplementary material, which is available to authorized users.

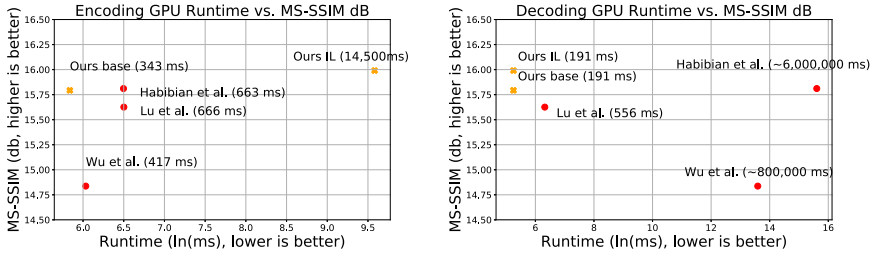


Fig. 1. Plots indicating the GPU runtime vs. MS-SSIM of our model vs. other deep approaches at bitrate 0.2, averaged over a 1920×1080 UVG frame. Runtimes are shown independent of the entropy coding implementation. We interpolate to obtain the MS-SSIM estimate at the exact bitrate.




Fig. 2. An illustration of the explicit transformations used in removing redundant information in subsequent frames vs. probabilistic modeling used in entropy coding. A typical lossy compression algorithm will contain elements of both approaches.

Recently, there has been an explosion of deep-learning based image compressors that have been demonstrated to outperform BPG on a variety of evaluation datasets across both MS-SSIM and PSNR as evaluation metrics [9, 18, 19, 23]. This explosion has also recently happened in video compression on a somewhat smaller scale, with the latest advances being able to outperform H.265 on MS-SSIM and PSNR in certain cases [12, 14, 20, 25]. Many of these approaches [12, 25, 39] involve learning-based generalizations of the traditional video compression techniques of motion-compensation, frame interpolation and residual coding.

While achieving impressive distortion-rate curves, there are several major facts blocking the wide adoption of these approaches for real-world, generic video compression tasks. First, most aforementioned approaches are still slower than standard video codecs at both encoding and decoding stage; moreover, due to the fact that they explicitly perform interpolation and residual coding between frames, a majority of the computations cannot be parallelized to accelerate coding speed; finally, the domain bias of the training dataset makes it difficult to generalize well to a wide range of different type of videos.

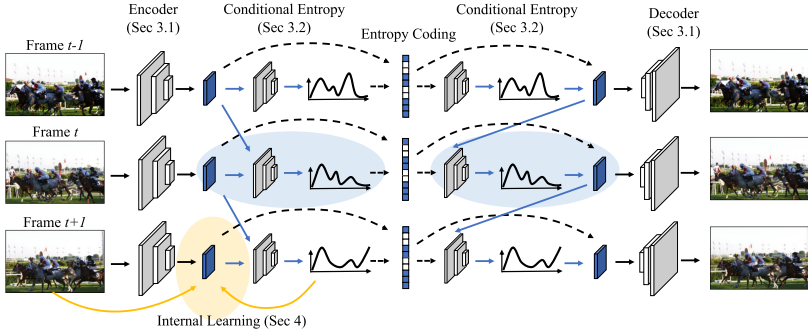


Fig. 3. Overview of the architecture of our approach. We highlight our key contribution, namely the conditional entropy model and internal learning in blue and yellow, respectively. (Color figure online)

In this paper, we address these issues by creating a remarkably simple *entropy-focused* video compression approach that is not only competitive with prior state-of-the-art learned compression, but also *significantly faster* (see Fig. 1), rendering it a practical alternative to existing video codecs. Such an entropy-focused approach focuses on better capturing the correlations between frames during entropy coding rather than performing explicit transformations (e.g. motion compensation). Our contributions are two-fold (illustrated in Fig. 3). First, we propose a base model consisting only of a **conditional entropy model** fitted on top of the latent codes produced by a deep single-image compressor. The intuition for why we don't need explicit transformations can be visualized in Fig. 2: given two video frames, prior works would code the first frame to store the full frame information while coding the second frame to store explicit motion information from frame 1 as well as residual bits. On the other hand, our approach encodes both frames as independent image codes, and reduces the joint bitrate by fitting probability model (an *entropy* model) to maximize the probability of the second image code given the first. We can thus extend this to a full video sequence by still encoding every frame independently, and simply considering every adjacent pair of frames for the probability model. While entropy modeling has been a subcomponent of prior works [14, 16, 20, 25, 39], they have tended to be very simple [25], only dependent on the image itself [12, 20], or use costly autoregressive models that are intractably expensive during decoding [14, 39]; here our conditional entropy model provides a viable means for video compression purely within itself.

Our second contribution is to propose **internal learning** of the latent code during inference. Prior works in video compression operate by using a fixed encoder during the inference/encoding stage. As a result, the latent codes of the video is not optimized towards reconstruction/entropy estimation for the specific test video. We observe as long as the decoder is fixed, we can trade off *encoding runtime* to further optimize the latent codes along the rate-distortion curve, while not affecting *decoding runtime* (Fig. 1, right).

