

CALLIC: Content Adaptive Learning for Lossless Image Compression

Daxin Li^{1*}, Yuanchao Bai^{1*}, Kai Wang¹, Junjun Jiang¹, Xianming Liu^{1†}, Wen Gao²

¹Faculty of Computing, Harbin Institute of Technology, Harbin

²Department of Computer Science and Technology, Peking University, Beijing
hahalidaxin@stu.hit.edu.cn, yuanchao.bai@hit.edu.cn, cswk@stu.hit.edu.cn, jiangjunjun@hit.edu.cn,
csxm@hit.edu.cn, wgao@pku.edu.cn

Abstract

Learned lossless image compression has achieved significant advancements in recent years. However, existing methods often rely on training amortized generative models on massive datasets, resulting in sub-optimal probability distribution estimation for specific testing images during encoding process. To address this challenge, we explore the connection between the Minimum Description Length (MDL) principle and Parameter-Efficient Transfer Learning (PETL), leading to the development of a novel content-adaptive approach for learned lossless image compression, dubbed CALLIC. Specifically, we first propose a content-aware autoregressive self-attention mechanism by leveraging convolutional gating operations, termed Masked Gated ConvFormer (MGCF), and pretrain MGCF on training dataset. Cache then Crop Inference (CCI) is proposed to accelerate the coding process. During encoding, we decompose pre-trained layers, including depth-wise convolutions, using low-rank matrices and then adapt the incremental weights on testing image by Rate-guided Progressive Fine-Tuning (RPFT). RPFT fine-tunes with gradually increasing patches that are sorted in descending order by estimated entropy, optimizing learning process and reducing adaptation time. Extensive experiments across diverse datasets demonstrate that CALLIC sets a new state-of-the-art (SOTA) for learned lossless image compression.

Introduction

In the realm of digital imaging, lossless image compression stands as a pivotal technology, playing a crucial role in efficiently managing and transmitting vast amounts of image data without sacrificing quality. This form of compression is essential in fields where precision and detail are paramount, such as medical imaging, remote sensing, and digital archiving. Furthermore, the pursuit of higher compression ratios in lossless encoding is of paramount importance. Achieving greater compression rates without loss of information not only conserves valuable storage space and reduces transmission time but also addresses the growing demand for handling larger and more complex images, such as those in high-resolution photography and 3D medical scans. Thus, advancements in lossless image compression techniques are

instrumental in driving forward the capabilities and efficiency of various data-intensive sectors.

Traditional image compression methods like PNG, JPEG-LS (Weinberger, Seroussi, and Sapiro 2000), and JPEG-XL (Alakuijala et al. 2019) use handcrafted algorithms to exploit the statistical properties of images. However, these methods often struggle to capture the complex and diverse distributions in raw images, limiting their compression performance. In recent years, deep learning-based techniques have revolutionized neural compression (Ballé, Laparra, and Simoncelli 2017; Cheng et al. 2020; He et al. 2022; Bai et al. 2022; Li et al. 2024b,a), and lossless compression methods (Mentzer et al. 2019; Kingma, Abbeel, and Ho 2019; Townsend et al. 2019; Hoogeboom et al. 2019; Ho, Lohn, and Abbeel 2019; Zhang et al. 2021b,a; Ryder et al. 2022; Rhee et al. 2022; Wang et al. 2023; Bai et al. 2024; Zhang et al. 2024; Wang et al. 2024) have achieved state-of-the-art (SOTA) results by learning the intricate distribution of raw images. These approaches use likelihood-based generative models, such as autoregressive models, flow models, and variational autoencoders (VAEs), to predict probabilities for entropy coding, converting data into compact bit-streams. Despite their success, these models often rely on training with large datasets, which can lead to sub-optimal probability estimates for specific testing images. Each image’s unique details, such as textures, lighting, or structures, pose challenges that need to be addressed to improve compression performance.

To address the aforementioned challenge, we propose CALLIC, a content-adaptive learned lossless image compression method that leverages the connection between the Minimum Description Length (MDL) principle and Parameter-Efficient Transfer Learning (PETL). Specifically, we introduce an effective and lightweight pre-trained autoregressive model with a masked gating mechanism to simulate content-aware self-attention, termed the Masked Gated ConvFormer (MGCF). The pre-trained MGCF identifies patterns common across various images, known in advance to both the encoder and decoder. To make MGCF a practical codec, we propose Cache then Crop Inference (CCI), which caches intermediate activations and convolves only on the cropped patches surrounding the activated positions during coding. During encoding, we enhance the pre-trained MGCF by incorporating additional parameter-efficient weights, fine-

*Equal contribution.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tuning it to capture the unique characteristics of each individual image and encoding this as a model prompt into a bitstream. We begin by decomposing the linear and depth-wise convolution layers in MGCF using low-rank matrices. Then, we introduce Rate-guided Progressive Fine-Tuning (RPFT), which fine-tunes low-rank weights with progressively increasing patches that are sorted in descending order by estimated entropy for more efficient adaptation. The MDL principle is employed to jointly optimize the bitrates for incremental weights and the compressed image. Finally, the adapted MGCF is used to encode the image, which is then sent to the decoder along with the compressed incremental weights.

Our contributions are encapsulated as follows:

- By bridging the MDL principle with PETL, we propose CALLIC, a novel content-adaptive lossless image compression method, effectively addressing the amortization gap between training data and testing images.
- We introduce MGCF, a content-aware autoregressive self-attention mechanism with convolutional gating operations. CCI is proposed to speed up the coding process.
- We decompose the pre-trained weights, including depth-wise convolutions, using low-rank matrices and adapt the model to testing image with RPFT, efficiently fine-tuning weights with progressively increasing patches sorted in descending order by estimated entropy.
- Extensive experimental results demonstrate that CALLIC achieves a new SOTA performance in lossless image compression, significantly outperforming existing learned methods across various datasets.

Related Work

Learned Lossless Image Compression

Lossless image compression converts images into the fewest possible bits while maintaining lossless reconstruction. Based on different types of likelihood models, these methods are categorized into autoregressive models, flow models, and VAE models.

Autoregressive models decompose data probability into conditional distributions via the chain of probabilities, using learnable models for these distributions. Oord *et al.* (Van Den Oord, Kalchbrenner, and Kavukcuoglu 2016) introduced PixelCNN, estimating the conditional distribution with masked convolution. Bai *et al.* (Bai *et al.* 2021) developed an end-to-end framework using lossy compression for prediction and a serial autoregressive model for residual image compression. They later advanced this with a parallel model in (Bai *et al.* 2024) for faster coding.

