

PILC: Practical Image Lossless Compression with an End-to-end GPU Oriented Neural Framework

Ning Kang^{1*} Shanzhao Qiu^{2*} Shifeng Zhang¹ Zhenguo Li^{1†} Shutao Xia^{2†}

¹ Huawei Noah's Ark Lab ² Tsinghua University

kang.ning2@huawei.com, qiusz20@mails.tsinghua.edu.cn

Abstract

Generative model based image lossless compression algorithms have seen a great success in improving compression ratio. However, the throughput for most of them is less than 1 MB/s even with the most advanced AI accelerated chips, preventing them from most real-world applications, which often require 100 MB/s. In this paper, we propose PILC, an end-to-end image lossless compression framework that achieves 200 MB/s for both compression and decompression with a single NVIDIA Tesla V100 GPU, 10× faster than the most efficient one before. To obtain this result, we first develop an AI codec that combines auto-regressive model and VQ-VAE which performs well in lightweight setting, then we design a low complexity entropy coder that works well with our codec. Experiments show that our framework compresses better than PNG by a margin of 30% in multiple datasets. We believe this is an important step to bring AI compression forward to commercial use.

1. Introduction

Lossy compression have shown great success in recent research [3, 6, 22, 25–28, 30]. In this paper, we focus on lossless compression. The basic idea of a lossless compression algorithm is to represent more likely appeared data with shorter codewords, while less frequent data with longer ones, such that the codeword is shorter than the original data in expectation. For example, an image with each pixel randomly generated is rarely seen in the real world, while a real picture taken by a camera occurs much more often, so the latter compresses better for almost all image compression algorithms.

According to Shannon's source coding theorem [34], no algorithm can compress data shorter than its entropy. On the other hand, if the distribution of data is known in ad-

*Equal contribution.

†Correspondence to: Zhenguo Li (li.zhenguo@huawei.com) and Shutao Xia (xiast@sz.tsinghua.edu.cn).




Figure 1. Comparison of compression throughput and ratio on CIFAR10. Our framework achieves high compression throughput with a competitive compression ratio.

vance, an entropy coder can be applied so that the code length would be very close to its entropy, thus obtaining the optimal compression ratio.

Unfortunately, in most cases, the data distribution is unknown. Traditional algorithms make use of some prior information to infer this distribution. For example, LZ77 [47] and LZ78 [48] used by most archive formats assume that data with lots of repeated segments occur more often than those without. For image compression algorithms, PNG [4] assumes most adjacent pixels are similar to each other, and JPEG-like algorithms apply the fact that lower frequency components are more significant than higher frequency ones.

Recently, deep generative models have shown great success in probabilistic modelling and lossless compression. However, the low throughput makes them difficult to be used in real scenarios. An end user does not want to wait for too long to compress/decompress a file, and a file server needs to manipulate tons of files every day. In many cases, a 100+ MB/s throughput is needed. For example, streaming a 1080p video with 30 fps requires 187 MB/s for decom-

pression. However, it is difficult for a learned algorithm to achieve this speed due to the following bottlenecks:

Network inference. Barely any AI compression algorithms have inference speed of 100+ MB/s because:

- A large network is usually needed for a good density estimation [5].
- Multiple network inferences are required to decompress with an auto-regressive [32,41] model.
- Bits-Back scheme [12] prevents algorithms from applying a large batch size [37].

Coder. Generative models only cooperate well with a dynamic entropy coder, which is far slower than static variants used in traditional algorithms. For most AI algorithms, the coder decodes only about 1 MB/s single-threaded.

Data transfer. Data transfer between CPU memory and GPU memory can also be a bottleneck for AI algorithms.

- Too many transfers are needed for decompression with an auto-regressive or hierarchical AE/VAE model.
- Transfer amount is too large for models predicting with a complicated distribution including lots of parameters, such as a distribution mixed by 10 logistic ones.

Besides, other problems exist in current AI algorithms. One is single image decompression. One would expect to decompress only this image, instead of having to decompress many unrelated ones at the same time, but it is not the case for some VAE [37] and Flow algorithms [13] using bits-back [12]. Another is compression with different image sizes. An algorithm that can only compress images with a fixed size hardly gets used in real applications, however, those with a fully connected layer or transformer layer in the network are not straight-forward to change the input size. Table 1 briefly show the limitation of current algorithms.

1.1. Our contributions

In this paper, we focus on practical image lossless compression. To solve above issues, we make the following contributions:

- We build an end-to-end framework with compression/decompression throughput of about 200 MB/s in one Tesla V100, and compression ratio 30% better than PNG, which also enables single image decompression and image of different sizes.
- We develop a very lite auto-regressive + Vector Quantized Variational Auto-encoder (VQ-VAE) [42] model, whose inference speed is about 300 MB/s, but log-likelihood is similar to L3C [23], whose inference speed is only about 30 MB/s.
- We design an AI compatible semi-dynamic entropy coder that is efficient in GPU besides CPU. As a result, only one data transfer with a minimal amount is needed between CPU memory and GPU memory. As far as we know, this is the first GPU coder applied in

Table 1. Limitations for some generative model based algorithms, where the columns mean: whether single image decompression is supported; whether images of different sizes are supported; whether inference time, coder, transfer time are faster than 100 MB/s, respectively.

