

Learned Video Compression via Joint Spatial-Temporal Correlation Exploration

Haojie Liu,¹ Han Shen,² Lichao Huang,² Ming Lu,¹ Tong Chen,¹ Zhan Ma^{1*}

¹Nanjing University, ²Horizon Robotics

{haojie, luming, tong}@smail.nju.edu.cn, {han.shen, lichao.huang}@horizon.ai, mazhan@nju.edu.cn

Abstract

Traditional video compression technologies have been developed over decades in pursuit of higher coding efficiency. Efficient temporal information representation plays a key role in video coding. Thus, in this paper, we propose to exploit the temporal correlation using both first-order optical flow and second-order flow prediction. We suggest an one-stage learning approach to encapsulate flow as quantized features from consecutive frames which is then entropy coded with adaptive contexts conditioned on joint spatial-temporal priors to exploit second-order correlations. Joint priors are embedded in autoregressive spatial neighbors, co-located hyper elements and temporal neighbors using ConvLSTM recurrently. We evaluate our approach for the low-delay scenario with High-Efficiency Video Coding (H.265/HEVC), H.264/AVC and another learned video compression method, following the common test settings. Our work offers the state-of-the-art performance, with consistent gains across all popular test sequences.

Introduction

Video content occupied more than 70% Internet traffic, and it became a big challenge for transmission and storage along its explosive growth. Researchers, engineers, etc, continuously pursue the (next-generation) high-efficiency video compression for wider application enabling and larger market adoption. Conventional video compression approaches usually follow the hybrid coding framework over decades (Sullivan and Wiegand 2005) with hand-crafted tools for individual components. It is not efficient to jointly optimize the system in an end-to-end manner, especially for the inter tool efficiency exploration despite its great success of H.264/AVC (Wiegand et al. 2003) and H.265/HEVC (Sullivan et al. 2012).

Recently, image compression algorithms (Ballé, Laparra, and Simoncelli 2016; Li et al. 2017; Ballé et al. 2018; Mentzer et al. 2018; Liu et al. 2019) based on machine learning have shown great superiority in coding efficiency for *spatial redundancy* removal, compared with conven-

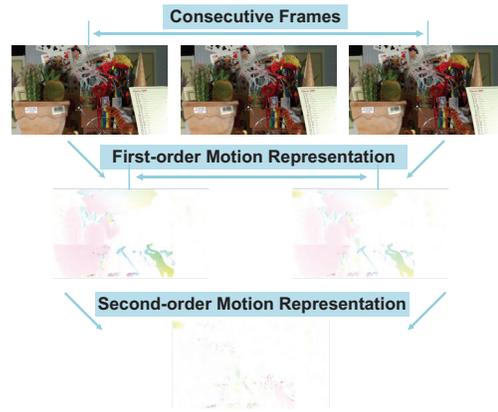


Figure 1: Temporal correlation exploration using both first-order and second-order motion representation. Frame-to-frame redundancies can be easily removed by accurate flow estimation and compensation (e.g., first-order); Then flow-to-flow correlation can be further eliminated or reduced using predictions from joint spatial-temporal priors (e.g., second-order).

tional codecs. These methods benefit from non-linear transforms, deep neural network (DNN) based conditional entropy model, and joint rate-distortion optimization (RDO), under an end-to-end learning strategy. Learned video compression can be extended from image compression by further exploiting the *temporal redundancy or correlation*. We proposed an end-to-end video compression framework using joint spatial-temporal priors to generate compact latent feature representations for intra texture, inter motion and sparse inter residual signals. Intra textures are well represented by spatial priors for both reconstruction and entropy context modeling using a variational autoencoder (VAE) structure (Minnen, Ballé, and Toderici 2018; Liu et al. 2019). We directly use NLAIC method proposed in (Liu et al. 2019) for our intra texture compression because of its state-of-the-art efficiency, and primarily investigate learned inter coding with the focus on efficient temporal motion representation in this paper.

*Zhan Ma is the corresponding author.

We represent temporal information or correlation using its both first-order and second-order statistics. The first-order temporal information is referred to as the motion fields (e.g., intensity, orientation) between consecutive frames. Motion fields can be described by either optical flow or block-based motion vectors. Here, we suggest an one-stage unsupervised flow learning approach, where first-order motions are quantized temporal features learned from the consecutive frames directly. Our unsupervised flow learning does not rely on a well pre-trained optical flow estimation network, such as FlowNet2 (Ilg et al. 2017; Sun et al. 2018), and can derive the compressed optical flow from quantized features directly.

The second-order temporal information is flow-to-flow correlations, describing the object acceleration. Flow can be further predicted for energy reduction. Thus, we fuse priors from spatial, temporal and hyper information to predict flow elements and compress the predictive difference. It turns out that this can be realized by entropy coding with adaptive contexts conditioned on fused priors.

Inter residual is derived between original frame and flow warped prediction. We reuse the intra texture coding network here directly for residual coding. It is worth to point out that VAE structures are applied for all components (e.g., intra, inter, residual) in this paper.

Contributions.

