

# Split Hierarchical Variational Compression

Tom Ryder\*, Chen Zhang\*<sup>†</sup>, Ning Kang<sup>†</sup>, Shifeng Zhang<sup>†</sup>  
<sup>†</sup>Huawei Noah's Ark Lab

{chenzhang10, kang.ning2, zhangshifeng4}@huawei.com

## Abstract

*Variational autoencoders (VAEs) have witnessed great success in performing the compression of image datasets. This success, made possible by the bits-back coding framework, has produced competitive compression performance across many benchmarks. However, despite this, VAE architectures are currently limited by a combination of coding practicalities and compression ratios. That is, not only do state-of-the-art methods, such as normalizing flows, often demonstrate out-performance, but the initial bits required in coding makes single and parallel image compression challenging. To remedy this, we introduce Split Hierarchical Variational Compression (SHVC). SHVC introduces two novelties. Firstly, we propose an efficient autoregressive prior, the autoregressive sub-pixel convolution, that allows a generalisation between per-pixel autoregressions and fully factorised probability models. Secondly, we define our coding framework, the autoregressive initial bits, that flexibly supports parallel coding and avoids – for the first time – many of the practicalities commonly associated with bits-back coding. In our experiments, we demonstrate SHVC is able to achieve state-of-the-art compression performance across full-resolution lossless image compression tasks, with up to 100x fewer model parameters than competing VAE approaches.*

## 1. Introduction

The volume of data, measured in terms of IP traffic, is currently witnessing an exponential year-on-year growth [9]. Fuelled by the demand for high-resolution media content, it is estimated that 80% of this data is in the form of images and video [9]. Data service providers, such as cloud and streaming platforms, have consequently seen costs associated with transmission and storage become prohibitively expensive. For example, an increased demand for streaming services forced major providers to throttle the maximum resolution of video content to 720p during the coronavirus pandemic. As such, these challenges have renewed the need for the development of high-performance data compression codecs.

\*co-first author. The work of Tom Ryder is conducted during his employment at Huawei Technologies R&D UK.

One solution to this problem has been the development of approaches using likelihood-based generative models capable of discrete density estimation [3, 6, 13, 14, 18, 22, 24, 35, 36, 42, 43]. Such methods operate by learning a deep probabilistic model of the data distribution, which, in combination with entropy coders, can be used to compress data. Here, according to Shannon's source coding theorem [21], the minimal required average codelength is bounded by the expected negative log-likelihood of the data distribution.

From this family of generative models, there have emerged three dominant modes for data compression: *normalizing flows* [3, 14, 42, 43], *variational autoencoders* [18, 24, 36] and *autoregressive models* [15, 31, 37]<sup>1</sup>. In fact, each of these approaches can be thought of as a traversal on the Pareto frontier of inference speed and compression performance. With broad generality, autoregressive models can often be the most powerful but the slowest; variational autoencoders are often the weakest but the fastest; and normalizing flows – depending on the variant – sit somewhere in between.

In this paper, we consider data compression with VAEs, and focus on extending the efficient frontier; obtaining solutions faster than popular VAEs that achieve state-of-the-art compression ratios. Use of VAEs, however, poses two outstanding challenges. Firstly, we should achieve competitive coding ratios without greatly sacrificing time complexity. For example, best iterates currently require one of two ingredients to improve performance: building either a deep hierarchy of latent variables [8] or use of autoregressive priors [11, 29]. The latter idea, especially popular in the codecs of the *lossy* compression community [25], posits a model that flexibly learns both local (via autoregression) and global (via hierarchical latent representation) data modalities (e.g. low-frequency information). Whilst these approaches, such as MS-PixelCNN [29] and PixelVAE [11], have had some success in achieving more efficient trade-offs, generation of even moderately sized images is still to the order of minutes [22].

Secondly, there should exist a practical means by which to efficiently perform single-image compression. Single-image compression then permits parallel coding, which is highly

<sup>1</sup>Most recently, score-based generative models have been adapted for data compression [16], but current approaches require an impractical number of operations at inference time.

desirable. However, translating a VAE into a lossless codec is currently achieved using the *bits-back* coding framework (predominantly, bits-back ANS), which requires a large number of *initial bits* [12, 36, 39] (see Section 3.1). Whilst this is a trivial number of bits on large image datasets (where we can amortize this cost), it renders bits-back an impractical approach for single-image compression. Furthermore, even large datasets are often coded such that images are interlinked. Access to a single image in the middle of a sequence would therefore require all prior images in the bitstream to be additionally decompressed.

To that end, we propose two novelties for use in VAE-based compression designed to address these challenges. The first, our *autoregressive sub-pixel convolution*, introduces a simple autoregressive factorisation – not dissimilar from the transformations used in normalizing flows [3, 42, 43] – designed to present an efficient interpolation between fully-factorised probability distributions and the impractical per-pixel autoregressions. Built from a modified space-to-depth convolution operator, we losslessly downsample data variables before performing a computationally efficient autoregression along the channel dimension. Our autoregressive operator is then advantaged by a number of network evaluations invariant to data dimensions, with each autoregression crucially performed on a downsampled version on the input tensor. More broadly, we view this framework as a generalisation of many popular autoregressive “context” models used in data compression [11, 26, 31, 41].

Our second contribution, *autoregressive initial bits*, presents a general framework for avoiding the impracticalities of bits-back ANS, allowing for eminently parallelizable coding. This technique, highly compatible with our autoregressive model, partitions the data variable into two *splits* such that the second partition is conditionally independent of the latent variable(s), given the first. In this way, we illustrate how we can use the entropy coding of the conditionally independent partition to both *supply* and *remove* the initial bits necessitated by bits-back ANS. We demonstrate that this approach reduces the bit overhead on a per-image basis by close to 20x.