For *flow models*, Ho *et al.* (Ho, Lohn, and Abbeel 2019) presented a local bit-back coding scheme to modify a continuous normalizing flow model for lossless compression. In contrast, discrete normalizing flows handle discrete data and employ invertible transformations for integer latent space, which is encoded with a simple prior distribution. Hoogeboom *et al.* (Townsend *et al.* 2019) developed an integer flow (IDF) model which learned invertible transformations for compression. Zhang *et al.* (Zhang *et al.* 2021b) introduced an invertible volume-preserving flow (iVPPF) model,

and later, the iFlow model in (Zhang *et al.* 2021a), featuring modular scale transforms and uniform base conversion systems.

VAE models are divided into deterministic and stochastic posterior sampling methods. Deterministic methods, like L3C (Mentzer *et al.* 2019), impose constraints on the latent codes’ posterior distribution. Stochastic methods, exemplified by BB-ANS (Townsend, Bird, and Barber 2019), employ stochastic latent code sampling, using the evidence lower bound (ELBO) for training. Kingma *et al.* (Kingma, Abbeel, and Ho 2019) later proposed the Bit-Swap scheme based on hierarchical VAE models. Townsend *et al.* (Townsend *et al.* 2019) introduced HiLLoC, another hierarchical latent lossless compression method.

Parameter-Efficient Transfer Learning

Presently, PETL strategies can be categorized into three predominant groups.

Adapter Modules, introduced by Houlsby *et al.* (Houlsby *et al.* 2019), are small, trainable networks added to existing models. These modules adjust the model’s internal representations to new tasks, keeping most original parameters unchanged.

Prompt Tuning uses task-specific prompts added to input data. This approach guides the pre-trained model to produce task-relevant outputs without changing its parameters (Li and Liang 2021; Lester, Al-Rfou, and Constant 2021).

Low-Rank Adaptation (Hu *et al.* 2021) modifies a pre-trained network’s weights with low-rank matrices, enabling task adaptation with few additional parameters.

There have been some works that combine learned lossy image compression with PETL, while they insert adapters into transform network (Tsubota, Akutsu, and Aizawa 2023; Shen, Yue, and Yang 2023), mainly to solve performance degradation for out-of-domain images in transform coding framework.

Minimum Description Length

The MDL principle suggests that the most effective model for data description is the one that maximizes the data compression efficiency (Grünwald 2005; Barron, Rissanen, and Yu 1998). A notable variant of this principle is the crude MDL or two-stage MDL principle, which involves a two-pass coding process. First, the process entails selecting and encoding the optimal model into a bitstream. Second, data is compressed using the chosen model. Given random variable \mathbf{x} ’s drawn from a probability distribution $p(\mathbf{x})$, the probability density function of $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is defined as $p(\mathbf{x}) = p(x_1, x_2, \dots, x_n)$. The model (a probability density function) q is selected from a countable collection \mathcal{Q} by minimizing the total codelength:

$$\mathcal{L}_{2\text{-stage}}(\mathbf{x}) = \min_{q \in \mathcal{Q}} L(q) + L_q(\mathbf{x}). \quad (1)$$

In this framework, $L(q) = \log \frac{1}{w(q)}$ denotes the codelength for the model (reflecting model complexity) and $w(q)$ represents a prior probability mass function on $q \in \mathcal{Q}$. $L_q(\mathbf{x}) = \log \frac{1}{q(\mathbf{x})}$ quantifies the codelength for data \mathbf{x} under model q .

In learned lossless image compression, the image \mathbf{x} is decomposed into a sequence of variables, enabling its representation as a set of non-independent data points x_1, x_2, \dots, x_n . Each variable's unknown real distribution is estimated with a generative model, and entropy coding tools are used to encode each variable based on the estimated distribution. The generative model is trained by minimizing the negative log-likelihood on a training dataset. During encoding, both the encoder and decoder are aware of the trained model (ignoring $L(q)$), and only the bitstream of image (of length $L_q(\mathbf{x})$) needs to be transmitted. Nevertheless, the generative model with learned amortized parameters on the whole training dataset usually results in sub-optimal probability model q for specific testing images. Each testing image manifests its own unique details, such as textures, lighting or structures, stemming from diverse shooting targets and imaging devices. This leaves room for improvement by closing the gap between training and testing images.

To address this challenge, this paper incorporates the two-stage MDL principle into learned lossless image compression. Initially, we propose to fine-tune the pre-trained generative model to a single image and encode the fine-tuned model into a bitstream. However, fine-tuning and encoding the entire generative model is extremely inefficient in terms of both computation and storage, necessitating a more efficient and scalable approach for the image-specific adaptation. Inspired by recent PETL, we fix the pre-trained generative model but add a small, trainable incremental weights to adjust the model's internal representations to each testing image. Instead of encoding the entire fine-tuned model, the pre-trained generative model is known by both the encoder and decoder, and only the incremental weights need to be encoded to a bitstream. We then compress the testing image using the adapted generative model, and transmit the compressed image together with the incremental weights to the decoder. In order to maximize the coding efficiency, we leverage the MDL principle to jointly optimize the bitrates for both the incremental weights and the compressed image:

$$\min_{\phi} L(\phi) + \log \frac{1}{q(\mathbf{x}; \theta, \phi)}, \quad (2)$$

where $L(\phi) = \log \frac{1}{w(\phi)}$ indicates the codelength required to encode the incremental weights ϕ , assumed to follow a prior distribution $w(\phi)$. The $\log \frac{1}{q(\mathbf{x}; \theta, \phi)}$ quantifies the codelength for current image \mathbf{x} under the adapted model $q(\mathbf{x}; \theta, \phi)$ incorporating both the pre-trained parameter θ and the incremental weights ϕ . The expected length of the two-stage code in relation to the actual data distribution p is:

$$L(\phi) + \mathbb{E}_p \left[\log \frac{1}{q(\mathbf{x}; \theta, \phi)} \right] = L(\phi) + D_{KL}(p||q) + H(p), \quad (3)$$

where $D_{KL}(p||q)$ represents the Kullback-Leibler divergence between p and q , and $H(p)$ is the entropy of p .