1) We propose an end-to-end video compression method which offers the state-of-the-art performance against traditional video codecs and recent learned approaches with consistent gains across a variety of common test sequences;

2) High-efficient inter coding is achieved by representing temporal correlation using both first-order optical flow and second-order flow predictive difference;

3) First-order flow is offered using an one-stage unsupervised learning and represented by quantized features derived from consecutive frames;

4) Second-order flow predictive difference compression is efficiently solved by entropy coding with adaptive contexts conditioned on fused priors (e.g., autoregressive spatial priors, hyperpriors, and temporal priors propagated using a ConvLSTM (Xingjian et al. 2015)).

Related Work

Learned Image Compression

DNN based image compression approaches generally rely on autoencoders. (Toderici et al. 2016) first proposed to use recurrent autoencoders to progressively encode bits for image compression. Recent years, convolutional autoencoders are studied extensively, including non-linear transforms (e.g., generalized divisive normalization (Ballé, Laparra, and Simoncelli 2016) and non-local attention transforms (NLAIC) (Liu et al. 2019)), differentiable quantization (e.g., soft-to-hard quantization (Mentzer et al. 2018) and uniform noise approximation (Ballé, Laparra, and Simoncelli 2016)), and adaptive entropy model using the Bayesian generative rules (e.g., PixelCNNs (Oord, Kalchbrenner, and Kavukcuoglu 2016) and variational autoencoders (Ballé et al. 2018)). RDO (Sullivan and Wiegand 1998) is applied by

minimizing Lagrangian cost $J = R + \lambda D$, in end-to-end training. Here, R is referred to as entropy rate, and D is the distortion measured by either mean squared error (MSE) or multiscale structural similarity (MS-SSIM).

These approaches demonstrated better coding efficiency against traditional image coders, both objectively and subjectively. In addition, extreme compression are under exploration with adversarial training (e.g., conditional GANs (Agustsson et al. 2018), multi-scale discriminators (Rippel and Bourdev 2017) for satisfied subjective quality at very low bit rates.

Learned Video Compression

Learned video compression (Djelouah et al. 2019; Habibiyan et al. 2019) is a relatively new area. (Chen et al. 2017) first proposed the DeepCoder where DNNs were used for intra texture and inter residual, and block motions were applied using traditional motion estimation for temporal information representation. Inspired by temporal interpolation and prediction (Niklaus, Mai, and Liu 2017; Niklaus and Liu 2018; Huang, Wang, and Wang 2015), (Wu, Singhal, and Krähenbühl 2018) introduced a RNN based video compression framework through frame interpolation, offering comparable performance with H.264/AVC. However, interpolation typically brings structural delay. Recently, unsupervised flow estimation methods (Jason, Harley, and Derpanis 2016; Meister, Hur, and Roth 2018) are introduced to utilize end-to-end learning for predicting optical flow between two frames without leveraging groundtruth flow. The proposed brightness constancy, motion smoothness and bidirectional census loss are proven to be efficient for flow generation. Robust and reliable optical flow derivation methods are emerged, such as FlowNet2 (Ilg et al. 2017) and PWC-Net (Sun et al. 2018). FlowNet2 stacks multiple simple flownet and applies small displacements to correct flow. PWC-Net extends traditional pyramid flow reconstructed rules to learn and generate precise flow faster.

(Lu et al. 2019) replaced the block based motion estimation with a pre-trained FlowNet2 followed by a cascaded autoencoder for flow compression. U-Net alike processing network was used to enhance the quality of predicted frame. In addition, they directly used models in (Ballé et al. 2018) for their intra and residual coding. The entire framework is so called DVC. DVC outperformed H.265/HEVC mainly at high bit rates but the coding efficiency dropped unexpectedly at low bit rates as reported.

All these attempts in learned video compression were trying to represent the temporal information more efficiently.

Proposed Method

Framework Overview

Fig. 2 sketches our learned video compression. Given a group of pictures (GOP) $\mathbb{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_t\}$, we first encode \mathbf{X}_1 using NLAIC-based intra coding in (Liu et al. 2019), having reconstructed frame as $\hat{\mathbf{X}}_1$.

For \mathbf{X}_2 , we first learn the first-order flow for the predicted frame $\hat{\mathbf{X}}_2^p$ using $\hat{\mathbf{X}}_1$. Corresponding $\mathbf{r}_2 = \mathbf{X}_2 - \hat{\mathbf{X}}_2^p$ is encoded using residual coding sharing the same architecture as

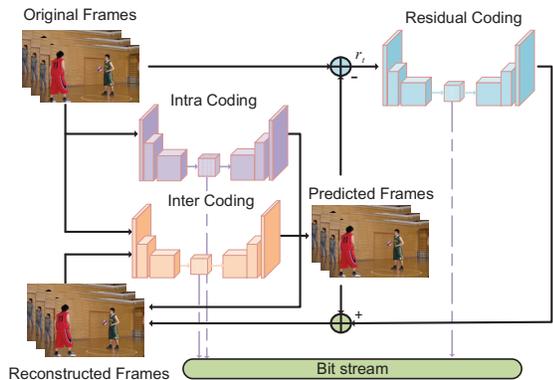


Figure 2: Learned Video Compression. Intra & residual coding using NLAIC in (Liu et al. 2019), and inter coding by exploring one-stage unsupervised flow learning, and quantized flow feature encoding using context adaptive entropy models.

NLAIC-based intra coding, and reconstructed as \hat{r}_2 . First-order flow is represented using quantized features that are entropy-coded with adaptive contexts conditioned on priors, by exploiting the second-order temporal correlation. Final reconstruction \hat{X}_2 is given by $\hat{X}_2 = \hat{X}_2^p + \hat{r}_2$. Subsequent frames follow the same process as of X_2 .