Algorithm	Single	Size	Inference	Code	Trans
PixelCNN [41]	✓	✓	✗	✗	✗
DistAug [16]	✓	✗	✗	✗	✗
L3C [23]	✓	✓	✗	✗	✓
BB-ANS [37]	✗	✗	✗	✗	✗
Bit-Swap [17]	✗	✗	✗	✗	✗
IDF [14]	✓	✓	✗	✗	✓
LBB [13]	✗	✓	✗	✗	✗
PILC (Ours)	✓	✓	✓	✓	✓

an AI compression algorithm.

- We implement the end-to-end version in pure Python and PyTorch, making it easier for future works to test and improve on our implementation.

2. Related Work

Generative model based lossless compression. While human designed algorithms are sometimes very powerful, deep-learning models can usually do better (See Fig. 1). For example, to compress a portrait, many traditional algorithms can make use of the fact that the skin color is similar everywhere, but few know a human has two eyes, one nose, etc., which can be usually be captured by a generative model. Denote the real probability of x as $p(x)$, the predicted probability of x as $q(x)$, and the expected codeword length as L . Then

$$L(X) = - \sum_{x \in X} p(x) \log_2 q(x) = H(X) + KL(P||Q) \quad (1)$$

Therefore, for a deep-learning algorithm, the compression ratio depends on how accurate the prediction is. Generative model based image lossless compression algorithms can be classified into 3 types.

- Auto-regressive (AR) algorithms predict each symbol from the previous context. Some of them are PixelCNN [41], PixelCNN++ [32], DistAug [16], etc.. They have great compression ratios with accurate modelling. However during decompression, since one pixel can be predicted only when all previous ones are known, one inference of the network can only decode one symbol, resulting in a very slow decompression speed.
- Autoencoder (AE) and Variational autoencoder (VAE) algorithms make use of latent variables to help with the prediction. Since the whole image can be predicted with only one network inference, an auto-encoder algorithm such as L3C [23] can decompress thousands of times faster than PixelCNN++ [32], but with a much

worse compression ratio. After bits-back [12] scheme is introduced to VAE algorithm [37], the compression ratio is improved significantly. Thus later works such as BB-ANS [37], Bit-Swap [17] and HiLLoC [38] all apply this scheme. Algorithms such as NVAE [39] and VDVAE [5] can even get ratios comparable to auto-regressive ones. However, bits-back introduces a dependency between different images, making it impossible to decompress a single one. Furthermore, the cost of random initial bits can only be leveraged by compressing/decompressing a large number of images sequentially, affecting the batch size and transfer speed greatly.

- Flow algorithms apply learnable invertible transforms to make data easier to compress. IDF [14] and IDF++ [40] that have only integer transforms are named as integer flows. On the other hand, general flows such as LBB [13] and iFlow [45] use transforms in real numbers. There is also one type in between called volume preserving flow, which is discovered in iVPF [46]. Integer flows are the only type that do not require bits-back, hence the only one that supports single image decoding. But as the expressive power of a single layer is weaker, the model needs to be large, affecting the inference time. Besides, general flow models require one coding step for each layer, making data transfer another bottleneck.

Entropy Coder. A static entropy coder codes data from a fixed distribution. It is used by most commercial algorithms. A dynamic entropy coder, on the other hand, can code many distributions. It is usually used to code one kind of distribution with parameters as input, such as a normal distribution with variable means and variances. The most widely used entropy coders such as Huffman Code [15], Arithmetic Coder (AC) [24, 31], and Asymmetric Numerical System (ANS) [10] all have both static and dynamic variants. Dynamic entropy coders are much slower than static ones, as the density mass needs to be calculated on the fly, unlike the latter one, which can precalculate all of them beforehand. For example, in a single-threaded setting, a static variant of ANS named FSE [8] decodes faster than 300 MB/s, while another variant named rANS [10] decodes at only about 1 MB/s with a distribution mixed by 10 logistic ones. Despite this fact, a dynamic coder is required for an AI algorithm to do fine-grained density estimation, which is also a key point that makes AI outperform traditional methods.

3. Methods

In this section, we introduce PILC, a practical image lossless compression framework. The proposed framework consists of a generative model and a semi-dynamic entropy coder. The model is composed of Three-Way Auto-

regressive module and VQ-VAE module. We demonstrate how the model interact with the coder to achieve high throughput.

3.1. Model Architecture

Overall Model. Among all current algorithms, AE have the highest throughput but the lowest compression ratio. As suggested by many previous works, AE models cannot capture local features very well, while a local AR model is capable of doing so [33, 44]. Therefore an intuitive idea would be to use a small AR to capture local features, and AE for more global ones. Since AR model hurts the parallelism and affect the throughput heavily, it must be very small. In this work, a Three-Way Auto-regressive (TWAR) model is designed, where each pixel is predicted with only 3 other ones, reserving the ability to decompress in parallel. For the AE model, we replace it with VQ-VAE [42], which is proven to be successful in image generation [9, 11, 29, 42], while the inference time is still similar to AE.

As is shown in Fig. 2a, TWAR predicts the original image and output image residual, then VQ-VAE estimates the distribution of the residual. Distribution of values in the residual is much simpler than those of original images, making it easier to be predicted by VQ-VAE and coded with our entropy coder.