Finally, we combine the above contributions to present our codec, *Split Hierarchical Variational Compression* (SHVC). SHVC posits a hierarchical VAE of general-form autoregressive priors that permits parallel coding. Using our framework, we outperform all other VAE-based compression approaches with fewer latent variables and a comparable number of neural network evaluations. We further illustrate the effectiveness of our architecture by training a small model which is able to outperform similar VAE approach Bit-Swap [18] – but with 100x fewer model parameters.

## 2. Related Work

**Compression with VAEs** can be separated into those *with* [18, 36] and those *without* [22] stochastic posterior sampling (the latter uses a discrete distribution of one symbol, assumed to have a probability of one). Whilst obtaining

theoretically superior compression ratios, approaches adopting stochastic posteriors, such as HiLLOc [36] and Bit-Swap [18], must entropy code using derivatives of the bits-back argument [12]. These approaches, under the umbrella of bits-back ANS (bb-ANS), require access to an *initial bitstream*. Although it is possible to amortize the cost of the initial bitstream across a large dataset, single and parallel data compression has presently proven challenging (see Section 3.1). In HiLLOc [36] the authors propose use of a conventional codec to compress and send parts of a dataset, which is then used as the basis for the initial bitstream. Whilst this partially solves some of the coding challenges, it requires the implementation of sub-optimal traditional techniques and still does not permit practical coding of a single image. In contrast, our approach avoids all of these challenges, requiring fewer latent variables, for a negligible additional bit cost. Similarly, approaches such as L3C [22] that leverage deterministic posteriors can also avoid the challenges associated with bb-ANS with use of arithmetic (or adaptive arithmetic) coding (AC) [40]. Closely related approach RC [24] – which can be loosely thought of as a VAE – uses a lossy compressed image as a de facto latent variable to condition the data distribution. However these techniques pay for their comparative practicality with a penalty in compression ratio as they must explicitly code the joint distribution of data and latent variables (see Section 3.2). Since this work focuses on VAEs based codecs, we refer readers to [3, 14, 42, 43] and references therein for compression with alternative deep generative models.

**Autoregressive Models** are a popular means to extend the independence assumptions of fully factorised models to high-dimensional multivariate densities. They are popular as both stand-alone models [31, 38] or in application with VAEs [11, 25, 29]. In their most computationally expensive forms, such as PixelCNN++ (and variants) [31], they posit a pixel-by-pixel autoregression, which then codes in raster-scan order. Whilst of broad academic interest, the  $\mathcal{O}(n^2)$  time complexity makes them prohibitive for application in data compression. One proposed solution to this problem has been to parametrise the priors of VAEs with autoregressive densities. Here probability estimation proceeds by combining hierarchical latent representations with decoded autoregressive context. Supplementing autoregressive components with auxiliary latent features permits causality restrictions that reduce time complexity, without greatly diminishing performance. These restrictions include channel-wise autoregressions [26]; independent, block-based models [29]; “checkerboard” context [41]; and small neural networks [11], amongst others. In fact, these restrictions are likely of dual purpose: combining powerful autoregressive models with VAEs will likely expedite posterior collapse (see Section 5.3) [5, 11, 20, 28]. Like these techniques, our approach combines a VAE with a restricted autoregressive model. Our method can be thought as most similar to [26] and [41]. Like the former, we perform a channel-wise autoregression, but do so after our autoregressive operator downsamples the data tensor. As such, our causality more closely re-

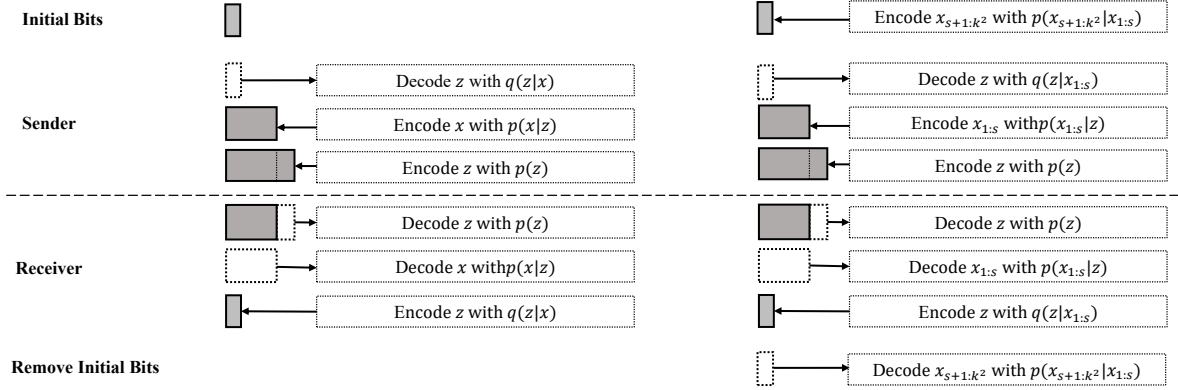


Figure 1. Left: The bb-ANS coding framework as discussed in Section 3.1. Right: The coding framework of ArIB as proposed in Section 4.2. Here we nest a bb-ANS coder inside of a block-based autoregressive structure that removes the need for initial bits.

sembles that of [41]. However, in contrast to the authors of [41], who enforce their causality with a binary mask, we do so using our sub-pixel convolution. This precipitates a greater degree of parallelism and presents the flexibility to efficiently recover a number of causal dependency schemes, such as PixelCNN++.