Adaptive entropy models for rate estimation and arithmetic encoding are embedded for intra, inter and residual components. Thanks to the VAE structures, intra and residual coding use the joint autoregressive spatial and hyper priors for the probability estimation; and inter coding applies priors from spatial, hyper and temporal information.

Intra Coding

We directly apply the NLAIC approach in (Liu et al. 2019) for our intra coding. Its state-of-the-art coding efficiency in image compression comes from the introduction of non-local attention transform that are embedded in both main and hyper encoder-decoder network in Fig. 3. Note that the main network is used to obtain the reconstructed frame and the hyperprior network is used for context modeling of adaptive entropy coding.

In NLAIC method, non-local attention modules (NLAM) are embedded to capture joint local and global correlations for both reconstruction and context probability modeling, by inheriting the advantages from both nonlocal processing and attention mechanism. NLAM applies joint spatial-channel attention masks for more compact feature representation. And, masked 3D convolutions are used to fuse hyperpriors and autoregressive priors for accurate context estimation of adaptive entropy coding,

$$(\mu_i, \sigma_i) = \mathbb{F}(\hat{x}_1, \dots, \hat{x}_{i-1}, \hat{z}_t). \quad (1)$$

\mathbb{F} represents cascaded 3D $1 \times 1 \times 1$ convolutions to fuse the priors. $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_{i-1}$ denote the causal (and possibly former reconstructed) pixels prior to current pixel \hat{x}_i obtained by a 3D $5 \times 5 \times 5$ masked convolution and \hat{z}_t are the hyperpriors. Probability of each pixel symbol in \hat{x}_t can be simply

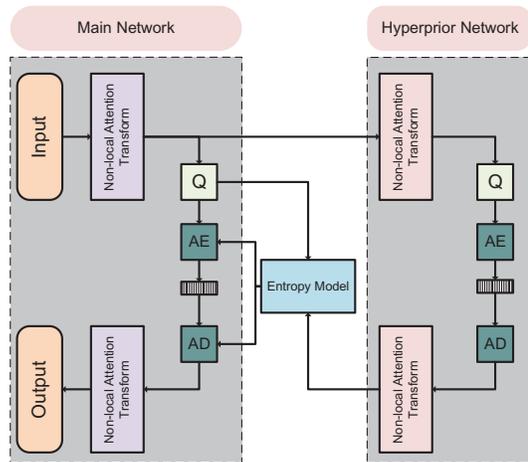


Figure 3: NLAIC-based Intra Coding. A VAE architecture with non-local attention transforms embedded. “AE” represents the arithmetic encoding to encode the quantized features with corresponding probability distribution, “AD” reverts binary strings to feature elements and “Q” is for quantization.

derived using

$$p_{\hat{x}_t | \hat{z}_t}(\hat{x}_t | \hat{z}_t) = \prod_i (\mathcal{N}(\mu_i, \sigma_i^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}))(\hat{x}_i), \quad (2)$$

with a Gaussian distribution assumption with mean (μ) and variance (σ).

Inter Coding

Video coding performance heavily relies on the efficient temporal information representation. It has two folds. One needs to have the most accurate first-order flow for compensation, and the other is to devise second-order statistics for flow prediction.

One-stage Unsupervised Flow Learning Previous work in (Lu et al. 2019) obtained decoded optical flow using typical two-stage methods shown in Fig. 4(a). It relied on a well pre-trained flow network to generate an uncompressed optical flow that was then compressed using a cascaded autoencoder. But, in our work, we leverage the quantized features between consecutive frames as compact motion representations and directly decode the compressed features for subsequent compensation. There is no need to explicitly derive raw and uncompressed flow in encoding process with supervised guidance as in FlowNet2 and PWC-Net. Thus, it is an one-stage unsupervised flow learning and compensation approach in Fig. 4(b).

We concatenated two consecutive frames as the input for feature fusion, and the network is consisted of stacked NLAM¹ and downsampling (e.g., 5×5 convolutions with stride 2), generating the fused features with $(H/16) \times$

¹Nonlocal attention is used to capture local and global correlations.

