

Microdosing: Knowledge Distillation for GAN based Compression

Leonhard Helming¹
leonhard.helminger@inf.ethz.ch
Roberto Azevedo²

¹ Department of Computer Science
ETH Zurich, Switzerland
² DisneyResearch|Zurich

Abdelaziz Djelouah²

Markus Gross^{1,2}

Christopher Schroers²

Abstract

Recently, significant progress has been made in learned image and video compression. In particular, the usage of Generative Adversarial Networks has led to impressive results in the low bit rate regime. However, the model size remains an important issue in current state-of-the-art proposals, and existing solutions require significant computation effort on the decoding side. This limits their usage in realistic scenarios and the extension to video compression. In this paper, we demonstrate how to leverage knowledge distillation to obtain equally capable image decoders at a fraction of the original number of parameters. We investigate several aspects of our solution including sequence specialization with side information for image coding. Finally, we also show how to transfer the obtained benefits into the setting of video compression. Altogether, our proposal allows to reduce a decoder model size by a factor of 20 and to achieve 50% reduction in decoding time.

1 Introduction

Initially, projections for 2022 have estimated that video content would reach 82% of internet traffic. While this would already make video the most prevalent form of media in terms of bandwidth by far, the current pandemic situation pushed video traffic even further than these expectations due to the need for social distancing and home officing [1]. As a result, compression techniques are more challenged than ever to handle visual data efficiently, and improvements can impact the daily lives of millions of people.

In contrast to the hand-crafted individual components of traditional codecs, learned image compression schemes aim to learn an optimal non-linear transform from data, ideally in an end-to-end fashion. At a high level, most of the methods can be understood as a sort of *generative model* that tries to reconstruct the input instance from a quantized latent

representation, coupled with a prior that is used to compress these latents losslessly [4]. Although providing good perceptual quality in the high bitrate target setting, it is the low bitrate setting in which neural image compression has shown most of its strength. In particular, GAN (Generative Adversarial Networks)-based architectures for image compression [2, 20] produce impressive results by generating an appropriate hallucination of detail in the output image. As a drawback, GAN-based compression frameworks usually have large decoder models that are many times trained on private datasets. Therefore, retraining these models to their original performance is not generally possible, and even if the data was present, it would not be straightforward and time consuming. In addition, the memory requirements and inference time make them less practical, especially in the context of video coding and mobile devices.

This paper proposes a knowledge distillation (KD) [8] approach that allows retaining good perceptual image quality while reducing the size of the decoder. The goal of KD is to transfer the learned knowledge of a *teacher network* on a smaller *student network* that remains competitive to the teacher network performance. By requiring less memory and computational power than the initial teacher network, the student network could, for instance, run on less powerful devices such as mobile phones or dedicated devices. Being able to compress the generator network or decoder in the auto-encoder setting is not only interesting in terms of memory requirements but also in terms of computational efficiency. This is especially important for image and video compression, where the majority of the computation should preferably be on the sender (encoder) side, while the decoding should be simple. Especially in the context of video streaming, an asset will typically be encoded once while it will be distributed and decoded millions of times. Traditional codecs do spend a lot of emphasis and compute on encoding while keeping decoding extreme lightweight. In contrast to that, deep learning based architectures are typically symmetric in terms of encoding and decoding or can even require considerably more compute for decoding when using a GAN [20].

Our proposal is based on: i) training a reduced *student decoder* with data generated from the *big decoder*. ii) overfitting the reduced *student decoder* model to a specific image or set of images; and iii) sending the specialized decoder weights alongside the image latents. To show the viability of our proposal, we incorporate it into state-of-the-art models for neural image and video compression targeting the low bitrate setting. First, we replace the **High-Fidelity Compression (HiFiC)** [20] decoder with a much smaller student decoder. *HiFiC* is the state-of-the-art in low bitrate neural image compression (~ 0.15 bpp) that produces extremely competitive results at the cost of a fairly big (~ 156 M parameters) decoder network. Our proposed KD approach allows for a much smaller decoder (~ 8 M parameters) and 50% faster decoding time while still producing visually similar output images. Second, we show how to apply our KD strategy in a neural video compression framework based on latent residuals [19]. In such a scenario, we overfit our reduced *student decoder* to a sequence so that we can provide a sequence specific decoder. We show that the additional bits needed for sending the weights amortize over the sequence.