### 3. Background

Suppose access to a dataset of size  $n$ ,  $\{x_1, x_2, \dots, x_n\}$ , drawn from some intractable  $p_{\text{data}}(x)$  that we wish to compress. In order to achieve this, we introduce a discrete probability distribution  $p(x)$  that, in combination with *entropy coding*, requires a *codelength* of  $\sum_{i=1}^n -\log_2 p(x_i)$  bits to represent. Ideally,  $p(x)$  should closely resemble  $p_{\text{data}}(x)$ . In such a case, the average codelength in the limit of  $n \rightarrow \infty$  is given by  $\mathbb{E}_{p_{\text{data}}}[-\log_2 p(x)] \rightarrow H(x)$ , where  $H(x)$  is the entropy of the data. Here, the compression scheme is said to be *optimal* under Shannon’s source coding theorem [32].

#### 3.1. Bits-back ANS for VAEs

As discussed in Section 1, variational autoencoders (VAEs) are one popular approach to estimating  $p_{\text{data}}(x)$ , which defines a latent variable model such that

$$p(x) = \int p(x, z) dz = \int p(x|z)p(z) dz, \quad (1)$$

where  $p(z)$  is the *prior* distribution over latent variable  $z$ . As  $p(x)$  is normally intractable, VAEs introduce an approximate *posterior*  $q(z|x)$ , which is optimised to maximise a lower bound on the marginal evidence, the Evidence Lower Bound (ELBO)

$$\log p(x) \geq \mathcal{L} = \mathbb{E}_q \left[ -\log q(z|x) + \log p(x|z)p(z) \right], \quad (2)$$

where low-variance estimates of the expectation in (2) are via Monte-Carlo integration and the *reparametrization trick* [17].

Entropy coding requires explicit probabilities of data symbols. However, in VAEs, the model is factorised into a prior and a likelihood, and therefore does not allow direct coding of the data. To remedy this, several authors have proposed coding variants of bb-ANS [18, 35, 36]. This process is outlined as follows, which, without loss of generality, we describe for a model with a single latent variable. During compression, one decodes  $z$  from some *auxiliary* initial bits with  $q(z|x)$ ; encodes  $x$  with  $p(x|z)$ ; and encodes  $z$  with  $p(z)$  to obtain the complete bitstream. In the decompression stage, one decodes  $z$  from the bitstream with  $p(z)$ ; decodes  $x$  with  $p(x|z)$ ; encodes  $z$  with  $q(z|x)$  and thus returns the initial bits (hence *bits-back* coding). This technique is visualised in Fig. 1 Left.

The first decoding step common in existing bb-ANS codecs requires access to an *initial bitstream*. This requirement leads to several disadvantages when compared to AC. Firstly, while the initial bits are returned *after* decompression, the same bits are occupied and not readily readable *beforehand*. Secondly, although we can amortize the cost of the initial bits across a large dataset by chaining the compressed data, access to any given data point requires decompression of all data points posterior to the one of our target in the original data sequence. As such, compression of a *single* image carries substantial overhead – and, by extension, so too do parallel coding implementations.

#### 3.2. Deterministic vs Stochastic Posterior Sampling

Within VAEs based lossless compression, the impracticalities associated with bits-back coding are not the only choice. Indeed, approaches that use “deterministic” posterior sampling [22] may eschew bb-ANS for AC. This approach is almost ubiquitous in the lossless codecs of lossy approaches [2, 7, 25, 26], where eminently parallelizable, low-latency codecs are especially preferable (e.g. streaming media). We note that when using deterministic posterior sampling, the likelihoods associated with the prior in Eq. (2) trivially cancel to zero, such that the objective resolves to maximum likelihood estimation of the




Figure 2. Average BPD savings against the number of latent variables ( $L = 1, \dots, 5$ ) used in three VAE lossless compression models. The baseline is the deterministic posterior with  $L = 1$  where the BPD saving is zero. Results are presented as an average for the three approaches trained across CIFAR10, ImageNet32 and ImageNet64.

joint distribution (see e.g. [2]).

Whilst gaining notable practical coding advantages, in sacrificing a stochastic posterior one also sacrifices the ability to minimize the cost of sending latent variables. To offset this limitation, models will repeatedly downsample (RDS) the number of symbols available to latent representations between each layer. Whilst this limits posterior expressiveness, there is little experimental evidence to support how much this matters in practice. In addition, models with stochastic posteriors require large  $L$  to excel (where  $L$  is the number of latent variables), which hinders run-time.

To that end, we display the results of a simple experiment in Fig. 2 designed to investigate this difference further. Here we train three VAEs across CIFAR10, ImageNet32 and ImageNet64: two with stochastic posteriors (one with and one without RDS) and a deterministic posterior (with RDS). The architectures in each model are identical (with the exception of downsampling operations) and we quantify compression performance in *bits per dimension* (BPD). Further experimental details can be found in the Appendix. Here we observe that, even with  $L \leq 3$ , both stochastic posteriors outperform by  $\sim 5\%$ . This difference grows as  $L$  increases – but it does not matter if RDS is used (at least, for  $L \leq 5$ ). This result is of important consequence: the best current approach to avoid the impracticalities of bb-ANS (ie. using a deterministic posterior) carries a 5% BPD penalty. For single-image compression with stochastic posteriors, the initial bits required would typically be much larger than this. Likewise, unless extending to a deep hierarchy of latent variables, RDS seems like a compute-efficient choice that does not limit performance.

## 4. Method

Our method posits a hierarchical VAE where we parameterize the priors using an autoregressive factorisation. We begin by defining a lossless downsampling convolution operator, before

describing its application to density estimation using both *weak* and *strong* autoregressive models. We then describe how this autoregressive structure can be leveraged to avoid many of the challenges associated with bb-ANS *without* sacrificing the performance of stochastic posteriors. Finally, we describe how these contributions can be combined to form our SHVC codec.