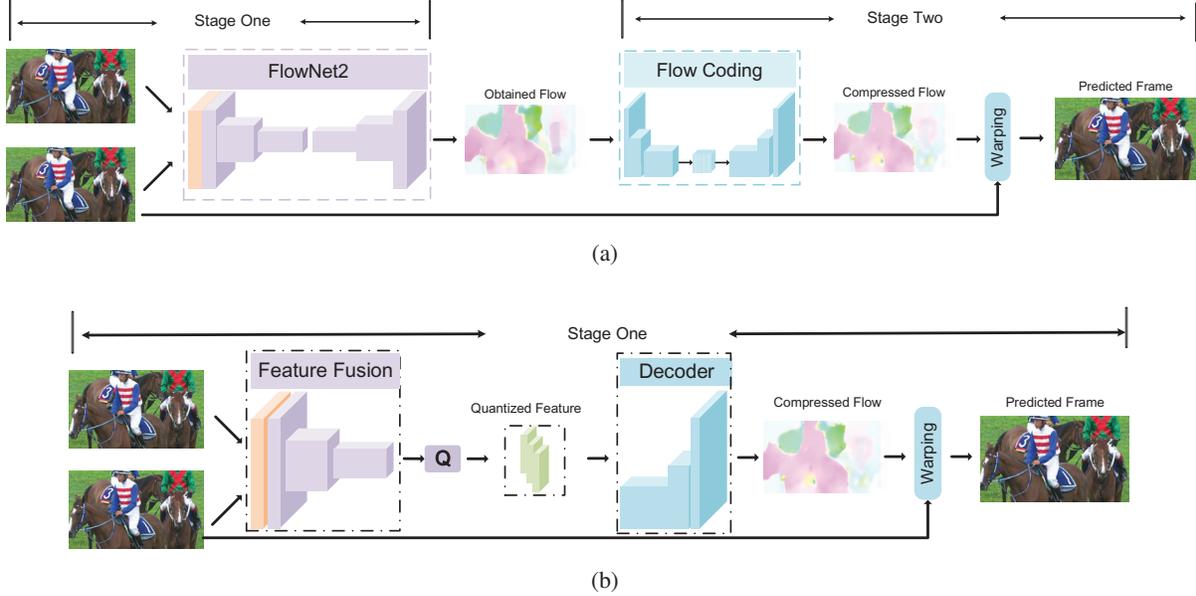


Figure 4: Flow Learning and Compensation. (a) Two-stage supervised approach using a pre-trained flow net (with explicit raw flow) and a cascaded flow compression autoencoder; (b) One-stage unsupervised approach with implicit flow represented by quantized features that will be directly decoded for compensation.

$(W/16) \times 64$ dimension. Quantization is then applied to obtain the quantized features \mathcal{F} for entropy coding. The decoder mirrors the stacked NLAM with upsampling (e.g., 5×5 deconvolutions with stride 2) in the feature fusion network, and derived the decoded flow \hat{f}_t^d (at a size of $H \times W \times 2$ for separable horizontal and vertical orientations) for compensation. Here, H, W denotes the height and width of the original frame, respectively.

To avoid quantization induced motion noise, we first pre-train the network with uncompressed consecutive frames \mathbf{X}_{t-1} and \mathbf{X}_t . Then we replace \mathbf{X}_{t-1} using its decoded correspondence $\hat{\mathbf{X}}_{t-1}$ as described in Eq. (3) and (4). Note that we only have $\hat{\mathbf{X}}_{t-1}$ not \mathbf{X}_{t-1} for inter coding in practice. And we have directly utilized the decoded flow \hat{f}_t^d for end-to-end training and do not need a flow explicitly at encoding (i.e., implicitly represented by quantized features).

A compressed flow representation of \hat{f}_t^d , i.e., quantized features \mathcal{F} , is encoded into the bitstream for delivery. \hat{f}_t^d is then used for warping with reference frame to have $\hat{\mathbf{X}}_t^p$ for compensation, i.e.,

$$\hat{f}_t^d = \mathbb{F}_d(\mathbb{F}_e(\hat{\mathbf{X}}_{t-1}, \mathbf{X}_t)), \quad (3)$$

$$\hat{\mathbf{X}}_t^p = \text{warping}(\hat{\mathbf{X}}_{t-1}, \hat{f}_t^d), \quad (4)$$

Here \mathbb{F}_e and \mathbb{F}_d represent the feature fusion network with quantization and decoder network, respectively. Note that $\mathcal{F} = \mathbb{F}_e(\hat{\mathbf{X}}_{t-1}, \mathbf{X}_t)$.

Context Adaptive Flow Compression Previous section introduced one-stage unsupervised flow learning targeting for the accurate first-order motion representation for compensation, where flow is represented implicitly using quan-

tized features \mathcal{F} . Efficient representation of \mathcal{F} is highly desirable

Generally, as shown in Fig. 1, there is not only the first-order correlation (frame-to-frame) that can be exploited by the flow, but also the second-order correlation both spatially and temporally. A way to predict flow efficiently could lead to much better compression performance. Ideally, flow element can be estimated by its spatial neighbors, temporal co-located element, and hyper priors for energy compaction. A duality problem for such flow prediction using neighbors, is flow element entropy coding using adaptive contexts.

Adaptive context modeling can leverage priors from spatial autoregressive neighbors, temporal and hyper information, shown in Fig. 5. For spatial autoregressive prior, we propose to apply the 3D masked convolutions on quantized features \mathcal{F} ; while for hyper priors, hyper decoder is used to decode corresponding information. Hyperpriors are widely used in VAE structured compression approaches. Note that, temporal correlations are exhibited in video sequence. Instead of applying the only pixel domain frame buffer in traditional video codecs, we propose to embed and propagate flow representation at a frame recurrent way using the ConvLSTM, which is also referred to as the temporal prior buffer. Temporal priors have the same dimension as the current quantized features \mathcal{F} .

