

Efficient Progressive Image Compression with Variance-aware Masking

Alberto Presta¹ Enzo Tartaglione² Attilio Fiandrotti^{1,2} Marco Grangetto¹ Pamela Cosman³

¹University of Turin, Italy

²LTCI, Télécom Paris, Institut Polytechnique de Paris

³Dept. of Electrical and Computer Engineering, UC San Diego, CA, USA

alberto.presta@unito.it

Abstract

Learned progressive image compression is gaining momentum as it allows improved image reconstruction as more bits are decoded at the receiver. We propose a progressive image compression method in which an image is first represented as a pair of base-quality and top-quality latent representations. Next, a residual latent representation is encoded as the element-wise difference between the top and base representations. Our scheme enables progressive image compression with element-wise granularity by introducing a masking system that ranks each element of the residual latent representation from most to least important, dividing it into complementary components, which can be transmitted separately to the decoder in order to obtain different reconstruction quality. The masking system does not add further parameters or complexity. At the receiver, any elements of the top latent representation excluded from the transmitted components can be independently replaced with the mean predicted by the hyperprior architecture, ensuring reliable reconstructions at any intermediate quality level. We also introduced Rate Enhancement Modules (REMs), which refine the estimation of entropy parameters using already decoded components. We obtain results competitive with state-of-the-art competitors, while significantly reducing computational complexity, decoding time, and number of parameters.

1. Introduction

In recent years, Learned Image Compression (LIC) has attracted significant interest, outperforming standardized codecs [5, 33, 39] in Rate-Distortion (RD) efficiency for natural images [26, 43]. In a learnable codec, an encoder on the transmitter side projects an image to a latent space that is quantized and entropy coded into a compressed bitstream. On the receiver side, the bitstream is processed by a decoder that reverses the encoding process, recovering (a distorted

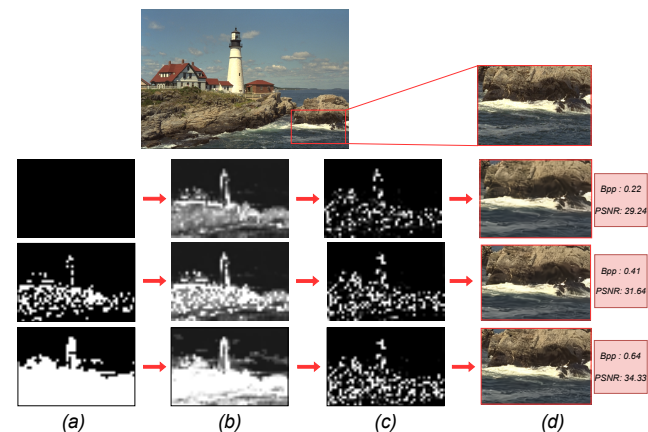


Figure 1. Compression results for three different qualities, which increase across rows. Adding details via the masking system (a) increases the standard deviation in the non-masked latent representation (b) to add details (c) for a better reconstruction (d).

version of) the original image.

However, LIC codecs still face the challenge of meeting the rate of channels whose capacity changes as a function of connection type and congestion at the nodes. Scalable coding [31] consists in encoding a content as one base bitstream enabling to recover a low-quality version of the content and a few enhancement bitstreams enabling improved quality when received. With progressive image compression, sometime known also as fine-grained scalability (FGS), this concept is further extended improving the quality of the reconstructed image as each extra bit from the same bitstream is received [23], allowing to truncate it at. A few learnable image compression schemes with progressive decoding properties have been known to exist to date; some allow adjusting the tradeoff between compressed bitrate and quality by exploiting a single rate-variable model [8, 18, 21]. However, a different bitstream must be encoded and delivered for each different bitrate target. Early models such as Diao *et*

al. [9] and Lu *et al.* [27] achieved progressive representation through distributed recurrent autoencoder and nested quantization, respectively. Lee *et al.* [22] introduced a novel way to represent each element of the latent representation based on trit-plane coding, achieving state-of-the-art RD efficiency, without exploiting any kind of context model, which were introduced by [15] for both rate and distortion reduction. A RD prioritized transmission system was introduced to be able to find the trit-planes with more information and give them priority for being encoded first. However, exploiting such modules for both encoding and decoding and finding the right priority transmission distribution make [15] and [22] costly in terms of computational resources and time.

In this work, we propose a learnable, efficient, and progressive image compression architecture. It has two initial levels, *base* and *top*; the first defines the lowest quality while the second allows the system to achieve all the other higher qualities. To achieve FGS, we first compute a residual latent representation by means of element-wise difference between the top and base ones, and then break it down into complementary parts, which form the final bitstream, which can be encoded and sent separately, resulting in reconstruction at multiple qualities, as shown in Fig. 1. Furthermore, we added in the main architecture learnable rate enhancement modules (REMs) to further improve the estimation of entropy parameters. Our progressive approach is built on a channel-wise entropy parameter module [30].

