
Image and Video Tokenization with Binary Spherical Quantization

Yue Zhao

UT Austin

yzhao@cs.utexas.edu

Yuanjun Xiong*

MThreads AI

bitxiong@gmail.com

Philipp Krähenbühl

UT Austin

philkr@cs.utexas.edu

Abstract

We propose a new transformer-based image and video tokenizer with Binary Spherical Quantization (BSQ). BSQ projects the high-dimensional visual embedding to a lower-dimensional hypersphere and then applies binary quantization. BSQ is (1) parameter-efficient without an explicit codebook, (2) scalable to arbitrary token dimensions, and (3) compact: compressing visual data by up to 100 \times with minimal distortion. Our tokenizer uses a transformer encoder and decoder with simple block-wise causal masking to support variable-length videos as input. The resulting BSQ-ViT achieves state-of-the-art visual reconstruction quality on image and video reconstruction benchmarks with 2.4 \times throughput compared to the best prior methods. Furthermore, by learning an autoregressive prior for adaptive arithmetic coding, BSQ-ViT achieves comparable results on video compression with state-of-the-art video compression standards. BSQ-ViT also enables masked language models to achieve competitive image synthesis quality to GAN- and diffusion-based methods.

1 Introduction

Learned discrete image and video tokenization allows for state-of-the-art visual compression [1, 2, 3], recognition [4, 5, 6, 7] and generation [8, 9, 10]. These models follow a proven recipe from large language modeling [11, 12, 13]: Tokenize input and outputs into discrete units and learn an autoregressive model to predict this tokenized stream one token at a time. The most widely used approach for image encoding is Vector-Quantized Variational Auto-Encoder (VQ-VAE) [8]. They encode inputs in continuous latent embeddings and map them to a learned codebook through nearest-neighbor lookup. However, VQ-VAE style approaches have two drawbacks: First, most image encoders are built upon convolutional networks (CNN) [9, 14]. Adapting spatial convolution for images to spatial-temporal convolution for videos requires non-trivial architectural changes [15, 16, 17] with increased computational cost. Treating videos as a sequence of images leads to a suboptimal quantization [16]. Second, vector quantization (VQ) scales poorly with the codebook size. The runtime scales linearly with the codebook size, and the codebook easily overfits on smaller datasets [17]. This is especially troubling for video inputs, as they rely on larger codebooks to represent both static visual patterns and dynamic motion patterns.

This paper proposes a unified visual tokenizer based on a Vision Transformer and Binary Spherical Quantization (BSQ). The Transformer-based encoder-decoder leverages a block-wise causal mask and uses only visual tokens from the current or past timestamps for reconstruction (Figure 3). BSQ first projects the high-dimensional visual embedding of the transformer encoder to a lower-dimensional hypersphere and then applies binary quantization. The transformer encoder, decoder, and BSQ are seamlessly integrated into the VQ-GAN [9] framework and trained end-to-end.

*Now at Predera.ai

Our proposed visual tokenizer features several advantages. First, the Transformer-based encoder-decoder shows a Pareto improvement in visual reconstruction quality and computational efficiency compared to standard CNNs. Second, the block-wise causal design unifies images and videos as input at training and supports variable-length videos at inference. BSQ constructs an implicit codebook whose effective vocabulary grows exponentially with the spherical dimension with no learned parameters. The increasing codebook size consistently yields better reconstruction results. Compared to Lookup-free Quantization (LFQ) [17], a recent technique that also builds an implicit codebook based on scalar quantization (SQ), BSQ has a bounded quantization error and is easier to train. Furthermore, we show that the soft quantization probability in BSQ reduces to a simple product of multiple channel-independent Bernoulli distributions, leading to efficient entropy regularization during training. Specifically, we show how a factorized approximation to the entropy for soft quantization of L bits reduces the theoretical computation complexity from $O(2^L \times L)$ to $O(L)$ with minimal approximation error, and negligible performance degradation in practice.

We validate the effectiveness of BSQ-ViT on visual reconstruction and compression benchmarks. On image reconstruction, our model achieves a state-of-the-art visual reconstruction quality by both pixel-level and semantic metrics. In particular, our best-performing BSQ-ViT achieves a reconstruction FID of 0.41 on ImageNet-1k val, a 43% reduction compared to the runner-up (SDXL-VAE [14]), while being $2.4\times$ faster. On video reconstruction, our best model reduces FVD on UCF-101 by more than half ($8.62 \rightarrow 4.10$). By further learning an autoregressive prior for adaptive arithmetic coding, BSQ-ViT achieves comparable results on video compression with state-of-the-art video compression standards, *e.g.* H.264 and HEVC. By learning a masked language model, BSQ-ViT enables image generation with similar quality to BigGAN [18] and ADM [19]. Code and models will be released at <https://github.com/zhaoyue-zephyrus/bsq-vit>.

2 Related Work

Visual Tokenization. VQ-VAE [8] introduced the concept of discrete tokenized bottlenecks in auto-encoder architectures. Recent improvements include better training objectives [20, 9], increasing VQ codebook usage [4, 21], replacing VQ with product quantization (PQ) [3] or scalar quantization (SQ) [22], and employing stronger generative models [9, 10]. Image tokenizers are trivially extended to video by tokenizing individual frames [23, 24]. However, this ignores dynamic motions and leads to suboptimal tokenization: The same visual information is compressed repeatedly across frames.

Video Tokenization. Dedicated video tokenizers make better use of temporal correlations in the input signal. VideoGPT [25] proposes 3D (de-)convolutions in VQ-VAE for video generation. TATS [15] replaces zero padding with replicate padding to mitigate the temporal corruption when video length varies. Yu *et al.* introduce central inflation of pretrained 2D convolutional filters to 3D [16] and further make them causal [17]. Phenaki [23] adopts a factorized causal video vision Transformer [26] (C-ViViT), which improves efficiency but sacrifices modeling complex motion across time.

Neural Compression. Since Shannon established the fundamental source coding theorem [27] it has formed the basis of lossless compression [28, 29, 30, 31] with probabilistic models including RNN [32, 33], CNN [34, 8], VAE [35, 36], and Transformers [37, 38]. L3C [39] presents a fast hierarchical probabilistic model for lossless image compression. LMIC [38] shows that LLMs trained primarily on text, *e.g.* Llama 2 [13] and Chinchilla [40], are general-purpose compressors for text, images, and audio. However, these LLMs are too big and slow to make this compression practical. Our tokenizer presents a lighter-weight alternative: Tokenization performs initial local lossy compression, while a lightweight and thus computationally efficient sequence model ($\sim 300M$) compresses the global video structure.

Video compression. Most high-performing modern video compression methods rely on hybrid coders that combine transform coding [41, 42] and motion compensation [43, 44]. Such belief continues in most of the recently popularized learning-based solutions [45, 46, 47, 48]. VCT [49] proposes a Transformer-based temporal entropy model to learn motion implicitly. However, VCT requires a heavily-engineered image compression model [50] and has a short temporal context window. In this work, we show that a learned video tokenizer combined with an arithmetic coder modeled by a sequence model achieves competitive compression results without explicitly modeling motion.

3 Preliminaries

A tokenization-based compression algorithm has three basic steps: A visual tokenizer, i.e. VQ-VAE [8] or LFQ [17], translates raw visual inputs to a discrete set of tokens and back. A sequence model then predicts an auto-regressive probability distribution over these discrete tokens. Finally, arithmetic coding translates this distribution into a compressed representation.

Visual Tokenization. VQ-VAE [8] introduced the concept of learning discrete visual representation with an auto-encoder architecture and a bottleneck module in between with vector quantization (VQ). Given a video $\mathbf{X} \in \mathbb{R}^{T \times H \times W \times 3}$, an encoder \mathcal{E} produces a set of d -dimensional latent embeddings $\mathbf{Z} = \mathcal{E}(\mathbf{X}) \in \mathbb{R}^{(\frac{T}{q} \times \frac{H}{p} \times \frac{W}{p}) \times d}$ with a spatial-temporal downsample factor of $q \times p \times p$. The bottleneck module q then transforms the real-valued latent embeddings into some discrete tokens $\hat{\mathbf{z}} = q(\mathbf{z})$.

In Vector Quantization (VQ) the quantizer q_{VQ} assigns each $\mathbf{z} \in \mathbf{Z}$ to the closest entry in a learnable code in a codebook $\mathbf{C} = [\mathbf{c}_1 \cdots \mathbf{c}_K] \in \mathbb{R}^{K \times d}$

$$\hat{\mathbf{z}} = q_{VQ}(\mathbf{z}) = \mathbf{c}_k = \arg \min_{\mathbf{c}_k \in \mathbf{C}} \|\mathbf{z} - \mathbf{c}_k\|_2. \quad (1)$$

Here, K is the vocabulary size of the codebook and the integer k is the discretized token representation of \mathbf{z} which can be stored in $\lceil \log(K) \rceil$ bits. A decoder \mathcal{G} maps the discretized tokens back into a visual representation $\hat{\mathbf{X}} = \mathcal{G}(\hat{\mathbf{z}})$. The entire network (\mathcal{E} , \mathcal{G} , and q) is end-to-end trainable and minimizes an MSE loss $\mathcal{L}_{\text{MSE}} = \|\hat{\mathbf{X}} - \mathbf{X}\|_2$ using straight-through estimator [51] to propagate gradients through the quantization bottleneck. More recent quantizers rely on a perceptual $\mathcal{L}_{\text{LPIPS}}$ and adversarial \mathcal{L}_{GAN} loss for better visual quality [9]

$$\underset{\mathcal{E}, \mathcal{G}, q}{\text{minimize}} \mathbb{E}_{\mathbf{X}} [\mathcal{L}_{\text{VQ}}(\mathcal{E}, \mathcal{G}, q) + \eta \mathcal{L}_{\text{LPIPS}}(\mathcal{E}, \mathcal{G}, q) + \lambda \mathcal{L}_{\text{GAN}}(\mathcal{E}, \mathcal{G}, q)], \quad (2)$$

where the quantization loss term \mathcal{L}_{VQ} emulates online clustering to learn \mathbf{c}_k . The main issue with VQ-VAE is that Vector Quantization scales poorly with increasing vocabulary size K [17]. Remedies include using a smaller code dimension [4], introducing stochasticity [52], reviving “dead” codevectors [21], and regularizing with a commitment loss [8]:

$$\mathcal{L}_{\text{commit}}(\hat{\mathbf{z}}, \mathbf{z}) = \|\text{sg}(\hat{\mathbf{z}}) - \mathbf{z}\|, \quad (3)$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operation.

Lookup-Free Quantization (LFQ) [17] uses a fixed implicit codebook $\mathbf{C}_{LFQ} = \{-1, 1\}^L$ as corners of a hypercube in L dimensional space. The best vector quantizer for this implicit codebook is the binary quantization $q_{LFQ}(\mathbf{z}) = \text{sign}(\mathbf{z})$. To optimize for an effective latent code and encourage usage of the implicit codebook, Yu *et al.* [17] use an additional entropy objective [53]:

$$\mathcal{L}_{\text{entropy}} = \mathbb{E} [H(q(\mathbf{z}))] - \gamma H [\mathbb{E} [q(\mathbf{z})]], \quad (4)$$

where both entropy terms rely on a soft quantization [2]

$$\hat{q}(\mathbf{c}|\mathbf{z}) = \frac{\exp(-\tau(\mathbf{c} - \mathbf{z})^2)}{\sum_{\mathbf{c} \in \mathbf{C}_{LFQ}} \exp(-\tau(\mathbf{c} - \mathbf{z})^2)} \quad (5)$$

to guarantee the loss is differentiable. The final loss \mathcal{L}_{LFQ} is a combination of \mathcal{L}_{MSE} , $\mathcal{L}_{\text{commit}}$, $\mathcal{L}_{\text{LPIPS}}$, \mathcal{L}_{GAN} , and $\mathcal{L}_{\text{entropy}}$. The main computational bottleneck in LFQ is the entropy optimization of a higher-dimensional codebook, as it involves summation over 2^L implicit codebook entries.