Explicitly, the main contributions of this paper are: i) proposing novel strategies for KD for neural image and video compression; ii) investigating KD in the low bitrate setting for GAN-based image compression; iii) investigating KD in the low bitrate GAN-based video compression setting with latent residuals.

2 Related Work

Neural image compression The first proposed neural image compression methods [8, 26, 27] showed improved results over JPEG or JPEG2000, while most recent approaches [9, 9, 13, 14, 22] are now on par or surpassing BPG [12]. Recent works [24, 23, 24, 25] also investigate the application of perceptual losses and how they relate to the human visual perception of decompressed images.

In particular, generative methods have been providing impressive results in the low bitrate setting [8, 24]. Generative models for image compression are able to synthesize details that would be very costly to store, resulting in visually pleasing results at bitrates in which previous methods show strong artifacts [8]. However, to synthesize those details, a powerful and potentially big decoder is necessary, which contradicts the general requirements on compression technologies, where the sender (encoder) should have the burden of computing a good compression, such that the receiver (decoder) can be as simple as possible. Our approach can be applied to any of the neural image compression methods above, whenever the decoder is too big and faster inference times are required. This is especially crucial in GAN-based neural compression methods. As a use case, we show an application of our proposal using *HiFiC* [24], which is the state-of-the-art GAN model for low bit rate image compression.

Neural video compression As an extension to the neural image compression methods, neural video compression approaches aim to leverage redundancy in both spatial and temporal information. Lu et al. [15] replace blocks in traditional video codecs with neural networks. Learning-based optical flow estimation is used to obtain the motion information and to reconstruct the current frames. Then they employ two auto-encoder networks to compress the corresponding motion and residual information. Veerabadrán et al. [29] show that minimizing an auxiliary adversarial distortion objective for neural video compression in the low bitrate setting creates distortions that better correlate with human perception. Djelouah et al. [11] propose compressing and sending residuals in the latent space instead of residuals in the pixel space, which allows the reuse of the same image compression network for both keyframes and intermediate frames. In this paper, we use such an approach as a starting point to show the viability of KD in neural video compression. Instead of explicitly computing residuals or differences, Ladune et al. [16] use feature space concatenation and train a decoder that operates on this joint information. Rozendaal et al. [28] present an extreme approach that fine tunes the full model to a single video, and sends model updates (quantized and compressed using a parameter-space prior) along with the latent representation. Such an approach is in line with our idea of sequence-specific information to be sent to the student decoder, and our work can be seen as complementary to [28], which does not include KD.

Knowledge distillation KD has been primarily used on vision tasks like object classification or segmentation. KD for generative models, however, is not well studied yet. In [8], the authors leveraged a teacher-student architecture to reduce the size of the BigGAN [7] architecture while still being competitive on Inception and FID scores. To the best of our knowledge, our work is the first to propose knowledge distillation to learn a smaller decoder for neural image and video compression frameworks.

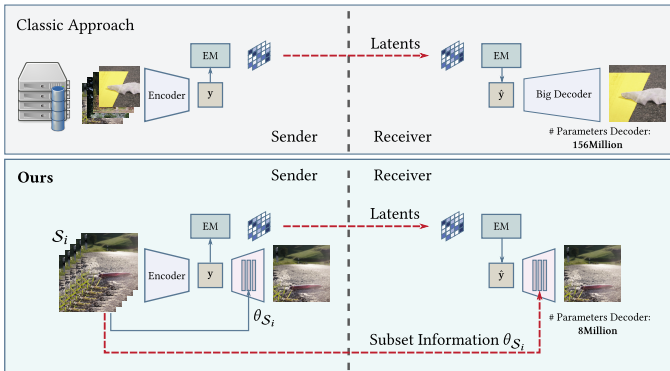


Figure 1: Overview of our proposal: we replace the Big Decoder of GAN compression models (top) with a smaller decoder plus a content-specific additional information (bottom).

3 Knowledge Distillation for Compression

Figure 1 illustrates an overview of our proposal. In the classic neural compression approach, the encoder-decoder pair is trained on a big dataset, to get an overall good performance on a variety of different content. Once the auto-encoder is fully trained, the decoder gets deployed and sent to the receiver. The potentially big decoder then allows to decode any type of content. In our approach, we enable the sender to partition the data into subsets \mathcal{S}_i , and learn a content-specific decoder with corresponding information $\theta_{\mathcal{S}_i}$ for each subset. This specialization allows us to train a model with less parameters, smaller memory footprint, and less computations. Once the decoder is fully trained, and the sender’s reconstruction quality requirement of the subset is fulfilled, the content-specific information is stored alongside the subset. If the receiver wants to decode an image $\mathbf{x} \in \mathcal{S}_i$, the subset specific information $\theta_{\mathcal{S}_i}$ has to be sent once per subset. Next, we detail how to apply our distillation process to image compression with GANs and extend this to video compression using latent space residuals.