The main contribution of this paper are the follows:

- We introduced a method where, to achieve FGS, complementary portions of a residual latent representation are added to the *base* one, which represents the lowest level in terms of bitrate.
- We exploited a lightweight masking policy which ranks the elements of the residual latent representation from most to least important, also identifying their positions, which allows for the creation of a progressive system where more elements are added incrementally.
- We achieved competitive results with respect to [15] in RD performance, significantly reducing complexity in terms of computational resources, decoding time, and number of parameters.

2. Related work

2.1. Learning-based image compression

While seminal works about learning-based image compression are based on recurrent neural networks [12, 36, 37], those based on variational autoencoders (VAEs) [19] have gained significant importance. Initially, the general architecture is formed by a simple autoencoder [3, 35] where the latent representation is modeled with a channel-wise probability distribution, extracted analytically [32] or with a

neural network. Quantization is replaced by uniform additive noise during training; this effective technique became the most common practice in such models. To further improve the entropy estimation of the latent elements, a hyperprior autoencoder is added to the main architecture [2, 29], with the aim of finding spatial correlation within the image. In such a scenario, the main latent representation is modeled with a Gaussian distribution. More effort has been put into optimizing the entropy estimation by introducing a parallelizable checkerboard context model [13] or a channel-wise approach [30]. The latter divides the latent space into blocks, called slices, along the channel dimension. For example, if the latent space has 320 channels and 10 slices are desired, each slice contains 32 channels. These slices are sequentially encoded, exploiting previous ones; this method can be seen as an autoregressive model along channels, improving efficiency of such models, and enhancing entropy parameter estimation. Other alternative techniques have been used, such as invertible modules [41], sparse representation spanned by learned visual codebooks [16], or graph-based attention [34].

Other works tried to exploit an attention mechanism to optimize the bits allocation, by integrating non-local attention blocks in the image compression architecture [7, 25, 42]. In [43], the authors replaced the non-local attention block introduced in [25] with a window block computed from spatially neighboring elements that focus on image details, proving that this can improve RD performance with less computation. Similarly, the model in [26] combines Transformer and CNN blocks to leverage both local and non-local modeling abilities. They also improved the entropy estimation module by adding a Swin-attention module.

2.2. Progressive image compression

In progressive image compression, reconstruction quality improves as more bits of the bitstream are received and decoded. To achieve this property, early works exploited recurrent neural networks (RNNs) [28]. The model in [36] exploits long-short memory [14] to transmit bits progressively. The encoder and decoder structures were improved in [37] for higher resolution images. In [12] a recurrent model is exploited, enhancing perceptual quality through a generative model; however, this method is suitable only for low resolution patches. In [9], a distributed recurrent auto-encoder was designed for scalable image compression, while [17] introduced an iterative process to generate the binary codec and supported spatially adaptive bit rates that dynamically adjust bit allocation based on context.

In [6], a new architecture with a single encoder and two decoders was introduced. This setup produces two representations: a preview image from the first decoder and a full-quality image from the second one. Although it uses two representations like our model, [6] relies on concatenation rather than residual addition, thus doubling the input

dimension of the second decoder. Furthermore, they did not introduce a technique for sending partial portions of the latent space. Due to these limitations, [6] supports only two levels of quality. These methods not only have worse performance than standard codecs, but also support only coarse granular scalability. In [40], several progressive encoders were introduced, one for each target quality, to achieve different levels of reconstruction; however, they do not provide FGS, and this method requires training separate encoders for each desired target quality.

To obtain FGS, [27] introduced a nested quantization, obtaining results comparable to conventional deep image codecs. [22] proposed a novel trit-plane coding method, which represents each element of the latent space with a ternary number. Specifically, they use L trits to express each element, ordered from the most significant to the least significant trit. In that sense, they also introduced an RD-prioritized transmission system that optimizes the order of trit transmission based on their priorities. This method was improved in [15] by introducing a context model capable of improving both the estimation of the entropy parameters and the quality of reconstruction. Although [15] obtains competitive results with standard codecs such as VVC [5], it still relies on complex context modules and iterative RD-prioritized transmission; this makes [15] extremely complex in terms of encoding/decoding time, computational resources, and the number of parameters required.

3. Proposed method

3.1. General architecture

Our progressive method is based on [43] as illustrated in Fig. 2. It uses a hyperprior-based architecture [2] and a channel-wise entropy model [30]. We generate two latent representations: the *base* provides the reconstructed image at the lowest bitrate, while the *top* yields the image at any higher bitrates; to obtain FGS we first obtain a residual representation by means of element-wise difference between the top and the base latents, and then we introduce a masking system that ranks its elements from the most to least important, and mask only the ones needed to achieve a specific target quality, replacing the others with the mean. Quality is indicated by $q \in [0, 100]$, where 0 and 100 represent the lowest and highest level, respectively.

Our approach has two advantages: It makes the first training step easier, since it is sufficient to introduce two different target qualities, i.e. *base* and *top*, in the loss function, and it enables choosing the distance between the lowest and highest levels, thus focusing on a specific bit range. Furthermore, the masking policy introduced in Sect. 3.3 is simple and efficient, with no parameters to train.