### 4.1. Autoregressive Sub-Pixel Convolutions

The space-to-depth and depth-to-space transformations are popular operations across image analysis, from generative modelling [3, 14] to super-resolution [33]. They define adjacent operations for efficient up and downsampling transformations by folding spatial dimensions into channel dimensions – and vice versa. Unlike learned operations, they greatly reduce computational complexity, allowing for greater parallelism by losslessly moving computation (and data) into the channels. Indeed, these operations have become an essential component in papers seeking real-time execution (e.g. [10, 19, 30, 33]). Specifically, given a tensor of  $C$  channels,  $H$  height and  $W$  width, we define the space-to-depth and depth-to-space transformations,  $f$  and  $f^{-1}$ , such that

$$f: \mathbb{R}^{C \times H \times W} \longrightarrow \mathbb{R}^{Ck^2 \times \frac{H}{k} \times \frac{W}{k}}, \quad (3)$$

$$f^{-1}: \mathbb{R}^{C \times H \times W} \longrightarrow \mathbb{R}^{\frac{C}{k^2} \times Hk \times Wk}, \quad (4)$$

where  $k$  is the scale factor.

As described in [33], these operations can be efficiently performed using *sub-pixel convolutions*, which are referred to as *pixel unshuffle* and *pixel shuffle*. In particular, their space-to-depth transformation, pixel unshuffle, is performed using a  $k$ -stride depthwise convolution where the  $n^{\text{th}}$  element of  $Ck^2$   $k \times k$  filters has one non-zero element such that

$$K_{h,w}^{(n)} = \begin{cases} 1 & \text{if } h = \lfloor n/k \rfloor \bmod k, w = n \bmod k \\ 0 & \text{else} \end{cases}, \quad (5)$$

where  $h, w$  are the indices over spatial dimensions. The result of this operation is visualised in Fig. 3 Centre.

Defining a channel-wise autoregression over the resulting tensor would posit a checkerboard autoregressive structure over each of the channels in the original tensor, sequentially. However, as identified in PixelCNN++ [31], sub-pixels in adjacent channels, sharing the same *spatial* location in the original tensor, have high correlation and therefore do not require complex models to describe the dependency structure. As such, the authors of PixelCNN++ use a linear model predicted by a single network evaluation, conditioned on decoded *context*, to define the joint distribution across channels. In this way, they obfuscate the need for separate RGB network evaluations. (We note that in our setting, context refers to previously decoded pixels in either the current or previous hierarchical latent variable.) From henceforth what we refer to as a *weak autoregression*, is then defined similarly to [31] according to






Figure 3. Left: a  $3 \times 4 \times 4$  input RGB image. Centre: the image  $x$  downsampled using the convolution operator of (5) with  $k = 2$ . Right: the image  $x$  downsampled using the convolution operator of (8) with  $k = 2$ .

$$p(x_{0:C,h,w}|D) = p(x_{0,h,w}|D) \prod_{c=1}^C p(x_{c,h,w}|x_{<c,h,w}, D) \quad (6)$$

where  $D$  is the decoded context and  $p$  is some parametric probability mass function (pmf), obtained via integrating a probability density function (pdf) over discretization bins, with mean at channel  $c$  location  $h,w$  given by

$$\mu_{c,h,w} = \alpha_{c,h,w} + \sum_{i=0}^{c-1} \beta_{c,h,w}^{(i)} x_{i,h,w}. \quad (7)$$

Here  $\alpha$  and  $\beta$  are scalars predicted for all channels and spatial locations by a single network evaluated on decoded context, and  $i$  is the index over channels in decoded context such that  $\beta_{c,h,w}^{(i)}$  is the scalar for prediction of the mean associated with pixel at channel  $i$ , spatial location  $h,w$ .

Inspired by this, we introduce a new space-to-depth convolution such that the resulting autoregression is alternatively re-ordered into  $k^2$  sub-blocks of  $C$  channels each. Crucially, the resulting channels in each sub-block share the same spatial index allowing application of the autoregression detailed in (6) and (7). We note that should  $k = H = W$  we return an equivalence to the per-pixel autoregression of PixelCNN++ but perform an autoregression exclusively in the channel dimension. Likewise should  $k < H$ , we define a block-based context model in raster scan order where, unlike MS-PixelCNN, adjacent blocks are dependent.

To achieve our desired downsampling operation, which we denote by  $g(\cdot)$ , we expand the depthwise convolutions of (5) into regular three-dimensional kernels where the  $n^{\text{th}}$  of  $Ck^2$   $C \times k \times k$  filters has one non-zero element such that

$$K_{c,h,w}^{(n)} = \begin{cases} 1 & \text{if } c = n \bmod C, h = \lfloor n/Ck \rfloor \bmod k, \\ & w = \lfloor n/C \rfloor \bmod k \\ 0 & \text{else} \end{cases}. \quad (8)$$

We further visualise this operation in Fig. 3 Right. The resulting density of the downsampled tensor for spatial location  $h,w$  is

then given by

$$p(g(x;k)_{0:Ck^2,h,w}|D) = \prod_{i=0}^{k^2} \left[ p(g(x;k)_{iC,h,w}|D) \times \prod_{j=iC+1}^{(i+1)C-1} p(g(x;k)_{j,h,w}|g(x;k)_{<j,h,w}, D) \right], \quad (9)$$

where  $i$  is the index over sub-blocks (i.e. a strong autoregression evaluated using neural networks).