3.1 Knowledge Distillation for Image Compression with GANs

High-Fidelity Generative Image Compression (HiFiC) Figure 2 shows the HiFiC architecture [1]. Its decoder can be divided into three sub-nets: *head* (~2M parameters), *res_blocks* (~149M parameters), and *tail* (~5.5M parameters). Through experimentation, and as shown in Figure 3, it is easy to conclude that the coarse information of the image is saved in the latent space, and the hallucination of the texture is generated by the residual network (*res_blocks*) of the decoder. In particular, the discrepancy in size of the the *res_blocks* is due to the fact that the model was trained on a big (private) dataset, thus such a big size is needed to capture all the textures seen during training. However, if we know in advance which images should be compressed (e.g. frames of a video with similar features), we can overfit to this data and sent only the necessary weights to properly decode these images. That is exactly what we propose with the Distilled-HiFiC architecture below.

¹The authors of [1] trained their architecture for different target bit rates and provide these models as: *HiFiC^{Hi}*, *HiFiC^{Mi}* and *HiFiC^{Lo}*.




Figure 2: Architecture of High-Fidelity Generative Image Compression (*HiFiC*)

Distilled-*HiFiC* Our proposed Distilled-*HiFiC* (See Figure 4(a)) reduces the size of the decoder by training a smaller sub-network, named *Micro-Residual-Network (Micro-RN)* that mimics the behavior of the residual network (*res_blocks*) for a specific subset, and therefore *microdosing* the capability of hallucinations. *Micro-RN* is based on the *degradation-aware (DA)* blocks introduced in [50]. While [50] utilizes a kernel prediction network to steer the weights according to a degradation vector, we learn a set of weights θ_{S_i} per subset S_i . *Micro-RN* is defined by two parameters: C_h , the number of hidden channels, and B , the number of DA Blocks. Similar as in [8], we train *Micro-RN* with the teacher-student architecture (Figure 4(a)), with the difference that our decoder borrows pre-trained layers (i.e., *head* and *tail*) from the teacher-decoder. Let $\mathbf{x} \in S_i$ be an image of subset S_i and $\tilde{\mathbf{x}}$ be the image compressed by the teacher network. We optimize the following loss:

$$\mathcal{L}(\mathbf{x}; \theta_S) = k_M \text{MSE}(\tilde{\mathbf{x}}, \hat{\mathbf{x}}) + k_p d_p(\hat{\mathbf{x}}, \mathbf{x}), \quad (1)$$

where $\hat{\mathbf{x}}$ is the output of the student network, MSE and d_p are the distortion losses, and k_M and k_p are their corresponding weights. Similar to [20] as the perceptual loss $d_p = \text{LPIPS}$ [51] is used. Hence, the loss forces the student-decoder to generate images that look similar to the teacher’s and further reduce the perceptual loss to the ground truth image. Note, that we freeze both the encoder and the entropy model which in this case is modeled using a hyperprior. Hence, θ_S only contains the weights of the *Micro-RN*. This allows us to leverage the powerful encoder and hyperprior of *HiFiC* as well as the models knowledge about the private training data set.




Figure 3: Even without residual blocks (RB), the *HiFiC* model outputs a coarse version of the image. The images show the output of *HiFiC* models **with** (a), (c) and **without** (b), (d) residual blocks.




Figure 4: Visualization of the proposed *Micro-RN*. If not stated explicitly, a 3×3 convolution is used. DConv denotes a depthwise convolution and C_h is the number of hidden channels.