For an image \mathbf{x} (Fig. 2) the two latent representations $\mathbf{y}^{b/t}$ are extracted with the encoders $g_a^{b/t}$, that is,




Figure 2. Overview of our proposed architecture. Green and red boxes represent encoder and decoder modules, respectively, while blue boxes must be stored at both encoder and decoder. q represents the target quality.

$\mathbf{y}^{b/t} = g_a^{b/t}(\mathbf{x}; \phi^{b/t})$, where $\phi^{b/t}$ are learnable parameters, and the superscript b/t represents both the *base* and the *top* representations in a more compact manner. The latents $\mathbf{y}^{b/t}$ are modeled as Gaussian distributed with mean $\mu^{b/t}$ and standard deviation $\sigma^{b/t}$. As in [2], a hyperprior autoencoder $\{h_a, h_s\}$ is introduced to find spatial correlations within the image, optimizing rate reduction. In particular, we have $\mathbf{z} = h_a(\mathbf{y}^b, \mathbf{y}^t; \phi_h)$, which is then quantized to $\hat{\mathbf{z}} = Q(\mathbf{z})$. Q is the quantization function, which is replaced by the addition of uniform noise (\mathcal{U}) during training. We then obtain $[d_\mu^{b/t}, d_\sigma^{b/t}] = h_s(\hat{\mathbf{z}}, \theta_h)$; these tensors are part of the input for the progressive channel-wise entropy estimation module (PCEEM), explained in Sec. 3.2, to extract both entropy parameters and $\hat{\mathbf{y}}^q$:

$$\hat{\mathbf{y}}^q = \text{PCEEM}(\mathbf{y}^{b/t}, q, [d_\mu^{b/t}, d_\sigma^{b/t}]). \quad (1)$$

We input q to obtain the latent representation for the target quality q , then encode and add them to the final bitstream.

The final reconstruction $\hat{\mathbf{x}}^q$ at quality q is obtained through the decoders: $\hat{\mathbf{x}}^q = g_s^{b/t}(\hat{\mathbf{y}}^q, \theta^{b/t})$, where $\theta^{b/t}$ are learnable parameters. We implemented two different decoders that are used depending on the target quality; If $q = 0$, we use g_s^b , while in any other case we use g_s^t .

3.2. Progressive channel-wise entropy estimation

To develop a *progressive channel-wise entropy estimation model* (PCEEM), both \mathbf{y}^b and \mathbf{y}^t are divided into s slices along the channels [30], for a total of $2 \times s$ slices. Fig. 3 illustrates PCEEM. To achieve the initial base quality reconstruction, only \mathbf{y}^b is necessary. For the other levels, we also consider \mathbf{y}^t .

To encode the *base* latent, the process is similar to previous works [26, 30, 43]. Given an index i such that $0 \leq i < s$ and the corresponding slice y_i^b , we obtain the entropy parameters $(\mu_i^b, \sigma_i^b) = \Psi_i^b(\mu_i^b, \sigma_i^b, \hat{y}_{<i}^b)$, where Ψ_i^b is a learnable module and $\hat{y}_{<i}^b$ are the previous slices. After the quantization step, $\hat{y}_i^b = Q(y_i^b)$, a module LRP_i^b is applied to reduce the quantization error; the resulting slice \hat{y}_i^b is concatenated to the previous slices. The final tensor $\hat{\mathbf{y}}^b$ can be used either as input to the decoder g_s^b to obtain the base reconstruction

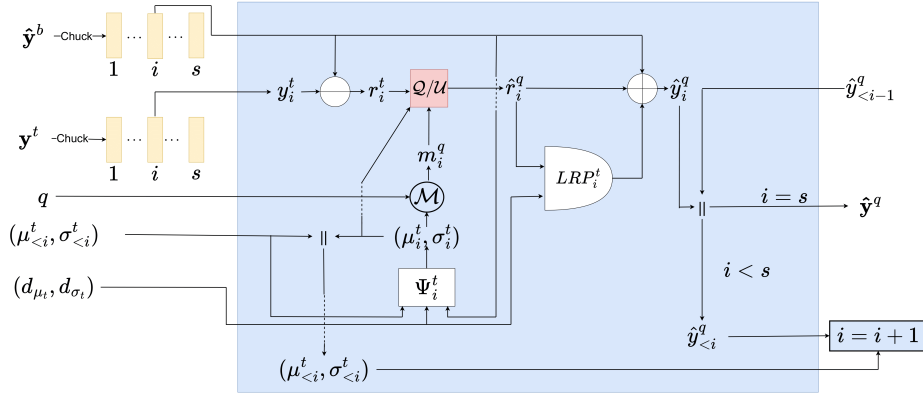


Figure 3. Progressive channel wise entropy estimation model (PCEEM) during the i -th slice, considering a general quality q . $\hat{\mathbf{y}}^b$ represents the base latent representation already obtained. \parallel represents concatenation along channels, while \odot represents element-wise operation, which can be summation (+) or subtraction (-).

image $\hat{\mathbf{x}}^b$, or as input for estimating the entropy of higher qualities.