We validate the performance of the proposed approach over several datasets across various framerates. We show that at standard framerates, our base model is much faster and easier to implement than most state-of-the-art deep video benchmarks, while matching or outperforming these benchmarks as well as H.265 on MS-SSIM. Adding internal learning provides additional $\sim 10\%$ bitrate gains with the same decoding time. Additionally, on lower framerates, our models outperform H.265 by a wide margin at higher bitrates. The simplicity of our method indicates that it is a powerful approach that is widely applicable across videos spanning a broad range of content, framerates, and motion.

2 Background and Related Work

2.1 Deep Image Compression

There is an abundance of work on learned, lossy image-compression [8, 9, 21, 23, 24, 31–33]. In general, these works follow a general autoencoder architecture minimizing the rate-distortion tradeoff. Typically, an encoder transforms the image into a latent space, quantizes the symbols, and applies entropy coding (typically arithmetic/range coding) on the symbols to output a compressed bitstream. During decoding, the recovered symbols are then fed through a decoder for image reconstruction.

Recent works approximate the rate-distortion tradeoff $\ell(\mathbf{x}, \hat{\mathbf{x}}) + \beta R(\hat{\mathbf{y}})$ in a differentiable manner by replacing the bitrate term R with the cross-entropy between the code distribution and a learned “prior” probability model: $R \approx \mathbb{E}_{\mathbf{x} \sim p_{data}} [\log p(E(\mathbf{x}); \boldsymbol{\theta})]$. Shannon’s source coding theorem [26] indicates that the cross-entropy is an asymptotic lower bound of the bitrate. One way to achieve this optimal bitrate during entropy coding is to use the learned “prior” model as the probability map during arithmetic coding or range coding to code the symbols. Hence, the smaller the cross-entropy term, the more the bitrate can be reduced. This then implies that the more expressive the prior model in modeling the true distribution of latent codes, the smaller the overall bitrate.

Sophisticated prior models have been designed for the quantized representation in order to minimize the cross-entropy with the code distribution. Autoregressive models [21, 23, 33], hyperprior models [9, 23], and factorized models [8, 9, 31] have been used to model this prior. [23] and [22] suggest that using an autoregressive model is intractably slow in practice, as it requires a pass through the model for every single pixel during decoding. [22] suggests that the hyperprior approach presents a good tradeoff between speed and performance.

A recent model by Liu *et al.* [17] presents an extension of [23] using residual blocks in the encoder/decoder, outperforming BPG and other deep models on both PSNR and MS-SSIM on Kodak.

2.2 Video Compression

Conceptually, traditional video codecs such as H.264/H.265 exploit temporal correlations between frames by categorizing the frames as follows [28, 38]:

- I-frames: compressed as an independent image
- P-frames: predicted from past frames using block-based flow estimate, then encode residual.
- B-frames: similar to P-frames but predicted from both past and future frames.

In order to predict P/B-frames, the motion between frames is predicted via block matching (and the flow is uniformly applied within blocks), and then the resulting difference is separately encoded as the “residual.” Generally, if neighboring frames are temporally correlated, encoding the motion and residual vectors requires fewer bits than recording the subsequent frame independently.

Recently, several deep-learning based video compression frameworks [12, 15, 18, 20, 25, 39] have been developed. Both Wu *et al.* [39] and Lu *et al.* [20] attempt to generalize various parts of the motion-compensation and residual learning framework with neural networks, and get close to H.265 performance (on *very-fast* setting). Rippel *et al.* [25] achieved state-of-the-art results in MS-SSIM by generalizing flow/residual coding with a global state, spatial multiplexing, and more. Djelouah *et al.* [12] jointly decode motion and blending coefficients from references frames, and represent residuals in latent space. Habibian *et al.* [14] utilizes a 3D convolutional architecture to avoid motion compensation, as well as an autoregressive entropy model to outperform [20, 39].

These prior works generally require specialized modules and explicit transformations, with the entropy model being an oftentimes intractable autoregressive subcomponent [12, 14, 39]. A more closely related work is that of Han *et al.* [15], who propose to model the entropy dependence between codes with an LSTM: $p(\mathbf{y}_i | \mathbf{y}_{<i})$. In contrast to these prior works, we focus on a *entropy-only* approach, with no explicit transformations across time. More importantly, our base approach carefully exploits the parallel nature of frame encoding/decoding, rendering it orders of magnitude faster than other state-of-the-art while being just as competitive.

2.3 Internal Learning

The concept of internal learning is not new. It is similar to the sample-specific nature of transductive learning [29, 35]. “Internal Learning” is a term proposed in [6, 13], which exploits the internal recurrence of information within a single-image to train an unsupervised super-resolution algorithm. Many related works have also trained deep networks on a single example, from DIP [34] to GANs [27, 30, 40]. Also related is Sun *et al.* [29] who propose “test-time training” on an auxiliary function for each test instance on supervised classification tasks. Concurrently and independently from our work, Campos *et al.* propose context adaptive optimization in image compression [10], which has demonstrated promising results on finetuning each latent code towards its test image.

In our setting, we leverage the fact that in video compression the ground-truth is simply the video itself, and we apply internal learning in a way that obeys codebook consistency while decreasing the conditional entropy between video frames during decoding. There are unique advantages to using internal learning

in our entropy-only video compression setting: it can optimize for conditional entropy between codes in a way that an independent frame encoder cannot (see Sect. 4).