Compress: Given an image x , we use the Three-Way Auto-regressive module to predict the image, and get the reconstruction image \hat{x} , image residual $r = x - \hat{x}$. Then we adopt VQ-VAE to estimate the residual distribution $p(r|x)$, which we model as a simple logistic distribution. With r and $p(r|x)$, we could compress r to bit stream with our proposed dynamic entropy coder. We also compress the vector index to bit stream using the coder.

Decompress: At decompress time, we decompress vector index from bit stream. Then we select corresponding vectors from the codebook Z to formulate the latent feature z_q . The feature is fed to the decoder and outputs the residual distribution $p(r|x)$. The coder subsequently decode the residual r using $p(r|x)$. The AR module then decodes the original image x from the residual r .

Learning Residual with Three-Way Auto-regressive. Inspired by the Paeth filter which is utilized in the traditional lossless compression method PNG [4] to transform the image to make it more efficiently compressible, we adopt a similar rule that utilizes three predicted points to predict the current one, which is shown in Fig. 2b. Our AR module is formulated by:

$$\hat{x}_{ruv} = W_r^T \cdot (x_{r(u-1)(v-1)}, x_{r(u-1)v}, x_{ru(v-1)}) + b_r \quad (2)$$

$$\hat{x}_{guv} = W_g^T \cdot (x_{gu(v-1)}, x_{ru(v-1)}, x_{ruv}) + b_g \quad (3)$$

$$\hat{x}_{buv} = W_b^T \cdot (x_{bu(v-1)}, x_{gu(v-1)}, x_{guv}) + b_b \quad (4)$$

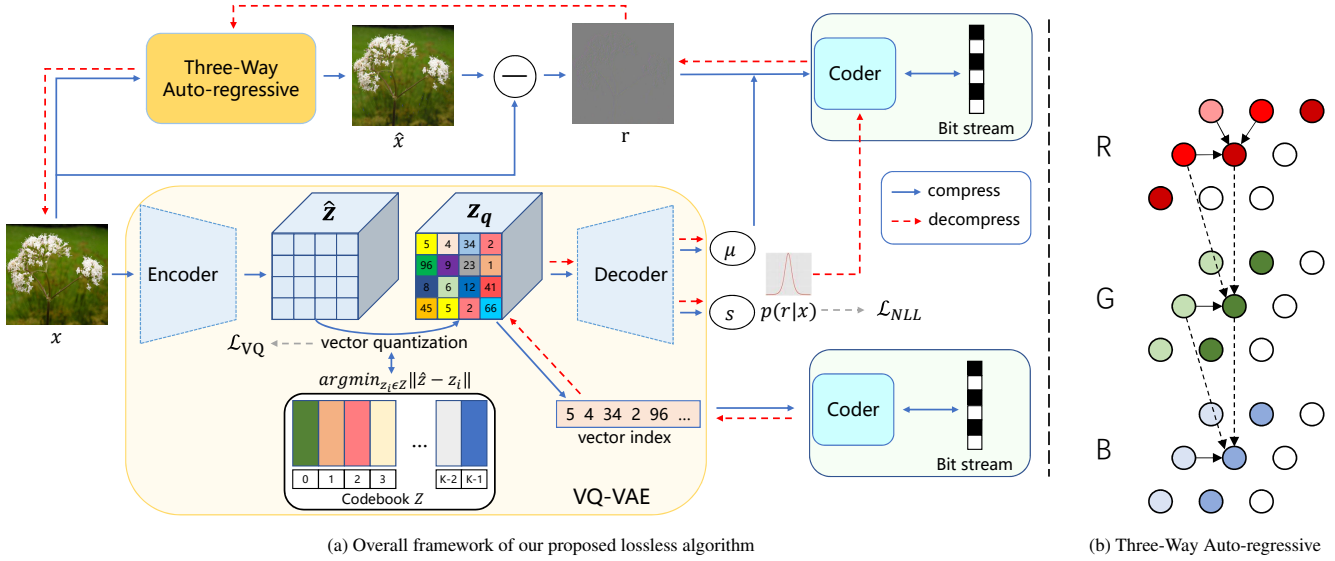


Figure 2. Left: Our framework consists of Three-Way Auto-regressive, VQ-VAE, and Coder. The AR module predicts the input image and obtains the residual. The VQ-VAE models the probability distribution of the residual with a codebook. The Coder compresses image residual and vector index to bit stream. When decompress, vector index is decompressed first and pushed into the VQ-VAE decoder to obtain the residual distribution. This distribution is subsequently adopted to decompress the residual and then recover the original image (blue solid arrows indicate compress process, red dotted ones indicate decompress, best viewed in color). Right: Three-Way Auto-regressive adopts three predicted points to predict the current point.

where r, g, b denote image channel, u, v denote the spatial location. We pad the input image with one row on the top and one column on the left to initiate the prediction. For the red channel, the current point is predicted by its up, left, upper-left points like Paeth filter. Inspired by YCoCg-R color space which is shown successful in FLIF [35] format and utilizes different channels to calculate the color space, for green and blue channel, we also use their previous channel to conduct the prediction. Different from [19], our AR module models the pictures channel-wise, predict the mean only, and only has 12 parameters (with a bias in each channel), but it could still generate a high-quality prediction. As shown in Fig. 2a, \hat{x} is quite close to x , which suggests the locality of natural images.