For any other possible quality q greater than 0, the *top* latent representation is handled differently. Given the index i and a target quality q , we have the initial slice y_i^t ; as in residual video coding [24], we first compute the latent residual representation $r_i^t = y_i^t - \hat{y}_i^b$, and then extract the entropy parameters (μ_i^t, σ_i^t) through a learnable module Ψ_i^t :

$$(\mu_i^t, \sigma_i^t) = \Psi_i^t(\hat{y}_i^b, d_\mu^t, \mu_{<i}^t, d_\sigma^t, \sigma_{<i}^t). \quad (2)$$

Note that Ψ_i^t is the same for every possible target quality. Unlike the *base* scenario, this module receives the entropy parameters $(\mu_{<i}^t, \sigma_{<i}^t)$ of previous slices instead of the slices $y_{<i}^t$, crucial for applying the masking policy described in Sect. 3.3, enabling progressive quality representation without altering the already obtained bitstream entropy parameters. σ_i^t is used to compute the mask $m_i^q = \mathcal{M}(\sigma_i^t, q)$, where \mathcal{M} represents the masking operator.

In summary, m_i^q masks certain elements of r_i^t , replacing them with the mean μ_i^t . In this way, r_i^t is effectively down-sampled, with only one encoded part and the rest replaced with reasonable values such as the mean; this process is described by the following equation:

$$\hat{r}_i^q = Q(r_i^t - \mu_i^t) \otimes m_i^q + \mu_i^t, \quad (3)$$

where \otimes represents element-wise multiplication. We first mask $Q(r_i^t - \mu_i^t)$, which is then encoded and added to the bitstream; At the decoder side, we then add back the mean μ_i^t (eq. 3); in this way, all the elements omitted by m_i^q are replaced by the mean already computed. This operation is represented by the red box in Fig. 3. We employed ANS coding [10] for entropy coding; since the residual representation elements are modeled as independent Gaussians with

parameters derived from Ψ_i^t , the bitstream can be split into complementary parts without loss of efficiency.

A network LRP_i^t is applied to \hat{r}_i^q to reduce the quantization error [30]. Finally, the resulting \hat{r}_i^q is added to the corresponding base slice \hat{y}_i^b , obtaining the final quantized slice \hat{y}_i^q ; the latter is further concatenated (symbol \parallel in Fig. 3) with the previous slices, obtaining $\hat{y}_{<i}^q$. If $i < s$, then the next slice $i + 1$ is encoded (blue box $i = i + 1$ in Fig. 3).

Once all slices have been computed, the resulting tensor $\hat{\mathbf{y}}^q$ is used as input for g_s^t to obtain the reconstructed image $\hat{\mathbf{x}}^q$ at a specific target quality q .

3.3. Progressive mask-based coding

To encode elements of r_i^t , we propose a progressive masking method based on the corresponding standard deviation σ_i^t ; it is worth pointing out that it is possible to compute σ_i^t on both the encoder and decoder sides, as they are required for the encoding / decoding of the bitstream, and also it does not change with respect to different qualities. The general idea is to add more and more elements of r_i^t to progressively improve reconstruction quality. In this context, it is reasonable to identify the most important elements in the residual latent representation as those with a larger variance, as they are more likely to change and cause a greater reconstruction error; this is the reason we use σ_i^t .

In particular, a progressive masking procedure for the i -th block r_i^t is described in Alg. 1. Each target quality value $q \in [0, 100]$ unambiguously determines the elements of r_i^t to be masked, i.e. skipped, in the encoded bitstream. A higher q means fewer masked elements. We compute the $(100 - q)$ -th percentile of σ_i^t and mask all elements of r_i^t whose corresponding standard deviation is below that value (lines 7-12 in Alg. 1). These elements are not encoded in the bitstream, but are replaced with the mean μ_i^t (line 16).

This approach first encodes elements with higher stan-

dard deviation, improving reconstruction quality. The latent representation is divided into complementary portions, prioritized by importance, which can be decoded separately. This masking policy has multiple advantages:

- It is an easy-to-compute progressive mechanism to obtain RD transmission priority, since it does not require any optimization procedure or further parameters;
- It is possible to compute such masks consistently both at the encoder and the decoder, providing also elements position: this allows the encoder to simply skip the masked values without any signaling overhead.
- Only the value q needs to be signaled within the encoded bitstream. However, q is represented by a floating point or an integer with negligible overhead.