3 Entropy-Focused Video Compression

Our base model consists of two components: we first encode each frame \mathbf{x}_i of a video \mathbf{x} with a straightforward, off-the-shelf image compressor consisting of a deep image encoder/decoder (Sect. 3.1) to obtain discrete image codes \mathbf{y}_i . Then, we capture the temporal relationships between our \mathbf{y}_i 's with a conditional entropy model that approximates the joint entropy of the video sequence (Sect. 3.2). The model is trained end-to-end with respect to the rate-distortion loss function (Sect. 3.3).

3.1 Single-Image Encoder/Decoder

We encode every video frame \mathbf{x}_i separately with a deep image compressor into a quantized latent code \mathbf{y}_i ; note that each \mathbf{y}_i contains *full* information to reconstruct each frame i and does not depend on previous frames. Our choice of architecture for single-image compression borrows heavily from the state-of-the-art model presented by Liu et al. [17], which has shown to outperform BPG on both MS-SSIM and PSNR. The architecture consists of the image encoder, quantizer, and image decoder. We simplify the model in two ways compared to the original paper: we remove all non-local layers for efficiency/memory reasons, and we remove the autoregressive context estimation due to its decoding intractability ([21, 23], also see Fig. 1).

More details about the image encoder/decoder architecture are found in supplementary material. In our video compression model, we use the image encoder/quantizer to produce the quantized code \mathbf{y}_i , and the image decoder to produce the reconstruction $\hat{\mathbf{x}}_i$. We do not use the existing entropy model (inspired from [9, 23]) which are only designed for modeling the intra-image entropy; instead we design our own conditional entropy model, as detailed next.

3.2 Conditional Entropy Model for Video Encoding

Our *entropy model* models the joint entropy of the video frame codes with a deep network in order to reduce the overall bitrate of the video sequence; this is because the cross-entropy between our entropy model and the actual code distribution is a tight lower bound of the bitrate [26]. Our goal is to design our entropy model to capture the temporal correlations as well as possible between the frames such that it can minimize the cross-entropy with the code distribution. Put another way, the bitrate for the entire video sequence code $R(\mathbf{y})$ is tightly approximated by the cross-entropy between the code distribution induced by the encoder $\mathbf{y} = E(\mathbf{x}), \mathbf{x} \sim p_{data}$ and our probability model $p(\cdot|\boldsymbol{\theta}): \mathbb{E}_{\mathbf{x} \sim p_{data}}[\log p(\mathbf{y}; \boldsymbol{\theta})]$.

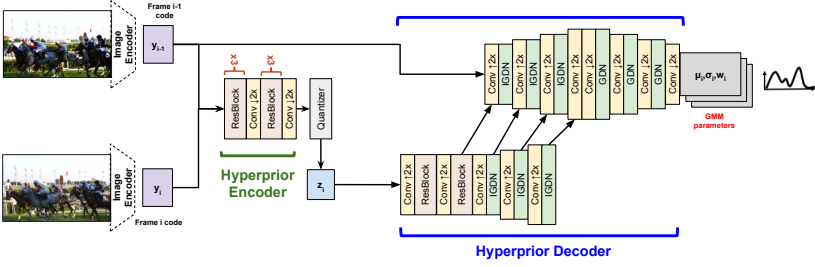


Fig. 4. Diagram of our conditional entropy model, consisting of both a hyperprior encoder (top) and decoder (bottom).

If $\mathbf{y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots\}$ represents the sequence of frame codes for the entire video sequence, then a natural factorization of the joint probability $p(\mathbf{y})$ would be to have every subsequent frame depend on the previous frames:

$$R(\mathbf{y}) \geq \mathbb{E}_{x \sim p_{data}} [\log p(\mathbf{y}; \boldsymbol{\theta})] = \mathbb{E}_{x \sim p_{data}} \left[\sum_{i=0}^n \log p(\mathbf{y}_i | \mathbf{y}_{<i}; \boldsymbol{\theta}) \right] \quad (1)$$

While other approaches (e.g. B-frames) model dependence in a hierarchical manner, our factorization makes sense in online and low-latency settings, where we want to decode frames sequentially. We further make a 1st-order Markov assumption such that each frame \mathbf{y}_i only depends on the previous frame \mathbf{y}_{i-1} ¹ and a small hyperprior code \mathbf{z}_i . Note that \mathbf{z}_i counts as side information, inspired from [9], and must also be counted in the bitstream. We encode it with a hyperprior encoder with \mathbf{y}_i and \mathbf{y}_{i-1} as input (see Fig. 4). We thus have

$$R(\mathbf{y}) \geq \mathbb{E}_{x \sim p_x} \left[\sum_{i=0}^n \log p(\mathbf{y}_i | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}) + \log p(\mathbf{z}_i; \boldsymbol{\theta}) \right]$$

