

Learning Better Lossless Compression Using Lossy Compression

Fabian Mentzer
ETH Zurich
mentzerf@vision.ee.ethz.ch

Luc Van Gool
ETH Zurich
vangool@vision.ee.ethz.ch

Michael Tschannen
Google Research, Brain Team
tschannen@google.com

Abstract

We leverage the powerful lossy image compression algorithm BPG to build a lossless image compression system. Specifically, the original image is first decomposed into the lossy reconstruction obtained after compressing it with BPG and the corresponding residual. We then model the distribution of the residual with a convolutional neural network-based probabilistic model that is conditioned on the BPG reconstruction, and combine it with entropy coding to losslessly encode the residual. Finally, the image is stored using the concatenation of the bitstreams produced by BPG and the learned residual coder. The resulting compression system achieves state-of-the-art performance in learned lossless full-resolution image compression, outperforming previous learned approaches as well as PNG, WebP, and JPEG2000.

1. Introduction

The need to efficiently store the ever growing amounts of data generated continuously on mobile devices has spurred a lot of research on compression algorithms. Algorithms like JPEG [51] for images and H.264 [53] for videos are used by billions of people daily.

After the breakthrough results achieved with deep neural networks in image classification [27], and the subsequent rise of deep-learning based methods, *learned* lossy image compression has emerged as an active area of research (e.g. [6, 45, 46, 37, 2, 4, 30, 28, 48]). In *lossy* compression, the goal is to achieve small bitrates R given a certain allowed distortion D in the reconstruction, i.e., the rate-distortion trade-off $R + \lambda D$ is optimized. In contrast, in *lossless* compression, no distortion is allowed, and we aim to reconstruct the input perfectly by transmitting as few bits as possible. To this end, a probabilistic model of the data can be used together with entropy coding techniques to encode and transmit data via a bitstream. The theoretical foundation for this idea is given in Shannon’s landmark paper [40], which proves a lower bound for the bitrate achievable by such a probabilistic model, and the




Figure 1. Overview of the proposed learned lossless compression approach. To encode an input image x , we feed it into the Q-Classifer (QC) CNN to obtain an appropriate quantization parameter Q , which is used to compress x with BPG. The resulting lossy reconstruction x_l is fed into the Residual Compressor (RC) CNN, which predicts the probability distribution of the residual, $p(r|x_l)$, conditionally on x_l . An arithmetic coder (AC) encodes the residual r to a bitstream, given $p(r|x_l)$. In gray we visualize how to reconstruct x from the bistream. Learned components are shown in *violet*.

overhead incurred by using an imprecise model of the data distribution. One beautiful result is that maximizing the likelihood of a parametric probabilistic model is equivalent to minimizing the bitrate obtained when using that model for lossless compression with an entropy coder (see, e.g., [29]). Learning parametric probabilistic models by likelihood maximization has been studied to a great extent in the generative modeling literature (e.g. [50, 49, 39, 34, 25]). Recent works have linked these results to learned lossless compression [29, 18, 47, 24].

Even though recent learned lossy image compression methods achieve state-of-the-art results on various data sets, the results obtained by the non-learned H.265-based BPG [43, 7] are still highly competitive, without requiring sophisticated hardware accelerators such as GPUs to run. While BPG was outperformed by learning-based approaches across the bitrate spectrum in terms of PSNR [30] and visual quality [4], it still excels particularly at high-PSNR lossy reconstructions.

In this paper, we propose a learned *lossless* compression system by leveraging the power of the lossy BPG, as

illustrated in Fig. 1. Specifically, we decompose the input image x into the lossy reconstruction x_l produced by BPG and the corresponding residual r . We then learn a probabilistic model $p(r|x_l)$ of the residual, conditionally on the lossy reconstruction x_l . This probabilistic model is fully convolutional and can be evaluated using a single forward pass, both for encoding and decoding. We combine it with an arithmetic coder to losslessly compress the residual and store or transmit the image as the concatenation of the bitstrings produced by BPG and the residual compressor. Further, we use a computationally inexpensive technique from the generative modeling literature, tuning the “certainty” (temperature) of $p(r|x_l)$, as well as an auxiliary shallow classifier to predict the quantization parameter of BPG in order to optimize our compressor on a per-image basis. These components together lead to a state-of-the-art full-resolution learned lossless compression system. All of our code and data sets are available on github.¹

In contrast to recent work in lossless compression, we do not need to compute and store any side information (as opposed to L3C [29]), and our CNN is lightweight enough to train and evaluate on high-resolution natural images (as opposed to [18, 24], which have not been scaled to full-resolution images to our knowledge).

In summary, our main contributions are:

- We leverage the power of the classical state-of-the-art lossy compression algorithm BPG in a novel way to build a conceptually simple learned lossless image compression system.
- Our system is optimized on a per-image basis with a light-weight post-training step, where we obtain a lower-bitrate probability distribution by adjusting the confidence of the predictions of our probabilistic model.
- Our system outperform the state-of-the-art in learned lossless full-resolution image compression, L3C [29], as well as the classical engineered algorithms WebP, JPEG200, PNG. Further, in contrast to L3C, we are also outperforming FLIF on Open Images, the domain where our approach (as well as L3C) is trained.

2. Related Work

Learned Lossless Compression Arguably most closely related to this paper, Mentzer *et al.* [29] build a computationally cheap hierarchical generative model (termed L3C) to enable practical compression on full-resolution images.