Algorithm 1 σ -masking procedure \mathcal{M}

Input: $\sigma_i^t, \mu_i^t, r_i^t, q$
Output: m_i^q, \hat{r}_i^q

- 1: $sh \leftarrow \text{Shapes}(\sigma_i^t)$ ▷ extract the shape of σ_i^t
- 2: $\sigma_i^t \leftarrow \text{Flatten}(\sigma_i^t)$ ▷ flatten σ_i^t to a 1-dim vector
- 3: $L \leftarrow \text{len}(\sigma_i^t)$ ▷ extract length of σ_i^t
- 4: q -th percentile $\leftarrow \text{Perc}(\sigma_i^t, 100 - q)$
- 5: $m_i^q \leftarrow \text{Zeros}(L)$
- 6: $j \leftarrow 0$
- 7: **while** $j < L$ **do**
- 8: **if** $\sigma_i^t < q$ -th percentile **then**
- 9: $m_i^q[j] = 0$
- 10: **else**
- 11: $m_i^q[j] = 1$
- 12: **end if**
- 13: $j \leftarrow j + 1$
- 14: **end while**
- 15: $m_i^q \leftarrow \text{Reshape}(m_i^q, sh)$ ▷ reshape m_i^q to sh
- 16: $\hat{r}_i^q \leftarrow Q(r_i^q - \mu_i^t) * m_i^q + \mu_i^t$




Figure 4. Blueprint of REM for a fixed checkpoint quality \bar{q} and for the slice i .

3.4. Slice-wise rate enhancement modules

Taking inspiration from [15], we introduce checkpoint-based learnable modules by means of *Slice-wise Rate En-*

hancement Modules (REMs) that aim to improve the estimation of the entropy parameters of each slice. The idea is to fix some checkpoint qualities along the bit range and use them as input for a REM to improve the encoding of higher quality. As explained in Sect. 3.2, regarding the *top* representation, we do not utilize slices \hat{y}_i^t as in [30]; REMs are meant to incorporate information coming from already encoded/decoded residual representation.

Consider a specific slice i , a checkpoint quality \bar{q} , and a target quality q^* , with $q^* > \bar{q}$. Fig. 4 shows how we inserted such a module, corresponding to the i -th slice and checkpoint quality \bar{q} , which we called $REM_i^{\bar{q}}$. It receives three elements as input: the quantized representation of the checkpoint $\hat{y}_i^{\bar{q}}$ (already decoded and present in memory), and (μ_i^k, σ_i^k) with $k = \{b, t\}$, and enhances the estimation of (μ_i^t, σ_i^t) only for residual latent elements that must be added to transition from checkpoint quality \bar{q} to q^* . This is represented by the mask m_i^Δ in Fig. 4, computed as:

$$m_i^\Delta = \mathcal{M}(\sigma_i^t, q^*) - \mathcal{M}(\sigma_i^t, \bar{q}). \quad (4)$$

which preserves only those points that need to be added progressively for the quality q^* ; this means that the $REM_i^{\bar{q}}$ will affect only the elements that have been added to the bitstream, leaving unchanged the parameters of the other elements. In practice, by fixing some checkpoint qualities \bar{q} 's for improving rate estimation, we divide the bit range into different subranges, denoted by the symbol L_b if only \hat{y}^b is used (*base* level), L_0 if no REMs are used, and $L_{\bar{q}_j}$ if the i -th $REM^{\bar{q}_j}$ is used; here j travels across the number of checkpoints added.

3.5. Training

Training occurs in multiple phases. The first phase is to train the general architecture without any REMs to reconstruct the image at the base and top qualities. The loss function is the sum of the *rate-distortion* trade-off for the two different levels:

$$\mathcal{L} = \sum_{k=\{b,t\}} L_k, \quad (5)$$

where each term of the summation is represented by

$$L_k = \lambda_k \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [d(\mathbf{x}, \hat{\mathbf{x}}^k)] + \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}} [-\log_2 p_{\hat{\mathbf{y}}^k | \hat{\mathbf{z}}}(\hat{\mathbf{y}}^k | \hat{\mathbf{z}}) - \log_2 p_{\hat{\mathbf{z}}}(\hat{\mathbf{z}})]. \quad (6)$$

The first term in (6) represents the quality of the reconstruction, where d represents a generic distortion metric, while the other two terms estimate the bitrate required for the two latent representations $\hat{\mathbf{y}}^k$ and $\hat{\mathbf{z}}$. Here, $\lambda = \{\lambda_b, \lambda_t\}$ denotes the Lagrangian parameters that balance the trade-off between rate and quality.

The second phase refines g_s^t to be able to reconstruct images with different qualities $0 < q \leq 100$. When $q=100$ we




Figure 5. Rate-distortion performance of our method compared with progressive image compression algorithms: Jeon [15], Lee [22], Lu [27] and JPEG2000 [33]. We tested our method on Kodak (left), JPEG-AI (center), and CLIC validation dataset (right).

consider the *top* latent representation without the mask. In this context, we only optimize the first term of (6), considering only $k=t$. In particular, in this stage, at each training step, we sample a value of q and refine g_s^t for that specific target quality; this enhances the decoder’s robustness to various values of q , representing different versions of the residual latent representation.

Finally, we insert a fixed number of REMs at certain checkpoint qualities and train them in the architecture, keeping other network parts frozen. In this phase, the training loss is the second term of (6) considering $k=t$. The number of REM modules is arbitrary, trading off the complexity of the entire architecture and the rate reduction.