Figure 5: Overview of knowledge distillation with latent residuals for video compression: Frame Prediction Network (FPN) and Latent Residual Network (LRN)

3.2 Knowledge Distillation for Video Compression with Latent Space Residuals

To show the application of KD in neural video compression scenarios, we use a similar approach as proposed by Djelouah et al. [14]² (see Figure 5). Such network is composed of two parts: a *Frame Prediction Network (FPN)* and a *Latent Residual Network (LRN)*. Given a sequence of frames (*group of pictures*, or *GOP*) to be encoded $\mathbf{x}_0, \dots, \mathbf{x}_{GOP}$, where \mathbf{x}_0 is a keyframe (I-frame) and $\mathbf{x}_1, \dots, \mathbf{x}_{GOP}$ are predicted frames (P-frames), it works as follows:

First, the I-frame (\mathbf{x}_0) is compressed using a neural image-only compression network to generate the encoded latent \mathbf{y}_0 . $\hat{\mathbf{x}}_0$ denotes the reconstructed frame from the latent \mathbf{y}_0 . Then, for each P-frame, $\mathbf{x}_{t+1}, 1 \leq t+1 \leq GOP$, we: (1) generate a temporal prediction, \mathbf{x}_{t+1}^{Pred} , of \mathbf{x}_{t+1} from the previous reconstructed frame, $\hat{\mathbf{x}}_t$, using *FPN*; *FPN* works by first computing the optical flow \mathbf{f}_{t+1} between \mathbf{x}_{t+1} and $\hat{\mathbf{x}}_t$, warping $\hat{\mathbf{x}}_t$, and then motion compensating it to generate \mathbf{x}_{t+1}^{Pred} (2) compress \mathbf{x}_{t+1}^{Pred} using the neural image compression network to generate the latent of the predicted frame, \mathbf{y}_{t+1}^{Pred} and (3) compute the latent residual, \mathbf{r}_{t+1} , between the latent of the compressed P-frame against the compressed predicted frame, $\mathbf{r}_{t+1} = \mathbf{y}_{t+1} - \mathbf{y}_{t+1}^{Pred}$, which is then quantized and entropy coded with EM_t . The final bitstream of a *GOP* is then composed of $\{\hat{\mathbf{y}}_0, \hat{\mathbf{f}}_1, \dots, \hat{\mathbf{f}}_{GOP}, \hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{GOP}\}$, i.e., latent of the I-frame and the computed flow fields and latent residuals for each of the P-frames (all quantized and entropy encoded)³.

In the low bitrate setting, HiFiC would seem like a suitable choice for the neural image compression architecture that could be used together with the above latent space residual

²The main differences of our implementation compared to Djelouah et al. [14] are: i) we only use P-Frames, while [14] also uses B-Frames; ii) Djelouah et al. [14] retrains the image encoder/decoder from scratch, while we reuse the pre-trained HiFiC encoder/decoder during training time, and replace the HiFiC decoder with our Distilled-HiFiC during inference time; iii) we start with a pre-trained optical flow and allow the flow to be fine tuned during the training of the latent residuals network.

³Details on the different network modules and training procedures are provided in the supplementary material.




Figure 6: Ablation Study: We show that increasing the number of hidden channels C_h or number of DA blocks B allows to achieve a similar performance as $HiFiC^{Lo}$.

framework. As previously mentioned, however, the size of the HiFiC decoder is a limiting factor. Also, inference time is even more important in the video setting that should be able to keep a decoding frame rate of ~ 30 frames per seconds (fps). Thus, we propose to use our Distilled-HiFiC in the latent residual framework above.

During encoding, we overfit our Distilled-HiFiC to a specific sequence so that we only need to send the θ_S once for all the frames of that sequence. Our decoding process starts then by receiving and loading sequence-specific Micro-RN weights on the Distilled-HiFiC decoder, which is then fixed during the decode of the whole sequence. As detailed next, with such a small overhead, we can reduce the decoding time by half while keeping similar visual quality to the latent residual framework using the original HiFiC decoder.

4 Experiments

In this section we first describe the experiment that leads to the minimal architecture that is powerful to mimic $HiFiC$'s residual network on a variety of subsets. In the second part of this section, we show that our student decoder can be used for video compression based on latent space residuals and compare the results with OpenDVC, a recent neural video compression method.

Dataset The models were trained on two types of subsets: \mathcal{S}_{avg} a subset of UVG [20] and \mathcal{S}_i a subset of sequence $i \in \{1, \dots, 7\}$. We created the subsets by taking every 10th frame of each clip. This allows Micro-RN to learn the image features present in the whole clip. \mathcal{S}_{avg} consists in total of 390 frames with resolution 1920×1080 and \mathcal{S}_i of up to 60 frames.

We evaluated the proposed model for qualitative comparisons on an in-house created movie. The subset \mathcal{S}_i corresponds to a specific clip, where \mathcal{S}_{lucid} denotes the union of four clips.