Townsend *et al.* [47] and Kingma *et al.* [24] leverage the “bits-back scheme” [17] for lossless compression of an image stream, where the overall bitrate of the stream is reduced by leveraging previously transmitted information. Motivated by recent progress in generative modeling using (continuous) flow-based models (e.g. [35, 23]), Hooge-

boom *et al.* [18] propose Integer Discrete Flows (IDFs), defining an invertible transformation for discrete data. In contrast to L3C, the latter works focus on smaller data sets such as MNIST, CIFAR-10, ImageNet32, and ImageNet64, where they achieve state-of-the-art results.

Likelihood-Based Generative Modeling As mentioned in Section 1, virtually every generative model can be used for lossless compression, when used with an entropy coding algorithm. Therefore, while the following generative approaches do not take a compression perspective, they are still related. The state-of-the-art PixelCNN [50]-based models rely on auto-regression in RGB space to efficiently model a conditional distribution. The original PixelCNN [50] and PixelRNN [49] model the probability distribution of a pixel given all previous pixels (in raster-scan order). To use these models for lossless compression, $O(H \cdot W)$ forward passes are required, where H and W are the image height and width, respectively. Various speed optimizations and a probability model amendable to faster training were proposed in [39]. Different other parallelization techniques were developed, including those from [34], modeling the image distribution conditionally on subsampled versions of the image, as well as those from [25], conditioning on a RGB pyramid and grayscale images. Similar techniques were also used by [9, 31].

Engineered Lossless Compression Algorithms The wide-spread PNG [33] applies simple autoregressive filters to remove redundancies from the RGB representation (e.g. replacing pixels with the difference to their left neighbor), and then uses the DEFLATE [11] algorithm for compression. In contrast, WebP [52] uses larger windows to transform the image (enabling patch-wise conditional compression), and relies on a custom entropy coder for compression. Mainly in use for lossy compression, JPEG2000 [41] also has a lossless mode, where an invertible mapping from RGB to compression space is used. At the heart of FLIF [42] is an entropy coding method called “meta-adaptive near-zero integer arithmetic coding” (MANIAC), which is based on the CABAC method used in, e.g., H.264 [53]. In CABAC, the context model used to compress a symbol is selected from a finite set based on local context [36]. The “meta-adaptive” part in MANIAC refers to the context model which is a decision tree learned *per image*.

Artifact Removal Artifact removal methods in the context of lossy compression are related to our approach in that they aim to make predictions about the information lost during the lossy compression process. In this context, the goal is to produce sharper and/or more visually pleasing images given a lossy reconstruction from, e.g., JPEG. Dong *et al.* [12] proposed the first CNN-based approach using a network inspired by super-resolution networks. [44] extends

¹<https://github.com/fab-jul/RC-PyTorch>

this using a residual structure, and [8] relies on hierarchical skip connections and a multi-scale loss. Generative models in the context of artifact removal are explored by [13], which proposes to use GANs [14] to obtain more visually pleasing results.

3. Background

3.1. Lossless Compression

We give a very brief overview of lossless compression basics here and refer to the information theory literature for details [40, 10]. In lossless compression, we consider a stream of symbols x_1, \dots, x_N , where each x_i is an element from the same finite set \mathcal{X} . The stream is obtained by drawing each symbol x_i independently from the same distribution \tilde{p} , i.e., the x_i are i.i.d. according to \tilde{p} . We are interested in encoding the symbol stream into a bitstream, such that we can recover the exact symbols by decoding. In this setup, the *entropy* of \tilde{p} is equal to the expected number of bits needed to encode each x_i :

$$H(\tilde{p}) = \text{bits}(x_i) = \mathbb{E}_{x_i \sim \tilde{p}} [-\log_2 \tilde{p}(x_i)].$$

In general, however, the exact \tilde{p} is unknown, and we instead consider the setup where we have an approximate model p . Then, the expected bitrate will be equal to the *cross-entropy* between \tilde{p} and p , given by:

$$H(\tilde{p}, p) = \mathbb{E}_{x_i \sim \tilde{p}} [-\log_2 p(x_i)]. \quad (1)$$

Intuitively, the higher the discrepancy between the model p used for coding is from the real \tilde{p} , the more bits we need to encode data that is actually distributed according to \tilde{p} .

Entropy Coding Given a symbol stream x_i as above and a probability distribution p (not necessarily \tilde{p}), we can encode the stream using *entropy coding*. Intuitively, we would like to build a table that maps every element in \mathcal{X} to a bit sequence, such that x_i gets a short sequence if $p(x_i)$ is high. The optimum is to output $\log_2 p(x_i)$ bits for symbol x_i , which is what entropy coding algorithms achieve. Examples include Huffman coding [19] and arithmetic coding [54].

In general, we can use a different distribution p_i for every symbol in the stream, as long as the p_i are also available for decoding. *Adaptive* entropy coding algorithms work by allowing such varying distributions as a function of previously encoded symbols. In this paper, we use adaptive arithmetic coding [54].

3.2. Lossless Image Compression with CNNs

As explained in the previous section, all we need for lossless compression is a model p , since we can use entropy coding to encode and decode any input losslessly given p .




Figure 2. Histogram of the marginal pixel distribution of residual values obtained using BPG and Q predicted from QC, on Open Images.