4. Experiments

4.1. Setup and training details

For training, we used 300k images from the OpenImages dataset [20] and the Adam optimizer, fixing the dimensions of $\mathbf{y}^{b/t}$ and \mathbf{z} to 320 and 192, respectively. In the first training phase, we trained the entire model for 100 epochs with a batch size of 16 and a learning rate of 10^{-4} , which was reduced by a factor of two after a plateau with patience of ten epochs. We set $\lambda = \{5 \times 10^{-3}, 5 \times 10^{-2}\}$ in (6).

Secondly, we refine g_s^t for 80 epochs with a learning rate of 10^{-4} , while in the last one we fixed the number of REMs to three to balance performance and complexity, empirically selecting quality checkpoints at $\bar{q} = \{0.5, 7.5, 20\}$ and refining them for 30 epochs. These values, approximately 5%, 40%, and 65% of the Kodak bitstream, ensure that a REM module covers the entire bit range while improving entropy parameter estimation in a sensible portion of the latent space, avoiding excessively high q .

We compare our method with several models (Jeon *et al.* [15], Lee *et al.* [22], Lu *et al.* [27], and JPEG2000 [33]), on three datasets: Kodak [11], CLIC validation dataset [38], and JPEG-AI [1]. Our implementation is built upon the `compressai` [4] library, and we used NVIDIA A40 gpu for training.

Table 1. BD-Rate and BD-PSNR on the three datasets considering Jeon *et al.* as reference and our method as proposed.

	<i>Kodak</i>	<i>JPEG-AI</i>	<i>CLIC</i>
<i>BD-Rate</i>	-1.05	0.47	-0.75
<i>BD-PSNR</i>	0.04	-0.01	0.01

4.2. Rate-distortion (RD) performance

Fig. 5 shows the PSNR results for the three datasets. Our method outperformed three of the other algorithms in the bit range and is competitive with Jeon *et al.*, as seen in the reconstruction in Fig. 6. In particular, we outperformed Jeon *et al.* for low and medium bit rates, with a slight deterioration for higher qualities. The competitiveness of our model is also seen in Table 1, which represents BD-Rate and BD-PSNR for the three different datasets considering Jeon *et al.* as the reference model. Our method is better for Kodak and CLIC, with a small deterioration for JPEG-AI.

Fig. 7 shows the RD trade-off with respect to some state-of-the-art fixed-rate (non-progressive) models. As these methods do not have the progressive property, our architecture is expected not to outperform them. Although we performed less well than the most recent models (for example, Zou *et al.* [43] and Liu *et al.* [26]), we obtained comparable results with other methods, demonstrating our competitiveness. We point out that the deep learning models in Fig. 7 are fixed-rate, i.e., there is a separate model for each target quality, and they do not achieve FGS.

4.3. Model Complexity

Figures 8a and 8b represent GFLOPs and the time to decode on GPU, respectively; we compared our results with the most complex version of Jeon *et al.*, which is competitive in terms of RD.

We analyze the different ranges defined in Sect. 3.4, to understand how complexity varies when adding additional portions of the bitstream. With our efficient masking system and parallelizable modules, our method is optimal in time and computational resources compared to Jeon *et al.*




Figure 6. Kodim14 reconstruction from Kodak dataset by different codecs: Proposed, Jeon *et al.*, JPEG.

Table 2. Complexity comparison between our proposed method and Jeon *et al.* on Kodak.

	GPU			CPU			Parameters (M)	
	Dec. time	Enc. time	~ GFLOPs	Dec. time	Enc. time	GFLOPs	Encoder	Decoder
Proposed	1.15	1.72	787.94	6.33	5.49	789.98	97.26	90.78
Jeon <i>et al.</i>	2.018	1.64	2012	10.9	4.16	1980.87	323.6	399




Figure 7. RD curve for some fixed-rate models on Kodak: Ballé *et al.* [2], Minnen *et al.* [29], Cheng *et al.* [7], Zou *et al.* [43], and Liu *et al.* [26]. Dotted lines represent that such models are not able to achieve FGS.

The complexity improvement of our method is substantial, especially for lower bitrates, i.e., where fewer modules, and so fewer iterations, are required. Our gain marginally decreases when increasing the bit rate while always being significantly faster than Jeon *et al.*. On average, we obtain an improvement of a factor of 2 both in terms of GFLOPs and decoding time, as shown in Figures 8.

Our method becomes less efficient at higher qualities due to REMs, as decoding the fixed checkpoint latent representations is necessary before achieving the target quality. We attribute our superior time performance compared to Jeon *et al.* to several factors: first, our general architecture, with also the addition of REMs, is smaller than the context model used by Jeon *et al.* and only needs to be applied once during decoding; furthermore, our masking policy is fast and efficient, whereas Jeon *et al.* relies on a complex optimization method for RD priority transmission; finally, our algorithm allows one to encode and decode the complemen-





Figure 8. GFLOPs and decoding time complexity on NVIDIA A40 GPU (a,b) vs. Jeon *et al.* for different subranges on Kodak.

tary portions of the final bitstream, making the computation simpler for a specific target quality. Tab. 2 compares the complexity of our method and Jeon *et al.* on the entire bit range of interest. Our method is better for both GPU and CPU usage, requires fewer parameters, and maintains similar performance at encoding.