Both VQ-VAE and LFQ lossily compress visual inputs \mathbf{X} into N discrete tokens $[k_1, \dots, k_N]$, where $k_i \in \{1, \dots, K\}$, in $N \lceil \log K \rceil$ bits. Neither tokenization strategy exploits the global image or video structure well. A sequence model with lossless arithmetic coding better fits this global structure.

Arithmetic Coding (AC) [29, 30, 54] offers a way of constructing a bitstream with near-optimal length by leveraging the statistical property of the coding distribution. Given a distribution over token streams $P_t : \{1, \dots, K\}^n \mapsto (0, 1]$, arithmetic coding looks to encode the token stream in $(-\lceil \log P_t(k_1, \dots, k_N) \rceil + 1)$ bits. The most common token distribution is an auto-regressive model

$$P_t(k_1, \dots, k_N) = P_t(k_1)P_t(k_2|k_1) \dots P_t(k_N|k_1, \dots, k_{N-1}) \quad (6)$$

for which efficient incremental encoding and decoding algorithms exist [49].

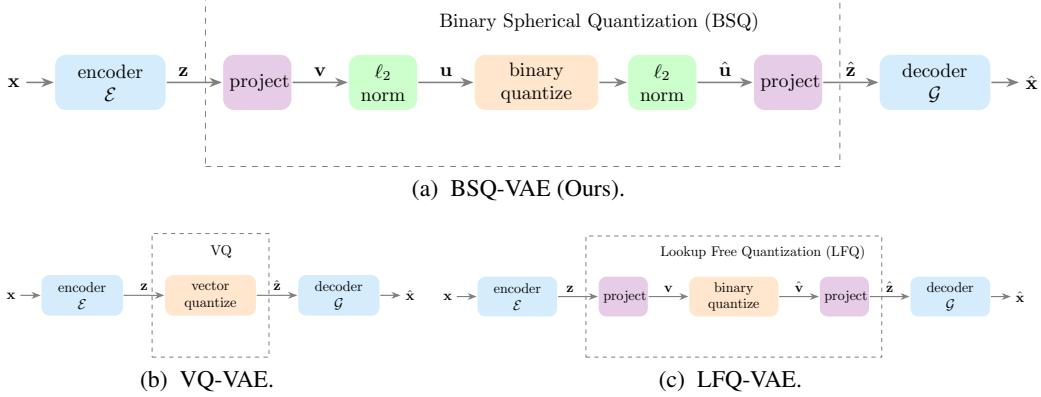


Figure 1: Variational Auto-Encoders (VAE) with different bottlenecks (BSQ, VQ, and LFQ).

4 Transformer-based Visual Tokenizer with Binary Spherical Quantization

Our video tokenizer follows an encoder-decoder architecture with a discretization bottleneck as illustrated in Figure 1a. It combines a transformer-based encoder, a transformer-based decoder, and a Binary Spherical Quantization (BSQ) layer. BSQ projects the latent code into a lower-dimensional spherical space, applies binary quantization, and then projects the result back up into the decoder’s latent space. This projection onto a low-dimensional spherical space has several theoretical advantages: The approximation error of the quantizer is bounded and much of the entropy computation factorizes along individual dimensions. These advantages result in experimental improvements as well. BSQ converges quicker and to a better tokenizer than other quantization schemes.

4.1 Binary Spherical Quantization

Binary Spherical Quantization (BSQ) optimizes over an implicit codebook $\mathbf{C}_{BSQ} = \{-\frac{1}{\sqrt{L}}, \frac{1}{\sqrt{L}}\}^L$, a hypercube projected onto a unit sphere. Each corner $\mathbf{c}_k \in \mathbf{C}_{BSQ}$ of a hypercube corresponds to a unique token k . The quantizer works as follows: it projects some high-dimensional latent embedding z to a lower-dimensional unit hypersphere u , applies binary quantization per axis $\hat{u} = \text{sign}(u)$, and back-projects to the quantized vector in the original latent space \hat{x} , as shown in Figure 1a. Specifically, we start with an encoded visual input $z = \mathcal{E}(x) \in \mathbb{R}^d$. We first linearly project the latent embedding to L dimensions $v = \text{Linear}(z) \in \mathbb{R}^L$, where $L \ll d$. Next, we obtain project v onto the unit sphere $u = \frac{v}{\|v\|}$, and perform binary quantization to each dimension of u independently $\hat{u} = \frac{1}{\sqrt{L}} \text{sign}(u)$, where $\text{sign}(x)$ is the sign function. To keep outputs on the unit sphere, we map $\text{sign}(0) \rightarrow 1$. We use a Straight-Through Estimator (STE) [51] to make the operator differentiable, $\text{sign}_{\text{STE}}(x) = \text{sg}(\text{sign}(x) - x) + x$, where $\text{sg}(\cdot)$ denotes the stop-gradient operation. Finally, we back-project the quantized \hat{u} to the d -dimensional space $\hat{z} = \text{Linear}(\hat{u}) \in \mathbb{R}^d$.

BSQ has a few appealing properties: As with LFQ, the implicit codebook entry is parameter-free and easy to compute. Unlike LFQ, a soft quantization of BSQ has a simple probabilistic interpretation, which leads to efficient entropy computation in an entropy loss $\mathcal{L}_{\text{entropy}}$. Finally, BSQ’s quantization error is bounded, which empirically leads to much faster and better convergence than LFQ.

Efficient implicit code assignment. At inference time, we map a projected embedding v to a token through simply binarization $k = \sum_{i=1}^L 1_{[v_i > 0]} 2^{i-1}$, where $1_{[\cdot]}$ is the indicator function. The inverse mapping uses the bitshift and the bitwise AND operations.

Soft BSQ and entropy. To best use the entire range of the implicit codebook \mathbf{C}_{BSQ} , we use the entropy loss $\mathcal{L}_{\text{entropy}} = \mathbb{E}_u [H(q(u))] - \gamma H [\mathbb{E}_u [q(u)]]$ [53]. To compute this entropy loss we first derive a soft quantization scheme [2]. Since both codebook entries and inputs to the quantizer are unit vectors, the soft quantization is a distribution

$$\hat{q}(\mathbf{c}|\mathbf{u}) = \frac{\exp(\tau \mathbf{c}^\top \mathbf{u})}{\sum_{\mathbf{c} \in \mathbf{C}_{BSQ}} \exp(\tau \mathbf{c}^\top \mathbf{u})} = \prod_{d=1}^L \sigma(2\tau c_d u_d), \quad (7)$$

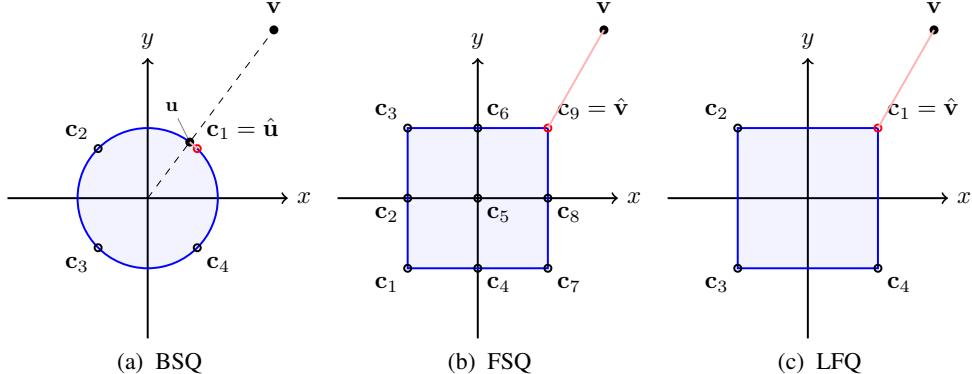


Figure 2: **Illustration of BSQ compared to LFQ and FSQ in the simplest case of $L = 2$.** In FSQ, we further consider each channel has 3 possible values $\{-1, 0, 1\}$. The Voronoi diagram for both FSQ and LFQ looks like hypercubes that partition the entire space while BSQ’s looks like a hypersphere evenly divided by 2^L centroids.

where σ is a sigmoid function, and the overall soft quantizer is independent along each dimension. See Sec. C.1 for a derivation. This form allows for an efficient computation of the first entropy term

$$\mathbb{E}_{\mathbf{u}} [H(\hat{q}(\mathbf{c}|\mathbf{u}))] = \mathbb{E}_{\mathbf{u}} \left[\sum_{d=1}^L H(\hat{q}(c_d|u_d)) \right]. \quad (8)$$

See Sec. C.2 for a derivation. Instead of reasoning over distributions over the entire codebook, which is exponentially large, we instead treat each dimension independently. The resulting entropy computation is linear to the dimension L of the bottleneck.

Unfortunately, the second entropy term cannot directly use the same independence assumption, as dimensions in the expected value $\mathbb{E}_{\mathbf{u}}[\hat{q}(\mathbf{c}|\mathbf{u})]$ are correlated through the distribution of \mathbf{u} . We find the closest factorized distribution $\tilde{q}(\mathbf{c}) = \prod_{d=1}^L \tilde{q}(c_d)$ to $\mathbb{E}_{\mathbf{u}}[\hat{q}(\mathbf{c}|\mathbf{u})]$, and instead minimize the entropy of the approximate distribution. As we will show in Sec. C.3 the best approximation in terms of the KL-divergence $\tilde{q}(c_d) = \mathbb{E}_{\mathbf{u}_d}[\hat{q}(c_d|\mathbf{u}_d)]$. The final approximate entropy term to maximize is

$$H(\mathbb{E}_{\mathbf{u}} [\hat{q}(\mathbf{c}|\mathbf{u})]) \approx H(\tilde{q}(\mathbf{c})) = \sum_{d=1}^L H(\mathbb{E}_{\mathbf{u}_d} [\hat{q}(c_d|\mathbf{u}_d)]). \quad (9)$$

As we will show in Sec. C.3 this approximation is an upper bound to the true entropy, but empirically closely tracks the true entropy. This entropy term is again efficient for evaluation.

Quantization error in BSQ. Most quantizers use pass-through gradient estimates during training [17, 8, 9]. Though simple to implement, it assumes that the gradients for an unquantized \mathbf{u} and quantized $\hat{\mathbf{u}}$ bottleneck are almost the same, which only holds if the quantization error $d(\mathbf{u}, \hat{\mathbf{u}}) = \|\mathbf{u} - \hat{\mathbf{u}}\|$ is small. As we show in Sec. C.4, this is true for BSQ

$$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] < \sqrt{2 - \frac{2}{\sqrt{L}}} < \sqrt{2}. \quad (10)$$

Relation to other quantization methods. BSQ is closely connected to many concepts introduced in information and coding theories. LFQ [17] uses the same binarization technique as BSQ but does not normalize its output. This leads to an unbounded quantization error and does not allow for as simple of a soft quantization for entropy computation. A pictorial comparison between LFQ and BSQ is shown in Figure 2 and a summary is provided in Table 7. Spherical Vector Quantization (SVQ) [55] also ensures all code vectors have a pre-defined radius. However, SVQ assumes a variety of radii, which have to be encoded by an additional gain quantizer. In our case, the source code is the output of a learned encoder \mathcal{E} . Therefore, the unit radius assumption is sound, and the gain quantizer can be avoided. Pyramid Vector Quantization (PVQ) [56] assumes all code vectors have a constant ℓ_1 norm, but the ℓ_1 normalized centroids partition the hypersphere less uniformly than ℓ_2 .

