TOWARDS REAL-TIME NEURAL IMAGE COMPRESSION WITH MASK DECAY

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural image compression has surpassed state-of-the-art traditional codecs (H.266/VVC) for rate-distortion (RD) performance, but suffers from large complexity and separate models for different rate-distortion trade-offs. In this paper, we propose an efficient single-model variable-bit-rate network, which is able to run at 30 FPS with 768x512 input images and still outperforms VVC for the RD performance. By further reducing both encoder and decoder complexities, our small model even achieves 30 FPS with 1920x1080 input images. To bridge the performance gap between our different capacities models, we meticulously design the mask decay, which transforms the large model's parameters into the small model automatically. And a novel sparsity regularization loss is proposed to mitigate shortcomings of L_p regularization. Our algorithm significantly narrows the performance gap by 50% and 30% for our medium and small models, respectively. At last, we advocate the scalable encoder for neural image compression. The encoding complexity is dynamic to meet different latency requirements. We propose decaying the large encoder multiple times to reduce the residual representation progressively. Both mask decay and residual representation learning greatly improve the RD performance of our scalable encoder. Our code will be public.

1 INTRODUCTION

The image compression based on deep learning has achieved extraordinary rate-distortion (RD) performance compared to traditional codecs (H.266/VVC) (Bross et al., 2021). However, two main issues limit its practicability in real-world applications. One is the large complexity. Most state-of-the-art (SOTA) neural image codecs rely on complex models, such as the large capacity backbones (Zhu et al., 2022; Zou et al., 2022), the sophisticated probability model (Cheng et al., 2020), and the parallelization-unfriendly auto-regressive model (Minnen et al., 2018). The large complexity easily results in unsatisfied latency in real-world applications. The second issue is the inefficient rate-control. Multiple models need to be trained and stored for different rate-distortion trade-offs. For the inference, it requires loading specific model according to the target quality. For practical purposes, a single model is desired to handle variable RD trade-offs.

In this paper, we try to solve both issues to make a single real-time neural image compression model, while its RD performance is still on-par with that of other SOTA models. Inspired by recent progress, we design an efficient framework that is equipped with Depth-Conv blocks (Liu et al., 2022) and the spatial prior (He et al., 2021; Li et al., 2022a; Qian et al., 2022). For variable RD trade-offs, we introduce an adjustable quantization step (Chen & Ma, 2020; Li et al., 2022a) for the representations. All modules within our framework are highly efficient and GPU friendly, different from recent Transformer based models (Zou et al., 2022). Encoding and decoding (including the arithmetic coding) achieves 30 FPS for the 768×512 inputs. Compared with other SOTA models, ours enjoys comparable RD performance, low latency, and a single model for all RD trade-offs.

To further accelerate our model, we reduce the complexity of both the encoder and decoder. Three series models are proposed: Large, Medium, and Small (cf. Appendix Tab. 3). Note that our small model even achieves 30 FPS for 1920×1080 inputs. Recent work (Zhu et al., 2022) points out that simply reducing the model's capacity results in significant performance loss. However, this crucial problem has not been solved effectively. In this paper, we advocate training small neural compression models with a teacher to mitigate this problem. In particular, *mask decay* is proposed

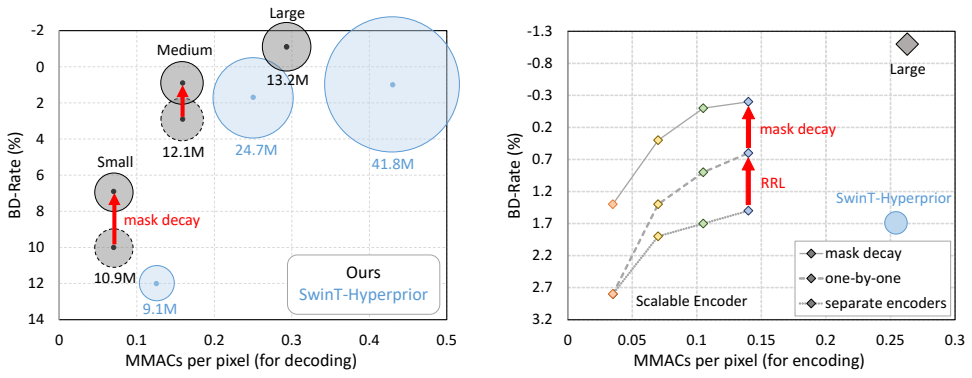

(a) (b)

Figure 1: The trade-off between BD-Rate and complexities on Kodak. The anchor is VTM. (a) and (b) show the performance improvement by our mask decay and residual representation learning (RRL). We cite results from SwinT-Hyperprior (Zhu et al., 2022) for comparison.

to transform the large model's parameters into the small model automatically. Specifically, mask layers are first inserted into a pretrained teacher model. Then we sparsify these masks until all layers become the student's structure. Finally, masks are merged into the raw layers meanwhile keeping the output unchanged. For the sparsity regularization, L_p-norm based losses are adopted by most previous works (Li et al., 2016; Ding et al., 2018; Zhang et al., 2021). But they hardly work for neural image compression (cf. Fig. 7). In this paper, we propose a novel sparsity regularization loss to alleviate the drawbacks of L_p so that optimize our masks more effectively. With the help of our large model, our medium and small models are improved significantly by 50% and 30%, respectively (cf. Fig. 1a). That demonstrates the effectiveness of our mask decay and reusing a large model's parameters is helpful for training a small neural image compression model.

In addition, considering the various device capabilities in real-world codec applications, the encoder scalability of supporting different encoding complexities while with only one decoder is also a critical need. To achieve this, we propose compressing the cumbersome encoder multi-times to progressively bridge the performance gap. Both the residual representation learning (RRL) and mask decay treat the cumbersome encoder as a reference implicitly, and encourage the diversity of different small encoders. Therefore ours achieves superior performance than training separate encoders (cf. Fig. 1b). And compared to SlimCAE (Yang et al., 2021) which is inspired by slimmable networks (Yu et al., 2019), ours enjoys a simple framework and better RD performance.

Our contributions are as follows.

- We propose an efficient framework for neural image compression. It enjoys a single model for different RD trade-offs. Our model is able to run at 30 FPS for the 768×512 inputs, while is on-par with other SOTA models for the RD performance. Our small model even achieves 30 FPS for the 1920×1080 inputs.
- We propose mask decay, an effective method to improve the student image compression model with the help of the teacher. A novel sparsity regularization loss is also introduced, which alleviates the shortcomings of L_p regularization. Thanks to mask decay, our medium and small models are significantly improved by 50% and 30%, respectively.
- We enable the encoding scalability for neural image compression. With residual representation learning and mask decay, our scalable encoder significantly narrows the performance gap from the teacher and achieves a superior RD performance than previous SlimCAE.

2 Related Works

Neural image compression is in a scene of prosperity. Ballé et al. (2017) proposes replacing the quantizer with an additive i.i.d. uniform noise, so that the neural image compression model enjoys the end-to-end training. Then, the hyperprior structure was proposed by Ballé et al. (2018). After that, numerous methods are proposed to improve the entropy model and the backbone (both encoder and decoder). Auto-regressive component (Minnen et al., 2018) and Gaussian Mixture Model (Cheng et al., 2020) are proposed to improve the probability estimation in entropy model. Recently, Transformer (Qian et al., 2022) and Swin (Zhu et al., 2022) are introduced to replace

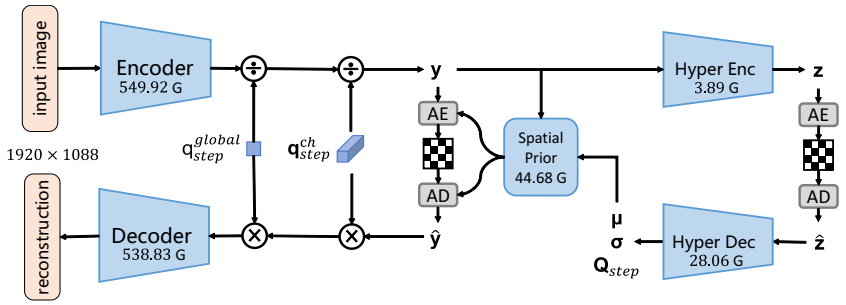

Figure 2: The overall framework of our model.