We assume that the hyperprior code distribution $p(\mathbf{z}_i; \boldsymbol{\theta})$ is modeled as a factorized distribution, $p(\mathbf{z}_i; \boldsymbol{\theta}) = \prod_j p(z_{ij} | \boldsymbol{\theta}_z)$, where j represents each dimension of \mathbf{z}_i . Since each z_{ij} is a discrete value, we design each $p(z_{ij} | \boldsymbol{\theta}_z) = c_j(z_{ij} + 0.5; \boldsymbol{\theta}_z) - c_j(z_{ij} - 0.5; \boldsymbol{\theta}_z)$, where each $c_j(\cdot; \boldsymbol{\theta}_z)$ is a cumulative density function (CDF) parametrized as a neural network similar to [9]. In the meantime, we also model each $p(\mathbf{y}_i | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta})$ as a conditional factorized distribution: $\prod_j p(y_{ij} | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta})$, with $p(y_{ij} | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}) = g_j(y_{ij} + 0.5 | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}_y) - g_j(y_{ij} - 0.5 | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}_y)$, where g_j is modeled as the CDF of a Gaussian mixture model: $\sum_k w_{jk} \mathcal{N}(\mu_{jk}, \sigma_{jk}^2)$. $w_{jk}, \mu_{jk}, \sigma_{jk}$ are all learned parameters depending on $\mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}_y$. Similar to [9, 19, 23], the GMM parameters are outputs of a deep hyperprior decoder.

Note that our entropy model is *not* autoregressive either at the pixel level or the frame level - mixture parameters for each latent “pixel” y_{ij} are predicted

¹ When $i = 0$, \mathbf{y}_{i-1} doesn’t exist and can be represented as a zeroed-out vector.

independently given $\mathbf{y}_{i-1}, \mathbf{z}_i$, hence requiring only one GPU pass per frame during decoding. Also, all \mathbf{y}_i 's are produced independently with our image encoder, removing the need to specify keyframes. All these aspects are advantageous in designing a fast, online video compressor. Yet we also aim to make our model expressive such that our prediction for each pixel $p(y_{ij}|\mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta})$ can incorporate both local and global structure information surrounding that pixel.

We illustrate this architecture in Fig. 4. Our hyperprior encoder encodes our hyperprior code \mathbf{z}_i as side information given \mathbf{y}_i and \mathbf{y}_{i-1} as input. Then, our hyperprior decoder takes \mathbf{z}_i and \mathbf{y}_{i-1} as input to predict the Gaussian mixture parameters for \mathbf{y}_i : $\boldsymbol{\sigma}_i, \boldsymbol{\mu}_i$, and \mathbf{w}_i . We can effectively think of \mathbf{z}_i as providing supplemental information to \mathbf{y}_{i-1} to better predict \mathbf{y}_i . The hyperprior decoder first upsamples \mathbf{z}_i to the spatial resolution of \mathbf{y}_{i-1} with residual blocks; then, it uses deconvolutions and IGDN nonlinearities [7] to progressively upsample both \mathbf{y}_{i-1} and \mathbf{z}_i to different resolution feature maps, and fuses the \mathbf{z}_i feature to the \mathbf{y}_{i-1} at each corresponding upsampled resolution. This helps to incorporate changes between \mathbf{y}_{i-1} to \mathbf{y}_i , encapsulated by \mathbf{z}_i , at multiple resolution levels from more global features at the lower resolution to finer features at higher resolutions. Then, downsampling convolutions and GDN nonlinearities are applied to match the original spatial resolution of the image code and produce the mixture parameters for each pixel of the code.

3.3 Rate-distortion Loss Function

We train our base compression models end-to-end to minimize the rate-distortion tradeoff objective used for lossy compression:

$$L(x) = \underbrace{\mathbb{E}_{x \sim p_{data}} \left[\sum_{i=0}^n \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|^2 \right]}_{\text{Distortion}} + \lambda \underbrace{\mathbb{E}_{x \sim p_{data}} \left[\sum_{i=0}^n \log p(\mathbf{y}_i | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}) + \log p(\mathbf{z}_i; \boldsymbol{\theta}) \right]}_{\text{Rate}} \quad (2)$$

where each $\mathbf{x}_i, \hat{\mathbf{x}}_i, \mathbf{y}_i, \mathbf{z}_i$ is a full/reconstructed video frame and code/hyperprior code respectively. The first term describes the reconstruction quality of the decoded video frames, and the second term measures the bitrate as approximated by our conditional entropy model. Each $\mathbf{y}_i, \hat{\mathbf{x}}_i$ is produced via our image encoder/decoder, while our conditional entropy model captures the dependence of \mathbf{y}_i on $\mathbf{y}_{i-1}, \mathbf{z}_i$. We can additionally clamp the rate term to enforce a target bitrate R_a : $\max(\mathbb{E}_{x \sim p_{data}} [\sum_{i=0}^n \log p(\mathbf{y}_i | \mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}) + \log p(\mathbf{z}_i; \boldsymbol{\theta})], R_a)$.

4 Internal Learning of the Frame Code

We additionally propose an internal learning extension of our base model, which leverages every frame of a test video sequence as its own example for which we can learn a better encoding, helping to provide more gains in rate-distortion performance with our entropy-focused approach.