4.4. Ablation study

Fig. 9 shows how REMs and decoder retraining improve the performance of our method on Kodak. The refinement of g_s^t turns out to be important in improving overall performance (purple line versus pink line in Fig. 9); this is




Figure 9. Results on Kodak considering different configurations.

Table 3. BD-RATE, BD-PSNR, and the number of parameters for different configurations considering Jeon *et al.* as the base model.

	BD-Rate	BD-PSNR	# Dec. Pars (M)
3 REMs + g_s^t Refine	-1.05	0.04	90.8
2 REMs + g_s^t Refine	-0.56	0.02	79.4
1 REM + g_s^t Refine	0.94	-0.04	67.9
g_s^t Refine	1.24	-0.05	56.6
No g_s^t Refine	5.53	-0.25	56.6

reasonable since we make the decoder robust to change of qualities. Furthermore, adding more REMs enhances coding efficiency across various subranges of the bitrate. The module operating in the central portion of the bit range offers the most significant improvement because at this stage in progressive coding, a large portion of the bitstream is already available, meaning the checkpoint quality has more information compared to lower-quality stages, but there is still room for enhancement in the remaining bitstream.

Table 3 shows results in terms of BD-Rate and BD-PSNR, using Jeon *et al.*, as reference model; The proposed method produces a gain in terms of general RD performance only from 2 REMs, always exploiting less parameters.

4.5. Masking policy

Fig. 10 shows how our masking policy works for different q 's, considering a sample image from the CLIC validation dataset. As expected, the base level is characterized by very low std values (Fig. 10 (b)) for less variable areas, i.e., the background, with the model focusing only on the main part of the image, such as the foreground. As q increases, more elements are added. Following the policy introduced in Sec. 3.3, the most relevant elements, which are part of the subject in the foreground, are prioritized (Fig. 10 (d-g)), and then those that have less impact on the final reconstruction in terms of standard deviation (10 (e,h)). In this way we obtain latent representations with increasing details (Fig. 10 (f,i)), resulting in reconstructions at multiple quality levels.

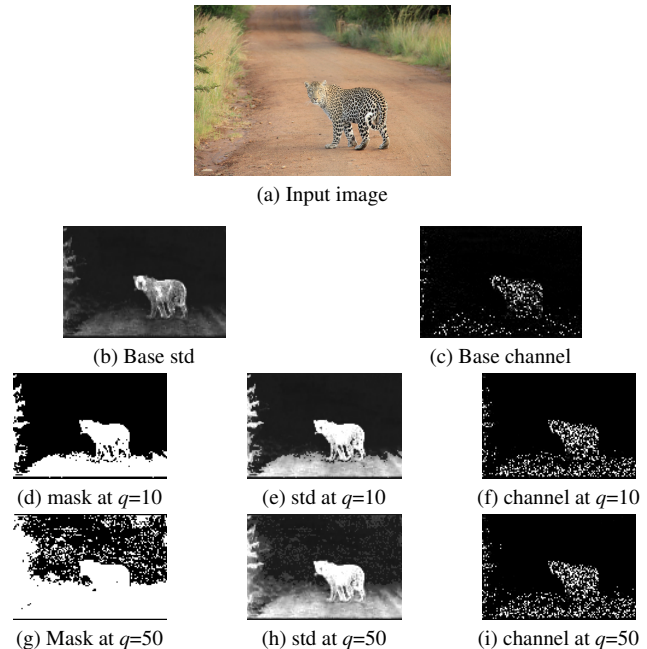


Figure 10. Original image (a), base latent channel and estimate standard deviation (b,c) obtained mask, estimate standard deviation, and resulting latent channel for two different q (d,e,f,g,h,i).

5. Conclusions

We proposed a novel learning-based model for progressive image compression based on two latent representations: *base* for the lowest quality reconstruction and *top* for all higher qualities. FGS is achieved by a lightweight masking system that decomposes the residual between the top and the base latent representations into complementary parts that can be encoded/decoded separately. We also introduced REM modules in the proposed architecture, improving entropy estimation.

Despite simplicity, our method achieves competitive, if not better, RD results with SOTA methods Jeon *et al.* [15], Lee *et al.* [22], reducing computational complexity, decoding time, and required parameters.

Given the promising results, this method could be further improved, for example, by introducing a system for calculating entropy parameters that takes advantage of the lower quality levels continuously, without the need to set checkpoints or REMs, or by introducing a scalable system for the *base* layer, to obtain both granularity on it and a larger bitrange.

Acknowledgements

This research was partially funded by Hi!PARIS Center on Data Analytics and Artificial Intelligence.