ResNet (He et al., 2016) in the encoder and decoder. However, recent SOTA neural image codec models suffer from high latency and separate models for each RD trade-off. In this paper, we propose an efficient neural image compression framework that aims at encoding and decoding in real-time. Especially, our framework can also handle variable rate-distortion trade-offs by the adjustable quantization step (Chen & Ma, 2020; Cui et al., 2021; Li et al., 2022a). The RD performance of our models is on-par with other SOTA neural image compression models, and surpasses the SOTA traditional image codec (H.266/VVC).

Model compression and knowledge distillation (KD) are two prevalent techniques for accelerating neural networks. Model compression, especially filter level pruning (He et al., 2017), aims at obtaining an efficient network structure from a cumbersome network. Some works (Li et al., 2016; Luo & Wu, 2020; He et al., 2019; Liebenwein et al., 2019; Molchanov et al., 2019) calculate the importance of each filter, and prune the unimportant filters. Other methods directly penalize filters during training. They treat some parameters as the importance implicitly, such as the scale in batch normalization (Liu et al., 2017), or the inserted redundancy layers (Ding et al., 2021; Wang & Wu, 2022). L_p-norm based sparsity losses are used in these methods (Ding et al., 2018; Wen et al., 2016; Zhang et al., 2021). For a pruned network, KD (Hinton et al., 2015) is utilized to boost the performance by mimicking the teacher's knowledge (e.g., the logit (Zhao et al., 2022), or the feature (Wang et al., 2021)). All these techniques are mainly studied on semantic tasks, but rarely advocated for neural image compression. This paper firstly tries to improve a small neural image compression model with the help of a large model. And we claim that, for neural image compression models, reusing the teacher's filters is an effective way to improve the student. To achieve this, we propose mask decay with a novel sparsity regularization loss. With all proposed techniques, our student models are greatly improved.

3 Methodology

We first introduce our efficient framework briefly. Then, three models with different complexities are proposed: Large, Medium, and Small. Next, we describe how to improve our small model (student) with the help of our large model (teacher). Mask decay is proposed to transform the teacher's parameters into the student automatically. At last, the scalable encoder is advocated. With residual representation learning and mask decay, several encoders are trained progressively to bridge the performance gap from the teacher.

3.1 Our framework for image compression

Our framework is shown in Fig. 2. Given an input image, the encoder generates its representation $\boldsymbol{y}$. Following Chen & Ma (2020); Li et al. (2022a), to achieve variable bit-rates within a single model, we introduce adjustable quantization steps for $\boldsymbol{y}$. $q_{\text{step}}^{\text{global}}$ and $\boldsymbol{q}_{\text{step}}^{\text{ch}}$ are two parameters that quantifying $\boldsymbol{y}$ globally and channel-wisely, respectively. And different RD trade-offs can be achieved by adjusting the magnitude of $q_{\text{step}}^{\text{global}}$. More details of our framework are introduced in Appendix Sec. A.

Fig. 2 also presents MACs of each module with the 1920×1088 input image. Obviously, most computational costs come from the encoder and the decoder. Hence, we focus on accelerating these two modules, while keeping other parts' architectures unchanged. Fig. 3 shows the detail structure of these two modules. Different from recent SOTA models (Zhu et al., 2022; Qian et al., 2022; Zou et al., 2022), we adopt residual blocks and depth-wise convolution blocks (Liu et al., 2022)

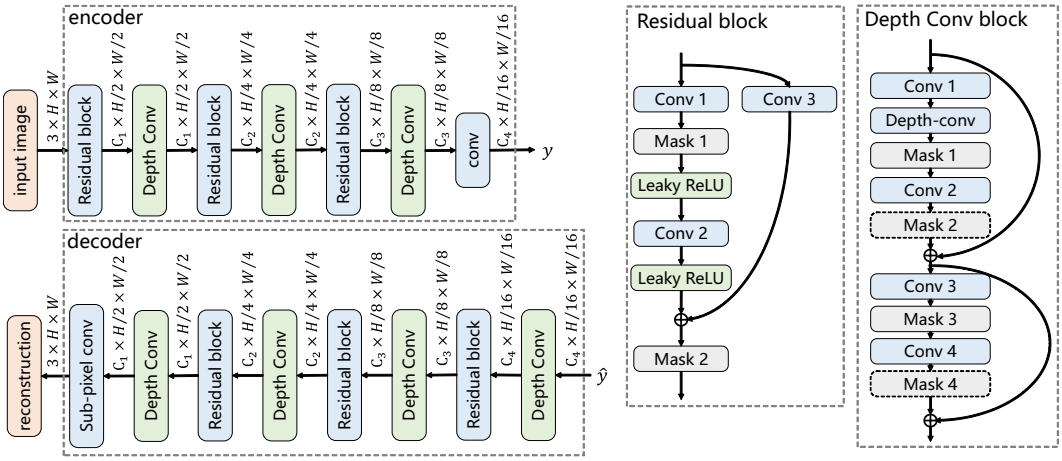

Figure 3: The architectures of our encoder and decoder. Note that all mask layers will be merged into conv. layers after training. There are Leaky ReLUs behind conv. layers within the Depth-Conv block. We omit them in this figure for simplification.

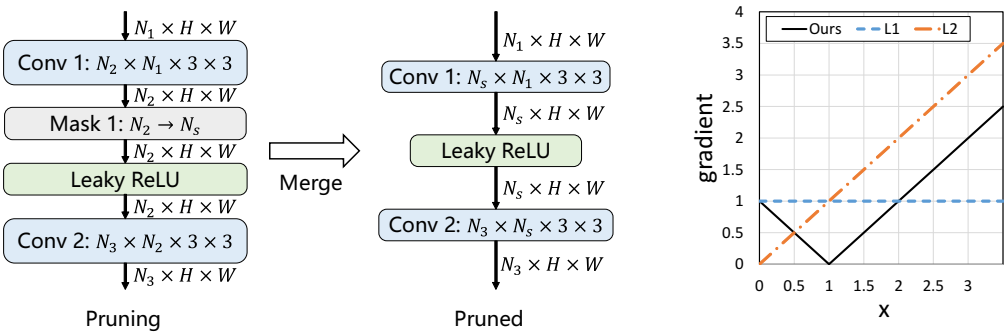

(a) (b)

Figure 4: (a) presents our pruning method. The mask layer is inserted between two conv. layers. These two cumbersome conv. layers will be transformed into efficient ones by merging the mask layer into them. (b) compares the gradient $\frac{\partial \mathcal{L}_{sparse}}{\partial x}$ of different sparsity loss functions.

rather than Transformer. Hence, ours is more GPU-friendly and enjoys lower latency. C_1, C_2, C_3, and C_4 are pre-defined numbers for different computational budgets. Three models are proposed with channels $[64, 64, 128, 192]$ (Small), $[128, 128, 192, 192]$ (Medium), and $[192, 192, 192, 192]$ (Large). Tab. 3 in the appendix summarizes models' #Params and MACs. Note that Large is our cumbersome teacher model while Medium and Small are efficient student models.

3.2 Improve the Student by Mask Decay

Due to the limited capacity, the student performs badly if training from scratch (Zhu et al., 2022). How to improve the student is a crucial but rarely studied problem for neural image compression. In this paper, we claim that reusing the teacher's parameters is helpful. To achieve it, we propose mask decay to transform parameters in an end-to-end manner.

Compress one layer's representation by mask decay. As showed in Fig. 4a, a mask layer is inserted between two layers to reduce the representation's channels (from N_2 to N_s). This mask layer contains N_2 parameters to prune corresponding channels. In the forward computation, the representation multiplies these parameters channel-wisely. During training, these parameters will become sparse by optimizing our sparsity loss. After training, indicated by the mask, useless parameters in Conv #1 and #2 are pruned. And the mask is eliminated by multiplying itself into Conv #2.