The goal of a compression algorithm is to find codes that can later be *decoded* according to a codebook that does not change during encoding. This is also intuitively why we can not overfit our entire compression architecture to a single frame in a video sequence; this would imply that every video frame would require a separate decoder to decode. However, we make the observation that in our models, the trained decoder/hyperprior decoder represent our codebook; hence as long as the decoder and hyperprior decoder parameters remain fixed, we can actually optimize the encoder/hyperprior parameters or the latent codes themselves, \mathbf{y}_i and \mathbf{z}_i , for every frame during inference. In practice we do the latter to reduce the number of parameters to optimize.

One benefit of internal learning in our video compression setting is similar to that suggested by Campos et al. [10]: the test distribution during inference is oftentimes different than the training distribution. This is especially true in videos, where the test distribution may have different artifacts, framerate, etc. Our base conditional entropy model may predict a higher entropy for test videos due to distributional shift - internal learning might help account for the shortcomings of out-of-distribution prediction by the encoder/hyperprior encoder.

The second benefit is unique to our video compression setting: we can optimize each frame code to reduce the joint entropy in a way that the base approach cannot. In the base approach, there is a restriction of assuming that \mathbf{y}_i is produced by an independent single-image compression model without accounting for past frames as input. Yet there exist configurations of \mathbf{y}_i with the same reconstruction quality that are more easily predictive from \mathbf{y}_{i-1} in our entropy model $p(\mathbf{y}_i|\mathbf{y}_{i-1}, \mathbf{z}_i)$. Performing internal learning allows us to more effectively search for a more optimal configuration of the frame code and hyperprior code \mathbf{z}_i^* , \mathbf{y}_i^* such that \mathbf{y}_i^* can be more easily predicted by \mathbf{y}_{i-1}^* , \mathbf{z}_i^* in the entropy model. As a result, internal learning helps open up a wider search space of frame codes that can potentially have a lower joint entropy.

To perform internal learning during inference, we optimize against a similar rate-distortion loss as in Eq. 2:

$$L_{\text{internal}}(\mathbf{x}) = \sum_{i=0}^n \ell(\mathbf{x}_i, \hat{\mathbf{x}}_i) + \lambda \left[\sum_{i=0}^n \log p(\mathbf{y}_i|\mathbf{y}_{i-1}, \mathbf{z}_i; \boldsymbol{\theta}) + \log p(\mathbf{z}_i; \boldsymbol{\theta}) \right] \quad (3)$$

where \mathbf{x} denotes the test video sequence that we optimize over, and ℓ represents the reconstruction loss function. We first initialize \mathbf{y}_i and \mathbf{z}_i as the output from the trained encoder/hyperprior encoder. Then we backpropagate gradients from (Eq. 2) to \mathbf{y}_i and \mathbf{z}_i for a set number of steps, while keeping all decoder parameters fixed. We can additionally customize λ in Eq. (3) depending on whether we want to tune more for bitrate or reconstruction. If the newly optimized codes are denoted as \mathbf{y}_i^* and \mathbf{z}_i^* , then we simply store \mathbf{y}_i^* and \mathbf{z}_i^* during encoding and discard the original \mathbf{y}_i and \mathbf{z}_i .

We do note that internal learning during inference prevents the ability to perform parallel frame encoding, since \mathbf{y}_i^* , \mathbf{z}_i^* now depend on \mathbf{y}_{i-1}^* as an output of internal learning rather than the image encoder; the gradient steps also increase

the encoding runtime per frame. However, after \mathbf{z}_i , \mathbf{y}_i are optimized, they are fixed during decoding, and hence decoding runtime does not increase. We analyze the tradeoff of increased computation vs. reduced bitrates in the next section.

5 Experiments

We present a detailed analysis of our video compression approach on numerous datasets, varying factors such as frame-rate and video codec quality.

5.1 Datasets, Metrics, and Video Codecs

Kinetics, CDVL, and UVG, and others: We train on the Kinetics dataset [11]. Then, we benchmark our method against standard video test sets which are commonly used for evaluating video compression algorithms. Specifically, we run evaluations on video sequences from the Consumer Digital Video Library (CDVL) [2] as well as the Ultra Video Group (UVG) [4]. UVG consists of 7 video sequences of 3900 frames total, each 1920×1080 and 120fps. Our CDVL dataset consists of 78 video sequences, each 640×480 and either 30fps or 60fps. The videos span a wide range of natural image settings as well as motion. For further analysis of our approach, we benchmark on video sequences from MCL-JVC [36] and Video Trace Library (VTL) [5], shown in supplementary material.

NorthAmerica: We collect a video dataset by driving our self-driving fleet in several North American cities and collecting monocular, frontal camera data. The framerate is 10 Hz. Our training set consists of 1160 video sequences of 300 frames each, and our test set consists of 68 video sequences of 300 frames each. All frames are 1920×1200 in resolution, and we train on 240×150 crops. We focus both on full street driving sequences as well as only on sequences where ego-vehicle is moving (so no red-lights or stop signs).