**Masked 3D Convolutions** Whilst we are restricted to  $k^2$  evaluations per latent variable at inference time, the same does not have to be true during training. One efficient parallel training scheme is to use 3-dimensional convolutions applied to the downsampled tensors by expanding them into a  $d \times Ck^2 \times H \times W$  volume [23, 27], where  $d$  is some auxiliary dimension. Here we can apply zero-masking along the channel dimension of the kernels to enforce the causality condition, along with  $k$ -stride channel convolutions on the input. Full details are available in the Appendix.

**Choice of Distribution** For our choices of  $p$  and  $q$ , we use a discretized mixture of logistic distributions for  $x$  and a discretized univariate logistic distribution for all  $z^{(l)}$  [31]. That is, given some mean  $\mu$ , scale  $s$  and uniform discretization bin-width  $b$ , one can obtain the univariate pmf by integrating the logistic pdf over the discretization bin. For  $x$ , we typically use a mixture of 5 discrete logistic distributions as defined above.

## 4.2. Autoregressive Initial Bits

As discussed in Section 3.1, bb-ANS is able to achieve efficient codelengths, but can lead to several shortcomings. Fortunately, our proposed autoregressive model naturally accommodates the possibility to bypass the auxiliary bits needed in other bb-ANS methods. We achieve this by exploiting the block-based autoregressive structure on the data variable. We outline this process below, which we refer to as *autoregressive initial bits* (ARiB). Different from the models considered in

existing VAE-based codecs, we remove the direct causality between the latent variable  $z$  and some partition of the data variable  $x$ . In practice, we simply remove  $z$  from  $D$  in Eq. (9) for the final  $n$  sub-blocks in  $p(x|z)$ , along with the partition from  $x$  in  $q(z^{(1)}|x)$ . As a result, we factorise the likelihood as  $p(x|z) = p(x_{s+1:k^2}|x_{1:s})p(x_{1:s}|z)$  with the approximate posterior as  $q(z|x) = q(z|x_{1:s})$ , where  $s$  is our ‘split’ index. Instead of conducting the first step by decoding  $z$  from  $q(z|x)$ , one can encode  $x_{s+1:k^2}$  with  $p(x_{s+1:k^2}|x_{1:s})$  and thus obtain the bitstream from which to decode  $z$ . Then one decodes  $z$  with  $q(z|x_{1:s})$ , encodes  $x_{1:s}$  with  $p(x_{1:s}|z)$  and encodes  $z$  with  $p(z)$ . At the decompression stage, one decodes  $z$  with  $p(z)$ , decodes  $x_{1:s}$  with  $p(x_{1:s}|z)$ , encodes  $z$  with  $q(z|x_{1:s})$  and decodes  $x_{s+1:k^2}$  with  $p(x_{s+1:k^2}|x_{1:s})$ . We illustrate this technique in Fig. 1 Right.

For this approach to be valid, we require the satisfaction of two criteria:

1. There exists some  $s$ ,  $k$  and  $z$  such that imposing  $(x_{s+1:k^2} \perp z|x_{1:s})$  does not greatly hinder performance.
2. The *entropy* of  $p(x_{s+1:k^2}|x_{1:s})$  and  $q(\hat{z}|x_{1:s})$ , where  $\hat{z}$  is the discretized analogue of  $z$ , should be such that  $\mathcal{H}_{p(x_{s+1:k^2}|x_{1:s})} \geq \mathcal{H}_{q(\hat{z}|x_{1:s})}$ .

In our experiments, we demonstrate that the performance costs associated with criteria one are negligible. Crucially, we demonstrate that it is both orders of magnitude less than initial bits required of vanilla bb-ANS and a parameterisation of our approach using deterministic posteriors.

For criteria two, we formulate the optimization of (2) as a constrained problem subject to  $\mathcal{H}_{p(x_{s+1:k^2}|x_{1:s})} \geq \mathcal{H}_{q(\hat{z}|x_{1:s})}$ , where we estimate the respective expectations during training using Monte-Carlo integration. Whilst a variety of techniques from optimization theory may be applied, we found it sufficient to simply penalise (2) according to

$$\mathcal{L}_{pen} = \mathcal{L} + \lambda \max\left(0, \mathcal{H}_{q(\hat{z}|x_{1:s})} - \mathcal{H}_{p(x_{s+1:k^2}|x_{1:s})}\right), \quad (10)$$

where  $\lambda$  is some Lagrange multiplier. We find that this further presents flexibility when choosing  $s$ , with a variety of choices yielding the same result.

### 4.3. Split Hierarchical Variational Compression

SHVC formulates a hierarchical VAE built from the components described above. Here we partition the latent variable into a simple disjoint hierarchy of  $L$  layers, such that  $z = \{z^{(1)}, \dots, z^{(L)}\}$ . We define the prior and posterior according to

$$p(x, z^{(1:L)}) = p(x|z^{(1)})p(z^{(L)}) \prod_{i=1}^{L-1} p(z^{(i)}|z^{(i+1)}), \quad (11)$$

$$q(z^{(1:L)}|x) = q(z^{(1)}|x) \prod_{i=1}^{L-1} q(z^{(i+1)}|z^{(i)}), \quad (12)$$

where we parameterise every conditional density in Eq. (11) as per Eq. (9). While this factorisation naturally fits the coding scheme proposed in Bit-Swap [18], we additionally introduce a local *reverse* encoding to accommodate the autoregressive structure for factors in Eq. (11). In more detail, for the encoding of  $z^{(i)} = [z_1^{(i)}, \dots, z_{k^2}^{(i)}]$  with  $p(z^{(i)}|z^{(i+1)})$ , one needs to encode in the reserved order of  $z_{k^2}^{(i)}, \dots, z_1^{(i)}$ , to accommodate the first-in-last-out nature of ANS based codecs.