4.2 Tokenization Network with Causal Video Transformer

We propose to use Vision Transformer (ViT) [57] to model both the encoder and decoder due to its better computational efficiency and higher reconstruction quality.

Video Transformer. We start from ViT-VQGAN [4] and extend it to take videos as input. We divide an input video $\mathbf{X} \in \mathbb{R}^{T \times H \times W \times 3}$ into non-overlapping patches of size $1 \times p \times p$, $\mathbf{x}_i \in \mathbb{R}^{1 \times p \times p \times 3}$. The visual tokens are flattened into a 1D sequence, linearly projected, and passed through a stack of N Transformer Encoder layers to yield the latent representation, $(\mathbf{z}_1, \dots, \mathbf{z}_N)$. The decoder takes the same architecture, maps the latent embeddings $\hat{\mathbf{z}}$ back to the pixel space, and regroups them into the original shape. $(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N) = \text{MLP}(\text{TransformerDecoder}(\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_N))$, where MLP is a decoding head with a two-layer MLP, *i.e.* $\text{Linear} \circ \text{Tanh} \circ \text{Linear}$.

Blockwise Causal Attention. During training, we always assume the input video has T frames, which might not hold at inference. Padding shorter video segments to T frames works but wastes a lot of bits, especially in the context of compression. To handle variable-length videos, we propose a simple blockwise causal masked attention analogous to causal attention in language modeling [58]. It specifies that only those tokens at time t or earlier can be used for reconstructing the visual tokens at time $t \in \{1, \dots, T\}$.

$$(\mathbf{z}_{(t-1) \times \frac{H}{p} \times \frac{W}{p} + 1}, \dots, \mathbf{z}_{t \times \frac{H}{p} \times \frac{W}{p}}) = \text{TransformerEncoder}(\mathbf{x}_1, \dots, \mathbf{x}_{t \times \frac{H}{p} \times \frac{W}{p}}), \quad (11)$$

$$(\hat{\mathbf{z}}_{(t-1) \times \frac{H}{p} \times \frac{W}{p} + 1}, \dots, \hat{\mathbf{z}}_{t \times \frac{H}{p} \times \frac{W}{p}}) = q_{LFQ}(\mathbf{z}_{(t-1) \times \frac{H}{p} \times \frac{W}{p} + 1}, \dots, \mathbf{z}_{t \times \frac{H}{p} \times \frac{W}{p}}), \quad (12)$$

$$(\hat{\mathbf{x}}_{(t-1) \times \frac{H}{p} \times \frac{W}{p} + 1}, \dots, \hat{\mathbf{x}}_{t \times \frac{H}{p} \times \frac{W}{p}}) = \text{MLP}(\text{TransformerDecoder}(\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_{t \times \frac{H}{p} \times \frac{W}{p}})). \quad (13)$$

This can be efficiently implemented with a blockwise causal attention mask written in a blockwise lower triangle matrix in Figure 3. When $T = 1$, the proposed encoder-decoder reduces to a ViT with a full attention mask. Therefore, we can easily train it using a mixture of images and videos.

We use factorized spatial-temporal position embedding to encode the temporal position information. Specifically, we add a set of zero-initialized temporal position embeddings $\text{PE}_t \in \mathbb{R}^{T \times d}$ and add it to the original spatial position embedding $\text{PE}_s \in \mathbb{R}^{N \times d}$ in the image tokenizer, *i.e.* $\text{PE}[i, :, :] = \text{PE}_t[i, \text{None}, :] + \text{PE}_s[\text{None}, :, :]$.

Training the Video Tokenizer from an Image Tokenizer. Due to the lack of diversity in existing video datasets, we first train an image tokenizer on image data and then fine-tune it to be a video tokenizer. Though previous works [7, 24] argue that a pre-trained image tokenizer can be used for videos as is, we observe that the video tokenizer after fine-tuning demonstrates much higher reconstruction quality on video benchmarks, see Sec. 5.1. The gain is further magnified when the effective vocabulary size becomes larger. We hypothesize that such increased vocabulary size, enabled by the proposed BSQ, is handy for learning video-specific motion and blur. In contrast, vanilla VQ methods fail to maintain high codebook usage when the codebook size exceeds 16K.

Optimizing the Visual Tokenizer. Following VQGAN [9], we use a perceptual loss [59] and an adversarial loss [60]. We use StyleGAN [61] as the discriminator since ViT-VQGAN [4] reports it is much easier to train than PatchGAN [62]. When we fine-tuned the tokenizer on videos, unlike MAGVIT or TATS, we did not inflate StyleGAN to be a 3D discriminator. Instead, we pass all reconstructed frames individually to the vanilla StyleGAN and sum up the losses.

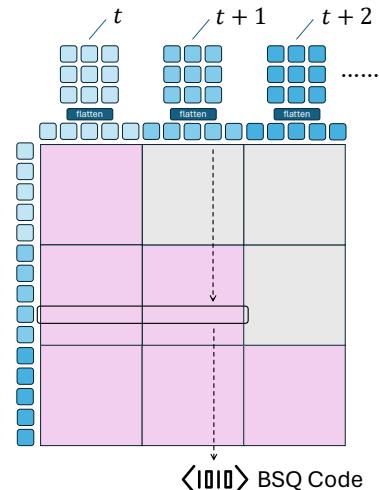


Figure 3: Given an input video, blockwise causal masked attention enables the Transformer encoder to only use the flattened patches from current or past timestamps to encode each visual patch and later translate it into a BSQ code.

Table 1: **Image reconstruction results on COCO2017 and ImageNet-1K** (256×256). The “data” column refers to the training data: CC for CC3M, YF for YFCC100M, OImg for OpenImages, LAION for LAION-5B, IN for ImageNet, and “?” for unknown source. The “arch.” column shows the encoder/decoder architecture: C for ConvNets with Self-Attention, and T-B for ViT-Base. The “# bits” column refers to the effective number of bits per token defined in Sec. 5.1. $\#$ is obtained by multiplying the latent dimension with the precision. The “TP” column means the inference throughput (images/second) per GPU. \dagger The number is taken from the paper. Note that STDs of PSNR, SSIM, and LPIPS are computed across samples instead of multiple runs.

Method	Data	Arch.	Quant.	Param.	# bits	TP \uparrow	COCO2017 val				ImageNet-1k val			
							PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow
DALL-E dVAE [20]	CC+YF	C	VQ	98M	13	34.0	25.15	.7497	.3014	55.07	25.46	.7385	.3127	36.84
MaskGIT [10]	IN-1k	C	VQ	54M	10	37.6	17.52	.4194	.2057	8.90	17.93	.4223	.2018	2.23
ViT-VQGAN [4]	IN-1k	T-B	VQ	182M	13	\dagger 7.5	-	-	-	-	-	-	-	\dagger 1.55
SD-VAE 1.x [72]	OImg	C	VQ	68M	10	22.4	21.78	.6139	.1042	6.79	22.12	.6046	.1039	1.52
SD-VAE 1.x [72]	OImg	C	VQ	68M	14	22.4	22.54	.6470	.0905	6.07	22.82	.6354	.0912	1.23
SD-VAE 1.x [72]	OImg	C	KL	68M	#64	22.4	21.68	.6375	.0985	5.94	21.99	.6275	.0980	1.35
SD-VAE 2.x [14]	OImg+LAION	C	KL	84M	#64	18.9	24.82	.7202	.0694	4.63	25.08	.7054	.0731	0.78
SDXL-VAE [14]	OImg+LAION+?	C	KL	84M	#64	18.9	25.11	.7433	.0623	4.23	25.38	.7276	.0666	0.72
Ours	IN-1k	T-B	BSQ	174M	18	45.1	25.08	.7662	.0744	5.81	25.36	.7578	.0761	1.14
Ours	IN-1k	T-B	BSQ	174M	36	45.1	27.64	.8485	.0412	3.42	27.88	.8410	.0432	0.41
Ours (w/. EMA)	IN-1k	T-B	BSQ	174M	36	45.1	27.92	.8526	.0380	3.34	28.14	.0814	.0400	0.45
							± 3.78	± 0.0698	± 0.187		± 4.32	± 0.0814	± 0.0237	

5 Experiments

We train the image tokenization model on the training set of ImageNet ILSVRC2012 [63] and evaluate the image reconstruction result on the validation set of MS-COCO [64] and ImageNet, denoted by **COCO 2017val** and **ImageNet-1k** respectively. We fine-tune the video tokenization model on **UCF-101** [65] and conduct video compression experiments on two standard benchmarks, *i.e.* MCL-JCV and UVG. We leave dataset statistics and implementation details in Sec. E.

Evaluation metrics. For image/video tokenization, we report perceptual metric (LPIPS-AlexNet) [59], PSNR, SSIM [66], and Fréchet Inception/Video Distance (FID/FVD) [67, 68]. To distinguish it from generation, we denote it as rFID/rFVD. For generation, we report FID, Inception Score (IS) [69], and improved precision and recall (IPR, Prec, and Rec) [70]. For compression, we report PSNR and MS-SSIM [71] under different levels of bits per pixel (bpp).

5.1 Main Results

Image Reconstruction. We first compare the image reconstruction result of BSQ on COCO and ImageNet (256×256) with state-of-the-art image tokenizers, including DALL-E dVAE [20], SD-VAE 1.x [72], SD-VAE 2.x, SDXL-VAE [14], MaskGIT [10], and ViT-VQGAN [4]. We observe that reconstruction metrics vary with many factors, especially preprocessing (*e.g.* interpolation), input resolution, and downsample scales (Sec. D.2 in [72]). To perform a comprehensive and fair comparison, we resize all images such that the smaller edge is 256 pixels using Lánczos interpolation², take the center crop ($H \times W$) = (256×256) , and ensure all models have the same spatial downsample ratio of $p = 8$ (except for MaskGIT, $p = 16$). We rerun all models on COCO 2017val and ImageNet-1k val except the undisclosed ViT-VQGAN. From Table 1, we can see that our model outperforms prior works on all metrics (PSNR, SSIM, LPIPS, and rFID), often by a big margin.

To compare the compression capability of different bottleneck modules, We study the **effective number of bits per token (# bits)**. For VQ-based models, # bits equals to $\log_2(K)$, where K is the

²The reconstruction result of bilinear interpolation is computed in Table 8 for reference. In short, the conclusion is that varying interpolation changes the values but unalts the order of all methods.

Table 2: Video reconstruction results on UCF-101 (split 1).