A bound $R_n(p)$ on the expected redundancy between the optimization target and the real entropy is derived as follows:

low:

$$\begin{aligned} & E_p \left[\min_{\phi} \{L(\phi) + D_{KL}(p||q)\} \right] \\ & \leq \min_{\phi} E_p [L(\phi) + D_{KL}(p||q)] \\ & = \min_{\phi} \{L(\phi) + D_{KL}(p||q)\} = R_n(p). \end{aligned} \quad (4)$$

The inequality holds using Jensen's inequality since minimum is a concave function. The bound $R_n(p)$ reflects the trade-off between estimation precision and model complexity in two-stage coding, as a more precise model requires more bits $L(\phi)$ to describe the model, but exhibits a smaller divergence $D_{KL}(p||q)$ with the true probability distribution.

This theoretical analysis paves the way for the following design of our content-adaptive learned lossless image compression, CALLIC. We approach content-adaptive lossless image compression from two aspects: **1) designing an effective and lightweight pre-trained generative model for content-aware structures** **2) developing an efficient fine-tuning method to adjust the pre-trained model to each testing image, optimizing the MDL in (2)**. The subsequent sections will explore these two aspects.

CALLIC Method

Pre-trained Content Adaptive Model

Masked Convolutional Gating Mechanism Transformers have empowered many fields (Vaswani 2017; Dosovitskiy 2020; Liang et al. 2024) due to their content adaptivity and global modeling capabilities but are limited by their quadratic time complexity. To effectively adapt to intricate image structures during pre-training, we propose a novel content-aware masked convolutional gating (MCG) mechanism to simplify attention mechanism, as depicted in Fig. 1(a). As suggested by Zhang *et al.* (Zhang, Zhang, and McDonagh 2021), our method initially concentrates on local features instead of global contexts to enhance efficiency and generalization. We employ masked depth-wise convolution with a $k \times k$ kernel to aggregate local contextual information. Furthermore, we simulate self-attention using gating mechanism, by modulating the value matrix through a Hadamard product with the output of the convolutional operation. The MCG mechanism can be expressed as:

$$\begin{aligned} \mathbf{A}_M &= \text{DWConv}_{k \times k}(\mathbf{W}_A \mathbf{X}, \mathbf{M}), \\ \mathbf{V} &= \mathbf{W}_V \mathbf{X}, \\ \text{MCG}(\mathbf{X}) &= \sigma(\mathbf{A}_M) \odot \mathbf{V}, \end{aligned} \quad (5)$$

where $\mathbf{W}_A, \mathbf{W}_V$ are two linear projections, σ is a non-linear activation, using swish function, \odot represents the Hadamard product, and \mathbf{M} is the convolutional mask that restricts information to the decoded contextual scope.

In MCG, the \mathbf{A}_M aggregates contextual information from local regions, capturing the characteristics of local structures and details. The resulting \mathbf{A}_M is then converted into gating values through a non-linear activation function. The value matrix \mathbf{V} is a linear projection of the input, retaining the projected spatial information at each position. Content