As explained in Sec. 1, auto-regressive models like PixelCNN [41] fail in decompression speed due to their dependencies nature because it is hard to parallelize the inference process. Fortunately, our AR module could be parallelized during inference. Figure 3 suggests one step in decoding. For the red channel, the current point is dependent on its up, left, upper-left points, so anti-diagonal points can be calculated in parallel. Once the red channel is recovered, for the green channel, we could calculate column by column to recover this channel, and the blue channel is the same. Thus, more points could be calculated at one time, which speeds up the inference process of the AR module. Because the padding rule is fixed, we could eventually losslessly recover the original image x using the residual r .

Learning Distribution with VQ-VAE. Vector quantization has been proved to be effective in lossy compression [1, 20]. Instead of designing a more sophisticated auto-encoder (AE) model, we utilize VQ-VAE [42] to learn residual distribution. Intuitively, given a fixed storage space, for VQ-VAE, we could just store the vector index and when decoding, with the codebook we could utilize more diversity features. But for vanilla AE, we only have the fixed size features. Experiments in Sec. 4.2 suggest VQ-VAE's superiority over vanilla AE. As the original image would contain much more spatial information, we choose to model the residual distribution given the original image, *i.e.*, $p(r|x)$.

As shown in Fig. 2a, given an image $x \in \mathbb{R}^{H \times W \times 3}$, we get its latent feature by:

$$\hat{z} = E(x) \in \mathbb{R}^{h \times w \times D} \quad (5)$$

where E is residual block based encoder, and D is feature depth dimension. Then, with a codebook $Z = \{z_i\}_{i=1}^K \in \mathbb{R}^D$, we could conduct vector quantization on \hat{z} :

$$z_q = q(\hat{z}) := \left(\arg \min_{z_i \in Z} \|\hat{z}_{uv} - z_i\| \right) \in \mathbb{R}^{h \times w \times D} \quad (6)$$

where q is the element-wise quantization for each $\hat{z}_{uv} \in \mathbb{R}^D$ to its closest code z_i , and u, v is the spatial location. Finally, we output the residual distribution $p(r|x)$ which we model as a logistic distribution and is formulated by:

$$(\mu, s) = D(z_q) \in \mathbb{R}^{H \times W \times 3 \times 2} \quad (7)$$




Figure 3. The decoding process of Red channel (left) and Green/Blue channel (right) for proposed AR model. For the red channel, anti-diagonal pixels can be decompressed in parallel, while for G and B, one column of pixels can be decoded in parallel.

where μ is the location parameter, s is the scale parameter, and D is residual block based decoder.

At compress time, the quantized vector index is encoded to bit stream by the coder. At decompress time, we could decode the index, select the corresponding vectors, and decode the distribution $p(r|x)$. In this way, we utilize the codebook Z to memorize more information.

Since there is no real gradient defined for Eq. (6), the gradient back-propagation is achieved by straight-through estimator, which just copies the gradient from the decoder to encoder. The VQ loss is formulated by:

$$\mathcal{L}_{VQ} = \|\text{sg}[\hat{z}] - z_q\|^2 + \beta \|\hat{z} - \text{sg}[z_q]\|^2 \quad (8)$$

where sg denotes the stop-gradient operator. The first term optimizes the codebook Z and the second term is a commitment loss to make sure \hat{z} commits to codebook and does not grow arbitrarily [42]. β is a weight parameter to balance the two terms.

Loss Function. During training, we could jointly optimize the AR module and VQ-VAE module by maximizing the negative log-likelihood of the residual r :

$$\mathcal{L}_{NLL} = \mathbb{E}_{x \sim p(x)} [-\log p(r|x)] \quad (9)$$

where $p(x)$ denotes the distribution of the training dataset. The lower the negative log-likelihood is, the fewer bits we need to compress the residual r . See Sec. 3.3 for the more implement detail.

Thus, the total objective function to optimize our neural compression model can be formulated by:

$$\mathcal{L}_{total} = \mathcal{L}_{NLL} + \alpha \mathcal{L}_{VQ} \quad (10)$$

where α balances the two terms.

Our lightweight model with the Three-Way Auto-regressive and VQ-VAE achieves a high inference speed, which plays an important role in our overall framework.

3.2. Coder Design

We design our entropy coder based on rANS [10], and save computation time by precalculating intermediate results. Previous works such as tANS [10] already adopt

this idea. However, tANS can only code a fixed distribution, which is inappropriate for AI algorithms. On the other hand, our proposed coder is more similar to dynamic ones, as it supports distributions with variable parameters. The only difference is that the parameters need to be quantized, and probability mass needs to be calculated for all of them. Thus we call it a semi-dynamic entropy coder.

To code with rANS, one first needs to choose a constant M , then for each symbol x , quantize its probability mass function (denoted as P_x), so that $\sum_x P_x = 2^M$. Let C_x be the cumulative distribution function of x . It satisfies that $C_x = \sum_{i < x} P_x$. To encode pixels within range $[0, 255]$, M is usually chosen in $[10, 15]$.

In this work, we modify rANS slightly, and propose Algorithm 1 and Algorithm 2 as the logic behind.