To sparsify parameters $\boldsymbol{m}$ in the mask layer, we design a novel sparsity regularization loss. Most previous works (Li et al., 2016; Wen et al., 2016; Zhang et al., 2021) adopt L1-norm ($|\boldsymbol{m}|$) or L2-norm ($\frac{1}{2}\|\boldsymbol{m}\|_2$) based functions. Note that with the gradient descend, $\boldsymbol{m}$ is updated by $sign(\boldsymbol{m})$ and

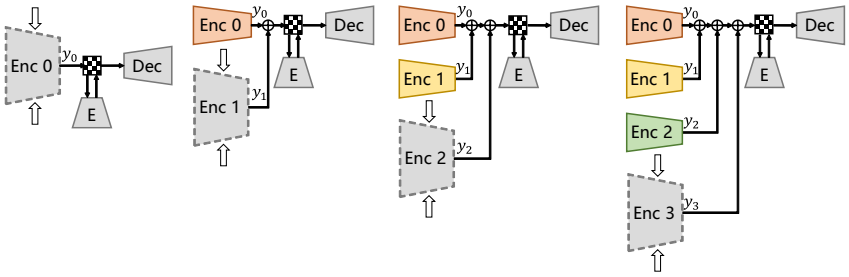

Figure 5: Illustration of compressing encoders multi-times to learn residual representations progressively. Note that "Enc", "Dec", and "E" denote the encoder, the decoder, and the entropy module, respectively.

$\boldsymbol{m}$, respectively. Miserably, both of them fail for neural image compression models (cf. Fig. 7). We argue that, for a parameter with a positive value, the gradient derived by the L1-norm is a constant without considering its own magnitude. And for the L2-norm, when the parameter is approaching zero, the gradient is too small to sparsify the mask. In addition, both of them have no stationary points when parameters are positive. To alleviate these problems, We design a novel sparsity regularization loss from the gradient view. As showed in Fig. 4b, the gradient is defined by

$$\frac{\partial \mathcal{L}_{sparse}(x)}{\partial x} = |x-1|\,, \tag{1}$$

where $x \geq 0$. For $x > 1$, the parameter with a large magnitude suffers a large gradient to decay, which likes the L2-norm. And for $0 \leq x < 1$, a smaller parameter enjoys a larger gradient to approach zero fast. In addition, there is a stationary point at $x = 1$, which means $\frac{\partial \mathcal{L}_{sparse}(x)}{\partial x}|_{x=1} = 0$. When it comes to L1 and L2, there always is a positive gradient for $x > 0$. To make the model converge to stable, it needs a negative gradient from the task loss to balance this positive gradient. Thanks to our sparsity loss, only a small negative gradient is needed when x is close to 1. $\mathcal{L}_{sparse}$ can be calculated by integrating Equation 1, which is

$$\mathcal{L}_{sparse}(x) = \begin{cases} -\frac{1}{2}x^2 + x, & \text{if } 0 \leq x \leq 1\,, \\ \frac{1}{2}x^2 - x + 1, & \text{if } x > 1\,. \end{cases} \tag{2}$$

For training, inspired by AdamW (Loshchilov & Hutter, 2018), we decouple $\mathcal{L}_{sparse}$ from the task loss $\mathcal{L}_{RD}$ as weight decay. Therefore, the task loss is not affected by our sparsity loss explicitly. So we name our algorithm as *mask decay*. For the mask $\boldsymbol{m}$, the updating formula is

$$\boldsymbol{m}_{t+1} = \boldsymbol{m}_t - \eta \frac{\partial \mathcal{L}_{sparse}}{\partial \boldsymbol{m}_t} - \gamma \frac{\partial \mathcal{L}_{RD}}{\partial \boldsymbol{m}_t} = \boldsymbol{m}_t - \eta |\boldsymbol{m}_t - 1| - \gamma \frac{\partial \mathcal{L}_{RD}}{\partial \boldsymbol{m}_t}\,, \tag{3}$$

where γ is the learning rate and η is the rate for mask decay.

Mask decay for our encoder and decoder. As illustrated in Fig. 3, to transform the teacher into the student, we insert mask layers in some positions if necessary. Our principle is to insert mask layers as few as possible. In one residual block, Mask #1 and #2 are inserted to prune channels inside and outside the residual connection, respectively. In the depth-conv block, Mask #1 determines channels of the depth convolution layer, while Mask #3 prunes representations between Conv #3 and #4. Note that the output channels of these two blocks are shared, so Mask #2 and #4 in the depth-conv block share weights with Mask #2 in the residual block. And we place Mask #2 and #4 behind the residual adding to guarantee the identity gradient back-propagation.

Training with mask decay. The training begins with a pretrained teacher model. Then, we insert mask layers and initialize them with $\mathbf{1}$. The whole network is optimized by minimizing $\mathcal{L}_{RD} = R + \lambda D$, where λ steers the trade-off between the rate R and the distortion D. Our masks are decayed before each iteration, likely weight decay (Loshchilov & Hutter, 2018). During training, we stop mask decay for one mask if it is sparse enough. When all masks meet the sparsity requirements, the whole sparse network can transform into the student by merging masks, meanwhile keeping the output unchanged. At last, the student will be finetuned for a few epochs.

3.3 THE SCALABLE ENCODER

At last, we advocate the scalable encoder in neural image codecs. A trivial solution for the scalable encoder is directly training separate small encoders with a uniform decoder. For a specific encoding

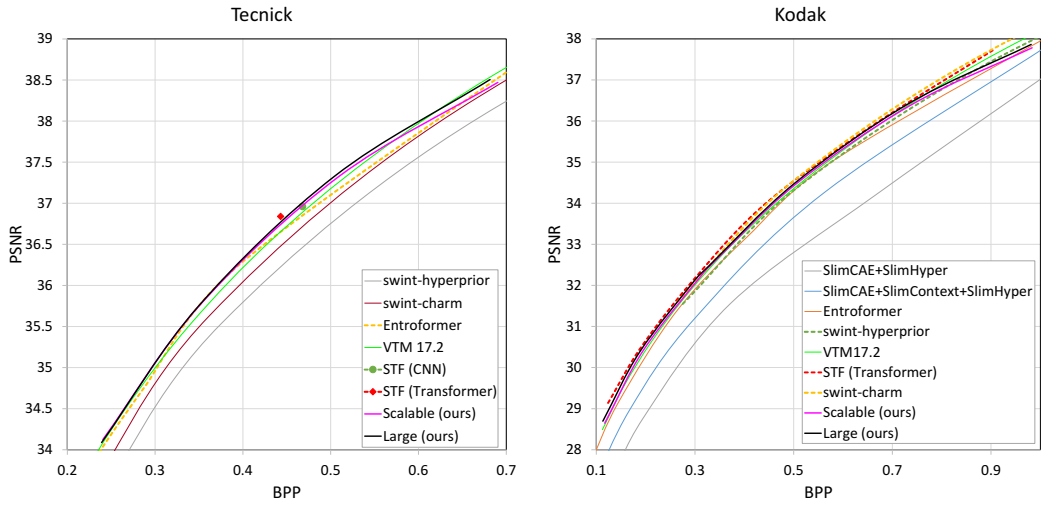

(a) (b)

Figure 6: Rate Distortion curves for Tecnick and Kodak. Our models are on-par with swint (Zhu et al., 2022), Entroformer (Qian et al., 2022), STF (Zou et al., 2022), and outperforms the traditional method VTM. Note that our models are dramatically faster than these methods (cf. Tab. 1).

complexity, the corresponding number of encoders are chosen to encode together, and then the best bitstreams will be sent to the decoder. However, these separate encoders suffer from the homogenization problem. We propose residual representation learning (RRL) to mitigate this problem. As illustrated in Fig. 5, we compress the cumbersome encoder several times to learn the residual representations progressively. Note that in each step, only parameters in the compressing encoder are optimized while other parts are frozen. That makes sure previous encoders generate decodable bitstreams. For a specific encoding complexity, the first k encoders are applied. Our RRL encourages the encoders' diversity. Both RRL and mask decay treat the cumbersome teacher as a reference, which makes the training more effective. Another related work SlimCAE (Yang et al., 2021) uses slimmable layers (Yu & Huang, 2019) to obtain an image compression model with dynamic complexities. Our framework is more simple to implement and achieves superior RD performance.

4 Experiments

Training. The training dataset is the training part of Vimeo-90K (Xue et al., 2019) septuplet dataset. During training, the raw image is randomly cropped into 256×256 patches. For a fair comparison, all models are trained with 200 epochs. For our method, it costs 60 epochs for the teacher to transform into the student with mask decay, then the student is finetuned by 140 epochs.