In particular, we can use a CNN to parametrize p . To this end, one general approach is to introduce (structured) side information z available both at encoding and decoding time, and model the probability distribution of natural images $x^{(i)}$ conditionally on z , using the CNN to parametrize $p(x|z)$.² Assuming that both the encoder and decoder have access to z and p , we can losslessly encode $x^{(i)}$ as follows: We first use the CNN to produce $p(x|z)$. Then, we employ an entropy encoder (described in the previous section) with $p(x|z)$ to encode $x^{(i)}$ to a bitstream. To decode, we once again feed z to the CNN, obtaining $p(x|z)$, and decode $x^{(i)}$ from the bitstream using the entropy decoder.

One key difference among the approaches in the literature is the factorization of $p(x|z)$. In the original PixelCNN paper [49] the image x is modeled as a sequence of pixels, and z corresponds to all previous pixels. Encoding as well as decoding are done autoregressively. In IDF [18], x is mapped to a z using an invertible function, and z is then encoded using a fixed prior $p(z)$, i.e., $p(x|z)$ here is a deterministic function of z . In approaches based on the bits-back paradigm [47, 24], while encoding, z is obtained by decoding from additional available information (e.g. previously encoded images). In L3C [29], z corresponds to features extracted with a hierarchical model that are also saved to the bitstream using hierarchically predicted distributions.

3.3. BPG

BPG is a lossy image compression method based on the HEVC video coding standard [43], essentially applying HEVC on a single image. To motivate our usage of BPG, we show the histogram of the marginal pixel distribution of the residuals obtained by BPG on Open Images (one of our testing sets, see Section 5.1) in Fig. 2. Note that while the possible range of a residual is $\{-255, \dots, 255\}$, we observe that for most images, nearly every point in the residual is in the restricted set $\{-6, \dots, 6\}$, which is indicative of the high-PSNR nature of BPG. Additionally, Fig. A1 (in the suppl.) presents a comparison of BPG to the state-of-the-art learned image compression methods, showing that BPG is still very competitive in terms of PSNR.

BPG follows JPEG in having a chroma format parameter to enable color space subsampling, which we disable by setting it to 4:4:4. The only remaining parameter to set

²We write $p(x)$ to denote the entire probability mass function and $p(x^{(i)})$ to denote $p(x)$ evaluated at $x^{(i)}$.

is the *quantization parameter* Q , where $Q \in \{1, \dots, 51\}$. Smaller Q results in less quantization and thus better quality (i.e., different to the *quality* factor of JPEG, where larger means better reconstruction quality). We learn a classifier to predict Q , described in Section 4.4.

4. Proposed Method

We give an overview of our method in Fig. 1. To encode an image x , we first obtain the quantization parameter Q from the *Q-Classifier* (QC) network (Section 4.4). Then, we compress x with BPG, to obtain the lossy reconstruction x_l , which we save to a bitstream. Given x_l , the *Residual Compressor* (RC) network (Section 4.1) predicts the probability mass function of the residual $r = x - x_l$, i.e.,

$$p(r|x_l) = \text{RC}(x_l).$$

We model $p(r|x_l)$ as a discrete mixture of logistic distributions (Section 4.2). Given $p(r|x_l)$ and r , we compress r to the bitstream using adaptive arithmetic coding algorithm (see Section 3.1). Thus, the bitstream B consists of the concatenation of the codes corresponding to x_l and r . To decode x from B , we first obtain x_l using the BPG decoder, then we obtain once again $p(r|x_l) = \text{RC}(x_l)$, and subsequently decode r from the bitstream using $p(r|x_l)$. Finally, we can reconstruct $x = x_l + r$. In the formalism of Section 3.2, we have $x = r, z = x_l$.

Note that no matter how bad RC is at predicting the real distribution of r , we can always do *lossless* compression. Even if RC were to predict, e.g., a uniform distribution—in that case, we would just need many bits to store r .

4.1. Residual Compressor

We use a CNN inspired by ResNet [15] and U-Net [38], shown in detail in Fig. 3. We first extract an initial feature map f_{in} with $C_f = 128$ channels, which we then downscale using a stride-2 convolution, and feed through 16 residual blocks. Instead of BatchNorm [20] layers as in ResNet, our residual blocks contain GDN layers proposed by [5]. Subsequently, we upscale back to the resolution of the input image using a transposed convolution. The resulting features are concatenated with f_{in} , and convolved to contract the $2 \cdot C_f$ channels back to C_f , like in U-Net. Finally, the network splits into four tails, predicting the different parameters of the mixture model, $\pi, \mu, \sigma, \lambda$, described next.

4.2. Logistic Mixture Model

We use a discrete mixture of logistics to model the probability mass function of the residual, $p(r|x_l)$, similar to [29, 39]. We closely follow the formulation of [29] here: Let c denote the RGB channel and u, v the spatial location. We define

$$p(r|x_l) = \prod_{u,v} p(r_{1uv}, r_{2uv}, r_{3uv}|x_l). \quad (2)$$

We use a (weak) autoregression over the three RGB channels to define the joint distribution over channels via logistic mixtures p_m :

$$p(r_1, r_2, r_3|x_l) = p_m(r_1|x_l) \cdot p_m(r_2|x_l, r_1) \cdot p_m(r_3|x_l, r_2, r_1), \quad (3)$$

where we removed the indices uv to simplify the notation. For the mixture p_m we use a mixture of $K = 5$ logistic distributions p_L . Our distributions are defined by the outputs of the RC network, which yields mixture weights π_{cuv}^k , means μ_{cuv}^k , variances σ_{cuv}^k , as well as mixture coefficients λ_{cuv}^k . The autoregression over RGB channels is only used to update the means using a linear combination of μ and the target r of previous channels, scaled by the coefficients λ . We thereby obtain $\tilde{\mu}$:

$$\begin{aligned} \tilde{\mu}_{1uv}^k &= \mu_{1uv}^k & \tilde{\mu}_{2uv}^k &= \mu_{2uv}^k + \lambda_{\alpha uv}^k r_{1uv} \\ \tilde{\mu}_{3uv}^k &= \mu_{3uv}^k + \lambda_{\beta uv}^k r_{1uv} + \lambda_{\gamma uv}^k r_{2uv}. \end{aligned} \quad (4)$$