4.1 Micro-Residual-Network

Training We trained our Micro-RN by minimizing Equation 1 on random crops of size 256×256 . We used the hyper parameters k_p and k_M proposed in [20]. We further used the Adam optimizer [15] with a learning rate of 10^{-4} and batch size 4. Training a model takes approximately 10 hours on a NVIDIA Titan Xp.

Ablation Study To find the minimum number of weights for Micro-RN, we trained our student decoder with various number of hidden channels $C_h = \{64, 128, 256\}$ and DA blocks $B = \{1, 2, 3, 4\}$. We are interested in the smallest possible architecture that is still capable to

mimic *HiFiC*'s residual network. *HiFiC^{Lo}* is selected as the teacher network, since it is the model that has the most difficult task of hallucinating details. Table 1 shows the number of parameters per configuration.

The experiments show that increasing either C_h or B results in a higher reconstruction quality, and better mimicking the *HiFiC* decoder. Another interesting aspect is that the required number of hidden channels C_h and number of DA Blocks B depend on the complexity and the details of a subset. For less complex sequences e.g. *Beauty*, the smallest tested configuration with $C_h = 64$ and $B = 1$ is sufficient to outperform *HiFiC^{Lo}* in terms of PSNR and LPIPS (see green curve in Figure 6(a)). Since the sequence consists primarily of black noise background with less features that need to be reconstructed, the Micro-RN do not need to save much details about the texture.

For sequences with many details and a high variation of features e.g. *YachtRide*, more parameters, i.e., more hidden channels and DA blocks are necessary. This can be seen in Figure 6(b). The green curve ($C_h = 64$) is always below the baseline of *HiFiC^{Lo}*. By increasing the number of channels to $C_h = 128$ (orange curve) and $B = 2$ our approach produces similar numbers to the baseline.

Architecture Based on these results and for simplicity, we chose to set the number of hidden channels to $C_h = 128$ and $B = 1$ for all other experiments. Increasing either B or C_h would lead to a double or almost quadratic increase of parameters for higher B , respectively.

Comparison to HiFiC We compare Distilled-*HiFiC* to the original *HiFiC* decoder on: number of parameters, decoding time, visual comparisons, and distortion metrics. Replacing *HiFiC*'s residual network with our proposed Micro-RN reduces the number of parameters of the decoder from 156M to 8M parameters. The proposed Micro-RN itself consists only of 600K parameters. This does not only reduce the memory footprint of the model, but it would also allow to decode images of higher resolution in a single pass. We also conducted a benchmark on the timings for decoding, for various implementations of *HiFiC*. We distinguish between the following models: *HiFiC*, *HiFiC*-(built) and ours. *HiFiC* is the original metagraph provided by [20] and *HiFiC*-(built) is the graph we built ourself during training. In addition we compare against an incomplete decoder which serves as a *Lower Bound*. It has the same architecture as *HiFiC*-(built) but without computing the output of the residual blocks.

Table 2 shows that decoding with our architecture is twice as fast as the original one. It further shows that our model is as fast the lower bound i.e. the added complexity by the Micro-RN is negligible in terms of decoding time. Since the encoder is the same for every setup and encoding time (~ 0.23 s) is similar for each model.

A visual comparison between the ground truth, a distilled decoder on both S_{avg} and S_i , and *HiFiC^{Lo}* is provided in Figure 8. Micro-RN is capable to mimic *HiFiC*'s residual network for both subsets, S_{avg} and S_i . Further, if we compare both models, the model which was trained only on the sequence, i.e. S_i , learns more details and better adapt to the specific

		C_h		
		64	128	256
B	1	211k	594k	1.88M
	2	294k	925k	3.19M
	3	378k	1.25M	4.51M
	4	461k	1.59M	5.82M

Table 1: Number of parameters per tested configuration.

		Decoding [s]
Lower Bound		0.20
HiFiC (pruned)		0.44
HiFiC (built)		0.45
Ours		0.20

Table 2: Timings for decoding.




Figure 7: Comparison to *HiFiC*: Distilled Decoders perform similar to *HiFiC* decoders. The quality of the models trained separately on each sequence is better, with the trade-off of requiring more bits for sending the weights of Micro-RN.




Figure 8: Image Compression: Visual comparison between *HiFiC*^{Lo} and our reduced Decoder.

sequence. It also seems that Micro-RN removes LPIPS specific patterns and has a slightly smoother reconstruction.