Testing. Testing datasets contain Kodak (Franzen, 1999), HEVC test sequences (Sharman & Suehring, 2017), Tecnick (Asuni & Giachetti, 2014). BD-Rate (Bjontegaard, 2001) for peak signal-to-noise ratio (PSNR) versus bits-per-pixel (BPP) is our main metric to compare different models.

More details about experimental settings are introduced in Appendix Sec. D.

4.1 Comparison with state-of-the-art

Fig. 6 compares our models with other SOTA models on Tecnick and Kodak. Our large model outperforms the SOTA traditional image codec ("VTM 17.2") and is on-par with other SOTA neural image codec models. Note that our model handles different RD trade-offs by only *one* model, while other SOTA methods train separate models for each trade-off. Tab. 1 compares latency of different models. For Entroformer (Qian et al., 2022) and STF (Zou et al., 2022), we measure the latency by running their official code. Note that the latency contains arithmetic coding, but ignores the disk I/O time. SwinT (Zhu et al., 2022) is also compared in Appendix Tab. 6. Our models surpass these SOTA models with a significant margin. Our large model runs at 30 FPS on A100 for the 768×512 inputs, while our small model even achieves 30 FPS for the 1080P inputs.

Table 1: Latency (ms) comparison. Lower means better. Entroformer (Qian et al., 2022) and STF (Zou et al., 2022) are the SOTA methods. 'OM' means out of memory. Note that 2080Ti and A100 have 12GB and 80GB memory, respectively. **Bold** denotes running over 30 FPS.

Resolution	GPU	Type	Entroformer	STF		Ours		
				Transformer	CNN	Large	Medium	Small
768×512	2080Ti	encoding	OM	176.3	158.5	63.0	44.7	**28.4**
		decoding	OM	202.3	210.2	41.1	**32.4**	**24.4**
	A100	encoding	816.8	115.9	96.4	**21.1**	**19.8**	**17.7**
		decoding	4361.9	143.2	118.0	**19.1**	**17.1**	**15.6**
1920×1080	2080Ti	encoding	OM	576.0	456.0	305.3	181.5	90.9
		decoding	OM	531.7	652.0	179.2	118.1	73.2
	A100	encoding	7757.4	355.6	278.1	84.2	56.3	**31.4**
		decoding	OM	354.8	281.7	60.2	46.5	**29.7**

Table 2: BD-Rate (%) comparison for PSNR. L, M, and S denote Large, Medium, and Small, respectively. The anchor is VTM. Lower BD-Rate means better. MACs is for 1920×1088 inputs.

Enc	Dec	MACs (G)	Method	HEVC					Kodak
				B	C	D	E	Avg.	
L	L	1165.39	Baseline	-5.2	-0.9	-0.9	-3.5	-2.63	-1.1
M	L	878.37	Baseline	-2.8	2.2	2.5	0.0	0.48	1.1
			Ours	**-4.5**	**-0.2**	**0.2**	**-3.0**	**-1.88** (76%↑)	**-0.4** (68%↑)
S	L	688.64	Baseline	0.3	7.1	7.0	2.8	4.30	4.0
			Ours	**-1.9**	**2.9**	**3.5**	**0.3**	**1.20** (45%↑)	**1.6** (47%↑)
L	M	885.12	Baseline	-2.1	3.1	3.6	1.1	1.43	2.1
			Ours	**-4.0**	**0.7**	**0.8**	**-1.6**	**-1.03** (60%↑)	**0.2** (59%↑)
L	S	700.38	Baseline	2.1	9.5	10.7	7.7	7.50	7.0
			Ours	**-0.4**	**5.6**	**6.9**	**3.3**	**3.85** (36%↑)	**4.0** (37%↑)
M	M	598.10	Baseline	-0.8	4.7	5.0	1.6	2.63	2.9
			Ours	**-3.1**	**1.6**	**1.5**	**-0.8**	**-0.20** (54%↑)	**0.9** (50%↑)
S	S	223.63	Baseline	6.3	14.7	15.6	10.6	11.80	10.0
			Ours	**2.6**	**9.5**	**10.6**	**6.5**	**7.30** (31%↑)	**6.9** (28%↑)

4.2 MASK DECAY

Previous work (Zhu et al., 2022) finds small models suffer from poor RD performance. This problem is alleviated by our mask decay. We accelerate the encoder and the decoder respectively to study more cases. Note that the large model is the teacher and the decay rate is $4e-5$. Fig. 1a and Tab. 2 summarize experimental results. In Tab. 2, we also calculate the relative improvement as a percentage. For example, if BD-Rate of the baseline student, our student, and the baseline teacher are 1.1%, -0.4%, and -1.1%, respectively. The relative improvement is $\frac{1.1-(-0.4)}{1.1-(-1.1)} = 68\%$. This metric indicates how much the performance gap between the student and the teacher our method narrows. RD curves are shown in Appendix Fig. 14. From experimental results, we conclude that:

- Students are improved by mask decay with a significant margin. For medium and small models, relative improvements are about 50% and 30%, respectively. This demonstrates the effectiveness of our method and the student benefits from reusing the teacher's parameters. Compared with SwinT-Hyperprior in Fig. 1a, ours achieve a superior trade-off between the RD performance and the complexity. In addition, our models enjoy small sizes.
- The encoder is more redundant than the decoder. Note that our encoder and decoder are symmetric. We surprisingly find compressing the encoder enjoys less RD performance loss, compared with the decoder. With the large decoder, compressing the encoder from Large to Small loses 2.7% BD-Rate on Kodak. On the other hand, only compressing the decoder loses 5.1% BD-Rate.

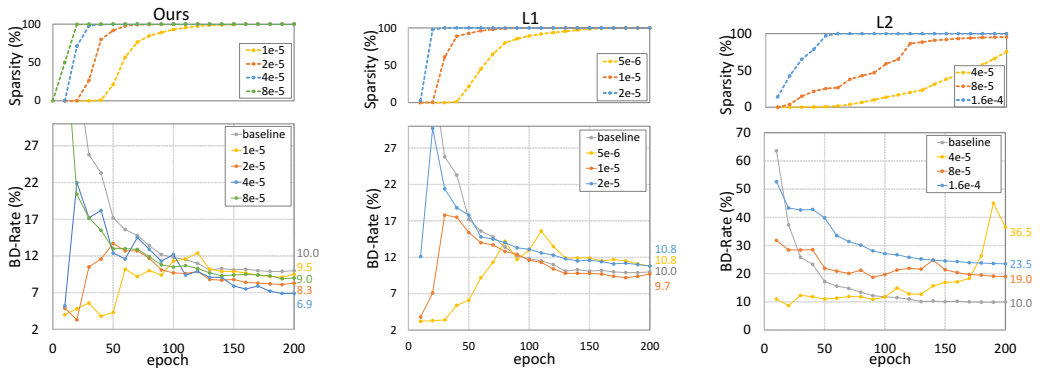

(a) (b) (c)

Figure 7: Ablation studies for different sparsity losses. We adjust the decay rate η for each criterion. This figure is best viewed in color and zoomed in.

Appendix Fig. 15 visualizes our models' reconstructions. Reconstructions of different capacities models are consistent. Our mask decay does not result in extra obvious artifacts. And our encoding and decoding time are dramatically shorter than VTM's.

4.3 THE SCALABLE ENCODER

Fig. 1b presents results for the scalable encoder. First, training separate encoders suffers from poor performance, which is due to the homogenization problem. Then, with our framework, training these encoders one-by-one to bridge the residual representations greatly improves the performance. Our RRL encourages these encoders' diversity. Finally, the mask decay treats the cumbersome encoder as a reference and makes training more efficient. More details can be found in Appendix Sec. C.

Fig. 6 shows RD curves for our scalable encoder ("Scalable"). Ours outperforms SlimCAE (Yang et al., 2021) significantly by 15.5% BD-Rate. Note that SlimCAE is a powerful model equipped with hyperprior (Ballé et al., 2018), conditional convolutions (Choi et al., 2019), and an autoregressive context model (Minnen et al., 2018). Overall, our scalable encoder is on-par other SOTA models.