For purposes of experimentation, we define two versions of our model: one with and one without the dependency structure permitting ArIB. From henceforth we shall refer to these models as SHVC and SHVC-ArIB, respectively. For SHVC, one can encode  $x$  with  $p(x|z^{(1)})$ , along with other variables in Eq. (11) as discussed above. For SHVC-ArIB, one performs encoding and decoding for  $x$  as discussed in Section 4.2; and applies local reverse encoding for slices in  $x_{s+1:k^2}$  and  $x_{1:s}$ , respectively. We note that the only difference in SHVC-ArIB is that, whilst  $p(x_{s+1:k^2}|x_{1:s})$  and  $p(x_{1:s}|z^{(1)})$  are both modelled using (9), the former evidently omits  $z^{(1)}$  from  $D$ . In addition, we restrict the posterior such that  $q(z^{(1)}|x_{1:s})$ . We visualise the overall architecture along with the coding scheme for SHVC-ArIB in the Appendix.

## 5. Experiments

In this Section, we perform a series of experiments to evaluate the effectiveness of compression with SHVC. We begin by discussing the architecture and training details in Section 5.1. In Section 5.2, we evaluate compression performance in terms on both low and full-resolution images. Here we additionally evaluate the effect of the ArIB constraints on compression performance and assess inference speed. Finally, in 5.3 we perform a series of ablation studies.

### 5.1. Architecture and Training Details

As discussed in Section 3, we observe an efficient trade-off between inference speed and performance where we repeatedly downsample and limit the number of latent variable layers. Inspired by this, we apply repeated downsampling, using three and four latent variables when training  $32 \times 32$  and  $64 \times 64$  images, respectively. With these settings, we require 3 and 16 network evaluations for posterior and prior inference on CIFAR10. In contrast, leading VAE approach HiLLOc [36] requires 24 prior *and* posterior evaluations.

For all of our experiments, we set  $k=2$  in our autoregressive model. Whilst this choice is discussed more extensively in Section 5.3, we found the  $k=2$  presented an effective compromise between inference speed and representational power. Indeed, increasing  $k$  can often result in *worse* performance. We observe that increasing  $k$  increases the possibility of posterior collapse [5, 11, 20, 28], making training a hierarchy of powerful autoregressive models challenging. We further detail comprehensive architectures for all of our experiments in the Appendix.

	Compression Model	CIFAR10	ImageNet32	ImageNet64	CLIC.mobile	CLIC.pro	DIV2K
<i>Generic</i>	PNG [4]	5.71	5.87	6.39	3.90	4.00	3.09
	FLIF [34]	4.19	4.19	4.52	2.49	2.78	2.91
	JPEG-XL [1]	5.74	5.89	6.39	2.36	2.63	2.79
<i>VAE-Based</i>	L3C [22]	-	4.76	4.42	2.64	2.94	3.09
	Bit-Swap [18]	3.82	4.50	-	-	-	-
	HiLLoC [36]	3.56 <sup>‡</sup>	4.20 <sup>‡</sup>	3.90 <sup>‡</sup>	-	-	-
	<b>SHVC</b>	<b>3.16/3.41<sup>‡</sup></b>	<b>3.98</b>	<b>3.68/3.71<sup>‡</sup></b>	<b>1.96*</b>	<b>2.02*</b>	<b>2.57*</b>
	<b>SHVC Lite</b>	<b>3.76</b>	<b>4.49</b>	<b>4.16</b>	-	-	-
<i>Flow-Based</i>	IDF [14]	3.34/3.60 <sup>‡</sup>	4.18	3.90/3.94 <sup>‡</sup>	-	-	-
	IDF++ [3]	3.26	4.12	3.81	-	-	-
	LBB [13]	3.12	3.88	3.70	-	-	-
	iVPF [43]	3.20/ 3.49 <sup>‡</sup>	4.03	3.75/3.79 <sup>‡</sup>	2.39*	2.54*	2.68*
	iFlow [42]	3.12/3.36 <sup>‡</sup>	3.88	3.70/3.65 <sup>‡</sup>	2.26*	2.44*	2.57*

Table 1. Compression results in BPD for SHVC and other popular codecs across three low-resolution datasets and three full-resolution datasets. Lower is better. Here <sup>‡</sup> and \* denote models trained on ImageNet32 and ImageNet64, respectively.

## 5.2. Compression Performance

**Low-Resolution Images** We begin by testing our method on three toy datasets: CIFAR10, Imagenet32 and Imagenet64. We compare our method against leading approaches from traditional codecs, normalizing flows and VAEs. (Given the time complexity associated with per-pixel autoregressive factorisations, we follow the broader compression community and eschew them from our comparisons). As an ad hoc test of generalizability, we additionally follow the authors of [14, 36, 42, 43] by training a model on Imagenet32 and testing it across all other datasets. We present our results in terms of bits per dimension (BPD), which we display in Table 1.

**Full-Resolution Images** We further compare our approach against full-resolution algorithms, i.e., L3C [22]. Here we follow [43] and adopt our Imagenet64 model using a patch-based evaluation protocol in which images are cropped to  $64 \times 64$ . From Table 1, we demonstrate reliable out-performance of every other considered codec. Further, when comparing the performance of SHVC to e.g. iFlow across small and full-resolution images we interestingly note that difference in out-performance becomes greater. Indeed, we hypothesise that autoregressive "context" becomes an increasingly important inductive bias as the resolution of the image increases – and likewise the extent of the spatial redundancy.