Algorithm 1 Modified rANS: Encode x

Input: M, P_x, C_x , State S , stream

- 1: $S \leftarrow S + 2^M$
 - 2: **while** $S \geq P_x * 2$ **do**
 - 3: $\text{push_stream}(S \bmod 2) \triangleright$ last bit of S to stream
 - 4: $S \leftarrow \lfloor S \div 2 \rfloor$
 - 5: **end while**
 - 6: $S \leftarrow S - P_x + C_x$
-

Algorithm 2 Modified rANS: Decode one symbol

Input: M, P_x, C_x for all x , State S , stream

- 1: $x \leftarrow \text{binary_search}(S): C_x \leq S < C_x + P_x$
 - 2: $S \leftarrow S - C_x + P_x$
 - 3: **while** $S < 2^M$ **do**
 - 4: $S \leftarrow S \times 2 + \text{pop_stream}() \triangleright$ last bit of stream
 - 5: **end while**
 - 6: $S \leftarrow S - 2^M$
 - 7: **return** x
-

Loops and binary search form the time-consuming steps for these algorithms, so tables are built to cache these results.

For encoding, one solution would be to create a table to store the number of bits to be pushed for each possible value of S and P_x . Each value can be represented as a 16-bit unsigned integer, so the memory needed is 2^{2M+1} bytes. In this work, we assume the number of distributions for different symbols is small. Denote the total number of distributions as D , and the distribution index $d \in [0, D)$ for each symbol can be calculated via the network, and let X be the number of colors. Since S is always in $[2^M, 2^{M+1})$ after Step 1, so for a fixed d , the number of loops can differ only 1. So we build a table δ with dimension $D \times X$ such that

$$\frac{S}{2^{\lfloor (\delta[d,x]+S)/2^M \rfloor}} \in [P_x, P_x * 2) \quad (11)$$

To further simplify the algorithm, Step 6 is merged with Step 1 of next symbol in this work, and table ϕ is built such

that $\phi[d, x] = 2^M - P_x + C_x$. S needs to be initialized in $[2^M, 2^{M+1})$ then. The encoding process now is shown in Algorithm 3. Both δ and ϕ use unsigned 16-bit integers, so the memory needed for tables are $4 \times D \times X$ bytes. As the total memory needed for all P_x 's and C_x 's are also $4 \times D \times X$ bytes, which is not needed any more, Algorithm 3 does not cost any additional memory.

Algorithm 3 ANS-AI: Encode x

Input: M, δ, ϕ, x, d , State S , stream
 1: $b = \lfloor (\delta[d, x] + S) / 2^M \rfloor$
 2: `push_stream(S, b)` ▷ last b bits of S to stream
 3: $S \leftarrow \lfloor S / 2^b \rfloor + \phi[d, x]$

In terms of decoding, for each distribution index d and state S , we maintain a table Θ to store the symbol to be decoded, a table B for number of bits to pop, and a table N to merge Step 2 to 6. Since Θ and B use 8-bit unsigned integers, and N use a 16-bit unsigned one, the total memory consumed for these tables are $D \times 2^{M+2}$ bytes. One alternative is to replace N and B with ϕ and δ , reducing the memory to $D \times 2^M + 4 \times D \times X$, but this will make the computation more complicated, and reduce the decoding speed by about 50% from our experiments.

Algorithm 4 ANS-AI: Decode one symbol

Input: M, Θ, B, N, d , State S , stream
 1: $x \leftarrow \Theta[d, S]$
 2: $S \leftarrow N[d, S] + \text{pop_stream}(B[d, S])$ ▷ last $B[d, S]$ bits
 3: **return** x

Our proposed coder not only simplifies the computation, but also makes it possible to implement with a machine learning framework such as PyTorch by removing all the logical units.

3.3. Integration of Model and Coder

To guarantee the effectiveness of the coder, the total number of distributions D needs to be small. In our model, each pixel of r is predicted with a truncated logistic distribution with location parameter μ and scale parameter s , truncated between 0 and 255. Denote the distribution as $r \sim L(\mu, s, 0, 255)$. Since the number of means is required to be quantized to at least 256 values, D would be over one thousand.

In this work, we predict $r - \mu + 128$ within `uint8` space, *i.e.*, the prediction of $r \sim L(\mu, s, 0, 255)$ is approximated as $(r - \mu + 128) \bmod 256 \sim L(128, s, 0, 255)$. To make it as accurate as possible, μ needs to be as close to 128 as possible. We apply the following two techniques for this:

1. In VQ-VAE module, instead of predicting residual r directly, we predict $(r + 128) \bmod 256$.
2. We add a sigmoid layer to predict the final value of μ in VQ-VAE module to make sure it is close to 128.




Figure 4. KL divergence between real truncated logistic distribution and approximated distribution, which indicates the BPD loss for the approximation. The figure shows that the approximated distribution is very close to the real one when $\mu \in [96, 160]$.

Fig. 4 and Fig. 6 show that this technique almost deals with no BPD loss.

4. Experiments

In this section, we evaluate the model and coder effectiveness with extensive experiments and we conduct experiments on both low resolution and high resolution images which demonstrate our framework obtains the fastest compression speed with a competitive compression ratio.

4.1. Experimental Settings

Implementation Details. The AR and VQ-VAE module are jointly trained. The AR module contains 12 parameters. In VQ-VAE, the codebook size K is set to 256 and the codebook dimension D is set to 32. The encoder and decoder both have 4 residual blocks, all with 32 channels. The encoder downsampling ratio is 2, and pixel shuffle is used when upsampling in the decoder. Thus, an image x with shape $32 \times 32 \times 3$ would be quantized to $16 \times 16 \times 1$ vectors. β in Eq. (8) is set to 0.25 as in [42], and α in Eq. (10) is set to 125.