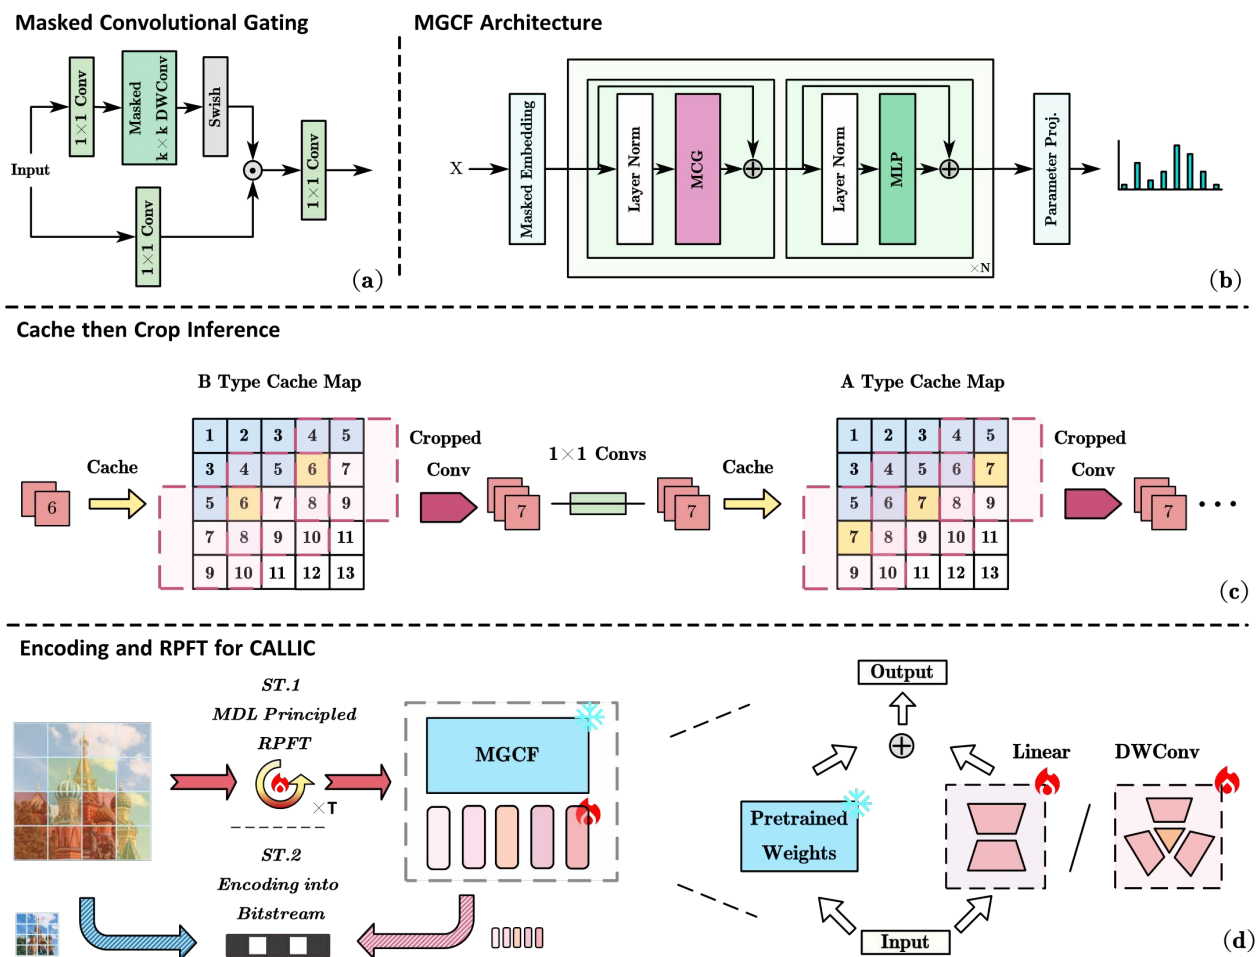


Figure 1: Overview of the proposed architectures and mechanisms in CALLIC: We first simplify the attention mechanism using a content-adaptive convolutional gating mechanism (a). Building on this, we introduce the Masked Gated ConvFormer (MGCF) (b). To accelerate the coding process, we propose Cache then Crop Inference (CCI) (c), which caches activations before masked convolution layers and then performs convolution on cropped features at coding positions. For instance-level content adaptivity, we fine-tune the pre-trained model on the test image. This involves adding learnable parameters through low-rank decomposition of MGCF layers (Right), fine-tuning these additional parameters with the proposed Rate-guided Progressive Fine-tuning (RPFT), and encoding them into the bitstream along with the image (Left).

awareness is primarily achieved through the modulation operation, which regulates the flow of information within the value matrix \mathbf{V} using locally adaptive gating values $\sigma(\mathbf{A}_M)$. This mechanism enables the network to learn efficient adaptation to the content on training dataset, enhancing the pre-trained model’s ability to respond to varying contexts.

Masked Gated ConvFormer Building on the MCG, we introduce an autoregressive architecture called the Masked Gated ConvFormer (MGCF) inspired by the MetaFormer architecture (Yu et al. 2022), as illustrated in Fig. 1(b). In MGCF, the input image first passes through an embedding layer, implemented via a 3×3 masked convolution. This is followed by N MGCF blocks, each consisting of a masked convolutional gated block, which includes MCG, and an MLP block, featuring two linear layers with GELU activa-

tion. The final output is transformed into parameters using a 1×1 convolutional layer. We model with a discrete logistic mixture model, as shown in (Salimans et al. 2016).

Cache then Crop Inference To make MGCF a practical codec, we introduce a strategy to accelerate our coding process, named Cache then Crop Inference (CCI), illustrated in Fig. 1 (c). CCI minimizes computations for decoded elements by caching their activations in masked convolutional layers. It then applies zero-padding convolution only to cropped windows surrounding the activated positions.

Specifically, image is first divided into patches and the patches are encoded in parallel. For each patch, pixels are initially grouped as $\mathbf{x} = \{\mathbf{x}_{G_1}, \dots, \mathbf{x}_{G_g}\}$ according to a parallel scan order, as in Fig. 1 (c). This parallel scan takes inspiration from DLPR (Bai et al. 2024) and entails $3P - 2$

autoregressive steps, where P is the patch size. Encoding is conducted group by group. At step i , the network processes activations from previous groups $x_{G_{\leq i-1}}$ and predicts the distribution for the current group x_{G_i} . In MGCF, types A and B masked convolutions are employed. For type B convolution, current group positions are masked out. Conversely, for type A convolution, positions for the current group and all previous groups are included in the receptive field. For type B convolution, which is used at the top as an embedding layer, pixels from the previous group $x_{G_{i-1}}$ are fed into the network and cached. The cached data then undergoes type B masked convolution at positions corresponding to the current group, aggregating contextual information. Instead of convolving over all positions, we crop windows around the current positions and perform zero-padding convolution on these windows in parallel. In the embedding layer, activations for the current group x_{G_i} are obtained and fed into the subsequent network layers. For type A convolution, used in subsequent masked convolution gating blocks, elements at positions for the current group are considered. We store activations at current group positions into cache map and perform cropped convolution to collect contextual information for these positions. The 1×1 convolution layer in the rest of the network transfers information solely across the channel dimension, and no caching is necessary.

MDL-Principled RPFT

Low-rank Decomposition for Weights We adapt the pre-trained weights using PETL to capture unique characteristics of each individual image. Based on the hypothesis that changes in weight matrices exhibit low-rank characteristics, LoRA proposes an incremental update of pre-trained weights in linear layers using the product of two low-rank matrices. For a pre-trained weight matrix $\mathbf{W} \in \mathbb{R}^{m \times n}$, the weight update is modeled as $\Delta \mathbf{W} \in \mathbb{R}^{m \times n}$. The fine-tuned weight matrix \mathbf{W}' is then given by:

$$\mathbf{W}' = \mathbf{W} + \Delta \mathbf{W} = \mathbf{W} + \mathbf{A}\mathbf{B}, \quad (6)$$

where $\mathbf{A} \in \mathbb{R}^{m \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times n}$, with $r \ll \min(n, m)$.

Additionally, aiming to fine-tuning MGCF, we extend this concept from two-dimensional linear layers to three-dimensional masked depth-wise convolutions. For a masked depth-wise convolution, the pre-trained kernel \mathbf{W}_{mc} is represented in $\mathbb{R}^{m \times 1 \times k \times k}$, where k denotes the kernel size. Utilizing Tucker decomposition (Tucker 1966), we decompose the pre-trained weights into four components: a core identity tensor $\mathbf{I} \in \mathbb{R}^{r_1 \times 1 \times r_2 \times r_3}$ and three mode matrices $\mathbf{A} \in \mathbb{R}^{m \times r_1}$, $\mathbf{C} \in \mathbb{R}^{k \times r_2}$, and $\mathbf{D} \in \mathbb{R}^{k \times r_3}$. The incremental update is then defined as $\Delta \mathbf{W}_{mc} = \mathbf{I} \times_1 \mathbf{A} \times_3 \mathbf{C} \times_4 \mathbf{D}$, where \times_n denotes the mode- n product, which is the multiplication of a tensor by a matrix along the n -th mode. The fine-tuned weight matrix is computed as:

$$\mathbf{W}'_{mc} = \mathbf{M} \odot (\mathbf{W}_{mc} + \Delta \mathbf{W}_{mc}). \quad (7)$$

Notably, the incremental weights for both linear and convolutional layers can be merged with the pre-trained weights, resulting in no additional computation time during inference. We incorporate low-rank decomposition into the masked convolutional gating blocks and MLP. Specifically,

we add incremental weights to the weight matrices \mathbf{W}_A and \mathbf{W}_V within the convolutional gating mechanism, as well as to the first linear layer \mathbf{W}_{up} in the MLP.