4.4 ABLATION STUDIES

Sparsity losses and decay rates. Fig. 7 summarizes experimental results for mask decay with different sparsity regularization losses and decay rates. For sparsity, 0% and 100% denote the teacher's and the student's structures, respectively. We adjust the decay rate η for each sparsity loss to get reasonable sparsity decay curves. BD-Rate is calculated over VTM on Kodak. For our sparsity loss, all decay rates achieve better RD performance compared with the baseline. A suitable decay rate $(4e-5)$ achieves the best BD-Rate. A large decay rate $(8e-5)$ results in fast decay which destroys the model's performance in early epochs. This model suffers from the insufficient decay process, and the performance is easily saturated. Small decay rates ($1e-5$ and $2e-5$) waste too many epochs for mask decay and the model lacks finetuning. For L1, the model suffers from poor performance with both large and small decay rates. With a suitable rate $(1e-5)$, it is still worse than ours with a significant margin. For L2, we find it hardly works because the gradient vanishes when the mask is close to zero.

Discussions for important channels. In the field of model pruning, many works (Li et al., 2016; He et al., 2019; Liebenwein et al., 2019; Molchanov et al., 2019) study how to identify unimportant channels for pruning. However, recent works (Liu et al., 2018; Li et al., 2022b) indicate randomly pruning channels results in comparable or even superior performance, which demonstrates there may be no importance difference among all channels. Previous works study mainly semantic tasks. We want to extend this conclusion for low-level tasks, i.e., image compression.

First, we notice that our mask decay selects similar channels with different decay rates in Figure 7a. We use $\mathbb{C}_i$ to represent the set of chosen channels, where $i \in F = \{1e-5, 2e-5, 4e-5, 8e-5\}$ means the decay rate η. We check masks in the first two blocks, and #channels of the teacher and the student are 192 and 64, respectively. Hence, about 33% channels need to be kept after pruning. However, we find $\cap_{i \in F} \mathbb{C}_i$ and $\cup_{i \in F} \mathbb{C}_i$ contain 24% and 44% channels, respectively, which means

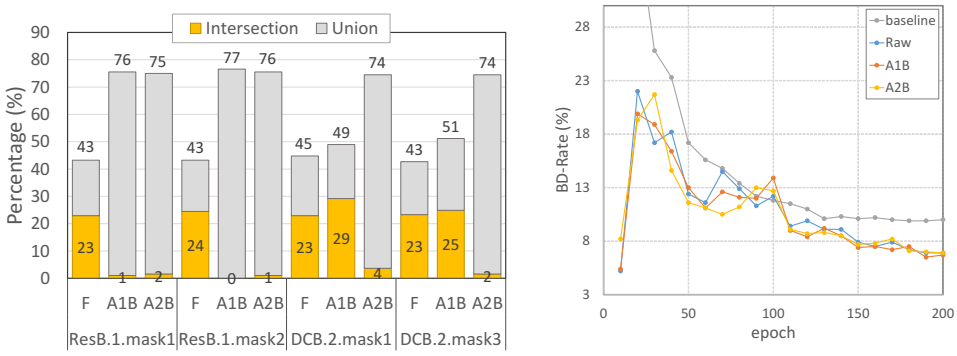

(a) (b)

Figure 8: (a) presents the percentage of channels w.r.t different mask decay manners in different layers. F, A1B, and A2B mean all channels are free, avoiding the first block, and avoiding the first 2 blocks, respectively. Numbers of the intersection of AkB ($k = 1, 2$) and F are small, which means our AMD indeed avoids channels in F. (b) presents the BD-Rate over VTM on Kodak w.r.t training epochs. Raw denotes the curve for $\eta = 4e - 5$ in Figure 7a.

all of them only consider 84 channels from 192 channels. Are these 44% channels important? Our answer is negative. To verify our opinion, we propose mask decay with avoiding specific channels. We divide all channels into two groups: freedom $\mathbb{F}$ and avoidance $\mathbb{A}$. To avoid channels in $\mathbb{A}$, we set a large decay rate for them. Therefore, the loss of our *Avoidable Mask Decay* (AMD) is

$$\mathcal{L}_{AMD}(x) = \eta_f \mathbf{1}_{\mathrm{x} \in \mathbb{F}} \mathcal{L}_{sparse}(x) + \eta_a \mathbf{1}_{\mathrm{x} \in \mathbb{A}} \mathcal{L}_{sparse}(x) \tag{4}$$

where $\mathbf{1}_{\text{condition}}$ is 1 if the condition is true, 0 otherwise. Defaultly, we set $\eta_a = 10\eta_f$. Hence, avoidable channels enjoy a large decay rate to approach zero quickly. We conduct two experiments that avoid selecting channels in $\cup_{i \in F} \mathbb{C}_i$ for the first one or two blocks, respectively. Figure 8a illustrates our AMD indeed avoid channels. Specifically, we calculate $\mathbb{C}_{AkB} \cap \mathbb{C}_F$ and $\mathbb{C}_{AkB} \cup \mathbb{C}_F$, where $k = 1, 2$ denotes avoiding channels in the first k blocks and $\mathbb{C}_F = \cup_{i \in F} \mathbb{C}_i$. At most 1% channels are chosen by both A1D and F in the first block, and A2B avoids channels of F in both two blocks. Figure 8b shows the BD-Rate for each epoch. From it, we find AMD results in comparable or a little bit superior performance compared to the raw mask decay.

Overall, if some channels are important, we can avoid these channels to achieve comparable performance. That demonstrates there are no important channels. For future works, this conclusion guides us to pay more attention to the optimization process rather than the importance of each channel.

5 Conclusions and Future Works

We proposed an efficient framework for image compression. With 768×512 inputs, our model is able to run at 30 FPS, while the RD performance is on-par with SOTA neural codecs and superior to the SOTA traditional codec (H.266/VVC). With an adjustable quantization step, our model handles variable RD trade-off within only one model. These three signs of progress make a new milestone in the field of image codec. By reducing the complexities of both encoder and decoder, a small model is proposed, which achieves 30 FPS for 1080P inputs. We proposed mask decay for training this small model with the help of the cumbersome model. And a novel sparsity regularization loss is designed to alleviate the drawbacks of L1 and L2 losses. Thanks to our algorithm, our medium and small models are improved significantly by 50% and 30%, respectively. At last, the scalable encoder is advocated. With only one decoder, there are several encoders for different complexities. With residual representation learning and mask decay, the scalable encoder narrows the performance gap between the teacher. And ours outperforms previous SlimCAE by a significant margin.

We believe this paper will inspire many future works. First, we point out the encoder is more redundant than the decoder. Hence, improving the decoder is promising. Second, this paper mainly focuses on the pre-defined student structure. How to determine the complexity of each layer is still an open problem for image compression models. Third, it is also encouraged to find a more powerful pruning algorithm. But according to our AMD, we advise paying more attention to the optimization process rather than the importance criterion for each filter.

REFERENCES

Nicola Asuni and Andrea Giachetti. TESTIMAGES: a large-scale archive for testing visual devices and basic image processing algorithms. In *Proceedings of the Smart Tools and Apps for Graphics*, pp. 63–70, 2014.

Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. In *International Conference on Learning Representations*, pp. 1–27, 2017.

Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, pp. 1–23, 2018.

Gisle Bjontegaard. Calculation of average PSNR differences between RD-curves. *VCEG-M33*, 2001.

Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021.

Tong Chen and Zhan Ma. Variable bitrate image compression with quality scaling factors. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 2163–2167, 2020.

Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7939–7948, 2020.

Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3146–3154, 2019.

Ze Cui, Jing Wang, Shangyin Gao, Tiansheng Guo, Yihui Feng, and Bo Bai. Asymmetric gained deep image compression with continuous rate adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10532–10541, 2021.

Xiaohan Ding, Guiguang Ding, Jungong Han, and Sheng Tang. Auto-balanced filter pruning for efficient convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pp. 6797–6804, 2018.

Xiaohan Ding, Tianxiang Hao, Jianchao Tan, Ji Liu, Jungong Han, Yuchen Guo, and Guiguang Ding. ResRep: Lossless cnn pruning via decoupling remembering and forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4510–4520, 2021.

Rich Franzen. Kodak lossless true color image suite. *source: http://r0k.us/graphics/kodak*, 4(2), 1999.

Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14771–14780, 2021.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 770–778, 2016.

Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4340–4349, 2019.

Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1389–1397, 2017.

Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

Jiahao Li, Bin Li, and Yan Lu. Hybrid spatial-temporal entropy modelling for neural video compression. *arXiv preprint arXiv:2207.05894*, 2022a.

Yawei Li, Kamil Adamczewski, Wen Li, Shuhang Gu, Radu Timofte, and Luc Van Gool. Revisiting random channel pruning for neural network compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 191–201, 2022b.

Lucas Liebenwein, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. Provable filter pruning for efficient neural networks. *arXiv preprint arXiv:1911.07412*, 2019.

Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2736–2744, 2017.

Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations*, pp. 1–21, 2018.

Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11976–11986, 2022.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, pp. 1–19, 2018.

Jian-Hao Luo and Jianxin Wu. Neural network pruning with residual-connections and limited-data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1458–1467, 2020.

David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, volume 31, pp. 1–10, 2018.

Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation for neural network pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11264–11272, 2019.

Yichen Qian, Ming Lin, Xiuyu Sun, Zhiyu Tan, and Rong Jin. Entroformer: A transformer-based entropy model for learned image compression. In *International Conference on Learning Representations*, pp. 1–14, 2022.

Karl Sharman and Karsten Suehring. Common test conditions for HM. *JCTVC-Z1100*, 2017.

Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1874–1883, 2016.

Guo-Hua Wang and Jianxin Wu. Practical network acceleration with tiny sets. *arXiv preprint arXiv:2202.07861*, 2022.

Guo-Hua Wang, Yifan Ge, and Jianxin Wu. Distilling knowledge by mimicking features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. doi: 10.1109/TPAMI.2021.3103973.

Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, volume 29, pp. 1–9, 2016.

Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019.

Fei Yang, Luis Herranz, Yongmei Cheng, and Mikhail G. Mozerov. Slimmable compressive autoencoders for practical neural image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4998–5007, 2021.

Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1803–1811, 2019.

Jiahui Yu, L. Yang, N. Xu, Jianchao Yang, and Thomas S. Huang. Slimmable neural networks. In *International Conference on Learning Representations*, pp. 1–12, 2019.

Yulun Zhang, Huan Wang, Can Qin, and Yun Fu. Aligned structured sparsity learning for efficient image super-resolution. In *Advances in Neural Information Processing Systems*, volume 34, pp. 2695–2706, 2021.

Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11953–11962, 2022.

Yinhao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International Conference on Learning Representations*, pp. 1–35, 2022.

Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17492–17501, 2022.

Towards Real-Time Neural Image Compression With Mask Decay
-Appendix-

A OUR FRAMEWORK

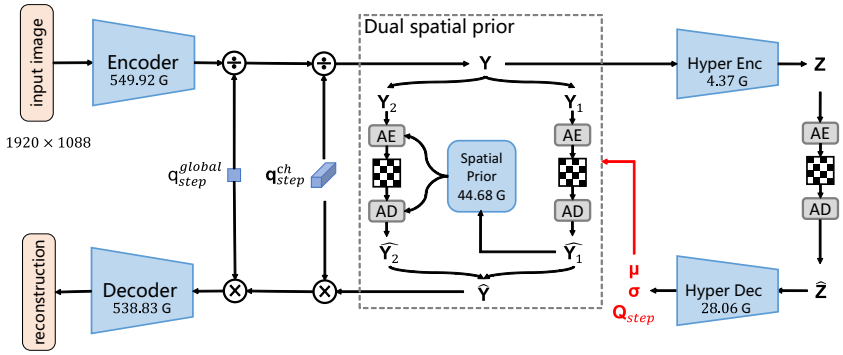

Figure 9: The overall framework of our model.

Our whole framework for image compression is illustrated in Fig. 9. Next, we introduce each module in detail.

The encoder and the decoder. The structures of these two modules are illustrated in Fig. 3. The block output dimensionalities are predefined by C_1, C_2, C_3, and C_4, respectively. Downsampling is implemented by setting the stride as 2 in specific conv. layers, while upsampling is achieved by the sub-pixel convolution (Shi et al., 2016). In each residual block, the kernel size of both Conv #1 and #2 is 3×3, while Conv #3's is 1×1. And Conv #1 enlarges the channel number to the output dimensionality. In each depth-conv block, only the depth conv. layer uses 3×3 conv., while others adopt 1×1 conv. Conv #3 enlarges the channels by 4 times, while Conv #4 converts it back.

The hyper encoder and the hyper decoder. Fig. 10 describes the structures of our hyperprior. Note that we use 3×3 conv. in the hyper encoder.

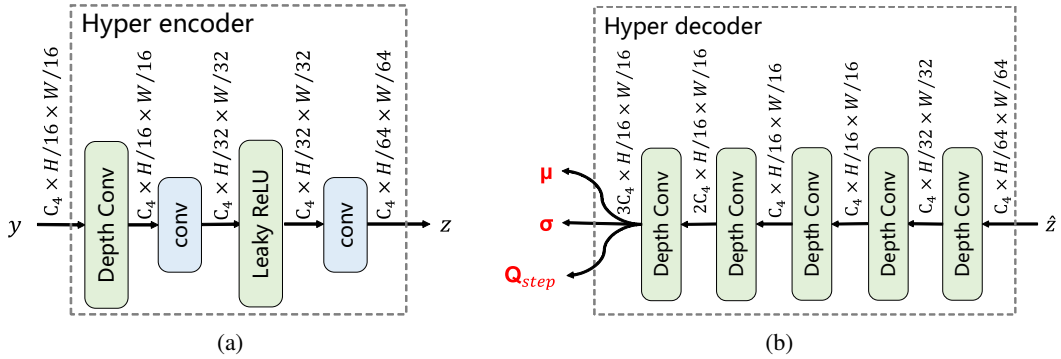

Figure 10: (a) and (b) present the structures of our hyper encoder and hyper decoder, respectively.

The dual spatial prior. Following He et al. (2021); Li et al. (2022a); Qian et al. (2022), we adopt dual spatial prior to reduce the spatial redundancy. The detail structure is illustrated in Fig. 11. The representation $\mathbf{Y}$ is split into two parts $\mathbf{Y}_1$ and $\mathbf{Y}_2$ spatially. $\mathbf{Y}_1$ is first coded according to the side information from the hyperprior. Then, given the reconstruction $\hat{\mathbf{Y}}_1$, $\mathbf{Y}_2$ is estimated more

effectively. Our dual spatial prior shares a similar idea of auto-regressive prior (Minnen et al., 2018) that predicts the target conditioning on its neighbor. But ours enjoys parallel computation and a fast inference speed.

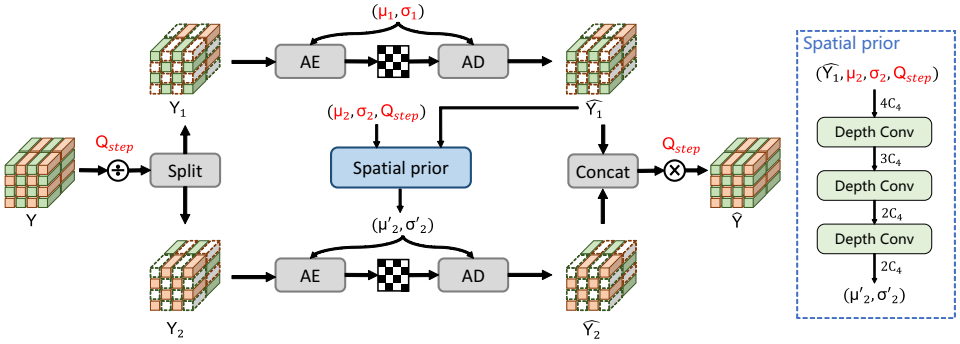

Figure 11: The structure of our dual spatial prior.

Schemes of models with different capacities. Tab. 3 summarizes three models used in this paper. We also compare the model size and MACs for 1920×1088 inputs.

Table 3: Architecture schemes with statistics of models' #Params (MB) and MACs (G). Note that Others include of the hyperprior model, the spatial prior model, and the entropy model. The resolution of the input image is 1920×1088.