We do all experiments in a Linux docker with Intel(R) Xeon(R) Gold 6151 CPU (3.0 GHz, 72 threads), and one NVIDIA Tesla V100 32GB. For speed tests, we ignore the time for disk I/O, since it has nothing to do with compression algorithms. But for fair comparison between AI and traditional methods, we do count the time for data transfer between CPU memory and GPU memory. That means for all AI algorithms, timer starts when data is loaded into RAM and ends when compressed data is stored in RAM.

Different from previous AI compression works which focus on density estimation, this paper targets practical compression. Therefore, for BPD values reported in this paper, all storage required to recover the original image is counted, including meta data, Huffman tables, auxiliary bits, etc.. For AI algorithms, real BPD is reported instead of theoret-

Table 2. Ablation study on the model architecture. Theoretical BPD on CIFAR10 is reported.

Model	BPD
Vanilla Auto-encoder (AE)	5.53
VQ-VAE	5.13
Three-Way Auto-regressive (AR)	4.92
AR + AE	4.35
AR + VQ-VAE	4.17

Table 3. Ablation study on the receptive field of AR model and the parallel mechanism. Theoretical BPD and decompress throughput on CIFAR10 (red channel) are reported. Receptive field means the current point is predicted by how many previous points.

Receptive Field	BPD	Throughput (MB/s)
3 (with parallel)	5.7714	382.5
3	5.7714	48.5
4	5.7737	47.5
5	5.7708	46.0
6	5.7701	44.8
7	5.7449	44.0

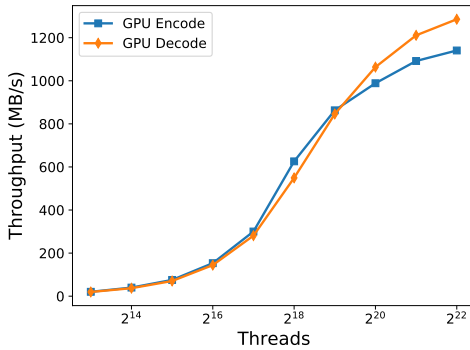


Figure 5. Encoding/decoding speed of the GPU coder. When the number of threads (batch size) is small, the throughput of the GPU coder is low, but it can benefit from large number of threads.

ical one unless specified otherwise, which means for each image, we calculate the number of bytes (instead of bits) needed for each image, then take the average.

Datasets. For low resolution experiments, we train and evaluate on CIFAR10, ImageNet32/64. For high resolution experiments, we train our model on 50000 images from Open Image Train Dataset [18] which are randomly selected and preprocessed as [23] to prevent overfitting on JPEG artifacts. Then, we evaluate on natural high resolution datasets, CLIC.mobile, CLIC.pro [7] and DIV2K [2] to test the model generalization ability. For high resolution setting, we crop 32×32 patches for evaluating.

4.2. Model Effectiveness

Effectiveness of model architecture. We analyze single and combinations of each component to evaluate the effectiveness of model architecture. In Tab. 2, the vanilla Auto-encoder (AE) contains the same encoder and decoder architecture as VQ-VAE with latent feature shape $16 \times 16 \times 1$. In AE and VQ-VAE, we directly use the model to predict the distribution of the original image. In the single AR, we compress the residual with a static entropy coder to get the BPD. AR+AE follows the same process as AR+VQ-VAE. As shown in Tab. 2, VQ-VAE outperforms AE which demonstrates the effectiveness of the learned codebook Z . We could intuitively explain that $16 \times 16 \times 1$ vector indexes would contain more diversity features than $16 \times 16 \times 1$ numbers in AE’s latent. Our proposed AR+VQ-VAE model achieves the best theoretical BPD on CIFAR10 which suggests the effectiveness of predicting residual.

Effectiveness of Three-Way AR and parallel mechanism. We compare different receptive fields of the AR model and the parallel mechanism. Experiments are conducted on red channel of CIFAR10 images. The receptive field of AR model means the current point is predicted by how many previous points. In our model, the receptive field is three. As shown in Tab. 3, as the receptive field increases, the BPD indeed descends, but the decline is very limited. With the parallel mechanism we designed in Sec. 3.1, the Three-Way AR achieves a throughput of 382.5 MB/s when recovering original images from residual in decompression, and outperforms other settings which cannot be paralleled.

4.3. Coder Effectiveness

We implement the GPU coder based on Algorithm 3 and Algorithm 4. We test the coder on the randomly generated data (1000000 pieces, each with length 1000, and range 0-255). The signals have the same dimension with each value in 0-7. For each setting, we run 10 times and take the average speed. From Fig. 5, we see that GPU coder is slow for the small number of threads. Fortunately, the number of threads can easily reach one million with a V100 32G. Therefore, we can inference on million batches, and save the result in GPU. Then we can code all of the batches at once. This is beneficial for the GPU coder, in the sense that the data transfer times and the amount is smaller for both compression and decompression.