These priors are fused together for context adaptive flow coding, i.e.,

$$(\mu_{\mathcal{F}}, \sigma_{\mathcal{F}}) = \mathbb{F}(\mathcal{F}_1, \dots, \mathcal{F}_{i-1}, \hat{\mathbf{z}}_t, \mathbf{h}_{t-1}), \quad (5)$$

$$p_{\mathcal{F}}(\mathcal{F}_1, \dots, \mathcal{F}_{i-1}, \hat{\mathbf{z}}_t, \mathbf{h}_{t-1}) (\mathcal{F}_i | \mathcal{F}_1, \dots, \mathcal{F}_{i-1}, \hat{\mathbf{z}}_t, \mathbf{h}_{t-1}) \\ = \prod_i (\mathcal{N}(\mu_{\mathcal{F}}, \sigma_{\mathcal{F}}^2) * \mathcal{U}(-\frac{1}{2}, \frac{1}{2}))(\mathcal{F}_i). \quad (6)$$

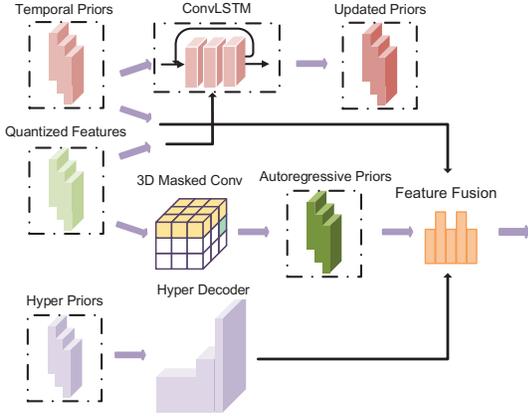


Figure 5: Adaptive Contexts Using Fused Priors. Spatial autoregressive priors are applied using 3D masked convolutions, temporal priors are propagated using convLSTM, and hyperpriors are hyper decoder in a VAE setup. All priors are then fused together for probability modeling. Note that contexts update at a frame recurrent fashion.

$\mathcal{F}_i, i = 0, 1, 2, \dots$ are elements of quantized features for implicit flow \hat{f}_t^d representation, \mathbf{h}_{t-1} is aggregated temporal priors from previous flow representations, which is updated using standard ConvLSTM:

$$(\mathbf{h}_t, \mathbf{c}_t) = \text{ConvLSTM}(\mathcal{F}_t, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}), \quad (7)$$

where $\mathbf{h}_t, \mathbf{c}_t$ are updated state at t and prepared for the next time step slot with \mathbf{c}_{t-1} as a memory gate.

Residual Coding

For the sake of simplicity, we encode the residual signals \mathbf{r}_t using the identical networks as the NLAIC-based intra coding. \mathbf{r}_t is obtained by $\mathbf{r}_t = \mathbf{X}_t - \mathbf{X}_t^p$. Here, we do not calculate the loss between \mathbf{r}_t and $\hat{\mathbf{r}}_t$ but directly target for overall reconstruction loss $\mathbb{D}(\mathbf{X}_t, \mathbf{X}_t^p + \hat{\mathbf{r}}_t)$ for optimizing the residual coding.

Experimental Studies

We proceed to the details about training strategy and evaluation in this section. More ablation studies are given to verify the effectiveness of our work.

Implementation Details

Training & Testing Datasets. We choose COCO (Lin et al. 2014) dataset to pre-train NLAIC-based intra coding network. And then we joint train video compression framework on Vimeo 90k (Xue et al. 2019) which is a widely used dataset for low-level video processing tasks. All images from these the datasets are randomly cropped into 192×192 patches for training.

We gave evaluations on NLAIC using Kodak dataset in ablation studies to understand the efficiency of nonlocal attention transforms. We then evaluated our video compression approaches on standard HEVC dataset and ultra video

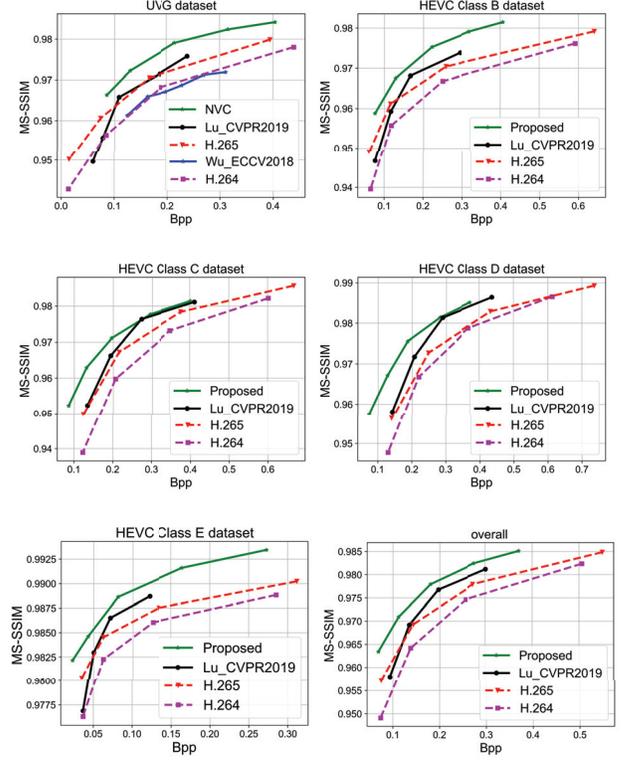


Figure 6: Rate-Distortion Performance Comparison. Experiments are performance for HEVC comment test classes and UVG videos for a variety of content distributions. Our method gains consistently across all test videos.

group (UVG) dataset with different classes, resolution and frame rate.