References

- [1] Leao Ascenso et al. Learning-based image coding: early solutions reviewing and subjective quality evaluation. In *Optics, Photonics and Digital Technologies for Imaging Applications VI*, 2020. 6
- [2] Johannes Ballé, , et al. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018. 2, 3, 7
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*, 2016. 2
- [4] Jean Bégaint et al. Compressai: a pytorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029*, 2020. 6
- [5] Benjamin Bross et al. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 2021. 1, 3
- [6] Chunlei Cai et al. A novel deep progressive image compression framework. In *Picture Coding Symposium*, 2019. 2, 3
- [7] Zhengxue Cheng et al. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *CVPR*, 2020. 2, 7
- [8] Ze Cui et al. Asymmetric gained deep image compression with continuous rate adaptation. In *CVPR*, 2021. 1
- [9] Enmao Diao et al. Drasic: Distributed recurrent autoencoder for scalable image compression. In *DCC. IEEE*, 2020. 2
- [10] Jarek Duda. Asymmetric numeral systems: entropy coding combining speed of huffman coding with compression rate of arithmetic coding. *arXiv preprint arXiv:1311.2540*, 2013. 4
- [11] Rich Franzen. Kodak lossless true color image suite. *source: http://r0k.us/graphics/kodak*, 4(2), 1999. 6
- [12] Karol Gregor, Frederic Besse, Danilo Jimenez Rezende, Ivo Danihelka, and Daan Wierstra. Towards conceptual compression. *NeurIPS*, 29, 2016. 2
- [13] Dailan He et al. Checkerboard context model for efficient learned image compression. In *CVPR*, June 2021. 2
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8), 1997. 2
- [15] Seungmin Jeon, Kwang Pyo Choi, Youngo Park, and Chang-Su Kim. Context-based trit-plane coding for progressive image compression. In *CVPR*, 2023. 2, 3, 5, 6, 8
- [16] Wei Jiang et al. Neural image compression using masked sparse visual representation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024. 2
- [17] Nick Johnston et al. Improved lossy image compression with priming and spatially adaptive bit rates for recurrent networks. In *CVPR*, 2018. 2
- [18] Fatih Kamisli et al. Variable-rate learned image compression with multi-objective optimization and quantization-reconstruction offsets. In *DCC. IEEE*, 2024. 1
- [19] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [20] Alina Kuznetsova et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 128(7), 2020. 6
- [21] Jooyoung Lee, Seyoon Jeong, and Munchurl Kim. Selective compression learning of latent representations for variable-rate image compression. In *NeurIPS*, 2022. 1
- [22] Jae-Han Lee, Seungmin Jeon, Kwang Pyo Choi, Youngo Park, and Chang-Su Kim. Dpict: Deep progressive image compression using trit-planes. In *CVPR*, 2022. 2, 3, 6, 8
- [23] Weiping Li. Overview of fine granularity scalability in mpeg-4 video standard. *IEEE Transactions on circuits and systems for video technology*, 11(3), 2001. 1
- [24] Dong Liu et al. Deep learning-based video coding: A review and a case study. *ACM Computing Surveys*, 53(1), 2020. 4
- [25] Haojie Liu et al. Non-local attention optimized deep image compression. *arXiv preprint arXiv:1904.09757*, 2019. 2
- [26] Jinming Liu et al. Learned image compression with mixed transformer-CNN architectures. In *CVPR*, 2023. 1, 2, 3, 6, 7
- [27] Yadong Lu, Yin hao Zhu, Yang Yang, Amir Said, and Taco S Cohen. Progressive neural image compression with nested quantization and latent ordering. In *ICIP*, 2021. 2, 3, 6
- [28] Larry R Medsker, Lakhmi Jain, et al. Recurrent neural networks. *Design and Applications*, 5(64-67):2, 2001. 2
- [29] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. *NeurIPS*, 31, 2018. 2, 7
- [30] David Minnen et al. Channel-wise autoregressive entropy models for learned image compression. In *ICIP*, 2020. 2, 3, 4, 5
- [31] J-R Ohm. Advances in scalable video coding. *Proceedings of the IEEE*, 93(1), 2005. 1
- [32] Alberto Presta et al. A differentiable entropy model for learned image compression. In *ICIAP*. Springer, 2023. 2
- [33] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5), 2001. 1, 6
- [34] Gabriele Spadaro et al. Gabic: Graph-based attention block for image compression. In *ICIP*, 2024. 2
- [35] Lucas Theis et al. Lossy image compression with compressive autoencoders. In *ICLR*, 2017. 2
- [36] George Toderici et al. Variable rate image compression with recurrent neural networks. *ICLR*, 2015. 2
- [37] George Toderici et al. Full resolution image compression with recurrent neural networks. In *CVPR*, 2017. 2
- [38] George Toderici et al. Workshop and challenge on learned image compression (clic2020). In *CVPR*, 2020. 6
- [39] Gregory K Wallace. The jpeg still picture compression standard. *Communications of the ACM*, 34(4), 1991. 1
- [40] Guo-Hua Wang, Jiahao Li, Bin Li, and Yan Lu. Evc: Towards real-time neural image compression with mask decay. In *ICLR*, 2023. 3
- [41] Yueqi Xie et al. Enhanced invertible encoding for learned image compression. In *ACM MM*, 2021. 2
- [42] Lei Zhou et al. End-to-end optimized image compression with attention mechanism. In *CVPRW*, 2019. 2
- [43] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *CVPR*, 2022. 1, 2, 3, 6, 7