4.4. Distribution Effectiveness

We count the predicted location parameter μ for the distribution specified in Sec. 3.3. It can be shown in Fig. 6 that for all the three datasets, the predicted values nearly always fall in [96, 160]. As stated in Sec. 3.3, this cause almost no loss in BPD.

Table 4. Compression performance in BPD and throughput in compression and decompression, compared to both engineered codec PNG [4], WebP [43], FLIF [35], JPEG2000 [36] and the learned L3C [23].

BPD	CIFAR10	ImageNet32	ImageNet64	DIV2K	CLIC.pro	CLIC.mobile	Throughput (MB/s)	
							Compress	Decompress
PNG [4] (fastest)	6.44	6.78	6.09	4.64	4.23	4.39	55.9	118.2
PNG [4] (best)	5.91	6.41	5.77	4.23	3.90	3.80	3.0	83.5
WebP [43] (-z 0)	4.77	5.44	4.92	3.43	3.22	3.03	29.8	99.1
FLIF [35] (-effort 0)	4.27	5.06	4.70	3.24	3.03	2.82	6.2	4.2
JPEG2000 [36]	6.75	7.50	6.08	4.11	3.79	3.94	7.6	9.1
L3C [23]	4.55	5.19	4.57	3.13	2.96	2.65	12.3	6.3
PILC (Ours)	4.23	5.10	4.76	3.41	3.23	3.00	180.3	217.2




Figure 6. The probability mass of the predicted μ on multiple datasets. They roughly follow the double exponential distribution with location parameter 128.

4.5. Compression performance

Compare with exists methods. We reproduce all methods in Tab. 4 on the same platform explained in Sec. 4.1 for fair comparison. We use Python CV2 lib for PNG and JPEG2000. We adopt FLIF version 1.0.1 and use command ‘flif -e -effort=0’ for fastest setting. For WebP, we use version 1.2.1 and the command ‘cwebp -lossless -z 0’. For L3C, we reproduce the results using the official model [21] and we replace the original coder in L3C with ours to test the speed as the original coder is very slow. The throughput results are reported based on high resolution settings. Table 4 shows our framework achieves the fastest compress and decompress throughput. Also, the framework achieves a competitive compression ratio on low resolution images and we still outperform PNG, WebP and JPEG2000 on high resolution images.

End-to-end throughput. To fully use the GPU coder and make the measured speed more accurate, we duplicate the CIFAR10 validation set 100 times (1 million images). Table 5 shows the details of the throughput and time taken for each step.

Table 5. Throughput & decomposition of time for the whole compression/decompression process on CIFAR10. Throughput is calculated with respect to the size of the original data. Time is measured as μ s per image.

	Phase	Throughput (MB/s)	Time (μ s)
Compress	RAM \rightarrow GPU	9246	0.33
	Model Inference	276	11.11
	Coder Encode	675	4.55
	GPU \rightarrow RAM	2985	1.03
	Total	180	17.02
Decompress	RAM \rightarrow GPU	11101	0.28
	Coder Decode	11091	0.28
	VQ-VAE Decode	721	4.26
	Code Decode	672	4.57
	AR Decode	869	3.53
	GPU \rightarrow RAM	2521	1.20
	Total	217	14.12

4.6. Limitation Analysis

Our framework, although very light-weighted, requires an AI chip such as Tesla V100 to work efficiently, which is inaccessible for most end-users. For example, one image with 2K resolution compresses only about 1.3 MB/s in the CPU of our server machine. Therefore, our framework is only applicable to file/cloud servers.

5. Conclusion

In this work, we develop the first efficient end-to-end generative model based image lossless compression framework with throughput of about 200 MB/s, and compression ratio 30% better than PNG. To achieve this result, we develop a model that combines AR model and VQ-VAE, and a semi-dynamic coder that is computation efficient.

To further improve PILC, future works could focus on algorithms that are more friendly in everyday devices, such as a PC, or mobile phone. This requires a network with much less FLOPS than the current VQ-VAE one.