Loss Function & Training Strategy. It is difficult to train multiple networks on-the-fly at one shot. Thus, we pre-train the intra coding and flow learning and coding networks first, followed by the jointly training with pre-trained network models for an overall optimization, i.e.,

$$L = \frac{\lambda_1}{n} \sum_{t=0}^n \mathbb{D}_1(\hat{\mathbf{X}}_t, \mathbf{X}_t) + \frac{\lambda_2}{n} \sum_{t=0}^n \mathbb{D}_2(\hat{\mathbf{X}}_t^p, \mathbf{X}_t) + R_s + \frac{1}{n-1} \sum_{t=1}^n R_t, \quad (8)$$

where \mathbb{D}_1 is measured using MS-SSIM, and \mathbb{D}_2 is the warping loss evaluated using $L1$ norm and total variation loss. R_s represents the bit rate of intra frame and R_t is the bit rate of inter frames including bits for residual and flow respectively. Currently, λ_1 and λ_2 will be adapted according to the specified overall bit rate and bit consumption percentage of flow information in inter coding.

Besides, entropy rate loss R is approximated by conditional probability p using Eq. (9), with main payload for context adaptive feature elements, and payload overhead for image height (H), width (W), number of frames (N) and

GOP length (n), e.g.,

$$R = -\sum_i \log_2(p) + \text{overhead}. \quad (9)$$

Note that bit consumption for “overhead” is less than 0.1% of the entire compressed bitstream, according to our extensive simulations.

To well balance the efficiency of temporal information learning and training memory consumption, we have enrolled 5 frames to train the video compression framework and shared the weights for the subsequent frames. The initial learning rate (LR) is set to $10e-4$ and is clipped by half for every 10 epochs. The final models are obtained using a LR of $10e-5$. We apply the distributed training on 4 GPUs (Titan Xp) for 5 days.

Evaluation Criteria. For fair comparison, we have applied the same setting as DVC in (Lu et al. 2019) for our method and traditional H.264/AVC, and HEVC codecs. We use GOP of 10 and encode 100 frames on HEVC test sequences and use GOP of 12 with 600 frames on UVG dataset. The reconstructed quality are measured in RGB domain using MS-SSIM. Bits per pixel (Bpp) is used for bit rate measure which can be easily translated to the kbps by scaling it with $H \times W \times n/\tau$. τ is the video duration in seconds.

Performance Comparison

Rate distortion Performance. Our approach outperforms all the existing methods as shown in Fig.6. Here, the distortion is measured by MS-SSIM which is proven to be a more relevant to human visual system, and used widely in learned compression methodologies (Minnen, Ballé, and Toderici 2018).

To the best of our knowledge, our work is the first end-to-end method that outperforms H.265/HEVC consistently across a variety of bit rates for all test sequences. In contrast, algorithm in (Wu, Singhal, and Krähenbühl 2018) only presents a similar performance as H.264/AVC. DVC (Lu et al. 2019) improves (Wu, Singhal, and Krähenbühl 2018) with better coding efficiency against HEVC at high bit rates. However, a cliff fall of performance is revealed for DVC at low bit rate (e.g., some rates having performance even worse than H.264/AVC). We have also observed that DVC’s performance varies for different test sequences. But, our approach shows consistent gains, across contents and bit rates, leading to the conclusion that our model presents better generalization for practical applications.

We use H.264/AVC as the anchor for BD-Rate calculation as shown in Table 1. Our approach reports 50.36% and 51.67% BD-rate reduction on HEVC test sequences and UVG dataset compared with H.264/AVC, respectively, offering a significant performance improvement margin in contrast to the HEVC or DVC over the H.264/AVC.

Visual Comparison We provide the visual quality comparison with H.264/AVC and H.265/HEVC as shown in Fig. 8.

Traditional codecs usually suffer from blocky artifacts, especially at low bit rate, because of its block based coding strategy. Our results eliminate this phenomenon and provide

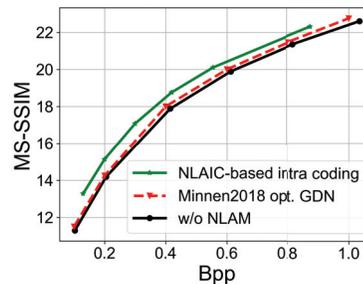


Figure 7: Efficiency of NLAM. NLAM is used in non-local attention transforms for intra & residual coding. Performance is reduced by removing NLAM, but still close to the work in (Minnen, Ballé, and Toderici 2018).

Table 1: BD-Rate Gain of Our Method, HEVC and DVC against the H.264/AVC

Sequences	H.265/HEVC	DVC	Ours
ClassB	-28.31%	-29.09%	-54.17%
ClassC	-20.50%	-28.11%	-39.17%
ClassD	-8.89%	-27.35%	-44.84%
ClassE	-29.52%	-33.91%	-63.28%
Average	-21.73%	-29.73%	-50.36%
UVG dataset	-37.25%	-28.28%	-51.67%

more visually satisfying quality of reconstructed frames. Meanwhile, we need less bits for similar visual quality.