Method	Backbone	Quantizer	Param.	# bits	UCF-101 train			UCF-101 val		
					PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
(IMAGE TOKENIZER, w/o ADAPTING TO VIDEOS)										
Ours	ViT	VQ	174M	14	25.64	.8142	.1120	357	25.58	.8120
Ours	ViT	BSQ	174M	18	25.86	.8273	.1089	326	25.83	0.8259
(IMAGE TOKENIZER \rightarrow VIDEO TOKENIZER)										
MaskGIT [10]	2D CNN	VQ	53M	10	21.5	.685	0.114	216	-	-
TATS [15]	3D CNN	VQ	32M	14	-	-	-	162	-	-
MAGVIT-L [16]	3D CNN	VQ	158M	10	22.0	.701	.0990	25	-	-
MAGVIT-v2 [17]	C-3D CNN	LFQ	158M	18	-	-	.0694	16.12	-	-
MAGVIT-v2 [17]	C-3D CNN	LFQ	N/A (>158M)	18	-	-	.0537	8.62	-	-
Ours	non-BC ViT	VQ	174M	14	33.06	.9518	.0223	9.16	32.92	.9506
Ours	BC ViT	VQ	174M	14	32.81	.9496	.0236	10.76	32.68	.9484
Ours	BC ViT	BSQ	174M	18	32.08	.9421	.0244	8.08	31.49	.9357
Ours	BC ViT	BSQ	174M	36	33.80	.9606	.0159	4.10	33.55	.9588
									.0167	6.21

codebook size; For KL-regularized models (SD-VAE 2.x and XL), since the latent is continuous, we count # bits as the latent dimension multiplied by the numeric precision (here we use 16 since the checkpoint is stored in FP16). For our BSQ, # bits is L because each latent channel is binary. We summarize the key observations as follows. (1) BSQ efficiently compresses image patches into a small amount of bits. It reconstructs images better in all metrics using fewer bits per token than the second-best method (SDXL-VAE). (2) BSQ is also computationally efficient. Although the ViT-based backbone doubles the parameters, our method yields a $2.4 \times$ higher throughput than SDXL-VAE. MaskGIT runs at a comparable speed but reconstructs significantly worse because of a small codebook size (1024) and more spatial downsampling ($16 \times$). (3) BSQ is generalizable across different domains of images. ImageNet is relatively object-centric while COCO is more scene-centric. Though trained on ImageNet only, our method does well on the scene-centric COCO too. It even works better than SD-VAE 1.x/2.x trained on the similarly scene-centric OpenImages dataset [73].

Video Reconstruction. We present the video reconstruction on both UCF-101 training and validation subsets in Table 2. First, we use the image tokenizer to reconstruct the video frame by frame. BSQ works slightly better than VQ but neither is comparable to the specialized video tokenizers fine-tuned on video data shown in the lower half of Table 2. Second, we finetune the image tokenizer on videos and see significant improvements. For example, our 18-bit BSQ with causal ViT reduces rFVD from 342 to 11.62 and improves PSNR from 25.83 to 31.49 dB. The compared prior methods include: (1) MaskGIT [10] which is a fine-tuned 2D-CNN based tokenizer, (2) TATS [15] which uses a 3D CNN with replicated padding, (3) MAGVIT [4] whose 3D CNN is initialized by zero-inflating a 2D filter, and (4) MAGVIT-v2 [17] which makes 3D CNN causal. Since most methods do not release checkpoints, we take their reported numbers directly. Our models with all configurations outperform MAGVIT-v2 with a comparable number of parameters (174M vs. 158M) by a large margin. The best-performing MAGVIT-v2 uses a larger backbone and achieves a rFVD of 8.62. Our causal BSQ-ViT with $L = 18$ achieves an 8.08 rFVD and halves the LPIPS. For BSQ with $L = 36$, our method further improves the reconstruction metrics.

We also show the effect of using block-wise causal masks. The non-causal variant (non-BC) works slightly better on all metrics because now the model can look at all visual patches within the temporal context window. This result resembles the observations in video compression that using bidirectional predicted pictures (B-frames) benefits compression quality given the same group of pictures (GoP).

Image Generation. Our BSQ-ViT tokenizer can be seamlessly integrated into existing generative models for visual generation. We follow MaskGIT [10], a masked language modeling approach. At training time, the underlying masked language model (masked LM) learns to predict the masked tokens given a random proportion of unmasked tokens like BERT [74]. At inference time, the model repeats decoding in a non-

Table 3: Image generation results on ImageNet-1K (128 \times 128).

[†]The number is taken from the paper.

Category	Method	# steps	FID \downarrow	IS \uparrow	Prec \uparrow	Rec \uparrow
GAN	BigGAN [18]	1	6.02	145.8	0.86	0.35
Diffusion	ADM [19]	1,000	5.91	93.3	0.70	0.65
	VQ	12	[†] 9.4	-	-	-
Masked LM	FSQ [22]	12	[†] 8.5	-	-	-
	BSQ (Ours)	32	5.44	139.6	0.80	0.50

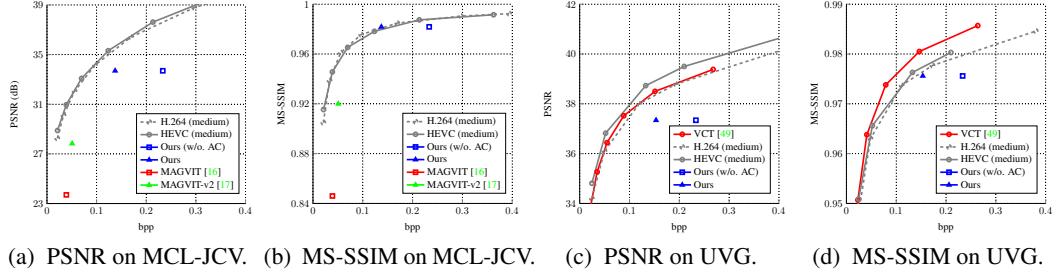


Figure 4: **Video compression results on MCL-JCV 640×360 and UVG 1920×1080 .**

Table 5: **Ablation studies on ImageNet-1k val 128×128 .**

Method	ℓ_2 -norm	# bits ($K \times d$ or L)	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow	Code usage
VQ	✓	10 (1024 \times 32)	23.61 \pm 3.21	.6873 \pm .1211	.1214 \pm .0434	7.05	57.5%
	✓	14 (16384 \times 8)	25.76 \pm 3.46	.7834 \pm .0988	.0669 \pm .0282	4.27	100.0%
	✓	16 (65536 \times 8)	25.67 \pm 3.36	.7851 \pm .0962	.0706 \pm .0283	6.61	100.0%
BSQ	✓	10	24.11 \pm 3.25	.7250 \pm .1121	.0919 \pm .0338	4.51	100.0%
	✓	14	25.26 \pm 3.31	.7710 \pm .0992	.0784 \pm .0293	4.60	99.8%
	✓	18	25.97 \pm 3.37	.7990 \pm .0906	.0629 \pm .0261	2.66	93.8%
LFQ	✗	18	18.58 \pm 2.10	.4828 \pm .1340	.2951 \pm .0806	30.7	0.6%

autoregressive way [75] for several steps and progressively decodes from an all-masked canvas to visually plausible contents following a pre-defined unmasking schedule. Unlike MaskGIT with a VQ-VAE with $K = 1024$, BSQ-ViT has an effective vocabulary size of 2^L and $L = 18$, resulting in a slow embedding lookup. We fix it by dividing each token into groups and treating sub-tokens independently with a similar rationale in Sec. 4.1. We increase the number of decoding steps accordingly. Table 3 shows that the masked LM with BSQ outperforms those with VQ and FSQ reported in [22]. Our method achieves comparable results with other generation paradigms such as GAN-based [18] and diffusion-based [19] approaches. We leave qualitative results in Sec. F.

Video Compression. We compare the compression result on MCL-JCV and UVG in Figure 4. Simply flattening the video token sequence to a bitstream achieves an MS-SSIM of 0.9818 at 0.2333 bpp. Although this is not great, we use an auto-regressive model to predict the conditional probability such that the bpp is reduced by 41%. This leads to a better tradeoff than standard video codecs including both H.264 and HEVC.

On UVG 1080P, our model is comparable to H.264 while being worse than HEVC and VCT [49]. Note that our model trains on UCF-101 which only has 9K 320×240 video clips encoded in MPEG-4 while VCT has been trained on a million high-resolution Internet video clips. We hypothesize that the gap will be mitigated by adding more diverse videos and removing compression artifacts from the training videos. Nevertheless, we show the potential advantage of our method in encoding and decoding speed in Table 4. Due to the simplicity of the Transformer-based encoder and decoder, our method runs faster than VCT.

5.2 Ablation Studies

For ablation studies, we train an ImageNet image tokenizer with resolution 128×128 with $p = 8$, although our conclusions generally hold for higher resolution, *e.g.* 256×256 in Sec. 5.1.

BSQ vs VQ. Table 5 shows that BSQ and VQ follow a similar trend: better reconstruction for increased L . Since $K = 2^{18}$ results in an out-of-memory issue, we try a smaller $K = 2^{16} = 65536$ for VQ. The gain for VQ already diminishes even though the small bottleneck dimension of 8 still guarantees full code usage. In contrast, BSQ consistently works better on all metrics when $L = 18$.

Method	Resolution	Encode	EC	Decode	FPS
VCT [49]	1920 \times 1080	[†] 494 ms	30.5 ms	168 ms	1.4
H.264	1920 \times 1080	-	-	-	2.6
Ours	1920 \times 1080	55.8 ms	42.2 ms	64.8 ms	6.1
VCT [49]	640 \times 360	[†] 22.2 ms	4.24 ms	10.1 ms	27.3
H.264	640 \times 360	-	-	-	22.4
Ours	640 \times 360	6.2 ms	4.69 ms	7.2 ms	55.2

Table 6: **Ablation studies of the loss design.**(a) Leave-one-out ablations for training losses. (b) Group size. ($L = 18$)

$\mathcal{L}_{\text{commit}}$	$\mathcal{L}_{\text{entropy}}$	$\mathcal{L}_{\text{LPIPS}}$	rFID	Code usage	group size	rFID	Code usage	Speed (ms)
✓	✓	✓	✓	2.95 45.6%	$g = 18$	(OOM)	70.0	
✗	✓	✓	✓	2.83 93.8%	$g = 9$	2.83	93.8%	0.335
✓	✗	✓	✓	2.44 78.3%	$g = 6$	2.76	95.2%	0.232
✓	✓	✗	✓	13.8 13.3%	$g = 3$	3.32	96.0%	0.233
✓	✓	✓	✗	19.2 6.9%	Ours	2.86	95.1%	0.212

Importance of ℓ_2 normalization in BSQ. We remove the ℓ_2 normalization in BSQ, which is equivalent to LFQ, and show results in the last rows of Table 5. We see much lower code usage and worse rFID, indicating that LFQ does not work well with a ViT-based tokenization encoder.

Contribution of losses. We study the effect of each loss in Table 6a. Although it is computationally prohibitive to enumerate all combinations of loss terms and their associative weights, we conduct a simple “leave-one-out” setting where one of the losses is removed at a time. BSQ works slightly better after removing $\mathcal{L}_{\text{commit}}$ and $H(p(\mathbf{c}|\mathbf{u}))$. However, the code usage varies greatly. The best configuration is to keep the minimal entropy term while dropping the commitment loss. The commit loss may be unnecessary because the quantization error in BSQ is already strictly bounded. On the contrary, the dataset entropy maximization term and perceptual term do matter. Without $-H(\mathbb{E}_{\mathbf{u}}[p(\mathbf{c}|\mathbf{u})])$, rFID increases to 13.8 while the code usage in the validation set significantly drops to 13.3%. We also observe that the perceptual loss is important for low FID and high code usage. However, a deeper look into its role is beyond the scope of this paper.

Approximating the dataset entropy term. We now show the efficacy of approximating the dataset entropy term using Eq (9). We compare with the approximation method in [17] that computes entropy in sub-groups of dimensions with varying group size $g \in \{9, 6, 3\}$. Our approximation method can also be interpreted as a group size of $g = 1$. From Table 6b, we conclude that our approximation achieves a similar level of rFID and code usage compared to other setups while running the fastest.

6 Conclusions

We present a new transformer-based image and video tokenizer with Binary Spherical Quantization (BSQ). The transformer-based architecture effortlessly integrates image and video tokenization over an arbitrary time horizon. The Binary Spherical Quantization allows for efficient and effective training of the quantized bottleneck. Our results indicate that the proposed tokenizer runs at a faster speed, reconstructs with higher fidelity, and in combination with a sequence model offers a strong baseline for lossy video compression and image synthesis.