With these parameters, we can define

$$p_m(r_{cuv}|x_l, r_{\text{prev}}) = \sum_{k=1}^K \pi_{cuv}^k p_L(r_{cuv}|\tilde{\mu}_{cuv}^k, \sigma_{cuv}^k), \quad (5)$$

where r_{prev} denotes the channels with index smaller than c (see Eq. 3), used to obtain $\tilde{\mu}$ as shown above, and p_L is the logistic distribution:

$$p_L(r|\mu, \sigma) = \frac{e^{-(r-\mu)/\sigma}}{\sigma(1 + e^{-(r-\mu)/\sigma})^2}.$$

We evaluate p_L at discrete r , via its CDF, as in [39, 29], evaluating

$$p_L(r) = \text{CDF}(r + 1/2) - \text{CDF}(r - 1/2). \quad (6)$$

4.3. Loss

As motivated in Section 3.1, we are interested in minimizing the cross-entropy between the real distribution of the residual $\tilde{p}(r)$ and our model $p(r)$: the smaller the cross-entropy, the closer p is to \tilde{p} , and the fewer bits an entropy coder will use to encode r . We consider the setting where we have N training images $x^{(1)}, \dots, x^{(N)}$. For every image, we compute the lossy reconstruction $x_l^{(i)}$ as well as the corresponding residual $r^{(i)} = x^{(i)} - x_l^{(i)}$. While the true distribution $\tilde{p}(r)$ is unknown, we can consider the empirical distribution obtained from the samples and minimize:

$$\mathcal{L}(\text{RC}) = - \sum_{i=1}^N \log p(r^{(i)}|x_l^{(i)}). \quad (7)$$

This loss decomposes over samples, allowing us to minimize it over mini-batches. Note that minimizing Eq. 7 is the same as maximizing the likelihood of p , which is the perspective taken in the likelihood-based generative modeling literature.




Figure 3. The architecture of the residual compressor (RC). On the left, we show a zoom-in of the *Residual Block* and the *Tail* networks. Given x_l , the lossy reconstruction of the image x , the network predicts the probability distribution of the residual, $p(r|x_l)$. This distribution is a mixture of logistics parametrized via $\mu, \sigma, \pi, \lambda$.

4.4. Q-Classifer

A random set of natural images is expected to contain images of varying “complexity”, where complex can mean a lot of high frequency structure and/or noise. While virtually all lossy compression methods have a parameter like BPG’s Q , to navigate the trade-off between bitrate and quality, it is important to note that compressing a random set of natural images with the same fixed Q will usually lead to the bitrates of these images being spread around some Q -dependent mean. Thus, in our approach, it is suboptimal to fix Q for all images.

Indeed, in our pipeline we have a trade-off between the bits allocated to BPG and the bits allocated to encoding the residual. This trade-off can be controlled with Q : For example, if an image contains components that are easier for the RC network to model, it is beneficial to use a higher Q , such that BPG does not waste bits encoding these components. We observe that for a fixed image, and a trained RC, there is a single optimal Q .

To efficiently obtain a good Q , we train a simple classifier network, the Q-Classifer (QC), and then use $Q = \text{QC}(x)$ to compress x with BPG. For the architecture, we use a light-weight ResNet-inspired network with 8 residual blocks for QC, and train it to predict a class in $Q = \{11, \dots, 17\}$, given an image x (Q was selected using the Open Images validation set). In contrast to ResNet, we employ no normalization layers (to ensure that the prediction is independent of the input size). Further, the final features are obtained by average pooling each of the final $C_f = 256$ channels of the $C_f \times H' \times W'$ -dimensional feature map. The resulting C_f -dimensional vector is fed to a fully connected layer, to obtain the logits for the $|Q|$ classes, which are then normalized with a softmax. Details are shown in Section A.1 in the supplementary material.

While the input to QC is the full-resolution image, the network is shallow and downsamples multiple times, making this a computationally lightweight component.

4.5. τ -Optimization

Inspired by the temperature scaling employed in the generative modeling literature (e.g. [22]), we further optimize the predicted distribution p with a simple trick: Intuitively, if RC predicts a μ_{cuv} that is close to the target r_{cuv} , we can make the cross-entropy in Eq. 7 (and thus the bitrate) smaller by making the predicted logistic “more certain” by choosing a smaller σ . This shifts probability mass towards r_{cuv} . However, there is a breaking point, where we make it “too certain” (i.e., the probability mass concentrates too tightly around μ_{cuv}) and the cross-entropy increases again.