References

- [1] Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc Van Gool. Soft-to-hard vector quantization for end-to-end learning compressible representations. In *NIPS*, pages 1141–1151, 2017. 4
- [2] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017. 7
- [3] Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 675–685. PMLR, 2019. 1
- [4] Thomas Boutell. PNG (portable network graphics) specification version 1.0. *RFC*, 2083:1–102, 1997. 1, 3, 8
- [5] Rewon Child. Very deep vaes generalize autoregressive models and can outperform them on images. In *ICLR*. OpenReview.net, 2021. 2, 3
- [6] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Variable rate deep image compression with a conditional autoencoder. In *ICCV*, pages 3146–3154. IEEE, 2019. 1
- [7] Workshop and Challenge on Learned Image Compression. <https://www.compression.cc/challenge/>. 7
- [8] Yann Collet. Finite State Entropy. <https://github.com/Cyan4973/FiniteStateEntropy>. 3
- [9] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, and Jie Tang. Cogview: Mastering text-to-image generation via transformers. *CoRR*, abs/2105.13290, 2021. 3
- [10] Jarek Duda. Asymmetric numeral systems. *CoRR*, abs/0902.0271, 2009. 3, 5
- [11] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883. Computer Vision Foundation / IEEE, 2021. 3
- [12] Brendan J. Frey and Geoffrey E. Hinton. Efficient stochastic source coding and an application to a bayesian network source model. *Comput. J.*, 40(2/3):157–165, 1997. 2, 3
- [13] Jonathan Ho, Evan Lohn, and Pieter Abbeel. Compression with flows via local bits-back coding. In *NeurIPS*, pages 3874–3883, 2019. 2, 3
- [14] Emiel Hoogeboom, Jorn W. T. Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. In *NeurIPS*, pages 12134–12144, 2019. 2, 3
- [15] David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952. 3
- [16] Heewoo Jun, Rewon Child, Mark Chen, John Schulman, Aditya Ramesh, Alec Radford, and Ilya Sutskever. Distribution augmentation for generative modeling. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 5006–5019. PMLR, 2020. 2
- [17] Friso H. Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 3408–3417. PMLR, 2019. 2, 3
- [18] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020. 7
- [19] Mu Li, Kede Ma, Jane You, David Zhang, and Wangmeng Zuo. Efficient and effective context-based convolutional entropy modeling for image compression. *IEEE Trans. Image Process.*, 29:5900–5911, 2020. 4
- [20] Mu Li, Wangmeng Zuo, Shuhang Gu, Jane You, and David Zhang. Learning content-weighted deep image compression. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(10):3446–3461, 2021. 4
- [21] Fabian Mentzer. L3C-PyTorch. <https://github.com/fab-jul/L3C-PyTorch>. 8
- [22] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional probability models for deep image compression. In *CVPR*, pages 4394–4402. Computer Vision Foundation / IEEE Computer Society, 2018. 1
- [23] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *CVPR*, pages 10629–10638. Computer Vision Foundation / IEEE, 2019. 2, 7, 8
- [24] Richard C. Pasco. Source coding algorithms for fast data compression (ph.d. thesis abstr.). *IEEE Trans. Inf. Theory*, 23(4):548, 1977. 3
- [25] Yash Patel, Srikanth Appalaraju, and R. Manmatha. Deep perceptual compression. *CoRR*, abs/1907.08310, 2019. 1
- [26] Yash Patel, Srikanth Appalaraju, and R. Manmatha. Human perceptual evaluations for image compression. *CoRR*, abs/1908.04187, 2019. 1
- [27] Yash Patel, Srikanth Appalaraju, and R. Manmatha. Hierarchical auto-regressive model for image compression incorporating object saliency and a deep perceptual loss. *CoRR*, abs/2002.04988, 2020. 1
- [28] Yash Patel, Srikanth Appalaraju, and R. Manmatha. Saliency driven perceptual image compression. In *WACV*, pages 227–236. IEEE, 2021. 1
- [29] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In *NeurIPS*, pages 14837–14847, 2019. 3
- [30] Oren Rippel and Lubomir D. Bourdev. Real-time adaptive image compression. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 2922–2930. PMLR, 2017. 1
- [31] Jorma Rissanen and Glen G. Langdon. Arithmetic coding. *IBM J. Res. Dev.*, 23:149–162, 1979. 3
- [32] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *ICLR (Poster)*. OpenReview.net, 2017. 2

- [33] Robin Schirrmeister, Yuxuan Zhou, Tonio Ball, and Dan Zhang. Understanding anomaly detection with deep invertible networks through hierarchies of distributions and features. In *NeurIPS*, 2020. 3
- [34] Claude E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27(3):379–423, 1948. 1
- [35] Jon Sneyers and Pieter Wuille. FLIF: free lossless image format based on MANIAC compression. In *ICIP*, pages 66–70. IEEE, 2016. 4, 8
- [36] David S. Taubman and Michael W. Marcellin. *JPEG2000 - image compression fundamentals, standards and practice*, volume 642 of *The Kluwer international series in engineering and computer science*. Kluwer, 2002. 8
- [37] James Townsend, Thomas Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. In *ICLR (Poster)*. OpenReview.net, 2019. 2, 3
- [38] James Townsend, Thomas Bird, Julius Kunze, and David Barber. Hilloc: lossless image compression with hierarchical latent variable models. In *ICLR*. OpenReview.net, 2020. 3
- [39] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *NeurIPS*, 2020. 3
- [40] Rianne van den Berg, Alexey A. Gritsenko, Mostafa Dehghani, Casper Kaae Sønderby, and Tim Salimans. IDF++: analyzing and improving integer discrete flows for lossless compression. In *ICLR*. OpenReview.net, 2021. 3
- [41] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1747–1756. JMLR.org, 2016. 2, 4
- [42] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NIPS*, pages 6306–6315, 2017. 2, 3, 4, 5, 6
- [43] webmproject. libwebp. <https://github.com/webmproject/libwebp>. 8
- [44] Mingtian Zhang, Andi Zhang, and Steven McDonagh. On the out-of-distribution generalization of probabilistic image modelling. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [45] Shifeng Zhang, Ning Kang, Tom Ryder, and Zhenguo Li. iflow: Numerically invertible flows for efficient lossless compression via a uniform coder. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [46] Shifeng Zhang, Chen Zhang, Ning Kang, and Zhenguo Li. ivpf: Numerical invertible volume preserving flow for efficient lossless compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 620–629, 2021. 3
- [47] Jacob Ziv and Abraham Lempel. A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory*, 23(3):337–343, 1977. 1
- [48] Jacob Ziv and Abraham Lempel. Compression of individual sequences via variable-rate coding. *IEEE Trans. Inf. Theory*, 24(5):530–536, 1978. 1