Model	C_1, C_2, C_3, C_4		Encoder	Decoder	Others	Total
Large (**L**)	192, 192, 192, 192	#Params	3.19	3.38	10.82	17.38
		MACs	549.92	538.83	76.64	1165.39
Medium (**M**)	128, 128, 192, 192	#Params	2.08	2.33	10.82	15.23
		MACs	262.90	258.56	76.64	598.10
Small (**S**)	64, 64, 128, 192	#Params	0.82	1.14	10.82	12.78
		MACs	73.17	73.82	76.64	223.63

B MASK DECAY

Compress one layer's representation by mask decay. In Fig. 4, our mask contains N_2 parameters which are initialized by 1. After training, N_s parameters remain to be positive while $N_2 - N_s$ parameters become zeros. For these $N_2 - N_s$ channels, corresponding filters in Conv #1 and #2 are pruned safely. And for N_s positive parameters, Claim 1 guarantees that they can be multiplied into Conv #2 without changing the outputs. Hence, our sparse mask is able to be eliminated safely.

Claim 1. *Our mask can be merged into the next convolution layer while keeping the outputs unchanged.*

Proof. As showed in Fig. 4, assume $\boldsymbol{W}_1$ and $\boldsymbol{b}_1$ are parameters in Conv #1, and $\boldsymbol{W}_2$ and $\boldsymbol{b}_2$ are parameters in Conv #2. The LeakyReLU applies the element-wise function:

$$f(x) = \begin{cases} x, & \text{if } x \geq 0\,, \\ \text{negative_slope} \times x, & \text{otherwise}\,. \end{cases} \tag{5}$$

Then given the input $\mathbf{X}$, the forward computation of the pruning model is

$$\mathbf{Y}_1 = \boldsymbol{W}_1 \circledast \mathbf{X} + \boldsymbol{b}_1\,, \tag{6}$$

$$\mathbf{Y}_2 = \mathbf{Y}_1 \odot_c \boldsymbol{m}\,, \tag{7}$$

$$\mathbf{Y}_3 = f(\mathbf{Y}_2)\,, \tag{8}$$

$$\mathbf{Y}_4 = \boldsymbol{W}_2 \circledast \mathbf{Y}_3 + \boldsymbol{b}_2\,. \tag{9}$$

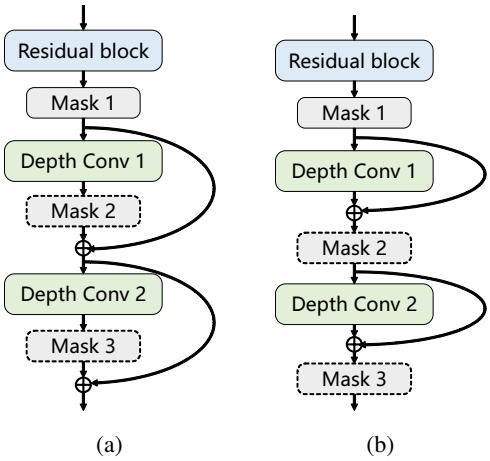

Figure 12: (a) and (b) present two manners to insert masks for the block-level dimensionality. (a) is better because it guarantees the identity gradient back-propagation.

where $\circledast$ denotes the convolution operator, and $\odot_c$ means the channel-wise product. Note that the element of $\boldsymbol{m}$ is positive, therefore

$$\mathbf{Y}_4 = \boldsymbol{W}_2 \circledast f(\mathbf{Y}_1 \odot_c \boldsymbol{m}) + \boldsymbol{b}_2 \,, \tag{10}$$

$$= \boldsymbol{W}_2 \circledast (f(\mathbf{Y}_1) \odot_c \boldsymbol{m}) + \boldsymbol{b}_2 \,, \tag{11}$$

$$= (\boldsymbol{W}_2 \odot_c \boldsymbol{m}) \circledast f(\mathbf{Y}_1) + \boldsymbol{b}_2 \,. \tag{12}$$

The last equation holds for the linearity property of the convolution operator. At last, keeping the outputs unchanged, the weights within Conv #2 become $\boldsymbol{W}_2 \odot_c \boldsymbol{m}$ by merging the mask. □

Mask decay for our encoder and decoder. How to prune the channels outside the residual connection is a crucial problem. Fig. 12 shows two manners. Note that Mask #1, #2, and #3 share weights. If we insert Mask #2 and #3 after the residual adding (as in Fig. 12b), the identity output will be degraded by these masks, because most masks are in $[0, 1]$. Inserting masks as in Fig. 12a mitigates this problem and guarantees the identity gradient back-propagation.

C THE SCALABLE ENCODER

Another option to achieve our scalable encoder is directly training our framework end-to-end (cf. Fig. 13). During training, all encoders generate bitstreams for decoding, and the corresponding loss is calculated. However, this baseline requires a large GPU memory for training and performs badly demonstrated by our experiments. Different from training one encoder each time, end-to-end training requires optimizing all parameters of the whole network. And it does not utilize the knowledge in the pretrained cumbersome model. Hence, the training is inefficient. Tab. 4 summarizes results for our scalable encoders with the large decoder. We also conduct experiments for the small decoder, and Tab. 5 summarizes results.

D DETAILS FOR EXPERIMENTS

Training. The training dataset is the training part of Vimeo-90K (Xue et al., 2019) septuplet dataset. During training, the raw image is randomly cropped into 256×256 patches. For a fair comparison, all models are trained with 200 epochs. For our method, it cost 60 epochs for the teacher to transform into the student with mask decay, then the student is finetuned by 140 epochs. AdamW is used as the optimizer with batch size 16. The initial learning rate is $2e-4$, and decays by 0.5 at 50, 90, 130, 170 epochs. The default decay rate η for our mask decay is $4e-5$. All models are trained on a computer with 8 V100 GPUs. Our large and small model cost about 33 and 16 hours for training, respectively.

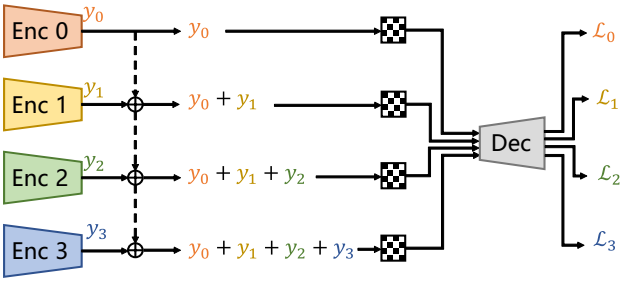

Figure 13: A naive baseline for the scalable encoder trained in an end-to-end manner. Dash lines means detaching gradients.

Table 4: BD-Rate (%) comparison for PSNR on Kodak. The encoder is scalable and the decoder is Large. MACs in the table are for encoding with 1920×1088 resolution, while MACs for other modules is 615.47 G in total. The anchor is VTM.

	S_0	$S_{0,1}$	$S_{0,1,2}$	$S_{0,1,2,3}$	L
MACs (G)	73.17	146.34	219.51	292.68	549.92
End-to-end	3.1	2.2	1.8	1.6	-1.1
Separate	2.8	1.9	1.7	1.5	
One-by-one	2.8	1.4	0.9	0.6	
Ours	**1.4**	**0.4**	**-0.1**	**-0.2**	

Testing. BD-Rate (Bjontegaard, 2001) for peak signal-to-noise ratio (PSNR) is our main metric to compare different models. We use VTM 17.2 as the anchor, while negative numbers indicate bitrate saving and positive numbers indicate bitrate increase. Testing datasets contain Kodak (Franzen, 1999), HEVC test sequences (Sharman & Suehring, 2017), Tecnick (Asuni & Giachetti, 2014). Kodak contains 24 images with the 768×512 resolution. And Tecnick consists of 100 images with the 1200×1200 resolution. HEVC is a video dataset that consists of 4 classes: B (1080P), C (480P), D (240P), E (720P). For testing image codecs on HEVC, only the first frame of each video sequence is used. We report BD-Rate on each HEVC class and average all for convenient comparison. For one testing image, we pad zeros on the boundary to make the resolution as multiples of 64. We test the latency on two computers. One is equipped with one 2080Ti GPU and two Intel(R) Xeon(R) E5-2630 v3 CPUs. Another is equipped with one A100 GPU and one AMD Epyc 7V13 CPU. Both encoding and decoding times are gathered. Note that the latency contains arithmetic coding, but ignores the disk I/O time.