While RC is already trained to learn a good σ , the prediction is only based on x_l . We can improve the final bitrate during encoding, when we additionally have access to the target r_{cuv} , by *rescaling* the predicted σ_{cuv}^k with a factor τ_c^k , chosen for every mixture k and every channel c . This yields a more optimal $\tilde{\sigma}_{cuv}^k = \tau_c^k \cdot \sigma_{cuv}^k$. Obviously, τ also needs to be known for decoding, and we thus have to transmit it via the bitstream. However, since we only learn a τ for every channel and every mixture (and not for every spatial location), this causes a completely negligible overhead of $C \cdot K = 3 \cdot 5$ floats = 60 bytes.

We find τ_c^k for a given image $x^{(i)}$ by minimizing the likelihood in Eq. 7 *on that image*, i.e., we optimize

$$\min_{\tau} \sum_{c,u,v} \log p_{\tau}(r^{(i)}|x_{l,cuv}^{(i)}), \quad (8)$$

where p_{τ} is equal to p predicted from RC but using $\tilde{\sigma}_{cuv}^k$. To optimize Eq. 8, we use stochastic gradient descent with a very high learning rate of 9E^{-2} and momentum 0.9, which converges in 10-20 iterations, depending on the image.

We note that this is also computationally cheap. Firstly, we only need to do the forward pass through RC once, to get $\mu, \sigma, \lambda, \pi$, and then in every step of the τ -optimization, we only need to evaluate $\tau_c^k \cdot \sigma_{cuv}^k$ and subsequently Eq. 8. Secondly, the optimization is only over 15 parameters. Finally, since for practical $H \times W$ -dimensional images, $15 \ll H \cdot W$, we can do the sum in Eq. 8 over a $4 \times$ spatially subsampled version of p_{τ} .

| [bpsp] | Open Images | CLIC.mobile | CLIC.pro | DIV2K |
|-----------|-------------|-------------|-------------|-------------|
| RC (Ours) | 2.790 | 2.538 | 2.933 | 3.079 |
| L3C | 2.991 +7.2% | 2.639 +4.0% | 2.944 +0.4% | 3.094 +0.5% |
| PNG | 4.005 +44% | 3.896 +54% | 3.997 +36% | 4.235 +38% |
| JPEG2000 | 3.055 +9.5% | 2.721 +7.2% | 3.000 +2.3% | 3.127 +1.6% |
| WebP | 3.047 +9.2% | 2.774 +9.3% | 3.006 +2.5% | 3.176 +3.2% |
| FLIF | 2.867 +2.8% | 2.492 -1.8% | 2.784 -5.1% | 2.911 -5.5% |

Table 1. Compression performance of the proposed method (RC) compared to the learned L3C [29], as well as the classical engineered approaches PNG, JPEG2000, WebP, and FLIF. We show the difference in percentage to our approach, using *green* to indicate that we achieve a better bpsp and *red* otherwise.

5. Experiments

5.1. Data sets

Training Like L3C [29], we train on 300 000 images from the Open Images data set [26]. These images are made available as JPEGs, which is not ideal for the lossless compression task we are considering, but we are not aware of a similarly large scale lossless training data set. To prevent overfitting on JPEG artifacts, we downscale each training image using a factor randomly selected from [0.6, 0.8] by means of the Lanczos filter provided by the Pillow library [32]. For a fair comparison, the L3C baseline results were also obtained by training on the exact same data set.

Evaluation We evaluate our model on four data sets: **Open Images** is a subset of 500 images from Open Images validation set, preprocessed like the training data. **CLIC.mobile** and **CLIC.pro** are two new data sets commonly used in recent image compression papers, released as part of the “Workshop and Challenge on Learned Image Compression” (CLIC) [1]. CLIC.mobile contains 61 images taken using cell phones, while CLIC.pro contains 41 images from DSLRs, retouched by professionals. Finally, we evaluate on the 100 images from **DIV2K** [3], a super-resolution data set with high-quality images. We show examples from these data sets in Section A.3.

For a small fraction of exceptionally high-resolution images (note that the considered testing sets contain images of widely varying resolution), we follow L3C in extracting 4 non-overlapping crops x_c from the image x such that combining x_c yields x . We then compress the crops individually. However, we evaluate the non-learned baselines on the full images to avoid a bias in favor of our method.

5.2. Training Procedures

Residual Compressor We train for 50 epochs on batches of 16 random 128×128 crops extracted from the training set, using the RMSProp optimizer [16]. We start with an initial learning rate (LR) of $5E^{-5}$, which we decay every 100 000 iterations by a factor of 0.75. Since our Q-Classifer is trained on the output of a trained RC network,

it is not available while training the RC network. Thus, we compress the training images with a random Q selected from $\{12, 13, 14\}$, obtaining a pair (x, x_l) for every image.

Q-Classifer Given a trained RC network, we randomly select 10% of the training set, and compress each selected image x once for each $Q \in \mathcal{Q}$, obtaining a $x_l^{(Q)}$ for each $Q \in \mathcal{Q}$. We then evaluate RC for each pair $(x, x_l^{(Q)})$ to find the optimal Q' that gives the minimum bitrate for that image. The resulting list of pairs $(x, x_l^{(Q')})$ forms the training set for the QC. For training, we use a standard cross-entropy loss between the softmax-normalized logits and the one-hot encoded ground truth Q' . We train for 11 epochs on batches of 32 random 128×128 crops, using the Adam optimizer [21]. We set the initial LR to the Adam-default $1E^{-4}$, and decay after 5 and 10 epochs by a factor of 0.25.

5.3. Architecture and Training Ablations

Training on Fixed Q As noted in Section 5.2, we select a random Q during training, since QC is only available after training. We explored fixing Q to one value (trying $Q \in \{12, 13, 14\}$) and found that this hurts generalization performance. This may be explained by the fact that RC sees more varied residual statistics during training if we have random Q 's.