Rate-guided Progressive Fine-Tuning To further accelerate the adaptation process, we propose a Rate-guided Progressive Fine-Tuning (RPFT) approach. This method is based on the assumption that it is unnecessary to use all image patches for model fine-tuning throughout the entire process. Two key considerations support this: 1) There is inherent redundancy within patches due to content similarities among them; 2) The distributions of different patches vary in their complexity, which affects the model’s learning process differently. These considerations guide our two design aspects: first, progressively fine-tuning on an increasing number of patches, and second, determining the focus of fine-tuning—patches, which contain more informative structures or details, should be trained with more steps.

Consequently, we select the estimated bitrate as an indicator of the informational content of each patch and employ a strategy of gradually increasing the number of training patches, which are sorted in descending order by pixel rate. We begin by estimating the bitrates for all image patches, organizing them from highest to lowest, and progressively incorporate them into our training samples for fine-tuning. Initially, we select only $b\%$ of the patches as training samples. As training progresses, we systematically include more patches, eventually utilizing the entire set in the final $d\%$ of the training steps. The tuning on full set ensures that our training and testing objectives are well-aligned. We use the following function to describe how the fine-tuning patches ratio $F(t)$ is increased:

$$t' = \frac{t}{T \cdot (1 - d)},$$

$$s(x) = \begin{cases} 0, & \text{if } x < 0, \\ 1, & \text{if } x > 1, \\ x^2(3 - 2x), & \text{if } 0 \leq x \leq 1, \end{cases} \quad (8)$$

$$F(t) = b + (1 - b) \cdot [s(t')]^e,$$

where t represents the current step, T denotes the total number of training steps, and e is a hyperparameter that controls the growth rate of the training sample ratio.

The progression of $F(t)$ is illustrated in Fig. 2. By modifying b , d , and e , we can adjust the number of patches involved in the fine-tuning process. Specifically, as both b and d increases and e decrease, the number of patches participating in training will increase. Through RPFT, we intensify our focus on learning the distribution of patches with higher informational content while effectively minimizing the number of patches involved in the training process. This strategy not only optimizes the learning process but also effectively reduces the time required for adaptation.

MDL-Principled Optimization The complete lossless image compression pipeline is detailed in Fig. 1 (d). Before encoding the test image into a bitstream, we first fine-tune the incremental weights for the current image as a model prompt while keeping the pre-trained MGCF parameters

Codec	Kodak	RS19	Histo24	DIV2K	CLIC.p	Params.
PNG	4.35 +71.3%	3.90 +124.1%	3.79 +38.3%	4.23 +72.0%	3.93 +70.9%	—
JPEG2000	3.19 +25.6%	2.57 +47.7%	3.36 +22.6%	3.12 +30.5%	2.93 +32.0%	—
FLIF	2.90 +14.6%	2.18 +25.3%	3.23 +17.9%	2.91 +18.3%	2.72 +18.3%	—
JPEG-XL	2.87 +13.4%	2.02 +16.1%	3.07 +12.0%	2.79 +13.4%	2.63 +14.3%	—
L3C (CVPR’19)	3.26 +28.3%	2.66 +53.5%	3.53 +28.8%	3.09 +25.6%	2.94 +27.8%	5M
RC (CVPR’20)	3.38 +32.8%	2.19 +25.9%	3.33 +21.5%	3.08 +25.2%	2.93 +27.4%	6.9M
iVPF (CVPR’21)	—	—	—	2.68 +8.9%	2.54 +10.4%	59.5M
LC-FDNet (CVPR’22)	2.98 +17.3%	2.15 +23.6%	3.07 +12.0%	2.72 +10.6%	2.63 +14.3%	23.7M
DLPR (TPAMI’24)	2.86 +14.3%	2.01 +15.5%	2.96 +8.0%	2.55 +3.7%	2.38 +3.5%	22.2M
ArIB-BPS (CVPR’24)	2.78 +9.8%	<u>1.92</u> +10.3%	2.92 +6.6%	2.55 +3.7%	2.42 +5.2%	146.6M
MGCF (Ours)	<u>2.77</u> +9.1%	1.94 +11.5%	<u>2.88</u> +5.1%	<u>2.49</u> +1.2%	<u>2.33</u> +1.3%	575K
CALLIC (Ours)	2.54	1.74	2.74	2.46	2.30	575K

Table 1: Compression performance of the proposed methods and other codecs in terms of bits per sub-pixel (bpsp). We show the difference in percentage to our approach, using green. The best is highlighted in bold, and the second is highlighted using underline. CALLIC introduces 25K additional trainable parameters compared to MGCF, which can be merged with the pre-trained weights during inference.




Figure 2: Illustration of incremental curves, $T = 50$. The shape of the curve is controlled by parameters b, d, e .

fixed. During encoding, these incremental parameters are quantized into discrete values before being incorporated into the bitstream. Given the relatively narrow numerical range of incremental weights, we employ a smaller quantization step size, $w < 1$. For the differentiable estimation of quantization during optimization, we utilize a mixed quantization approach, as detailed in (Tsubota, Akutsu, and Aizawa 2023). Let ϕ represent the incremental weights, and θ the pre-trained MGCF weights. This process quantizes incremental weights using a straight-through estimator (STE) (Theis et al. 2016) for network inference, where sg denotes stop gradient operation: $\hat{\phi} = \text{sg}(\lfloor \frac{\phi}{w} \rfloor * w - \phi) + \phi$. Additionally, uniform noise is added to the weights for entropy estimation: $\tilde{\phi} = \phi + U(-\frac{w}{2}, \frac{w}{2})$. We leverage the MDL principle to jointly optimize the bitrates for both the incremental weights and the image pixels:

$$\mathcal{L} = -\log p_s(\tilde{\phi}) + \sum_i -\log q(x_{G_i} | x_{G_{<i}}; \theta, \hat{\phi}), \quad (9)$$

Codec	Bpsp	Params.	Enc. Time	Dec. Time
LC-FDNet	2.98	23.7M	1.6s	1.6s
DLPR	2.86	22.2M	1.5s	2.0s
ArIB-BPS	2.78	146.6M	7.1s	7.0s
MGCF (Ours)	2.77	575K	1.4s	1.6s
-w/o CCI	—	—	15.2s	15.4s
CALLIC (Ours)	2.54	575K	9.7s	1.7s
-w/o RPFT	2.54	—	12.9s	—

Table 2: Coding speed analysis on Kodak. Bpsp denotes bits per sub-pixel. Our pre-trained model delivers fast coding speeds, leveraging the proposed CCI mechanism. Fine-tuning improves compression but increases encoding time, which is effectively reduced with the RPFT method.

where $p_s(\tilde{\phi})$ denotes the probability distribution of the incremental weights, modeled using a static logistic distribution with zero mean and a constant scale $s \in \mathbb{R}$.