References

- [1] Thomas J Daede, Nathan E Egge, Jean-Marc Valin, Guillaume Martres, and Timothy B Terriberry. Daala: A perceptually-driven next generation video codec. *arXiv preprint arXiv:1603.03129*, 2016.
- [2] Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc V Gool. Soft-to-hard vector quantization for end-to-end learning compressible representations. *NeurIPS*, 2017.
- [3] Alaaeldin El-Nouby, Matthew J Muckley, Karen Ullrich, Ivan Laptev, Jakob Verbeek, and Hervé Jégou. Image compression with product quantized masked image modeling. *TMLR*, 2023.
- [4] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved VQGAN. In *ICLR*, 2022.
- [5] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image transformers. In *ICLR*, 2022.
- [6] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. iBOT: Image BERT pre-training with online tokenizer. In *ICLR*, 2022.

- [7] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luwei Zhou, and Lu Yuan. BEVT: BERT pretraining of video transformers. In *CVPR*, 2022.
- [8] Aaron Van Den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017.
- [9] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021.
- [10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. MaskGIT: Masked generative image transformer. In *CVPR*, 2022.
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020.
- [12] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [13] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [14] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- [15] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic VQGAN and time-sensitive transformer. In *ECCV*, 2022.
- [16] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. MAGVIT: Masked generative video transformer. In *CVPR*, 2023.
- [17] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion–tokenizer is key to visual generation. In *ICLR*, 2024.
- [18] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [19] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *NeurIPS*, 2021.
- [20] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021.
- [21] Chuanxia Zheng and Andrea Vedaldi. Online clustered codebook. In *ICCV*, 2023.
- [22] Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: VQ-VAE made simple. In *ICLR*, 2024.
- [23] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual descriptions. In *ICLR*, 2022.
- [24] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [25] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. VideoGPT: Video generation using vqvae and transformers. *arXiv preprint arXiv:2104.10157*, 2021.
- [26] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A video vision transformer. In *ICCV*, 2021.
- [27] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [28] David A Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952.
- [29] Richard Clark Pasco. *Source coding algorithms for fast data compression*. PhD thesis, Stanford University CA, 1976.
- [30] Jorma Rissanen and Glen G Langdon. Arithmetic coding. *IBM Journal of research and development*, 23(2):149–162, 1979.

[31] Jarek Duda. Asymmetric numeral systems. *arXiv preprint arXiv:0902.0271*, 2009.

[32] Tomáš Mikolov. *Statistical language models based on neural networks*. PhD thesis, Brno University of Technology, 2012.

[33] Mohit Goyal, Kedar Tatwawadi, Shubham Chandak, and Idoia Ochoa. Deepzip: Lossless data compression using recurrent neural networks. In *DCC*, 2019.

[34] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *NeurIPS*, 2016.

[35] James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. In *ICLR*, 2019.

[36] James Townsend, Thomas Bird, Julius Kunze, and David Barber. Hilloc: Lossless image compression with hierarchical latent variable models. In *ICLR*, 2020.

[37] Fabrice Bellard. Lossless data compression with neural networks. URL: <https://bellard.org/nncp/nncp.pdf>, 2019.

[38] Grégoire Delétang, Anian Ruoss, Paul-Ambroise Duquenne, Elliot Catt, Tim Genewein, Christopher Mattern, Jordi Grau-Moya, Li Kevin Wenliang, Matthew Aitchison, Laurent Orseau, et al. Language modeling is compression. In *ICLR*, 2024.

[39] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *CVPR*, 2019.

[40] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. An empirical analysis of compute-optimal large language model training. In *NeurIPS*, 2022.

[41] Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Signal Processing Magazine*, 18(5):9–21, 2001.

[42] Johannes Ballé, Philip A Chou, David Minnen, Saurabh Singh, Nick Johnston, Eirikur Agustsson, Sung Jin Hwang, and George Toderici. Nonlinear transform coding. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):339–353, 2020.

[43] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *TCSVT*, 2003.

[44] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *TCSVT*, 2012.

[45] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In *CVPR*, 2019.

[46] Oren Rippel, Sanjay Nair, Carissa Lew, Steve Branson, Alexander G Anderson, and Lubomir Bourdev. Learned video compression. In *ICCV*, 2019.

[47] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *CVPR*, 2020.

[48] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. In *NeurIPS*, 2021.

[49] Fabian Mentzer, George Toderici, David Minnen, Sung-Jin Hwang, Sergi Caelles, Mario Lucic, and Eirikur Agustsson. VCT: A video compression transformer. In *NeurIPS*, 2022.

[50] Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. ELIC: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *CVPR*, 2022.

[51] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

[52] Yuhta Takida, Takashi Shibuya, WeiHsiang Liao, Chieh-Hsin Lai, Junki Ohmura, Toshimitsu Uesaka, Naoki Murata, Shusuke Takahashi, Toshiyuki Kumakura, and Yuki Mitsufuji. SQ-VAE: Variational bayes on discrete representation with self-annealed stochastic quantization. In *ICML*, 2022.

[53] Aren Jansen, Daniel PW Ellis, Shawn Hershey, R Channing Moore, Manoj Plakal, Ashok C Popat, and Rif A Saurous. Coincidence, categorization, and consolidation: Learning to recognize sounds with minimal supervision. In *ICASSP*, 2020.

[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.

[55] Jon Hamkins and Kenneth Zeger. Gaussian source coding with spherical codes. *IEEE Transactions on Information Theory*, 48(11):2980–2989, 2002.

[56] Thomas Fischer. A pyramid vector quantizer. *IEEE Transactions on Information Theory*, 32(4):568–583, 1986.

[57] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

[58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017.

[59] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

[60] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.

[61] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.

[62] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.

[63] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115:211–252, 2015.

[64] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[65] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

[66] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[67] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. *NeurIPS*, 2017.

[68] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. In *ICLR Workshop*, 2019.

[69] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. *NeurIPS*, 2016.

[70] Tuomas Kynkänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *NeurIPS*, 2019.

[71] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *ACSSC*, 2003.

[72] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.

[73] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallochi, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.

[74] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.

[75] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. In *ICLR*, 2018.

[76] Haiqiang Wang, Weihao Gan, Sudeng Hu, Joe Yuchieh Lin, Lina Jin, Longguang Song, Ping Wang, Ioannis Katsavounidis, Anne Aaron, and C-C Jay Kuo. MCL-JCV: a JND-based H. 264/AVC video quality assessment dataset. In *ICIP*, 2016.

[77] Alexandre Mercat, Marko Viitanen, and Jarno Vanne. Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 297–302, 2020.

[78] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

[79] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

[80] Jean Bégaint, Fabien Racapé, Simon Feltman, and Akshay Pushparaja. CompressAI: a pytorch library and evaluation platform for end-to-end compression research. *arXiv preprint arXiv:2011.03029*, 2020.

Table 7: Comparing BSQ and LFQ.

	LFQ [17]	BSQ (Ours)
Quantized output	$\hat{\mathbf{v}} = \text{sign}(\mathbf{v})$	$\hat{\mathbf{u}} = \frac{1}{\sqrt{L}} \text{sign}(\mathbf{u}) = \frac{1}{\sqrt{L}} \text{sign}(\frac{\mathbf{v}}{ \mathbf{v} })$
STE gradient	$\frac{\partial \hat{v}_i}{\partial v_i} = 1$	$\frac{\partial \hat{u}_i}{\partial v_i} = \frac{1}{\sqrt{L}} (1 - v_i^2 / \mathbf{v} ^2)$
Quantization Error	$\mathbb{E}_{\mathbf{v}} [d(\mathbf{v}, \hat{\mathbf{v}})] = \infty$ Unbounded	$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] < \sqrt{2 - \frac{2}{\sqrt{L}}} < \sqrt{2}$ Upper-bounded (See Sec. C.4)
Training objective	$\mathcal{L}_{\text{MSE}}, \mathcal{L}_{\text{commit}}, \mathcal{L}_{\text{LPIPS}}, \mathcal{L}_{\text{GAN}}, \mathcal{L}_{\text{entropy}} = H[p(\mathbf{c} \mathbf{v})] - H[\mathbb{E}_{\mathbf{u}}[p(\mathbf{c} \mathbf{v})]]$	$\mathcal{L}_{\text{MSE}}, \mathcal{L}_{\text{LPIPS}}, \mathcal{L}_{\text{GAN}}, \mathcal{L}_{\text{entropy}} = H[p(\mathbf{c} \mathbf{u})] - \hat{H}[\mathbb{E}_{\mathbf{u}}[p(\mathbf{c} \mathbf{u})]]$

A Arithmetic Coding Details

Starting from the initial interval $I_0 = [0, 1]$, the AC encoder recursively partitions the interval into a series of sub-interval $I_n = [l_n, u_n]$ such that $I_n \subset I_{n-1} \subset \dots \subset I_0$, and I_n is determined by I_{n-1} and $\rho(y|x_{<n})$.

$$I_n(y) = \left[l_{n-1} + (u_{n-1} - l_{n-1}) \sum_{y=1}^{x_n-1} \rho(y|x_{<n}), \quad l_{n-1} + (u_{n-1} - l_{n-1}) \sum_{y=1}^{x_n} \rho(y|x_{<n}) \right]. \quad (14)$$

Any number in the final interval I_N can sufficiently represent the encoded sequence. To obtain the final bit stream, we take a binary fraction $\lambda = \sum_{i=1}^C b_i \times 2^{-i}$, $b_i \in \{0, 1\}$ in I_N such that $l_N \leq \lambda < u_N$. The bit stream $\{b_0, \dots, b_C\}$ is the encoding result with a length of C bits.

The AC decoder takes in λ , starts with I_0 , and performs a similar interval partitioning process. At the n -th step, the decoder queries the model $\rho_n(y|x_{<n})$, calculate the sub-intervals for all possible values of y using Eq. (14), and decodes output x_n that leads to $\lambda \in I_n(x_n)$. The decoder can recover the encoded token sequence by continuing with I_{n+1} based on the decoded x_n and repeating for step $n+1$ for N steps.

In practice, the encoder and the decoder can be implemented efficiently with fixed-length integer numbers and operate incrementally for arbitrarily long input sequences.

B Comparison between BSQ and LFQ

In Sec. 4.1, we have introduced the mechanism of BSQ and briefly discussed the connections and differences with LFQ. We summarize them in Table 7. Note that STE gradient in BSQ is anisotropic and is more likely to be a good estimation because of an upper-bounded quantization error regardless of L . This property explains why a commitment loss like $\mathcal{L}_{\text{commit}}(\hat{\mathbf{u}}, \mathbf{u})$ is not needed in BSQ but useful for LFQ.