Effect of the Crop Size Using crops of 128×128 to train a model evaluated on full-resolution images may seem too constraining. To explore the effect of crop size, we trained different models, each seeing the same number of pixels in every iteration, but distributed differently in terms of batch size vs. crop size. We trained each model for 600 000 iterations, and then evaluated on the Open Images validation set (using a fixed $Q = 14$ for training and testing). The results are shown in the following table and indicate that smaller crops and bigger batch-sizes are beneficial.

| Batch Size | Crop Size | BPSP on Open Images |
|------------|------------------|---------------------|
| 16 | 128×128 | 2.854 |
| 4 | 256×256 | 2.864 |
| 1 | 512×512 | 2.877 |

GDN We found that the GDN layers are crucial for good performance. We also explored instance normalization, and conditional instance normalization layers, in the latter case conditioning on the bitrate of BPG, in the hope that this would allow the network to distinguish different operation modes. However, we found that instance normalization is more sensitive to the resolution used for training, which led worse overall bitrates.

6. Results and Discussion

6.1. Compression performance in bpsp

We follow previous work in evaluating *bits per subpixel* (Each RGB pixel has 3 subpixels), *bpsp* for short, sometimes called *bits per dimension*. In Table 1, we show the performance of our approach on the described test sets. On Open Images, the domain where we train, we are outperforming all methods, including FLIF. Note that while L3C was trained on the same data set, it does not outperform FLIF. On the other data sets, we consistently outperform both L3C and the non-learned approaches PNG, WebP, and JPEG2000.

These results indicate that our simple approach of using a powerful lossy compressor to compress the high-level image content and leverage a complementary learned probabilistic model to model the low level variations for lossless residual compression is highly effective. Even though we only train on Open Images, our method can generalize to various domains of natural images: mobile phone pictures (CLIC.mobile), images retouched by professional photographers (CLIC.pro), as well as high-quality images with diverse complex structures (DIV2K).




Figure 4. *Top*: Distribution of bpsp, on the 500 images from Open Images validation set. The images are sorted by the bpsp achieved using our approach. We show PNG and FLIF, as well as the bpsp needed to store the lossy reconstruction only (“ x_l only”). *Bottom*: Fraction of total bits used by our approach that are used to store x_l . Images follow the same order as on the top panel.

In Fig. 4 we show the bpsp of each of the 500 images of Open Images, when compressed using our method, FLIF, and PNG. For our approach, we also show the bits used to store x_l for each image, measured in bpsp on top (“ x_l only”), and as a percentage on the bottom. The percentage averages at 42%, going up towards the high-bpsp end of the figure. This plot shows the wide range of bpsp covered by a random set of natural images, and motivates our Q-Classifer. We can also see that while our method tends to outperform FLIF on average, FLIF is better for some high-bpsp images, where the bpsp of both FLIF and our method approach that of PNG.

6.2. Runtime

We compare the decoding speed of RC to that of L3C for 512×512 images, using an NVidia Titan XP. For our components: BPG: 163ms; RC: 166ms; arithmetic coding: 89.1ms; i.e., in a total 418ms, compared to L3C’s 374ms.

QC and τ -optimization are only needed for encoding. We discussed above that both components are computationally cheap. In terms of actual runtime: QC: 6.48ms; τ -optimization: 35.2ms.

6.3. Q-Classifer and τ -Optimization

In Table 2 we show the benefits of using the Q-Classifer as well as the τ -optimization. We show the resulting bpsp for the Open Images validation set (top) and for DIV2K (bottom), as well as the percentage of predicted Q that are ± 1 away from the optimal Q' (denoted “ ± 1 to Q' ”), against a baseline of using a fixed $Q = 14$ (the mean over QC’s training set, see Section 5.2). The last column shows the required number of forward passes through RC.

Q-Classifer We first note that even though the QC was only trained on Open Images (see Sec 5.2), we get similar behavior on Open Images and DIV2K. Moreover, we see that using QC is clearly beneficial over using a fixed Q for all images, and only incurs a small increase in bpsp compared to using the optimal Q' (0.18% for Open Images, 0.26% for DIV2K). This can be explained by the fact that QC manages to predict Q within ± 1 of Q' for 94.8% of the images in Open Images and 90.2% of the DIV2K images.

Furthermore, the small increase in bpsp is traded for a reduction from requiring 7 forward passes to compute Q' to a single one. In that sense, using the QC is similar to the “fast” modes common in image compression algorithms, where speed is traded against bitrate.

τ -Optimization Table 2 shows that using τ -Optimization on top of QC reduces the bitrate on both testing sets.