**ArIB** ArIB poses non-trivial constraints on the data variable, such that a sufficiently large partition should be: 1. conditionally independent of latent variables; 2. encoded with an entropy larger than that with which the first latent variable is decoded with. Using the objective of Eq. (10), we train models capable of single-image compression: SHVC, SHVC-ArIB and a

parameterisation of SHVC with deterministic posteriors using AC (henceforth known as *Deterministic SHVC*). We display results across CIFAR10, Imagenet32 and Imagenet64 in Fig. 4, and quantify our results as additional bits on a per-image basis against SHVC. In more detail, the overhead in SHVC comes from initial bits, whilst other models incur a performance cost. Here we see that our ArIB adds minimal additional bits – less than the number in SHVC by a factor of  $\sim 20$ . Crucially, our approach also outperforms Deterministic SHVC, which we hope may serve as motivation to adapt the codec to lossy compressors.

	SHVC	Bit-Swap	HiLLoC	IDF
BPD	<b>3.18</b>	3.82	3.32	3.34
Time (s)	<b>4.63</b>	5.86	10.20	20.58

Table 2. Inference time and BPD across CIFAR10. Inference time measures evaluation of 10,000 test images with a batch size of 100.

**Inference Speed** To better evaluate our approach in the context of VAEs, we compare SHVC against popular publicly available compression models, HiLLoC<sup>2</sup> [36], Bit-Swap [18] and Integer Discrete Flow (IDF) [14]<sup>3</sup>. We report achieved BPDs in Table 2 and measure inference time in seconds (s) to evaluate the 10,000 CIFAR10 test images with a batch size of 100. Here we observe that our model is faster than HiLLoC for a lower BPD and achieves lower BPD for the same speed as Bit-Swap. We further note that, unlike both considered approaches, we are able to easily achieve parallel coding with minimal overhead.

<sup>2</sup>In contrast to SHVC and Bit-Swap, HiLLoC is implemented in TensorFlow. As such, HiLLoC is likely advantaged in terms of run-time environment.

<sup>3</sup>We note follow-up paper IDF++ [3] has no publicly available code but uses largely the same architecture as IDF. As such IDF++ inference times can be loosely inferred from IDF speeds.




Figure 4. Additional bits per single image compressed across SHVC, SHVC-ArIB and Deterministic SHVC. The overhead in SHVC comes from initial bits, whilst other models incur a performance cost.




Figure 5. The theoretically minimal BPD against iteration, comparing the common space-to-depth operation vs ours.

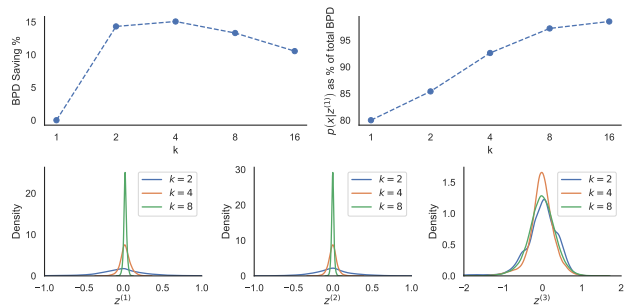


Figure 6. Top: (left) BPD savings against the choice of  $k$  (the baseline is  $k = 1$ ), and (right) the reconstruction loss as a percent of the total BPD. Bottom: marginal densities of latent variables for different  $k$ .

### 5.3. Ablation Studies

**Space-to-depth operation** One alternative in our approach is to replace our downsampling operator of (8) with the usual space-to-depth transformation described in Eq. (5). However, as discussed, this would instead define a weak autoregressive property over spatially adjacent pixels, channel-by-channel. In Fig. 5 we demonstrate the differences resulting from the choice of spatial downsampling transformation by training two models on CIFAR10. Here we observe that our convolutional operator provides non-trivial benefits over the vanilla space-to-depth transformation.

**Choice of  $k$**  As discussed, one important hyper-parameter choice is that of  $k$ . As  $k$  grows, the prior becomes more powerful but the time complexity grows. In Fig. 6 we visualise the effect of increasing  $k$  on the compression ratio, which is displayed as an average across three models trained on CIFAR10, ImageNet32 and ImageNet64. As discussed, we note that performance of the model peaks at  $k=4$ , before becoming worse as  $k$  increases. This non-intuitive behaviour can be better explained further in Fig. 6, where we see evidence of a posterior collapse common in hierarchical VAEs – especially those with hierarchical autoregressive priors [5, 11, 20, 28].

## 6. Conclusion

We have proposed and evaluated a new VAE model for data compression. SHVC is able to outperform existing VAE approaches in terms of both speed and compression ratios. Additionally, unlike competing approaches using bb-ANS, SHVC-ArIB is able to support parallel coding with minimal overhead. As such, we believe ArIB represents the most promising means to convert a VAE into a lossless codec. Motivated by this, one promising area of future work could consider the application of SHVC to lossy compression, which has traditionally ignored bits-back schemes.

Whilst our approach does not introduce any new negative societal consequences, we remain acutely aware of the issues precipitated by this area of research. These concerns are two-fold. Firstly, improving the capacity to store data could greatly increase the amount of non-essential personal data held by third-parties. This may be viewed as a challenge on broader digital liberties. Secondly, any class of generative model trained on sensitive data will learn to closely approximate the distribution of that data. As such, the model itself might extend to malignant use-cases beyond its intended purpose., e.g. classifiers.