4.2 Application: Video Compression with Latent Space Residuals

Training We used Adam optimizer [15] with a learning rate of 10^{-4} and a batch size of 4 on random crops of size 256×256 . The training of the full video compression pipeline is separated into four stages (the corresponding loss functions are defined in the supplementary materials): During the first 2 phases (100k steps, 50k steps) only the *FPN* is trained. Phase 3 includes the entropy model for the latent residuals for 150k steps. During the last phase we optimize for multiple frames. The back propagation of the reconstruction error through multiple frames in one optimization step, allows the model to alleviate error accumulation [18]. Due to memory limitations and runtime constraints, we chose $N = 3$.

Architecture Our *FPN* (Figure 5) is based on OpenDVC⁴. The flow field compression network (E_f , EM_f and D_f) uses the architecture proposed in [6] with 128 channels. Further, the pre-trained optical flow and motion compensation network are taken OpenDVC. For *LRN* (E_l , EM_l and D_l) we used the pre-trained components from *HiFiC*⁵.

Results We evaluated the learned models using PSNR, LPIPS and MS-SSIM and compared the latent residual compression (LRC) model to OpenDVC. Figure 9 shows that our LRC models capture details (e.g. textures of the street and trees), while OpenDVC over-smoothes the texture details at the same bitrate. To obtain the rate distortion curve (Fig-

⁴<https://github.com/RenYang-home/OpenDVC>

⁵<https://hific.github.io/>

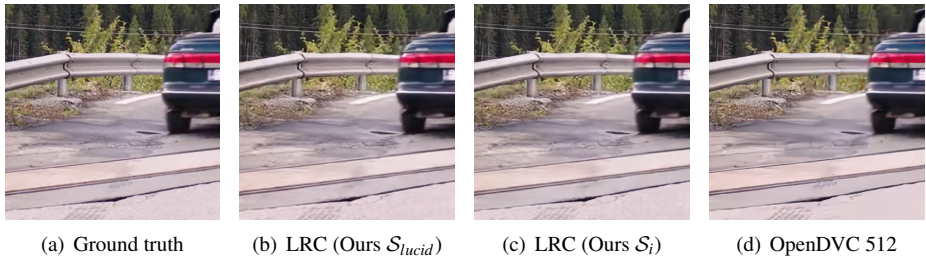


Figure 9: Video Compression: Visual Comparison between LRC models (trained on \mathcal{S}_{uvg} and \mathcal{S}_{uv}) and OpenDVC at the same bit rate.




Figure 10: Results for neural video codecs on UVG Dataset. The points on the curve correspond to (left) $HiFiC^{Lo}$ (mid) $HiFiC^{Mi}$ and (right) $HiFiC^{Hi}$. If not noted, GOP=10. The content-specific information per subset is included in the plots and amortized over the full dataset.

ure 10) we trained a video compression model for each of $HiFiC^{\{Hi,Mi,Lo\}}$. It can be seen that the LRC model leverages the motion information of previous frames. When compared to $HiFiC^{Lo}$, the LRC model decreases the bitrate by almost 50% (0.056bpp vs 0.11bpp) while keeping the reconstruction quality at a similar level. In case of $HiFiC^{Hi}$, the improvement in terms of percentage is less, namely ($\sim 30\%$). The results also show that using the distilled decoder $LRC(Ours \mathcal{S}_{avg})$ produces similar quality of reconstruction and only adds 0.005bpp if encoded on the full UVG dataset (3900 frames). The content-specific information that has to be additionally sent is already included in the numbers and amortized over the full dataset. Such information is sent uncompressed. Also, it can be seen that if we overfit to each sequence separately the model can better adapt to each sequence with the trade-off of sending the weights for each subset. This explains the increase of bpp for $LRC(Ours \mathcal{S}_1, \dots, \mathcal{S}_7)$.

5 Conclusion

In this paper, we showed how to leverage knowledge distillation in the context of neural image and video compression. More specifically we focused on GAN-based neural networks targeting low bitrates. Our proposal distills the information from a big decoder and replaces it with a much smaller one along with sequence-specific side information. Altogether this allows us to have only $\sim 5\%$ of the original model size and to achieve 50% reduction in decoding time. Future work could focus on appropriate encoding of the side information and additional methods to ensure temporal consistency.