C Proofs

C.1 Proof of Eq (7)

Before proving Eq (7), we will first prove the following identity:

Let $\mathbf{u} \in \mathbb{R}^L$, $\mathbf{C} = \Omega^L \in \mathbb{R}^{L \times 2^L}$ for $\Omega = \{-\frac{1}{\sqrt{L}}, \frac{1}{\sqrt{L}}\}$,

$$\sum_{\mathbf{c} \in \mathbf{C}} e^{\tau \mathbf{u}^\top \mathbf{c}} = \sum_{\mathbf{c} \in \mathbf{C}} \prod_{d=1}^L e^{\tau u_d c_d} = \prod_{d=1}^L \sum_{c_d \in \Omega} e^{\tau u_d c_d}. \quad (15)$$

Proof. With τ dropped for simplicity of notation.

$$\begin{aligned}
\sum_{\mathbf{c} \in \mathbf{C}} e^{\mathbf{u}^\top \mathbf{c}} &= \sum_{\mathbf{c} \in \mathbf{C}} \prod_{k=1}^L e^{u_k c_k} \\
&= \sum_{c_1 \in \Omega} \sum_{c_2 \in \Omega} \dots \sum_{c_L \in \Omega} \prod_{d=1}^L e^{u_d c_d} \\
&= \sum_{c_1 \in \Omega} \sum_{c_2 \in \Omega} \dots \sum_{c_{L-1} \in \Omega} e^{u_L c_L} \prod_{d=1}^{L-1} e^{u_d c_d} \\
&= \sum_{c_1 \in \Omega} \sum_{c_2 \in \Omega} \dots \sum_{c_{L-1} \in \Omega} \left(\prod_{d=1}^{L-1} e^{u_d c_d} \right) \left(\sum_{c_L \in \Omega} e^{u_L c_L} \right) \\
&= \dots \\
&= \left(\sum_{c_1 \in \Omega} e^{u_L c_L} \right) \left(\sum_{c_2 \in \Omega} e^{u_2 c_2} \right) \dots \left(\sum_{c_L \in \Omega} e^{u_L c_L} \right) = \prod_{d=1}^L \sum_{c_d \in \Omega} e^{u_d c_d}. \square
\end{aligned}$$

Therefore, the probability of \mathbf{u} being assigned to \mathbf{c}_i can be written as:

$$\begin{aligned}
\hat{q}(\hat{\mathbf{c}}|\mathbf{u}) &= \frac{e^{\tau \mathbf{u}^\top \hat{\mathbf{c}}}}{\sum_{\mathbf{c} \in \mathbf{C}} e^{\tau \mathbf{u}^\top \mathbf{c}}} = \frac{\prod_{d=1}^L e^{\tau u_d \hat{c}_d}}{\prod_{d=1}^L \sum_{c_d \in \{-\frac{1}{\sqrt{L}}, \frac{1}{\sqrt{L}}\}} e^{\tau u_d c_d}} \quad (\text{Using Eq (15)}) \\
&= \prod_{d=1}^L \frac{e^{\tau u_d \hat{c}_d}}{e^{\tau u_d \hat{c}_d} + e^{\tau u_d - \hat{c}_d}} \quad (\text{since } c_d = \pm \frac{1}{\sqrt{L}} = \pm \hat{c}_d) \\
&= \prod_{d=1}^L \sigma(2\tau u_d \hat{c}_d).
\end{aligned}$$

C.2 Proof of Eq (8)

Since $\hat{q}(\hat{\mathbf{c}}|\mathbf{u}) = \prod_{d=1}^L \sigma(2\tau u_d \hat{c}_d)$ each variable c_d is independent of each other. Thus by definition

$$H[\hat{q}(\mathbf{c}|\mathbf{u})] = \sum_{d=1}^L H(\sigma(2\tau u_d c_d)). \quad \square$$

C.3 Proof of Eq (9)

Now we look at $H[\mathbb{E}_{\mathbf{u}}[\hat{q}(\mathbf{c}|\mathbf{u})]]$. We first compute $Q(\mathbf{c}) = \mathbb{E}_{\mathbf{u}}[\hat{q}(\mathbf{c}|\mathbf{u})]$.

$$Q(\mathbf{c}) = \mathbb{E}_{\mathbf{u}}[\hat{q}(\mathbf{c}|\mathbf{u})] = \frac{1}{N} \sum_{\mathbf{u}} \hat{q}(\mathbf{c}|\mathbf{u}) = \frac{1}{N} \sum_{\mathbf{u}} \prod_k \sigma(2u_k c_k).$$

Unlike \mathbf{c} , \mathbf{u} does not factorize like Eq (15). This would require us to compute $Q(\mathbf{c})$ as a full distribution over 2^L states, which is slow ($O(L \times 2^L)$) and easily overfits. Instead, we approximate $Q(\mathbf{c})$ by a factorized distribution $\tilde{q}(\mathbf{c}) = \prod_{d=1}^L \tilde{q}_d(c_d)$, where $c_d \in \Omega$ for $\Omega = \{-\frac{1}{\sqrt{L}}, \frac{1}{\sqrt{L}}\}$, using

an M-projection. We again omit τ for notational brevity.

$$\begin{aligned}
D(Q\|\tilde{q}) &= H(Q, \tilde{q}) - H(Q) \\
&= -\sum_{i=1}^{2^L} Q(\mathbf{c}_i) \log \tilde{q}(\mathbf{c}_i) - H(Q) \\
&= -\sum_{i=1}^{2^L} Q(\mathbf{c}_i) \sum_d \log \tilde{q}_d(c_d) - H(Q) \\
&= -\sum_d \underbrace{\sum_{i=1}^{2^L} Q(\mathbf{c}_i) \log \tilde{q}_d(c_d)}_{= -\sum_d \left(\log \tilde{q}_d(\mathbf{c}_d = 1) \sum_{\mathbf{c}_{-d}} p(\mathbf{c}_i) + \log \tilde{q}_d(\mathbf{c}_d = -1) \sum_{\mathbf{c}_{-d}} Q(c_i) \right)} - H(Q) \\
&= -\sum_d \sum_{c_d \in \{-1, 1\}} \log \tilde{q}_d(c_d) \sum_{\mathbf{c}_{-d}} Q(\mathbf{c}) - H(Q), \text{ where } \mathbf{c}_{-d} \text{ sums over all dimensions except } d.
\end{aligned}$$

The minimizer of the above projection $\frac{\partial}{\partial \tilde{q}_d} D(Q\|\tilde{q}) = 0$

$$\begin{aligned}
\tilde{q}_d(c_d)^* &= \sum_{\mathbf{c}_{-d}} Q(\mathbf{c}) = \mathbb{E}_{\mathbf{u}} \left[\sum_{\mathbf{c}_{-d}} p(\mathbf{c}|\mathbf{u}) \right] \\
&= \mathbb{E}_{\mathbf{u}} \left[\sum_{\mathbf{c}_{-d}} \prod_k \sigma(2u_k c_k) \right] = \mathbb{E}_{\mathbf{u}} \left[\sum_{\mathbf{c}_{-d}} \sigma(2u_d c_d) \prod_{k \neq d} \sigma(2u_k c_k) \right] \\
&= \mathbb{E}_{\mathbf{u}} \left[\sigma(2u_d c_d) \sum_{\mathbf{c}_{-d}} \prod_{k \neq d} \sigma(2u_k c_k) \right] = \mathbb{E}_{\mathbf{u}} \left[\sigma(2u_d c_d) \underbrace{\prod_{k \neq d} \sum_{\mathbf{c}_{-d}} \sigma(2u_k c_k)}_{=1} \right] = \mathbb{E}_{\mathbf{u}} [\sigma(2u_d c_d)]
\end{aligned}$$

Therefore, the entropy term is simplified to:

$$H(\tilde{q}) = \sum_d H(\tilde{q}_d(c_d)) = \sum_d H(\mathbb{E}_{\mathbf{u}}[\sigma(2u_d c_d)]).$$

By the nature of the above derivation the cross entropy $H(Q, \tilde{q}) = H(\tilde{q})$ equals the entropy of the approximation. This means that $D(Q\|\tilde{q}) = H(\tilde{q}) - H(Q) \geq 0$, and the entropy of the approximation is an upper bound $H(\tilde{q}) \geq H(Q)$ to the true entropy.

In practice, this bound is relatively tight. The most adversarial distribution $P(\mathbf{u})$ is $P(\frac{1}{\sqrt{L}}\vec{1}) = \frac{1}{2}$ and $P(-\frac{1}{\sqrt{L}}\vec{1}) = \frac{1}{2}$, where all inputs are maximally correlated, but the factorized distribution is not. Figure 5 shows an empirical estimate of this approximation error for various values of τ . In practice, we use $\tau = \frac{1}{100}$, which has little to no approximation error.

C.4 Proof of the Quantization Error Bound of BSQ (Eq. (10))

We consider ℓ_2 -distance $d(\mathbf{u}, \hat{\mathbf{u}}) = \|\mathbf{u} - \hat{\mathbf{u}}\|$. A simple (but loose) bound is:

$$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] = \mathbb{E}_{\mathbf{u}} [d_{\max}(\mathbf{u}, \hat{\mathbf{u}})] < \sqrt{2 - \frac{2}{\sqrt{L}}} < \sqrt{2}, \quad (16)$$

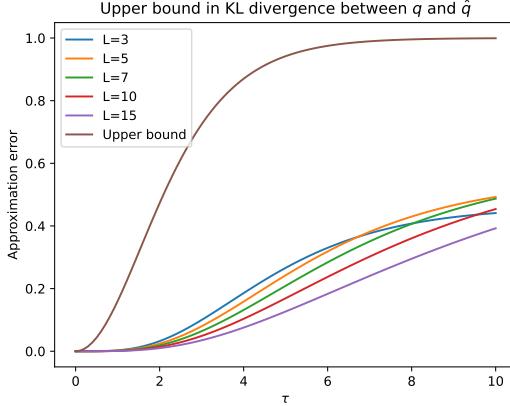


Figure 5: Empirical estimation of the approximation error with respect to τ at different bottleneck dimensions L .

where d_{\max} is attained if \mathbf{u} is at any axis, $\mathbf{u} = [\underbrace{0, \dots, 0}_n, 1, \underbrace{0, \dots, 0}_{L-1-n}]$. To achieve a tighter bound, we first expand the definition,

$$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] = \frac{\underbrace{\int \dots \int}_{S^{L-1}} d_{S^{L-1}} V d(\mathbf{u}, \hat{\mathbf{u}})}{\underbrace{\int \dots \int}_{S^{L-1}} d_{S^{L-1}} V}, \quad (17)$$

where $S^{L-1} = \{x \in \mathbb{R}^L : \|x\| = 1\}$ denotes the unit L -sphere of radius 1 and $d_{S^{L-1}} V$ denotes its surface area element. We further define a hyperspherical coordinate system that is analogous to the spherical coordinate system for 3D Euclidean space or the polar coordinate system for 2D space to represent the surface area element.

$$\begin{aligned} u_1 &= \cos(\varphi_1), \\ u_2 &= \sin(\varphi_1) \cos(\varphi_2), \\ u_3 &= \sin(\varphi_1) \sin(\varphi_2) \cos(\varphi_3), \\ &\dots \\ u_{L-1} &= \sin(\varphi_1) \sin \varphi_2 \dots \sin(\varphi_{L-2}) \cos(\varphi_{L-1}), \\ u_L &= \sin(\varphi_1) \sin \varphi_2 \dots \sin(\varphi_{L-2}) \sin(\varphi_{L-1}), \end{aligned}$$

(surface area element) $d_{S^{L-1}} V = \sin^{L-2}(\varphi_1) \sin^{L-3}(\varphi_2) \dots \sin(\varphi_{L-2}) d\varphi_1 \dots d\varphi_{L-1}$,

$$\text{(surface area)} S_{L-1} = \underbrace{\int \dots \int}_{S^{L-1}} d_{S^{L-1}} V = \frac{2\pi^{L/2}}{\Gamma(\frac{L}{2})}.$$

Due to symmetry, we assume the subarea A^{L-1} where $\forall i \in \{1, \dots, L\}$, $u_i > 0$, and it will be quantized to $\mathbf{c}_1 = \hat{\mathbf{u}}_1 = \frac{1}{\sqrt{L}} \vec{1}$. The unit hypersphere S^{L-1} has 2^L of such subareas interchangeably. Computing Eq (17) is equivalent to

$$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] = \frac{\underbrace{\int \dots \int}_{A^{L-1}} d_{S^{L-1}} V d(\mathbf{u}, \hat{\mathbf{u}})}{\underbrace{\int \dots \int}_{A^{L-1}} d_{S^{L-1}} V}. \quad (18)$$