## References

- [1] Jyrki Alakuijala, Ruud van Asseldonk, Sami Boukourt, Martin Bruse, Iulia-Maria Comşa, Moritz Firsching, Thomas Fischbacher, Evgenii Kliuchnikov, Sebastian Gomez, Robert Obryk, et al. Jpeg xl next-generation image compression architecture and coding tools. In *Applications of Digital Image Processing XLII*, volume 11137, page 111370K. International Society for Optics and Photonics, 2019. 7
- [2] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018. 3, 4
- [3] 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 *International Conference on Learning Representations*, 2021. 1, 2, 4, 7
- [4] Thomas Boutell and T Lane. Png (portable network graphics) specification version 1.0. *Network Working Group*, pages 1–102, 1997. 7
- [5] Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. 2016. 2, 6, 8
- [6] Sheng Cao, Chao-Yuan Wu, and Philipp Krähenbühl. Lossless image compression through super-resolution. *arXiv preprint arXiv:2004.02872*, 2020. 1
- [7] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [8] Rewon Child. Very deep VAEs generalize autoregressive models and can outperform them on images. In *International Conference on Learning Representations*, 2021. 1
- [9] Cisco. Cisco visual networking index: global mobile data traffic forecast update, 2018–2023. *Update*, 2018:2023, 2020. 1
- [10] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. Salsanext: fast, uncertainty-aware semantic segmentation of lidar point clouds for autonomous driving. *arXiv preprint arXiv:2003.03653*, 2020. 4
- [11] Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville. Pixelvae: A latent variable model for natural images. *International Conference on Learning Representations*, 2017. 1, 2, 6, 8
- [12] GE Hinton and Drew van Camp. Keeping neural networks simple by minimising the description length of weights. 1993. In *Proceedings of COLT-93*, pages 5–13. 2
- [13] Jonathan Ho, Evan Lohn, and Pieter Abbeel. Compression with flows via local bits-back coding. In *Advances in Neural Information Processing Systems*, pages 3879–3888, 2019. 1, 7
- [14] Emiel Hoogeboom, Jorn Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. In *Advances in Neural Information Processing Systems*, pages 12134–12144, 2019. 1, 2, 4, 7
- [15] Heewoo Jun, Rewon Child, Mark Chen, John Schulman, Aditya Ramesh, Alec Radford, and Ilya Sutskever. Distribution augmentation for generative modeling. In *International Conference on Machine Learning*, pages 5006–5019. PMLR, 2020. 1
- [16] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in Neural Information Processing Systems*, 2021. 1
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2014. 3
- [18] Friso H Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. *International Conference on Machine Learning*, 2019. 1, 2, 3, 6, 7
- [19] Haojie Liu, Tong Chen, Qiu Shen, Tao Yue, and Zhan Ma. Deep image compression via end-to-end learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018. 4
- [20] James Lucas, George Tucker, Roger B Grosse, and Mohammad Norouzi. Don’t blame the elbow! a linear vae perspective on posterior collapse. *Advances in Neural Information Processing Systems*, 32:9408–9418, 2019. 2, 6, 8
- [21] David JC MacKay and David JC Mac Kay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003. 1
- [22] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10629–10638, 2019. 1, 2, 3, 7
- [23] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional probability models for deep image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [24] Fabian Mentzer, Luc Van Gool, and Michael Tschannen. Learning better lossless compression using lossy compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6638–6647, 2020. 1, 2
- [25] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, 2018. 1, 2, 3
- [26] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *IEEE International Conference on Image Processing, ICIP*. 2, 3
- [27] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016. 5
- [28] Ali Razavi, Aaron van den Oord, Ben Poole, and Oriol Vinyals. Preventing posterior collapse with delta-VAEs. In *International Conference on Learning Representations*, 2019. 2, 6, 8
- [29] Scott Reed, Aäron van den Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov, and Nando de Freitas. Parallel multiscale autoregressive density estimation. In *Proceedings of the 34th International Conference on Machine Learning*, 2017. 1, 2
- [30] Oren Rippel, Alexander G. Anderson, Kedar Tatwawadi, Sanjay Nair, Craig Lytle, and Lubomir Bourdev. Elf-vc: Efficient

- learned flexible-rate video coding. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [4](#)
- [31] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *International Conference on Learning Representations*, 2017. [1](#), [2](#), [4](#), [5](#)
- [32] C. E. Shannon. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 2001. [3](#)
- [33] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016. [4](#)
- [34] Jon Sneyers and Pieter Wuille. Flif: Free lossless image format based on maniac compression. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 66–70. IEEE, 2016. [7](#)
- [35] James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. *International Conference on Learning Representations*, 2019. [1](#), [3](#)
- [36] James Townsend, Thomas Bird, Julius Kunze, and David Barber. Hilloc: Lossless image compression with hierarchical latent variable models. *International Conference on Learning Representations*, 2020. [1](#), [2](#), [3](#), [6](#), [7](#)
- [37] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in Neural Information Processing Systems*, pages 4790–4798, 2016. [1](#)
- [38] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, 2016. [2](#)
- [39] Chris S Wallace. Classification by minimum-message-length inference. In *International Conference on Computing and Information*, pages 72–81. Springer, 1990. [2](#)
- [40] Ian H. Witten, Radford M. Neal, and John G. Cleary. Arithmetic coding for data compression. *Commun. ACM*, 1987. [2](#)
- [41] Honglei Zhang, Francesco Cricri, Hamed R. Tavakoli, Nannan Zou, Emre Aksu, and Miska M. Hannuksela. Lossless image compression using a multi-scale progressive statistical model. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2020. [2](#), [3](#)
- [42] 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*, 2021. [1](#), [2](#), [7](#)
- [43] Shifeng Zhang, Chen Zhang, Ning Kang, and Zhenguo Li. ivpf: Numerical invertible volume preserving flow for efficient lossless compression. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. [1](#), [2](#), [7](#)