After optimization, we explicitly compute and merge the weights according to Eq. 6 and Eq. 7, and then proceed with network inference as usual, ensuring no additional time is required during this stage. Finally, the image is encoded into a bitstream using the final model. The total bitrate for compressing the image includes the bitrates for the quantized incremental weights and the image pixels. This bitstream is then transmitted to the decoder, where the incremental weights are first decoded and used for image decoding.

Experiments

Experimental Settings

Pre-trained Settings For pretraining, we make a collection of 3450 images from publicly available high-resolution datasets: DIV2K (Agustsson and Timofte 2017), with 800 training images, and Flickr2K (Lim et al. 2017), with 2650

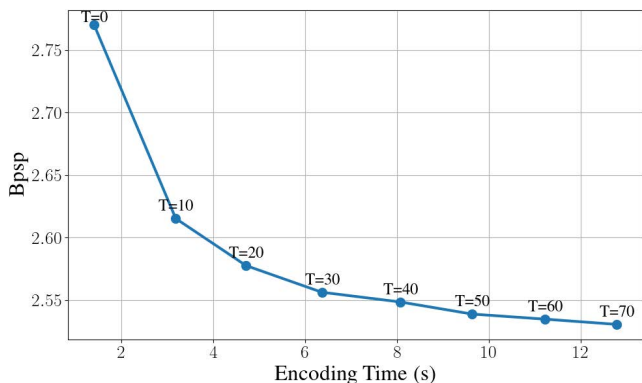


Figure 3: Steps and performance trade-off for CALLIC. Faster encoding can be achieved with minimal performance trade-offs.

training images. We crop these images into non-overlapping patches of size 64×64 , resulting in a training set of 612, 806 images. We train MGCF for 2M steps using the Adam optimizer with minibatches of size 32 and learning rate $5e - 4$.

Adaptation Settings The maximum number of optimization steps is set to $T = 50$, with a learning rate $1e - 2$. A relatively small scale for the prior weight distribution is chosen, set as $s = 0.05$. Quantization width is set to $w = 0.05$. For RPFT, we choose $b = 0.2$, $d = 0.1$, $e = 1$ as default.

Evaluation Settings We selected five high-resolution datasets: Kodak (Kodak 1993), WHU-RS19 validation (Xia et al. 2010), Histo24 (Bai et al. 2024), DIV2K (Agustsson and Timofte 2017) and CLIC2020 professional validation (CLIC.p) (Toderici et al. 2020) as our evaluation datasets. The Kodak, DIV2K, CLIC.p, and CLIC.m are natural images datasets. Histo24 is a dataset that includes 24768×512 histological images proposed by (Bai et al. 2024). We center-cropped 190 satellite images to 576×576 from the WHU-RS19 validation set, which were exported from Google Earth, to form our validation dataset, RS19. Our method is compared with traditional lossless codecs, including JPEG2000 (Skodras, Christopoulos, and Ebrahimi 2001), FLIF (Sneyers and Wuille 2016), and JPEG-XL (Alakuijala et al. 2019), and *open-sourced* learned methods like L3C (Mentzer et al. 2019), RC (Mentzer, Gool, and Tschannen 2020), iVPF (Zhang et al. 2021b), LC-FDNet (Rhee et al. 2022), DLPR (Bai et al. 2024) and ArIB-BPS (Zhang et al. 2024).

Experimental Results

Model Performance The compression performance results are summarized in Tab. 1. Our pre-trained model, MGCF, demonstrates superior performance compared to previous methods. Specifically, MGCF outperforms all other methods in four out of five datasets and performs slightly worse than ArIB-BPS on the RS19 dataset. By integrating RPFT into MGCF during encoding to create CALLIC, we achieve further improvements. CALLIC consistently surpasses all other methods across all datasets. For datasets

Depth	Dim.	Kernel	Bpsp	Enc./Dec. Time
1	128	7	2.84	1.0/1.1s
3	128	7	2.77	1.4/1.6s
5	128	7	2.73	2.0/2.1s
3	64	7	2.82	1.4/1.5s
3	192	7	2.74	1.8/1.8s
3	256	7	2.72	2.1/2.1s
3	128	5	2.82	1.4/1.5s
3	128	9	2.77	1.8/1.9s
3	128	11	2.76	2.2/2.3s

Table 3: Analysis about network architecture. We perform an ablation study on network depth, model channels, and the kernel size of masked convolutions.

Method	Bpsp	Params.	Enc. Time	Dec. Time
MGCF	2.77	575K	1.4s	1.6s
NAT	2.78	767K	7.6s	7.8s

Table 4: Comparison with neighborhood attention mechanism on Kodak. Masked convolutional gating mechanism performs better than neighborhood attention mechanism.

with gaps between the training set, such as Kodak, RS19, and Histo24, CALLIC achieves significant improvements of 9.1%, 11.5%, and 5.1% over MGCF, respectively. For other datasets, CALLIC also shows better, albeit less significant, improvements. This is attributed to the effectiveness of MGCF, which has learned similar distributions on the training dataset, making further gains challenging. We highlight that MGCF and CALLIC are lightweight models, with MGCF having only 575K parameters and CALLIC adding 25K *mergeable* weights. This is significantly smaller than other methods. Compared to ArIB-BPS, MGCF achieves better or comparable results with only 4% of the parameters. *It is worth noting that the lightweight characteristic is what makes further adaptation possible and efficient.*

Runtime We present the encoding speeds for recent learned methods on the Kodak dataset in Tab. 2. MGCF achieves competitive encoding and decoding times of 1.4 and 1.6 seconds, respectively, which are faster or comparable to other advanced methods. When the CCI is removed, encoding and decoding times increase to 15.2 and 15.4 seconds respectively. Content adaptive fine-tuning enhances performance while increasing the encoding time. The introduction of RPFT reduces the encoding time from 12.9 to 9.7 seconds while maintaining same performance, demonstrating its effectiveness. A trade-off between speed and performance can be precisely adjusted by modifying the training steps, as illustrated in Fig 3. For instance, with $T = 10$, the encoding time is reduced to 3.2 seconds, with a bpsp of 2.62.