Discussion While the gains of both components are small, their computational complexity is also very low (see Section 6.2). As such, we found it quite impressive to get the

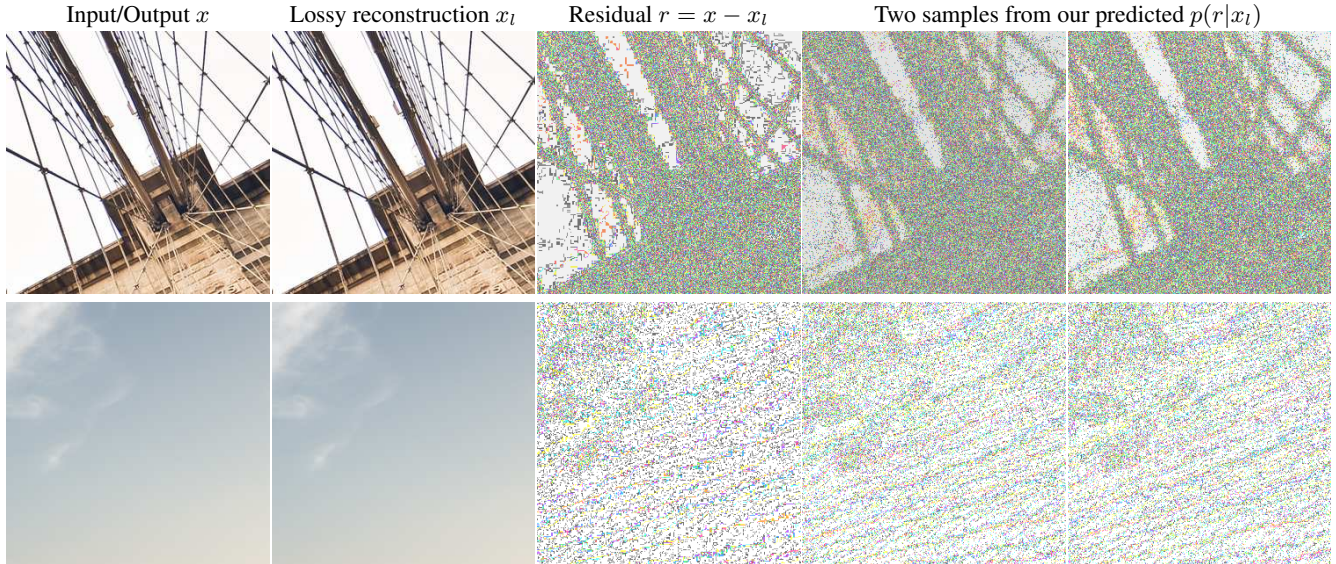


Figure 5. Visualizing the learned distribution $p(r|x_l)$ by sampling from it. We compare the samples to the ground-truth target residual r . We also show the image x that we losslessly compress as well as the lossy reconstruction x_l obtained from BPG. For easier visualizations, pixels in the residual images equal to 0 are set to white, instead of gray. *Best viewed on screen due to the high-frequency noise.*

reported gains. We believe the direction of tuning a handful of parameters post training on an instance basis is a very promising direction for image compression. One fruitful direction could be using dedicated architectures and including a tuning step end-to-end as in meta learning.

6.4. Visualizing the learned $p(r|x_l)$

While the bpsp results from the previous section validate the compression performance of our model, it is interesting to investigate the distribution predicted by RC. Note that we predict a mixture *distribution* per pixel, which is hard to visualize directly. Instead, we sample from the predicted

| Data set | Setup | bpsp | ± 1 to Q' | # forward |
|-------------|-----------------|-------|-----------------|---------------------|
| Open Images | Optimal Q' | 2.789 | 100% | $ \mathcal{Q} = 7$ |
| | Fixed $Q = 14$ | 2.801 | 82.6% | 1 |
| | Our QC | 2.794 | | 1 |
| | Our QC + τ | 2.790 | 94.8% | 1 |
| DIV2K | Optimal Q' | 3.080 | 100% | $ \mathcal{Q} = 7$ |
| | Fixed $Q = 14$ | 3.096 | 73.0% | 1 |
| | Our QC | 3.088 | | 1 |
| | Our QC + τ | 3.079 | 90.2% | 1 |

Table 2. On Open Images and DIV2K, we compare using the optimal Q' for encoding images, vs. a fixed $Q = 14$ and vs. using Q predicted by the Q-Classifer. For each data set, the last row shows the additional gains obtained from applying the τ -optimization. The forth column shows the percentage of predicted Q that are ± 1 away from the optimal Q' and the last column corresponds to the number of forward passes required for Q -optimization.

distribution. We expect the samples to be visually similar to the ground-truth residual $r = x - x_l$.

The sampling results are shown in Fig. 5, where we visualize two images from CLIC.pro with their lossy reconstructions, as obtained by BPG. We also show the ground-truth residuals r . Then, we show two samples obtained from the probability distribution $p(r|x_l)$ predicted by our RC network. For the top image, r is in $\{-9, \dots, 9\}$, for the bottom it is in $\{-5, \dots, 4\}$ (cf. Fig. 2), and we re-normalized r to the RGB range $\{0, \dots, 255\}$ for visualization, but to reduce eye strain we replaced the most frequent value (128, i.e., gray), with white.

We can clearly see that our approach i) learned to model the noise patterns discarded by BPG inherent with these images, ii) learned to correctly predict a zero residual where BPG manages to perfectly reconstruct, and iii) learned to predict structures similar to the ones in the ground-truth.

7. Conclusion

In this paper, we showed how to leverage BPG to achieve state-of-the-art results in full-resolution learned lossless image compression. Our approach outperforms L3C, PNG, WebP, and JPEG2000 consistently, and also outperforms the hand-crafted state-of-the-art FLIF on images from the Open Images data set. Future work should investigate input-dependent optimizations, which are also used by FLIF and which we started to explore here by optimizing the scale of the probabilistic model for the residual (τ -optimization). Similar approaches could also be applied to latent probability models of lossy image and video compression methods.