Ablation Study

Non-local Attention Transforms. Most existing image compressions apply Generalized Divisive Normalization (GDN) as non-linear transform to de-correlate spatial-channel redundancy (Minnen, Ballé, and Toderici 2018; Ballé 2018). Alternative non-local attention transform utilizes NLAM to capture both local and global correlations, leading to the state-of-the-art efficiency as reported in (Liu et al. 2019). Algorithms in (Minnen, Ballé, and Toderici 2018) ranks the second place for coding efficiency. Both methods in (Liu et al. 2019) and (Minnen, Ballé, and Toderici 2018) apply the VAE structures. Fig. 7 experiments the efficiency of NLAM, revealing that performance can be retained closely to (Minnen, Ballé, and Toderici 2018) even by removing all nonlocal operations.

Second-order Flow Correlations In Fig.1, we have shown that the second-order motion representations imply the further temporal redundancy between optical flows. Thus, we present a recurrent state (e.g., ConvLSTM) to aggregate temporal priors for inter coding which can effectively reduce the bits for flow compression.

Temporal priors are fused with autoregressive and hyper priors to improve the context modeling of flow element. Then we use ConvLSTM to combine temporal priors with current quantized features for the updated priors in a recurrent way. Flow prediction provides an effective means to exploit the redundancy between complex motion behaviors. Fig. 9 shows that efficiency variations when removing the

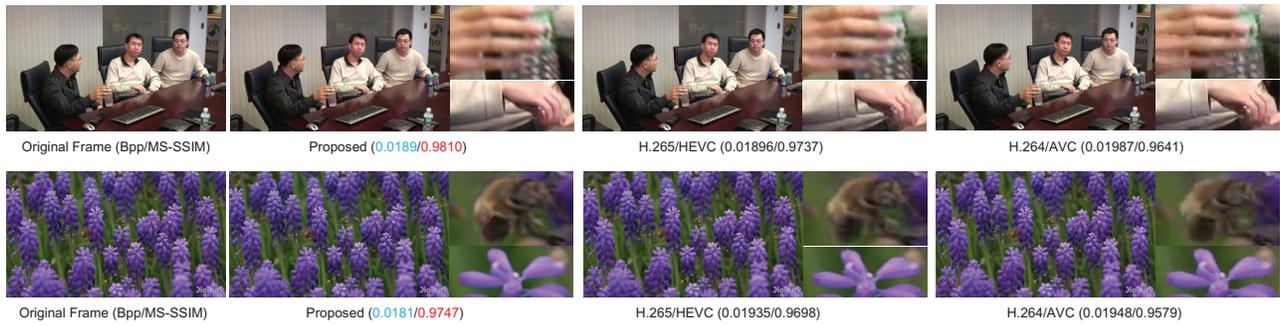


Figure 8: Visual Comparison. Reconstructed frames of our method, H.265/HEVC and H.264/AVC. We avoid blocky artifacts and provide better quality of reconstructed frame at low bit rate.

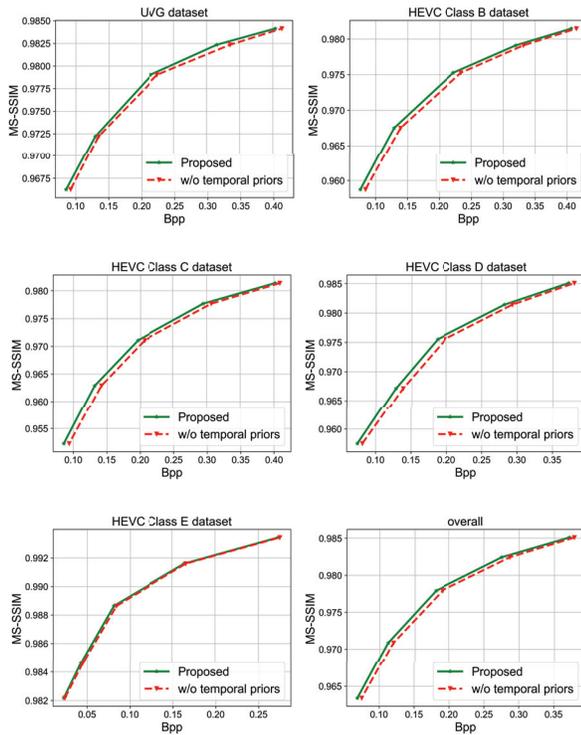


Figure 9: Efficiency of Temporal Priors. 2% - 10% loss captured at similar bit rate when removing temporal priors for context modeling.

ConvLSTM system from our entire framework without exploiting the temporal correlations, where 2% to 10% quality loss is captured.

Generally, bits consumed by motion information varies across different content and bit rates, leading to a variety of percentages to the total bits. More bit saving is revealed for low bit rates, and motion intensive content. For stationary content, such as HEVC Class E, spatial and hyper priors already give a good reference, thus temporal priors are less used.

Conclusions & Future Work

In this paper, we present an end-to-end video compression framework and fully exploit the spatial and temporal redundancies. Key novelty laid on the accurate motion representation for exploiting temporal correlation, via both first-order optical flow learning and second-order flow predictive coding. An one-stage unsupervised flow learning is applied with implicit flow representation using quantized features. These features are then compressed using joint spatial-temporal priors by which the probability model is conditioned adaptively.

We evaluate our methods and report the state-of-the-art performances among all the existing video compression approaches, including traditional H.264/AVC, H.265/HEVC, and learning-based DVC. Our approach offers the consistent gains over existing methods across a variety of contents and bit rates.