We expand the the numerator in Eq (18) as follows:

$$\begin{aligned}
&= \int_0^{\frac{\pi}{2}} \cdots \int_0^{\frac{\pi}{2}} d_{S^{L-1}} V \{ [\cos(\varphi_1) - \frac{1}{\sqrt{L}}]^2 + [\sin(\varphi_1) \cos(\varphi_2) - \frac{1}{\sqrt{L}}]^2 + \cdots \\
&\quad + [\sin(\varphi_1) \sin(\varphi_2) \cdots \sin(\varphi_{L-2}) \cos(\varphi_{L-1}) - \frac{1}{\sqrt{L}}]^2 \\
&\quad + [\sin(\varphi_1) \sin(\varphi_2) \cdots \sin(\varphi_{L-2}) \sin(\varphi_{L-1}) - \frac{1}{\sqrt{L}}]^2 \}^{\frac{1}{2}}
\end{aligned} \tag{19}$$

It is composed of L square terms. It is easy to see that the sum of constant terms leads to 1. Next, let's sum over all quadratic terms and keep on using $\sin^2(\theta) + \cos^2(\theta) = 1$:

$$\cos^2(\varphi_1) + \sin^2(\varphi_1) \cos^2(\varphi_2) + \cdots + \prod_{j=1}^{L-2} \sin^2(\varphi_j) \sin^2(\varphi_{L-1}) + \prod_{j=1}^{L-2} \sin^2(\varphi_j) \cos^2(\varphi_{L-1}) = 1$$

So the distance function to be integrated simplifies to

$$\begin{aligned}
&\left[2 - \frac{2}{\sqrt{L}} \cos(\varphi_1) - \frac{2}{\sqrt{L}} \sin(\varphi_1) \cos(\varphi_2) - \cdots - \frac{2}{\sqrt{L}} \prod_{j=1}^{L-2} \sin(\varphi_j) \cos(\varphi_{L-1}) - \frac{2}{\sqrt{L}} \prod_{j=1}^{L-2} \sin(\varphi_j) \sin(\varphi_{L-1}) \right]^{\frac{1}{2}} \\
&< \left(2 - \frac{2}{\sqrt{L}} \cos(\varphi_1) \right)^{\frac{1}{2}}.
\end{aligned}$$

Plug into the numerator in Eq (18) and continue simplifying:

$$\underbrace{\int \cdots \int}_{A^{L-1}} d_{S^{L-1}} V \left(2 - \frac{2}{\sqrt{L}} \cos(\varphi_1) \right)^{\frac{1}{2}} \tag{20}$$

$$\begin{aligned}
&= \underbrace{\int \cdots \int}_{A^{L-1}} \underbrace{d_{S^{L-2}} V \int_0^{\frac{\pi}{2}} \left(2 - \frac{2}{\sqrt{L}} \cos(\varphi_1) \right)^{\frac{1}{2}}}_{\frac{S_{L-2}}{2^{L-1}}} \sin^{L-2}(\varphi_1) d\varphi_1. \tag{21}
\end{aligned}$$

Therefore, we have

$$\mathbb{E}_{\mathbf{u}} [d(\mathbf{u}, \hat{\mathbf{u}})] < \frac{2\Gamma(\frac{L}{2})}{\sqrt{\pi}\Gamma(\frac{L-1}{2})} \int_0^{\frac{\pi}{2}} \left(2 - \frac{2}{\sqrt{L}} \cos(\varphi_1) \right)^{\frac{1}{2}} \sin^{L-2}(\varphi_1) d\varphi_1, \tag{22}$$

where RHS can be numerically computed and plotted in Figure 6.

D Dataset Overview

ImageNet-1k has 1.28M training images and 50,000 validation images; **COCO 2017val** has 5,000 images.

UCF101 has 13,320 video clips and three train-val splits. Following prior works [16], we consider split-1 which has 9,537 clips for training and 3,783 for validation.

The **MCL-JCV** dataset [76] consists of thirty 1080P ($1,920 \times 1,080$) video sequences with 24~30 FPS. The Open Ultra Video Group (**UVG**) dataset [77] consists of sixteen 4K ($3,840 \times 2,160$) test video sequences captured at 50/120 FPS. Following prior works [47], we report the performance on a subset of seven videos in YUV 8bit format at 120 FPS under the resolution of $1,920 \times 1,080$.

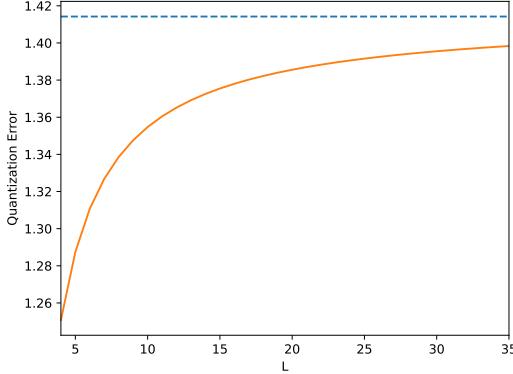


Figure 6: Quantization error with vocabulary size L .

E Implementation Details

Training Image Tokenizers. We train the image tokenizer with a batch size of 32 per GPU. We use AdamW optimizer [78] with $(\beta_1, \beta_2) = (0.9, 0.99)$ with 1×10^{-4} weight decay. The base learning rate is 4×10^{-7} (or a total learning rate of 1×10^{-4}) and follows a half-period cosine annealing schedule. The model is trained for 1M steps which amounts to 200 epochs over the entire ImageNet-1k training set. We did not heavily study the effect of loss weights. Instead, we keep $\gamma = 1$ in the entropy terms. We use a perceptual loss weight of 0.1 and an adversarial loss weight of 0.1 throughout the experiments.

Training Video Tokenizers. We finetune the video tokenizer with a batch size of 32 per GPU. The optimization schedule follows the image-based one but trains for fewer iterations. The network is initialized from the ImageNet-pretraining checkpoint and undergoes another 500K steps which amounts to 1600 epochs over UCF-101 split-1 train.

Training a Masked Language Model for Generation. The masked LM is a standard post-LN Transformer with 24 layers and a hidden dimension of 768 following MaskGIT [10]. We train the masked LM on 2 nodes of $8 \times$ GPUs (16 in total) with a total batch size of 1024 for 1M steps. We use AdamW optimizer with $(\beta_1, \beta_2) = (0.9, 0.96)$ with 0.045 weight decay. At inference time, we use a cosine unmasking schedule in MaskGIT [10] and set the sampling temperature to 15. We use classifier-free guidance [79]: At training, we replace 20% of the class condition labels with the mask token so that the model learns an unconditional distribution simultaneously. Let ℓ_c be class-conditioned logits and ℓ_\emptyset be unconditional logits. During inference, we interpolate logits using $\ell' = \ell_c + \alpha(\ell_c - \ell_\emptyset)$, where $\alpha = 0.5$.

Training an Auto-Regressive Model for Arithmetic Coding. The auto-regressive model is a Transformer with 24 layers and a hidden dimension of 768. We train this model on $8 \times$ GPUs with a total batch size being 64. We use AdamW optimizer with $(\beta_1, \beta_2) = (0.9, 0.96)$ with 0.045 weight decay.

Hardware. The hardware for training is $8 \times$ GPU-servers with NVIDIA A5000 (24GB). Pre-training an image tokenizer and fine-tuning a video tokenizer in the full schedule is done across two servers with distributed training and takes around 5 days. Training the AR model for AC is done on an $8 \times$ GPU server and takes around 1 week. When measuring the tokenizer’s throughput and the compression runtime, we use a server with $4 \times$ A5000 GPU and $1 \times$ AMD Ryzen Threadripper PRO 5975WX 32-Core CPU (64 threads).

F Qualitative Results

In Figure 7, we show reconstructed images produced by the proposed BSQ-ViT in comparison to the best prior work, SDXL-VAE [14]. We can see that our method is able to preserve more details about high-frequency texture and fine-grained shape/geometry. BSQ-ViT often shows better reconstruction results for characters.

Table 8: **Image reconstruction results on COCO2017 and ImageNet-1K** (256×256). The settings strictly follow Table 1 except that all images are resized with **bilinear** interpolation.

Method	Data	Arch.	Quant.	Param.	# bits	TP \uparrow	COCO2017 val				ImageNet-1k val			
							PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	rFID \downarrow
DALL-E dVAE [20]	CC+YF	C	VQ	98M	13	34.0	26.97 ±3.41	.0837 ±.0922	.2544 ±.1057	48.60 ±3.81	27.31 ±.1114	.7943 ±.1057	.2544	32.63
MaskGIT [10]	IN-1k	C	VQ	54M	10	37.5	18.21 ±2.74	.4596 ±0.1606	.1930 ±.0444	8.47 ±2.90	18.63 ±.1812	.4619 ±.0497	.1884	1.98
ViT-VQGAN [4]	IN-1k	T-B	VQ	182M	13	†7.5	-	-	-	-	-	-	-	†1.55
SD-VAE 1.x [72]	OImg	C	VQ	68M	10	22.4	23.29 ±3.34	.6705 ±.1316	.0949 ±.0313	6.49 ±3.69	23.65 ±.1540	.6615 ±.0367	.0940	1.40
SD-VAE 1.x [72]	OImg	C	VQ	68M	14	22.4	24.17 ±3.50	.7042 ±.1276	.0814 ±.0289	5.75 ±3.98	24.48 ±.1502	.6931 ±.0289	.0814	1.13
SD-VAE 1.x [72]	OImg	C	KL	68M	64	22.4	23.21 ±3.24	.6930 ±.1249	.0908 ±.04282	5.94 ±3.62	23.54 ±.1465	.6835 ±.0337	.0899	1.22
SD-VAE 2.x [14]	OImg+ LAION	C	KL	84M	64	18.9	26.62 ±3.64	.7722 ±.1086	.0584 ±.0273	4.26 ±4.09	26.90 ±.1300	.7592 ±.0349	.0609	0.70
SDXL-VAE [14]	OImg+ LAION+?	C	KL	84M	64	18.9	27.08 ±3.88	.7953 ±.1066	.0541 ±.0250	3.93 ±4.36	27.37 ±.1282	.7814 ±.0320	.0574	0.67
Ours	IN-1k	T-B	BSQ	174M	18	45.1	26.89 ±3.47	.8133 ±.0851	.0652 ±.0255	5.41 ±3.99	27.78 ±.0987	.8171 ±.0307	.0633	0.99
Ours	IN-1k	T-B	BSQ	174M	36	45.1	29.85 ±3.65	.8862 ±.0570	.0341 ±.0163	3.07 ±4.13	30.12 ±.0670	.8803 ±.0207	.0355	0.36
Ours (w/. EMA)	IN-1k	T-B	BSQ	174M	36	45.1	30.19 ±3.69	.8904 ±.0561	.0314 ±.0153	3.07 ±4.19	30.45 ±.0661	.8843 ±.0194	.0329	0.42

In Figure 8, we show sampled results produced by a Masked LM with the proposed BSQ-ViT in comparison to existing methods, BigGAN [18] and ADM [19]. We also plot the samples from the ground-truth ILSVRC2012 validation set for reference. Our method produces competitive results with state-of-the-art methods.

G Baselines for Video Compression

Following SSF [47], we used FFmpeg³ to produce the evaluation metrics for H.264 and HEVC. We use the commands provided in CompressAI [80].

```
ffmpeg -y -s:v $RESOLUTION -i $FILE.yuv -c:v h264 -crf $CRF -preset medium \
-bf 0 -pix_fmt yuv420p -threads 4 $FILE.mp4
```

where $\$Resolution \in \{1920 \times 1080, 640 \times 360\}$, and $\$CRF \in \{17, 20, 22, 27, 32, 37, 42, 47\}$.