Metrics: We measure runtime in milliseconds on a per-frame basis for both encoding and decoding. Moreover we plot the rate-distortion curve for multi-scale structural similarity (MS-SSIM) [37], which is a commonly-used perceptual metric that captures the overall structural similarity in the reconstruction. We report the MS-SSIM curve at log-scale similar to [9], where log-scale is defined as $-10 \log_{10}(1 - \text{MS-SSIM})$. Additionally, we report some curves using PSNR: $-10 \log_{10}(\text{MSE})$, where MSE is mean-squared error, and hence measuring the absolute error in the reconstructed image.

Note from Sect. 3.3 that all our base models are trained/optimized with mean-squared error (MSE), and we show that our base models are robust in both MS-SSIM and PSNR. However with internal learning, we demonstrate the flexibility of tuning to different metrics during test-time, so we optimize reconstruction loss towards MS-SSIM and MSE separately (see Sect. 4).




Fig. 5. Rate-distortion plot of our model against competing deep compression works [14, 20, 39]. Results are on full 1920×1080 UVG video.

Video Codecs and Baselines: We benchmark with both libx265 (HEVC/H.265) and libx264 (AVC/H.264). To the best of our knowledge all prior works on learned video compression [12, 14, 20, 25, 39] have artificially restricted codec performance either by using a faster setting or by imposing additional limitations on the codecs (such as removing B-frames). In contrast, we benchmark both H.265 and H.264 on the *veryslow* setting in ffmpeg in order to maximize the performance of these codecs. For the sake of illustration (and also to have a consistent comparison in Fig. 5 with other authors) we also plot H.265 with the *medium*




Fig. 6. Demonstration of our approach vs H.265/H.264 on 10 Hz 1920×1200 NorthAmerica video, 12 Hz 1920×1080 UVG video, and 6 Hz 640×480 CDVL video. Even at lower bitrates, our approach demonstrates significant reductions in bitrate and distortion on lower framerate video.

preset for benchmarking. We additionally include the official HEVC HM [3] and AVC JM [1] implementations in Fig. 5. We also incorporate corresponding numbers from the learned compression methods of [14, 20, 39]. Finally, we add our single-image compression model, inspired by [17], as a baseline.

In addition, we remove Group of Picture (GoP) restrictions when running H.265/H.264, such that the maximum GoP size is equivalent to the total number of frames of each video sequence. We note that neither our base approach nor internal learning require an explicit notion of GoP size: in the base approach, every frame code is produced independently with an image encoder, and with internal learning we optimize every frame sequentially.

Implementation Details: We use a learning rate of $7 \cdot 10^{-5}$ to $2 \cdot 10^{-4}$ for our models at different bitrates, and optimize parameters with Adam. We train with a batch size of 4 on two GPU’s. For test/runtime evaluations, we use a single Intel Xeon E5-2687W CPU and 1080Ti GPU. For internal learning we run 10-12 steps of gradient descent per frame. Our range coding implementation is written in C++ interfacing with Python; during encoding/decoding we compute the codes and distributions on GPU, then pass the data over to our C++ implementation.

5.2 Runtime and Rate-distortion on UVG

We showcase runtime vs. MS-SSIM plots of our method (both the base model and internal learning extension) against related deep compression works on UVG 1920×1080 video: Wu et al. [39], Lu et al. [20], and Habibian et al. [14].² Results are shown in Fig. 1, and detail the frame encoding/decoding runtimes on GPU excluding the specific entropy coding implementation.³

Overall our base approach is significantly faster than most deep compression works. During decoding, our base approach is *orders of magnitude* faster than approaches that use an autoregressive entropy model (Habibian et al. [14], Wu et al. [39]). We note that closest works in the GPU runtime and MS-SSIM is Lu et al., [20] who reported **666 ms/556 ms** for encoding/decoding. Nevertheless, our GPU-only pass is still faster (**340 ms/ 191 ms** for encoding/decoding). Our entropy coding implementation has room for optimization; the C++ algorithm itself is fast (**140 ms/139 ms** for range encoding/decoding of a 1080p frame) though the Python binding interfacing brings the time up to **1.19/0.65 s** for encoding/decoding. Additionally, we benchmark against prior works’ entropy coding runtime as well as codec runtime in supplementary.

While the optional internal learning extension improves the rate-distortion trade-off in all the benchmarks, it brings overhead in encoding runtime. We note that our implementation of internal learning is unoptimized with the backward operator in PyTorch. However, it brings no overhead in decoding runtime, meaning our approach is still faster than all other approaches during decoding.

² We don’t show the results of Djelouah et al. [12] and Rippel et al. [25] because we were unable to get consistent MS-SSIM metrics on the UVG dataset.

³ We thank the authors for providing us detailed runtime information.

In addition, we evaluate all the competing algorithms’ performance and plot the rate-distortion curve on UVG test dataset, as shown in Fig. 5. The results demonstrate that our approach is competitive or even outperforms existing approaches, especially on MS-SSIM.⁴ Between bitrate ranges 0.1–0.3, which is where other deep baselines present their numbers, our base approach is as competitive as a motion-compensation approach [20] or one that uses autoregressive entropy models [14]. At higher bitrates, the base approach outperforms H.265 *veryslow* in both MS-SSIM and PSNR. Internal learning further improves upon all bitrates by $\sim 10\%$.