References

- [1] The virus changed the way we internet. <https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html>. Accessed: 2021-10-20.
- [2] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [3] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression. *CoRR*, abs/1611.01704, 2016. URL <http://arxiv.org/abs/1611.01704>.
- [4] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *ICLR*, 2017.
- [5] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *ICLR*, 2018.
- [6] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. *arXiv preprint arXiv:2106.05237*, 2021.
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *CoRR*, abs/1809.11096, 2018. URL <http://arxiv.org/abs/1809.11096>.
- [8] Ting-Yun Chang and Chi-Jen Lu. Tinygan: Distilling biggan for conditional image generation. *CoRR*, abs/2009.13829, 2020. URL <https://arxiv.org/abs/2009.13829>.
- [9] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 3146–3154. IEEE, 2019. doi: 10.1109/ICCV.2019.00324. URL <https://doi.org/10.1109/ICCV.2019.00324>.
- [10] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers. Neural inter-frame compression for video coding. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6420–6428. IEEE, 2019. doi: 10.1109/ICCV.2019.00652. URL <https://doi.org/10.1109/ICCV.2019.00652>.
- [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.
- [12] *BPG Specification*. Fabrice Bellard, 2015. version 0.9.5.
- [13] Leonhard Helminger, Abdelaziz Djelouah, Markus Gross, and Christopher Schroers. Lossy image compression with normalizing flows. *arXiv preprint arXiv:2008.10486*, 2020.

- [14] Alexander Hepburn, Valero Laparra, Raúl Santos-Rodríguez, Johannes Ballé, and Jesús Malo. On the relation between statistical learning and perceptual distances. *CoRR*, abs/2106.04427, 2021. URL <https://arxiv.org/abs/2106.04427>.
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- [16] Théo Ladune, Pierrick Philippe, Wassim Hamidouche, Lu Zhang, and Olivier Déforges. Conditional coding for flexible learned video compression. *arXiv preprint arXiv:2104.07930*, 2021.
- [17] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In *CVPR*, 2019.
- [18] Guo Lu, Chunlei Cai, Xiaoyun Zhang, Li Chen, Wanli Ouyang, Dong Xu, and Zhiyong Gao. Content adaptive and error propagation aware deep video compression. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part II*, volume 12347 of *Lecture Notes in Computer Science*, pages 456–472. Springer, 2020. doi: 10.1007/978-3-030-58536-5_27. URL https://doi.org/10.1007/978-3-030-58536-5_27.
- [19] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional probability models for deep image compression. In *CVPR*, 2018.
- [20] Fabian Mentzer, George Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. *CoRR*, abs/2006.09965, 2020. URL <https://arxiv.org/abs/2006.09965>.
- [21] 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, MMSys '20*, page 297–302, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450368452. doi: 10.1145/3339825.3394937. URL <https://doi.org/10.1145/3339825.3394937>.
- [22] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *NeurIPS*. 2018.
- [23] Yash Patel, Srikar Appalaraju, and R. Manmatha. Deep perceptual compression. *CoRR*, abs/1907.08310, 2019. URL <http://arxiv.org/abs/1907.08310>.
- [24] Yash Patel, Srikar Appalaraju, and R. Manmatha. Human perceptual evaluations for image compression. *CoRR*, abs/1908.04187, 2019. URL <http://arxiv.org/abs/1908.04187>.
- [25] Yash Patel, Srikar Appalaraju, and R. Manmatha. Saliency driven perceptual image compression. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 227–236, January 2021.

- [26] 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. *ICLR*, 2016.
- [27] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. In *CVPR*, 2017.
- [28] Ties van Rozendaal, Iris AM Huijben, and Taco S Cohen. Overfitting for fun and profit: Instance-adaptive data compression. *arXiv preprint arXiv:2101.08687*, 2021.
- [29] Vijay Veerabadran, Reza Pourreza, AmirHossein Habibiian, and Taco Cohen. Adversarial distortion for learned video compression. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14-19, 2020*, pages 640–644. IEEE, 2020. doi: 10.1109/CVPRW50498.2020.00092. URL <https://doi.org/10.1109/CVPRW50498.2020.00092>.
- [30] Longguang Wang, Yingqian Wang, Xiaoyu Dong, Qingyu Xu, Jungang Yang, Wei An, and Yulan Guo. Unsupervised degradation representation learning for blind super-resolution. *CoRR*, abs/2104.00416, 2021. URL <https://arxiv.org/abs/2104.00416>.
- [31] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.