H Limitations

The proposed tokenizer has been tested on images with a 256×256 or 128×128 resolution and videos with a 128×128 resolution. Training a visual tokenizer on higher-resolution inputs and variable aspect ratio remains unexplored. Also, the training dataset is limited to ImageNet-1k and UCF-101. Scaling the proposed model to larger-scale visual contents remains an interesting problem to study.

I Broader Impacts

The video compression application illustrated in the paper may be useful to reduce the storage and transmission cost of video data. Also, the proposed visual tokenization model runs more efficiently than prior ones, resulting in potential energy savings. Both of them will ultimately benefit society.

³<https://ffmpeg.org/>

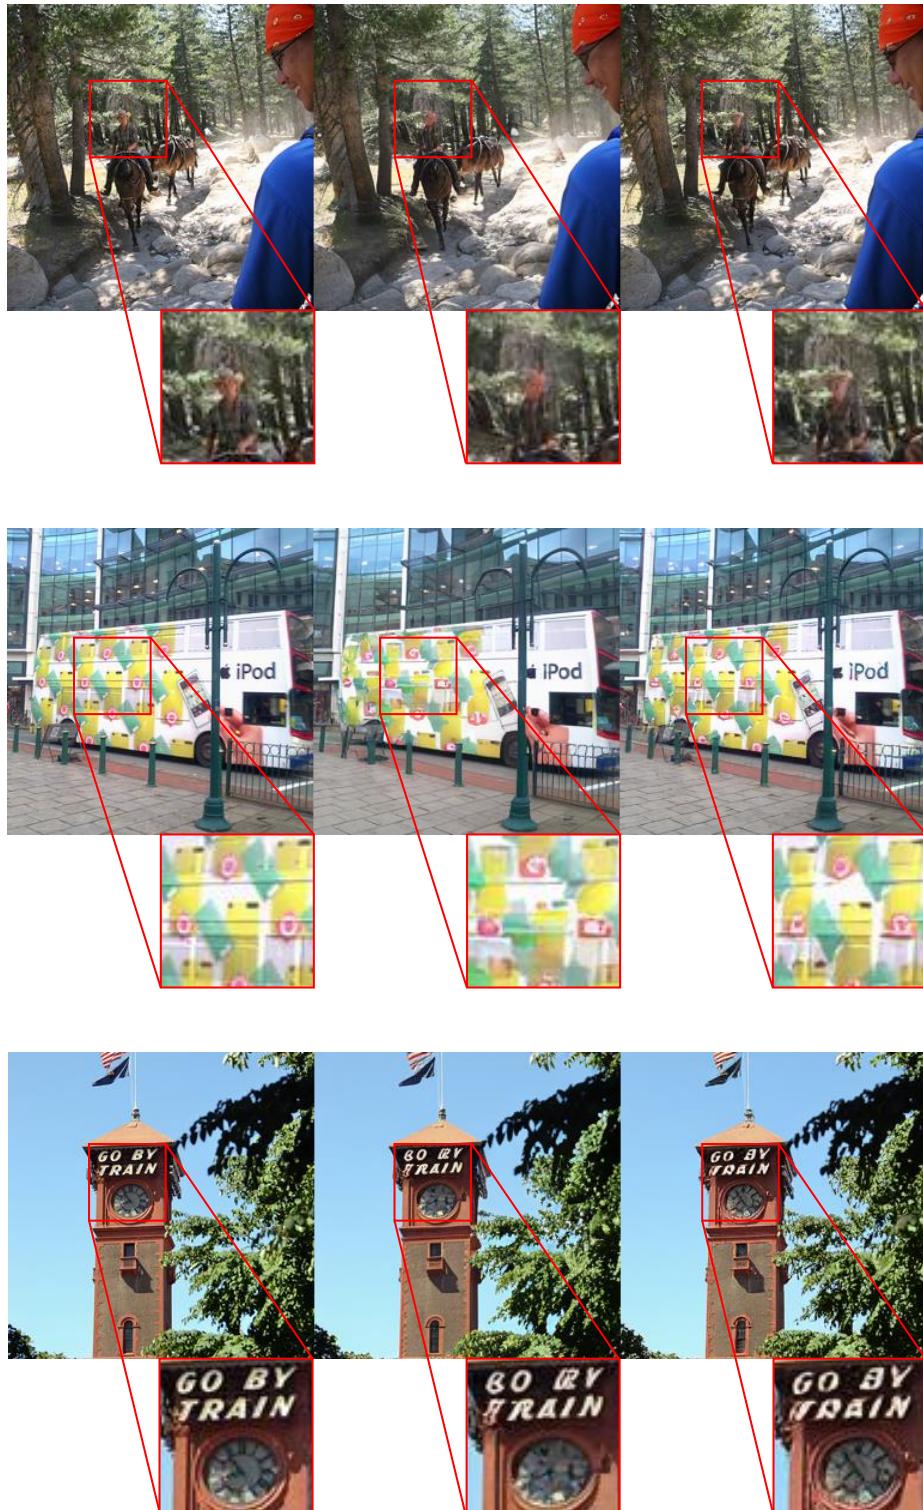


Figure 7: Reconstruction results of BSQ-ViT (right) compared to the original image (left) and SDXL-VAE [14] (middle). The three images are taken from COCO 2017val which are more scene-centric compared to ImageNet data that our model is trained on.

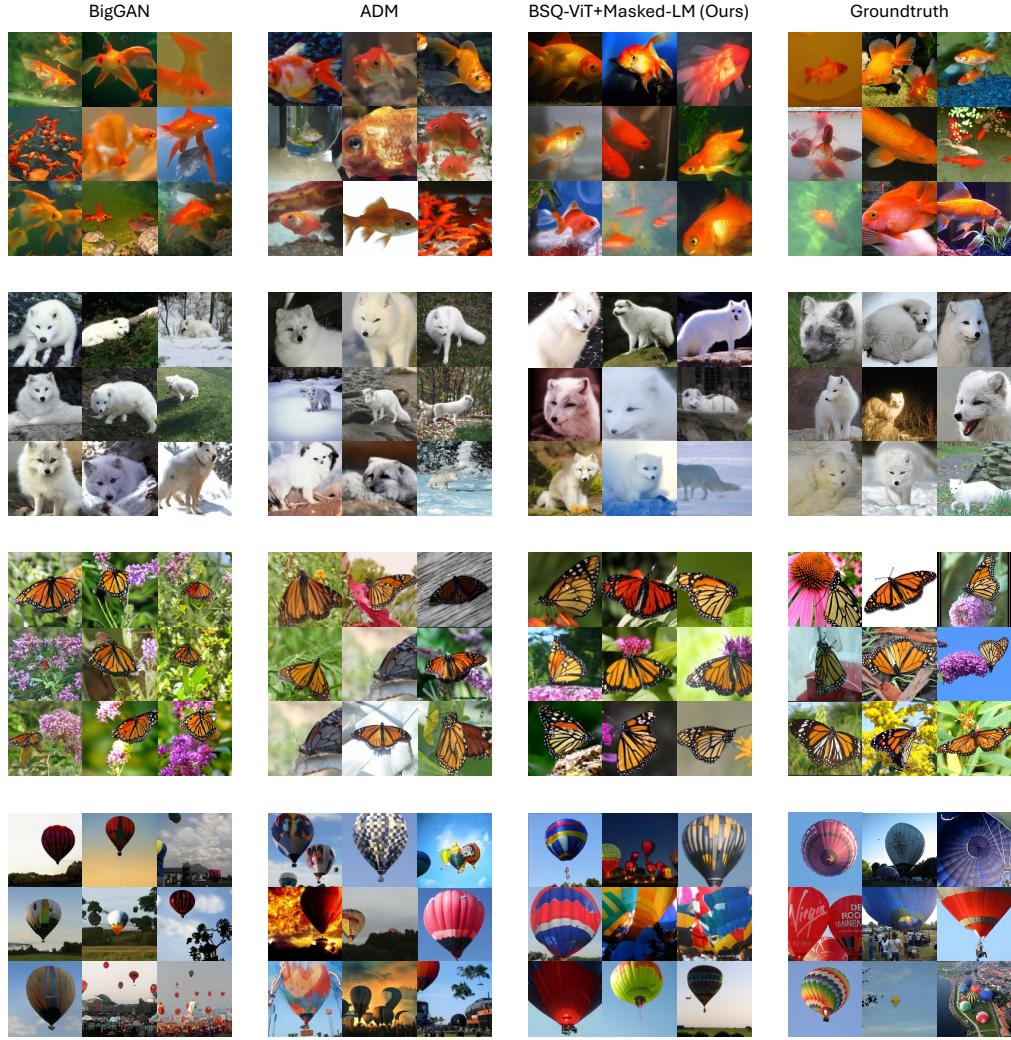


Figure 8: Sampled generation results of BSQ-ViT + Masked-LM (**second column from left**) compared to BigGAN [18] (**right**), ADM [19] (**second column from right**) and the original images (**left**). Classes are 1: goldfish, 279: arctic fox, 323: monarch butterfly, 417: balloon.

J Licenses for Existing Assets

J.1 Datasets

ImageNet The terms of service are available on <https://image-net.org/about>.

COCO The terms of service are available on <https://cocodataset.org/#termsosuse>.

MCL-JCV Copyright is available on the website <https://mcl.usc.edu/mcl-jcv-dataset/>.

UVG⁴ The dataset is licensed under a CC BY-NC 3.0 Deed license.

J.2 Evaluation Metrics

FID score is based on the PyTorch re-implementation⁵. The original implementation⁶ is based on TensorFlow. Both are licensed under an Apache-2.0 License.

⁴<https://ultravideo.fi/dataset.html>

⁵<https://github.com/mseitzer/pytorch-fid>

⁶<https://github.com/bioinf-jku/TTUR>

LPIPS is based on the implementation⁷ licensed under a BSD-2-Clause license.

SSIM and MS-SSIM are based on the PyTorch implementation⁸ licensed under an MIT license.

Generation Metrics (FID, Inception Score, Precision, and Recall) are reported using a TensorFlow implementation⁹ licensed under an MIT license.

J.3 Baseline Methods

DALL-E dVAE¹⁰ is licensed under a Modified MIT License.

SD-VAE 1.x¹¹ is licensed under an MIT License.

SD-VAE 2.x¹² is licensed under an MIT License.

SDXL-VAE¹³ is licensed under an MIT License.

ADM¹⁴ is licensed under an MIT License.

MaskGIT¹⁵ is licensed under an Apache-2.0 License.

CompressAI¹⁶ is licensed under a BSD-3-Clause-Clear license.

FFmpeg is licensed under the GNU LGPL version 2.1 or later. For more detail, please refer to <https://ffmpeg.org/legal.html>.

⁷<https://github.com/richzhang/PerceptualSimilarity>

⁸<https://github.com/VainF/pytorch-msssim>

⁹<https://github.com/openai/guided-diffusion/tree/main/evaluations>

¹⁰<https://github.com/openai/DALL-E>

¹¹<https://github.com/CompVis/latent-diffusion>

¹²<https://huggingface.co/stabilityai/sd-vae-ft-mse>

¹³<https://huggingface.co/stabilityai/sdxl-vae>

¹⁴<https://github.com/openai/guided-diffusion>

¹⁵<https://github.com/google-research/maskgit/tree/main>

¹⁶<https://github.com/InterDigitalInc/CompressAI>