5.3 Rate-distortion on NorthAmerica




Fig. 7. Plot of our approaches compared against compression baselines for NorthAmerica, both over the entire dataset as well as only when the ego-vehicle has positive velocity.

We show our conditional entropy model and internal learning extension on the NorthAmerica dataset, in Fig. 7. The graph shows that even our single-image Liu model baseline [17] outperforms H.265 on MS-SSIM at higher bitrates and approaches H.265 in PSNR. Our conditional entropy model demonstrates bitrate improvements of 20–50% across bitrates, and internal learning demonstrates an additional 10% improvement.

Figure 7 also shows graphs in which we only analyze video sequences where the autonomous vehicle is in motion, which creates a fairly large gap in H.265 performance. In this setting, both our video compression algorithm as well as the single-image model outperform H.265 by a wide margin on almost all bitrates in MS-SSIM and at higher bitrates in PSNR.

5.4 Varying Framerates on UVG and CDVL

We can additionally control the framerate by dropping frames for CDVL and UVG. We follow a scheme of keeping 1 out of every n frames, denoted as $/n$. We analyze UVG and CDVL video in $1/3$, $1/6$, and $1/10$ settings. Since all UVG videos are 120 Hz, the corresponding framerates are 40 Hz, 20 Hz, 12 Hz.

⁴ HEVC HM performs much better in PSNR, but as we see in supplementary, HEVC HM/AVC JM are *significantly slower* than ffmpeg codecs and our method.

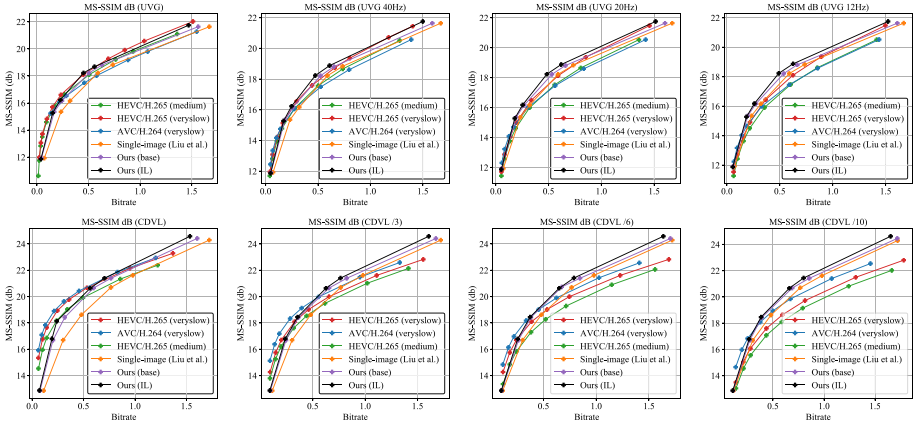


Fig. 8. Plot of our conditional entropy + internal learning adaptations against various baselines for UVG and CDVL. A separate graph is shown for each framerate.

The effects of our conditional entropy model and internal learning, evaluated at different framerates, are shown in separate graphs, in Fig. 8. The conditional entropy model is competitive with H.265 at the original framerate for UVG, and outperforms video codecs at lower framerates. In fact, we found that *single-image compression* matches H.265 *veryslow* on lower framerates! We find a similar effect on CDVL at lower framerates as well, where both single-image compression and our approach far outperform H.265 at lower framerates.

Our base conditional entropy model generally demonstrates a 20%–50% reduction of bitrate compared to the single-image model. The effect of internal learning on each frame code provides an additional 10–20% reduction in bitrate, demonstrating that internal learning of the latent codes during the inference stage provides additional gains.

5.5 Qualitative Results

We showcase qualitative outputs of our model vs H.265 and H.264 *veryslow* in Fig. 6, demonstrating the power of our model on lower framerate video. On 10Hz NorthAmerica, 12Hz UVG video, and 6Hz CDVL video, our model contains big reductions in bitrate compared to video codecs, while producing results that are more even and with fewer artifacts.

6 Conclusion

We propose a novel entropy-focused video compression architecture consisting of a base conditional entropy model as well as an internal learning extension. Rather than explicitly transforming information across frames as in prior work, our model aims to model the correlations between each frame code, as well as

perform internal learning of each frame code during inference to better optimize this entropy model. We show that our lightweight, entropy-focused method is competitive with prior work and video codecs as well as being much faster and conceptually easier to understand. With internal learning, our approach outperforms H.265 in numerous video settings, especially at higher bitrates and lower framerates. Our adaptations are anchored against single-image compression which is robust against varied framerates, whereas video codecs such as H.265/H.264 are not. Hence, we demonstrate that such a video compression approach can have wide applicability in a variety of settings.