Ablation Studies

Network Architecture of MGCF The Tab. 3 presents an analysis of the impact of various hyperparameters on the per-

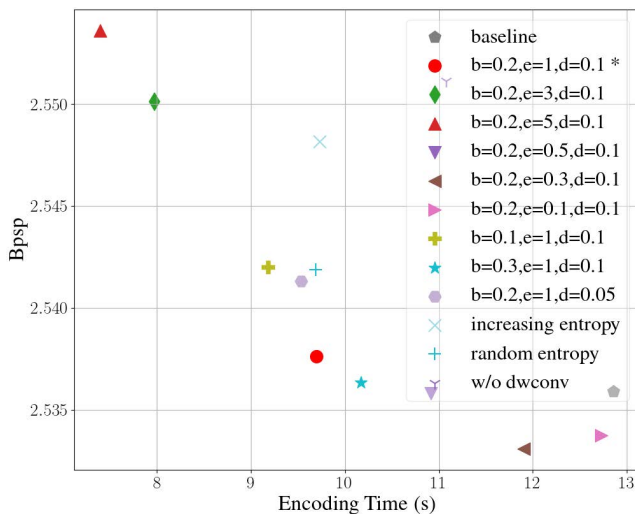


Figure 4: Different configurations for RPFT on Kodak. We ablate different increasing curves and sorting orders in this experiment.

formance of MGCF. The study explores changes in model complexity by varying the depth, dimension, and kernel size. The results demonstrate that increasing model complexity generally enhances performance. Notably, while increasing the kernel size beyond 7 shows diminishing returns, expanding the dimension and depth consistently leads to performance gains. The final model, with a depth of 3, a dimension of 128, and a kernel size of 7, was chosen to balance the trade-off between improved performance and competitive encoding/decoding complexity.

Convolutional Gating vs. Neighborhood Attention To compare the effectiveness of original self-attention and convolutional gating mechanisms, we adopt Neighborhood Attention Transformer (Hassani et al. 2023), which uses local attention on feature maps within a sliding window. In our experiments, we replaced masked convolutional gating with masked neighborhood attention, called “NAT”. The results, shown in Tab. 4, indicate that convolutional gating not only performs better but also works faster, mainly because it uses only convolution operations.

Different Configurations for RPFT The experimental data presented in Fig. 4 explores the effects of different configurations for RPFT. The baseline method, without RPFT, results in an encoding time of 12.9 seconds and a bpsp of 2.535. When employing RPFT with default configuration ($b=0.2, e=1, d=0.1$), the encoding time significantly decreases to 9.7 seconds while maintaining a comparable bpsp of 2.537. We found that the default configuration demonstrates robust performance across different image types, making it a reliable choice for various datasets.

Increasing the participation of image patches, characterized by higher b and d values and lower e values, generally leads to improved performance, albeit at the cost of longer encoding times. Notably, the configurations ($b=0.2, e=0.3, d=0.1$) and ($b=0.2, e=0.1, d=0.1$) outperform the baseline in

both encoding time and compression performance, underscoring the effectiveness of RPFT.

The results in Fig. 4 also indicate that focusing on higher rate patches (default) achieves better performance compared to starting from lower rate patches (“increasing entropy”) or employing random sampling (“random entropy”) strategies.

The absence of adaptation for depth-wise convolutions, as seen in the “w/o dwconv” condition, results in a higher bpsp of 2.551. This highlights the crucial role of depth-wise convolution adaptation in enhancing overall performance.

Conclusion

In this work, we propose CALLIC by leveraging the connection between the MDL principle and PETL. We introduce MGCF, an efficient pre-trained autoregressive model, combined with the CCI acceleration strategy to enhance its practicality as a codec. For efficient fine-tuning on test images, we decompose layers, including depth-wise convolutions, using low-rank matrices and introduce RPFT. This approach progressively fine-tunes additional parameters using gradually increasing patches prioritized by estimated entropy. Extensive experiments demonstrate that CALLIC establishes a new SOTA in learned lossless image compression.

Acknowledgments

This work was supported in part by National Key Research and Development Program of China under Grant 2022YFF1202104, in part by the National Key R&D Program of China under Grant 2023YFC2509100, in part by National Natural Science Foundation of China under Grants 62301188, 92270116 and U23B2009, in part by the Strategic Research, and Consulting Project by the Chinese Academy of Engineering under Grant 2023.XY-39, in part by China Postdoctoral Science Foundation under Grant 2022M710958, and in part by Heilongjiang Postdoctoral Science Foundation under Grant LBH-Z22156.

References

- Agustsson, E.; and Timofte, R. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Alakuijala, J.; Van Asseldonk, R.; Boukourt, S.; Bruse, M.; Comsa, I.-M.; Firsching, M.; Fischbacher, T.; Kliuchnikov, E.; Gomez, S.; Obryk, R.; et al. 2019. JPEG XL next-generation image compression architecture and coding tools. In *Applications of digital image processing XLII*, volume 11137, 112–124. SPIE.
- Bai, Y.; Liu, X.; Wang, K.; Ji, X.; Wu, X.; and Gao, W. 2024. Deep lossy plus residual coding for lossless and near-lossless image compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bai, Y.; Liu, X.; Zuo, W.; Wang, Y.; and Ji, X. 2021. Learning scalable l_∞ -constrained near-lossless image compression via joint lossy image and residual compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11946–11955.