As for the future study, an interesting topic is to devise implicit flow without actual bits consumption, such as the decoder-side flow derivation, or frame interpolation and extrapolation. Currently, residual shares the same network with intra coding, which may be worth for deep investigation for network simplification. It is also significant to generalize the whole system to more complex video data sets such as spectral video (Cao et al. 2016) and 3D video (Müller et al. 2013; Cao, Li, and Dai 2011).

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 61571215.

References

- Agustsson, E.; Tschannen, M.; Mentzer, F.; Timofte, R.; and Van Gool, L. 2018. Generative adversarial networks for extreme learned image compression. *arXiv preprint arXiv:1804.02958*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2016. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*.

- Ballé, J. 2018. Efficient nonlinear transforms for lossy image compression. *arXiv preprint arXiv:1802.00847*.
- Cao, X.; Yue, T.; Lin, X.; Lin, S.; Yuan, X.; Dai, Q.; Carin, L.; and Brady, D. J. 2016. Computational snapshot multispectral cameras: Toward dynamic capture of the spectral world. *IEEE Signal Processing Magazine* 33(5):95–108.
- Cao, X.; Li, Z.; and Dai, Q. 2011. Semi-automatic 2d-to-3d conversion using disparity propagation. *IEEE Transactions on Broadcasting* 57(2):491–499.
- Chen, T.; Liu, H.; Shen, Q.; Yue, T.; Cao, X.; and Ma, Z. 2017. Deepcoder: A deep neural network based video compression. In *Visual Communications and Image Processing (VCIP), 2017 IEEE*, 1–4. IEEE.
- Djelouah, A.; Campos, J.; Schaub-Meyer, S.; and Schroers, C. 2019. Neural inter-frame compression for video coding. In *Proceedings of the IEEE International Conference on Computer Vision*, 6421–6429.
- Habibian, A.; Rozendaal, T. v.; Tomczak, J. M.; and Cohen, T. S. 2019. Video compression with rate-distortion autoencoders. In *Proceedings of the IEEE International Conference on Computer Vision*, 7033–7042.
- Huang, Y.; Wang, W.; and Wang, L. 2015. Bidirectional recurrent convolutional networks for multi-frame super-resolution. In Cortes, C.; Lawrence, N. D.; Lee, D. D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 28. Curran Associates, Inc. 235–243.
- Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; and Brox, T. 2017. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, 1647–1655. IEEE.
- Jason, J. Y.; Harley, A. W.; and Derpanis, K. G. 2016. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *European Conference on Computer Vision*, 3–10. Springer.
- Li, M.; Zuo, W.; Gu, S.; Zhao, D.; and Zhang, D. 2017. Learning convolutional networks for content-weighted image compression. *arXiv preprint arXiv:1703.10553*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, H.; Chen, T.; Guo, P.; Shen, Q.; Cao, X.; Wang, Y.; and Ma, Z. 2019. Non-local attention optimized deep image compression. *arXiv preprint arXiv:1904.09757*.
- Lu, G.; Ouyang, W.; Xu, D.; Zhang, X.; Cai, C.; and Gao, Z. 2019. Dvc: An end-to-end deep video compression framework. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Meister, S.; Hur, J.; and Roth, S. 2018. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Mentzer, F.; Agustsson, E.; Tschannen, M.; Timofte, R.; and Van Gool, L. 2018. Conditional probability models for deep image compression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, 3.
- Minnen, D.; Ballé, J.; and Toderici, G. D. 2018. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, 10794–10803.
- Müller, K.; Schwarz, H.; Marpe, D.; Bartnik, C.; Bosse, S.; Brust, H.; Hinz, T.; Lakshman, H.; Merkle, P.; Rhee, F. H.; et al. 2013. 3d high-efficiency video coding for multi-view video and depth data. *IEEE Transactions on Image Processing* 22(9):3366–3378.
- Niklaus, S., and Liu, F. 2018. Context-aware synthesis for video frame interpolation. *arXiv preprint arXiv:1803.10967*.
- Niklaus, S.; Mai, L.; and Liu, F. 2017. Video frame interpolation via adaptive convolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, 3.
- Oord, A. v. d.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*.
- Rippel, O., and Bourdev, L. 2017. Real-time adaptive image compression. *arXiv preprint arXiv:1705.05823*.
- Sullivan, G. J., and Wiegand, T. 1998. Rate-distortion optimization for video compression. *IEEE Signal Processing Magazine* 15(6):74–90.
- Sullivan, G. J., and Wiegand, T. 2005. Video compression - from concepts to the h.264/avc standard. *Proceedings of the IEEE* 93(1):18–31.
- Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video technology* 22(12):1649–1668.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8934–8943.
- Toderici, G.; Vincent, D.; Johnston, N.; Hwang, S.-J.; Minnen, D.; Shor, J.; and Covell, M. 2016. Full resolution image compression with recurrent neural networks. *CoRR abs/1608.05148*.
- Wiegand, T.; Sullivan, G. J.; Bjontegaard, G.; and Luthra, A. 2003. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology* 13(7):560–576.
- Wu, C.-Y.; Singhal, N.; and Krähenbühl, P. 2018. Video compression through image interpolation. In *ECCV*.
- Xingjian, S.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-K.; and Woo, W.-c. 2015. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, 802–810.
- Xue, T.; Chen, B.; Wu, J.; Wei, D.; and Freeman, W. T. 2019. Video enhancement with task-oriented flow. *International Journal of Computer Vision* 127(8):1106–1125.