References

- [1] Workshop and Challenge on Learned Image Compression. <https://www.compression.cc/challenge/>. 6
- [2] 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*, 2017. 1
- [3] Eirikur Agustsson and Radu Timofte. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *CVPR Workshops*, 2017. 6
- [4] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative Adversarial Networks for Extreme Learned Image Compression. In *ICCV*, 2019. 1
- [5] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. Density modeling of images using a generalized normalization transformation. *arXiv preprint arXiv:1511.06281*, 2015. 4
- [6] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end Optimized Image Compression. *ICLR*, 2016. 1
- [7] Fabrice Bellard. BPG Image format. <https://bellard.org/bpg/>. 1
- [8] Lukas Cavigelli, Pascal Hager, and Luca Benini. Cas-cnn: A deep convolutional neural network for image compression artifact suppression. In *IJCNN*, pages 752–759, 2017. 3
- [9] Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. PixelSNAIL: An Improved Autoregressive Generative Model. In *ICML*, 2018. 2
- [10] Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012. 3
- [11] Peter Deutsch. DEFLATE compressed data format specification version 1.3. Technical report, 1996. 2
- [12] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *ICCV*, pages 576–584, 2015. 2
- [13] Leonardo Galteri, Lorenzo Seidenari, Marco Bertini, and Alberto Del Bimbo. Deep generative adversarial compression artifact removal. In *ICCV*, pages 4826–4835, 2017. 3
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 3
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4, 11
- [16] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural Networks for Machine Learning Lecture 6a Overview of mini-batch gradient descent. 6
- [17] Geoffrey Hinton and Drew Van Camp. Keeping neural networks simple by minimizing the description length of the weights. In *COLT*, 1993. 2
- [18] Emiel Hoogeboom, Jorn WT Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. In *NIPS*, 2019. 1, 2, 3
- [19] David A Huffman. A method for the construction of minimum-redundancy codes. *Proc. IRE*, 40(9):1098–1101, 1952. 3
- [20] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 2015. 4
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [22] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018. 5
- [23] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *NIPS*, 2016. 2
- [24] Friso H Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *ICML*, 2019. 1, 2, 3
- [25] Alexander Kolesnikov and Christoph H Lampert. PixelCNN Models with Auxiliary Variables for Natural Image Modeling. In *ICML*, 2017. 1, 2
- [26] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Shahab Kamali, Matteo Mallocci, Jordi Pont-Tuset, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. OpenImages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*, 2017. 6
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [28] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional Probability Models for Deep Image Compression. In *CVPR*, 2018. 1
- [29] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *CVPR*, 2019. 1, 2, 3, 4, 6
- [30] David Minnen, Johannes Ballé, and George D Toderici. Joint Autoregressive and Hierarchical Priors for Learned Image Compression. In *NeurIPS*. 2018. 1, 11
- [31] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Łukasz Kaiser, Noam Shazeer, and Alexander Ku. Image Transformer. *ICML*, 2018. 2
- [32] Pillow Library for Python. <https://python-pillow.org>. 6
- [33] Portable Network Graphics (PNG). <http://libpng.org/pub/png/libpng.html>. 2
- [34] Scott Reed, Aäron Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, and Nando Freitas. Parallel Multiscale Autoregressive Density Estimation. In *ICML*, 2017. 1, 2
- [35] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015. 2
- [36] Iain E Richardson. *H. 264 and MPEG-4 video compression: video coding for next-generation multimedia*. John Wiley & Sons, 2004. 2

- [37] Oren Rippel and Lubomir Bourdev. Real-Time Adaptive Image Compression. In *ICML*, 2017. 1
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 4
- [39] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. PixelCNN++: A PixelCNN Implementation with Discretized Logistic Mixture Likelihood and Other Modifications. In *ICLR*, 2017. 1, 2, 4
- [40] C. E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3):379–423, 1948. 1, 3
- [41] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The JPEG 2000 still image compression standard. *IEEE Signal Processing Magazine*, 18(5):36–58, 2001. 2
- [42] J. Sneyers and P. Wuille. FLIF: Free lossless image format based on MANIAC compression. In *ICIP*, 2016. 2
- [43] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12):1649–1668, 2012. 1, 3
- [44] Pavel Svoboda, Michal Hradis, David Barina, and Pavel Zemcik. Compression artifacts removal using convolutional neural networks. *arXiv preprint arXiv:1605.00366*, 2016. 2
- [45] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszar. Lossy Image Compression with Compressive Autoencoders. In *ICLR*, 2017. 1
- [46] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full Resolution Image Compression with Recurrent Neural Networks. In *CVPR*, 2017. 1
- [47] James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. In *ICLR*, 2019. 1, 2, 3
- [48] Michael Tschannen, Eirikur Agustsson, and Mario Lucic. Deep Generative Models for Distribution-Preserving Lossy Compression. In *NeurIPS*. 2018. 1
- [49] Aäron van den Oord, Nal Kalchbrenner, Lasse Espeholt, koray kavukcuoglu, Oriol Vinyals, and Alex Graves. Conditional Image Generation with PixelCNN Decoders. In *NIPS*, 2016. 1, 2, 3
- [50] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel Recurrent Neural Networks. In *ICML*, 2016. 1, 2
- [51] Gregory K Wallace. The JPEG still picture compression standard. *IEEE Transactions on Consumer Electronics*, 38(1):xviii–xxxiv, 1992. 1
- [52] WebP Image format. <https://developers.google.com/speed/webp/>. 2
- [53] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7):560–576, 2003. 1, 2
- [54] Ian H Witten, Radford M Neal, and John G Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, 1987. 3