References

1. Avc jm reference software. <http://iphone.hhi.de/suehring/>. Accessed 01 May 2020
2. Consumer digital video library. <https://www.cdvl.org/>. Accessed 01 Nov 2019
3. Hvc hm reference software. <https://vcgit.hhi.fraunhofer.de/jct-vc/HM>. Accessed 01 May 2020
4. Ultra video group. <http://ultravideo.cs.tut.fi/#testsequences>. Accessed 01 Nov 2019
5. Video trace library. <http://trace.kom.aau.dk/>. Accessed 01 Nov 2019
6. Shocher, A., Cohen, N., Irani, M.: “zero-shot” super-resolution using deep internal learning. In: CVPR (2018)
7. Ballé, J., Laparra, V., Simoncelli, E.P.: Density modeling of images using a generalized normalization transformation. ArXiv (2015)
8. Ballé, J., Laparra, V., Simoncelli, E.P.: End-to-end optimized image compression. In: ICLR (2017)
9. Ballé, J., Minnen, D., Singh, S., Hwang, S.J., Johnston, N.: Variational image compression with a scale hyperprior. In: ICLR (2018)
10. Campos, J., Meierhans, S., Djelouah, A., Schroers, C.: Content adaptive optimization for neural image compression (2019)
11. Carreira, J., Zisserman, A.: Quo vadis, action recognition? A new model and the kinetics dataset (2017)
12. Djelouah, A., Campos, J., Schaub-Meyer, S., Schroers, C.: Neural inter-frame compression for video coding. In: ICCV (2019)
13. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: ICCV (2009)
14. Habibian, A., van Rozendaal, T., Tomczak, J.M., Cohen, T.S.: Video compression with rate-distortion autoencoders. In: ICCV (2019)
15. Han, J., Lombardo, S., Schroers, C., Mandt, S.: Deep generative video compression (2019)
16. Lee, J., Cho, S., Beack, S.K.: Context-adaptive entropy model for end-to-end optimized image compression. In: ICLR (2019)
17. Liu, H., Chen, T., Guo, P., Shen, Q., Cao, X., Wang, Y., Ma, Z.: Non-local Attention Optimized Deep Image Compression. ArXiv (2019)
18. Liu, H., Chen, T., Lu, M., Shen, Q., Ma, Z.: Neural Video Compression using Spatio-Temporal Priors. ArXiv (2019)
19. Liu, J., Wang, S., Urtasun, R.: Dsic: deep stereo image compression. In: ICCV (2019)
20. Lu, G., Ouyang, W., Xu, D., Zhang, X., Cai, C., Gao, Z.: Dvc: An end-to-end deep video compression framework. In: CVPR (2019)

21. Mentzer, F., Agustsson, E., Tschannen, M., Timofte, R., Gool, L.V.: Conditional probability models for deep image compression. In: CVPR (2018)
22. Mentzer, F., Agustsson, E., Tschannen, M., Timofte, R., Gool, L.V.: Practical full resolution learned lossless image compression. In: CVPR (2019)
23. Minnen, D., Ballé, J., Toderici, G.: Joint autoregressive and hierarchical priors for learned image compression. In: NIPS (2018)
24. Rippel, O., Bourdev, L.: Real-time adaptive image compression. In: ICML (2017)
25. Rippel, O., Nair, S., Lew, C., Branson, S., Anderson, A.G., Bourdev, L.: Learned video compression. In: ICCV (2019)
26. Shannon, C.E.: A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948)
27. Shocher, A., Bagon, S., Isola, P., Irani, M.: Ingan: capturing and remapping the “DNA” of a natural image. In: ICCV (2019)
28. Sullivan, G.J., Ohm, J.R., Han, W.J., Wiegand, T.: Overview of the high efficiency video coding (HEVC) standard. *IEEE Trans Circuits Syst. Video Technol.* **22**, 1649–1668 (2012)
29. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A.A., Hardt, M.: Test-time training for out-of-distribution generalization (2019)
30. Tamar Rott Shaham, Tali Dekel, T.M.: Singan: Learning a generative model from a single natural image. In: ICCV (2019)
31. Theis, L., Shi, W., Cunningham, A., Huszar, F.: Lossy image compression with compressive autoencoders. In: ICLR (2017)
32. Toderici, G., O’Malley, S.M., Hwang, S.J., Vincent, D.: Variable rate image compression with recurrent neural networks. In: ICLR (2016)
33. Toderici, G., Vincent, D., Johnston, N., Hwang, S.J., Minnen, D., Shor, J., Covell, M.: Full resolution image compression with recurrent neural networks. In: CVPR (2017)
34. Ulyanov, D., Vedaldi, A., Lempitsky, V.S.: Deep image prior. In: CVPR (2018)
35. Vapnik, V.N.: *The Nature of Statistical Learning Theory*. Springer, Heidelberg (1995). <https://doi.org/10.1007/978-1-4757-2440-0>
36. Wang, H., et al.: MCL-JCV: a JND-based H.264/AVC video quality assessment dataset. In: ICIP (2016)
37. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: ACSSC (2003)
38. Wiegand, T., Sullivan, G.J., Bjontegaard, G., Luthra, A.: Overview of the H.264/AVC video coding standard. *IEEE Trans. Circuits Syst. Video Technol.* **13**, 560–576 (2003)
39. Wu, C.-Y., Singhal, N., Krähenbühl, P.: Video compression through image interpolation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018, Part VIII. LNCS, vol. 11212, pp. 425–440. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01237-3_26
40. Zhou, Y., Zhu, Z., Bai, X., Lischinski, D., Cohen-Or, D., Huang, H.: Non-stationary texture synthesis by adversarial expansion. In: SIGGRAPH (2018)