- Bai, Y.; Yang, X.; Liu, X.; Jiang, J.; Wang, Y.; Ji, X.; and Gao, W. 2022. Towards end-to-end image compression and analysis with transformers. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 104–112.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2017. End-to-end optimized image compression. In *5th International Conference on Learning Representations, ICLR 2017*.
- Barron, A.; Rissanen, J.; and Yu, B. 1998. The minimum description length principle in coding and modeling. *IEEE transactions on information theory*, 44(6): 2743–2760.
- Cheng, Z.; Sun, H.; Takeuchi, M.; and Katto, J. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7939–7948.
- Dosovitskiy, A. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Grünwald, P. 2005. Minimum description length tutorial.
- Hassani, A.; Walton, S.; Li, J.; Li, S.; and Shi, H. 2023. Neighborhood attention transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6185–6194.
- He, D.; Yang, Z.; Peng, W.; Ma, R.; Qin, H.; and Wang, Y. 2022. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5718–5727.
- Ho, J.; Lohn, E.; and Abbeel, P. 2019. Compression with flows via local bits-back coding. *Advances in Neural Information Processing Systems*, 32.
- Hoogeboom, E.; Peters, J.; Van Den Berg, R.; and Welling, M. 2019. Integer discrete flows and lossless compression. *Advances in Neural Information Processing Systems*, 32.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Kingma, F.; Abbeel, P.; and Ho, J. 2019. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *International Conference on Machine Learning*, 3408–3417. PMLR.
- Kodak, E. 1993. Kodak lossless true color image suite (PhotoCD PCD0992). <http://r0k.us/graphics/kodak>. Accessed: 2024-08-14.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045–3059.
- Li, D.; Bai, Y.; Wang, K.; Jiang, J.; and Liu, X. 2024a. Semantic Ensemble Loss and Latent Refinement for High-Fidelity Neural Image Compression. *arXiv preprint arXiv:2401.14007*.
- Li, D.; Bai, Y.; Wang, K.; Jiang, J.; Liu, X.; and Gao, W. 2024b. GroupedMixer: An Entropy Model with Group-wise Token-Mixers for Learned Image Compression. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4582–4597.
- Liang, P.; Jiang, J.; Liu, X.; and Ma, J. 2024. Image deblurring by exploring in-depth properties of transformer. *IEEE Transactions on Neural Networks and Learning Systems*.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Lee, K. M. 2017. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Mentzer, F.; Agustsson, E.; Tschannen, M.; Timofte, R.; and Gool, L. V. 2019. Practical full resolution learned lossless image compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10629–10638.
- Mentzer, F.; Gool, L. V.; and Tschannen, M. 2020. Learning better lossless compression using lossy compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6638–6647.
- Rhee, H.; Jang, Y. I.; Kim, S.; and Cho, N. I. 2022. LC-FDNet: Learned lossless image compression with frequency decomposition network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6033–6042.
- Ryder, T.; Zhang, C.; Kang, N.; and Zhang, S. 2022. Split hierarchical variational compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 386–395.
- Salimans, T.; Karpathy, A.; Chen, X.; and Kingma, D. P. 2016. PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications. In *International Conference on Learning Representations*.
- Shen, S.; Yue, H.; and Yang, J. 2023. Dec-adapter: Exploring efficient decoder-side adapter for bridging screen content and natural image compression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12887–12896.
- Skodras, A.; Christopoulos, C.; and Ebrahimi, T. 2001. The JPEG 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5): 36–58.
- Sneyers, J.; and Wuille, P. 2016. FLIF: Free lossless image format based on MANIAC compression. In *2016 IEEE international conference on image processing (ICIP)*, 66–70. IEEE.

- Theis, L.; Shi, W.; Cunningham, A.; and Huszár, F. 2016. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations*.
- Toderici, G.; Timofte, R.; Ballé, J.; Agustsson, E.; Johnston, N.; and Mentzer, F. 2020. Workshop and Challenge on Learned Image Compression (CLIC). <http://www.compression.cc>. Accessed: 2024-08-14.
- Townsend, J.; Bird, T.; and Barber, D. 2019. Practical lossless compression with latent variables using bits back coding. In *7th International Conference on Learning Representations, ICLR 2019*, volume 7. International Conference on Learning Representations (ICLR).
- Townsend, J.; Bird, T.; Kunze, J.; and Barber, D. 2019. HiLLOC: lossless image compression with hierarchical latent variable models. In *International Conference on Learning Representations*.
- Tsubota, K.; Akutsu, H.; and Aizawa, K. 2023. Universal deep image compression via content-adaptive optimization with adapters. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2529–2538.
- Tucker, L. R. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3): 279–311.
- Van Den Oord, A.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. In *International conference on machine learning*, 1747–1756. PMLR.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Wang, K.; Bai, Y.; Li, D.; Zhai, D.; Jiang, J.; and Liu, X. 2024. Learning Lossless Compression for High Bit-Depth Volumetric Medical Image. *IEEE Transactions on Image Processing*, 1–1.
- Wang, K.; Bai, Y.; Zhai, D.; Li, D.; Jiang, J.; and Liu, X. 2023. Learning lossless compression for high bit-depth medical imaging. In *2023 IEEE International conference on multimedia and expo (ICME)*, 2549–2554. IEEE.
- Weinberger, M. J.; Seroussi, G.; and Sapiro, G. 2000. The LOCO-I lossless image compression algorithm: Principles and standardization into JPEG-LS. *IEEE Transactions on Image processing*, 9(8): 1309–1324.
- Xia, G.-S.; Yang, W.; Delon, J.; Gousseau, Y.; Sun, H.; and Maître, H. 2010. Structural high-resolution satellite image indexing. Vienna, Austria.
- Yu, W.; Luo, M.; Zhou, P.; Si, C.; Zhou, Y.; Wang, X.; Feng, J.; and Yan, S. 2022. MetaFormer Is Actually What You Need for Vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10819–10829.
- Zhang, M.; Zhang, A.; and McDonagh, S. 2021. On the out-of-distribution generalization of probabilistic image modelling. *Advances in Neural Information Processing Systems*, 34: 3811–3823.
- Zhang, S.; Kang, N.; Ryder, T.; and Li, Z. 2021a. iflow: Numerically invertible flows for efficient lossless compression via a uniform coder. *Advances in Neural Information Processing Systems*, 34: 5822–5833.
- Zhang, S.; Zhang, C.; Kang, N.; and Li, Z. 2021b. ivpf: Numerical invertible volume preserving flow for efficient lossless compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 620–629.
- Zhang, Z.; Wang, H.; Chen, Z.; and Liu, S. 2024. Learned Lossless Image Compression based on Bit Plane Slicing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 27579–27588.