E IMPLEMENT DETAILS FOR OTHER IMAGE CODECS

VTM-17.2[1] is the reference software for the state-of-the-art traditional codec (H.266/VVC), and it is released on Jul 17, 2022. We use CompressAI-1.1.0[2] to gather evaluation results. We install these two softwares according to the official instructions. The default configuration of VTM is used. All results are gathered by running the following command:

```
python -m compressai.utils.bench vtm [path to image folder] -c [path to VTM folder]/cfg/
    encoder_intra_vtm.cfg -b [path to VTM folder]/bin -j 8 -q 24 26 28 30 32 34 36 38
    40 42
```

[1] https://vcgit.hhi.fraunhofer.de/jvet/VVCSoftware_VTM/-/releases/VTM-17.2

[2] https://github.com/InterDigitalInc/CompressAI/releases/tag/v1.1.0

Table 5: BD-Rate (%) comparison for PSNR on Kodak. The encoder is scalable and the decoder is Small. MACs in the table are for encoding with 1920×1088 resolution, while MACs for other modules is 150.46 G in total. The anchor is VTM.

	S_0	$S_{0,1}$	$S_{0,1,2}$	$S_{0,1,2,3}$	L
MACs (G)	73.17	146.34	219.51	292.68	549.92
End-to-end	9.1	8.6	8.2	7.9	
One-by-one	7.3	6.3	5.8	5.5	4.0
Ours	**6.3**	**5.6**	**5.0**	**4.8**	

Table 6: Latency (ms) comparison with SwinT. Lower means better. To compare with SwinT, we test ours on 2080Ti with 768×512 inputs, and directly cite their numbers. Note that SwinT does not reconstruct images for encoding, whereas ours reconstructs images.

Latency	SwinT				Ours		
	SwinT-ChARM	SwinT-Hyperprior	Conv-ChARM	Conv-Hyperprior	Large	Medium	Small
encoding	135.3	100.5	95.7	62.0	63.0	44.7	28.4
decoding	167.3	114.0	264.0	219.3	41.1	32.4	24.4

Entroformer[3] is proposed in Qian et al. (2022). Two series models are released: parallel, non-parallel. We evaluate the RD performance of the parallel model on Tecnick and Kodak, respectively. But when we measure its latency, we find it requires a large GPU memory for generating bitstreams.

STF[4] is proposed in Zou et al. (2022). The authors release two architectures: CNN-based model and Transformer-based model. But, only two pretrained models for two qualities are available. We evaluate one model on Tecnick and cite numbers in their paper for Kodak.

SwinT-ChARM is proposed in Zhu et al. (2022). Because the authors did not release their code, we cite numbers in their paper and the webpage[5]. The latency is reported in their paper, and Tab. 6 compares ours with theirs. Our models achieve lower latency.

F MORE EXPERIMENTAL RESULTS

Comparison for our different models. Fig. 14 presents RD curves on Kodak. Our mask decay "MD" improves our baseline models significantly. In these figures, "L", "M", and "S" denote Large, Medium, and Small, respectively. And "EncS-DecL" means the encoder is small and the decoder is large.

Disscusions for important channels. Tab. 7 presents the specific numbers in Fig. 8a. There are four masks in the first two blocks. We analyze the first mask (ResB.1.mask1) for example. Mask decay is required to choose 64 channels from the teacher's 192 channels. With four different decay rates, only 83 channels are considered and they choose 44 channels together. It seems like these channels are important for superior performance. We treat these 83 channels as $\mathbb{C}_F$ and let our algorithm avoid choosing these channels. A1B (avoid specific channels in the first block) only selects 2 channels in $\mathbb{C}_F$, while A2B selects 3. First, it demonstrates our AMD indeed avoids specific channels. Next, we compare the RD performance of these methods. Surprisingly, AMD results in comparable performance. That indicates that there are no important channels for superior performance.

[3] https://github.com/damo-cv/entroformer
[4] https://github.com/Googolxx/STF
[5] https://openreview.net/forum?id=IDwN6xjHnK8

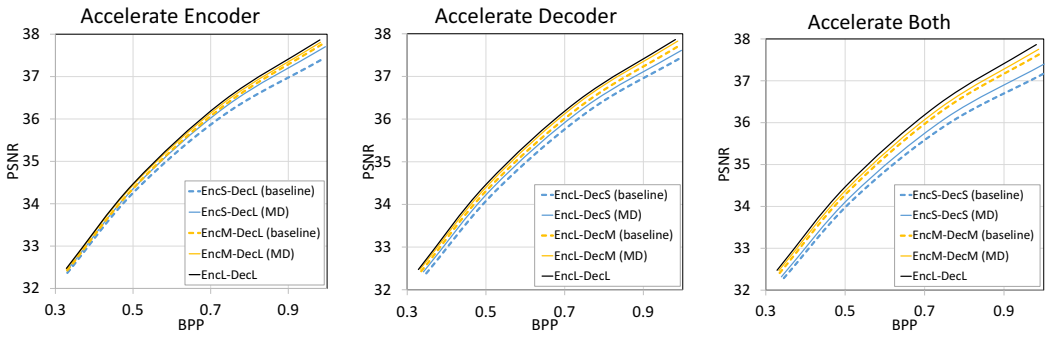

(a) (b) (c)

Figure 14: RD curves on Kodak for our models with different complexities. "MD" denotes mask decay. (a) and (b) show curves for accelerating only encoder or decoder, respectively. (c) presents curves for accelerating both encoder and decoder.

Table 7: The relationship for channels chosen by different settings. "ResB.1" and "DCB.2" means the Residual block and the Depth Conv block, respectively. Note that they are the first two blocks in our model. $\mathbb{C}$ denotes the set of chosen channels, and $F = \{1e-5, 2e-5, 4e-5, 8e-5\}$ which indicates different decay rates in Figure 7a. A1B and A2B means avoiding specific channels in the first 1 and 2 blocks, respectively. We report the cardinality of each set.

Mask Layer	Teacher	Student	$\cap_{i \in F} \mathbb{C}_i$	$\cup_{i \in F} \mathbb{C}_i$	$\mathbb{C}_{A1B} \cap \mathbb{C}_F$	$\mathbb{C}_{A1B} \cup \mathbb{C}_F$	$\mathbb{C}_{A2B} \cap \mathbb{C}_F$	$\mathbb{C}_{A2B} \cup \mathbb{C}_F$
ResB.1.mask1	192	64	44	83	2	145	3	144
ResB.1.mask2	192	64	47	83	0	147	2	145
DCB.2.mask1	192	64	44	86	56	94	7	143
DCB.2.mask3	768	256	179	328	191	393	12	572

G VISUALIZATIONS

Fig. 15 visualizes our models' reconstructions. Comparing our small model and large model, the reconstructions are almost the same. That demonstrates our mask decay does not result in extra artifacts. Compared with VTM, ours enjoy fast encoding and decoding.

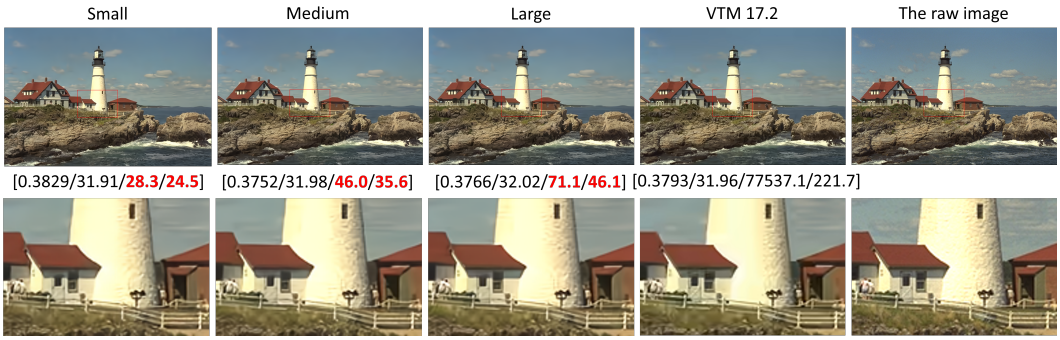

Figure 15: Visualization of our models' reconstruction. Numbers in the tuple are BPP, PSNR, the encoding time (ms), and the decoding time (ms), respectively. Note that the latency is measured on a computer with 2080Ti GPU. Our models are dramatically faster than VTM.