

Spectral Image Tokenizer

Carlos Esteves

Mohammed Suhail

Ameesh Makadia

{machc, suhailmhd, makadia}@google.com

Google Research

Abstract

Image tokenizers map images to sequences of discrete tokens, and are a crucial component of autoregressive transformer-based image generation. The tokens are typically associated with spatial locations in the input image, arranged in raster scan order, which is not ideal for autoregressive modeling. In this paper, we propose to tokenize the image spectrum instead, obtained from a discrete wavelet transform (DWT), such that the sequence of tokens represents the image in a coarse-to-fine fashion. Our tokenizer brings several advantages: 1) it leverages that natural images are more compressible at high frequencies, 2) it can take and reconstruct images of different resolutions without retraining, 3) it improves the conditioning for next-token prediction – instead of conditioning on a partial line-by-line reconstruction of the image, it takes a coarse reconstruction of the full image, 4) it enables partial decoding where the first few generated tokens can reconstruct a coarse version of the image, 5) it enables autoregressive models to be used for image upsampling. We evaluate the tokenizer reconstruction metrics as well as multiscale image generation, text-guided image upsampling and editing.

1. Introduction

In natural language processing, tokenization associates sets of characters to entries in a vocabulary, which is straightforward since language is discrete. The sequence of tokens is suitably modeled as a categorical distribution with autoregressive (AR) transformers, which is the foundation of modern large language models (LLMs) [33, 46, 47].

While natural images are represented by discrete pixel values, they exhibit high dimensionality, redundancies, and noise so it’s impractical to associate one token per pixel. This motivated a long line of research of learnable image tokenizers [15, 37, 37, 51, 57]. While there are successful autoregressive image generation models [15, 57, 58],

images are not sequential like language, which motivated developing alternatives such as denoising diffusion models [11, 38, 43] and masked transformers [3, 4, 60]. Nevertheless, most of these alternatives also operate on the latent space of tokenizers like VQGAN [15] instead of raw pixels.

In this work, we revisit AR transformer-based image generation. Our main contribution is a tokenizer operating on the image spectrum, specifically on DWT coefficients, where the coarse-to-fine representation lends itself more naturally to a sequential interpretation. Our Spectral Image Tokenizer (SIT) has several useful properties and enables different applications:

- P.1 Since the power spectrum of natural images decreases with frequency, high frequencies can be more heavily compressed with little effect in visual quality. SIT leverages this by associating tokens to larger patches at higher wavelet scales than at lower scales (see Fig. 2).
- P.2 SIT is transformer-based [52]; by using an attention mask where each scale depends on itself and lower scales (“Scale-Causal attention”), SIT can be trained at a single resolution and used to tokenize images of multiple resolutions (any number of scales up to the trained maximum), and detokenize partial token sequences (up to some scale), reconstructing a coarse image.
- P.3 Using SIT, we train an autoregressive generative transformer (AR-SIT) that models images coarse-to-fine. The next-token prediction is then conditioned on a coarse reconstruction of the image given by the partial token sequence, instead of the usual conditioning on the partial reconstruction of the previous rows of the image.
- P.4 AR-SIT can quickly generate only the first few tokens and reconstruct a coarse version of the image, enabling applications like quickly showing multiple coarse generations and letting the user select which ones to refine.
- P.5 AR-SIT can be used for text-based upsampling of an input low resolution image, by starting the decoding process with the few tokens output by SIT, and generating the rest of the sequence up to a desired resolution.
- P.6 AR-SIT can be used for text-guided image editing, by

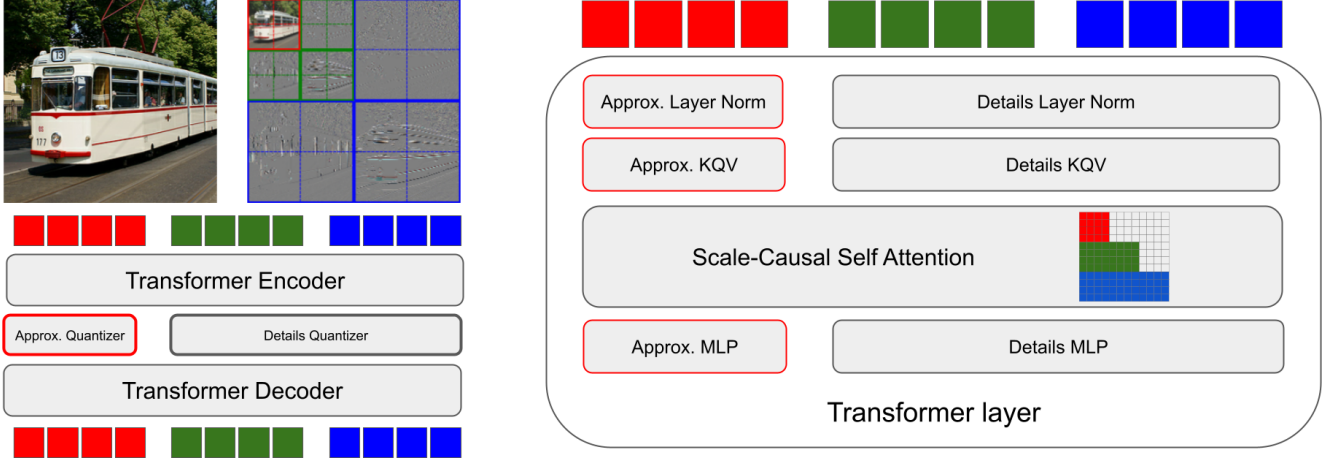


Figure 1. *Left*: we introduce a Spectral Image Tokenizer (SIT), that learns to encode and decode discrete wavelet transform (DWT) coefficients to and from a small set of discrete tokens, such that the sequence represents the image in a coarse-fine fashion. SIT is naturally multiscale and enables coarse-to-fine autoregressive image generation with our AR-SIT model. SIT also leverages the sparsity of high frequency coefficients in natural images. *Right*: details of the encoder/decoder transformer layers. The main architectural difference with respect to previous tokenizers is that the distributions of DWT approximation and details coefficients are distinct, hence we use specialized parameters for each in the quantizer codebooks and inner transformer layers. We also introduce a scale-causal attention where each token attends to its own scale and lower scales, which enables encoding, decoding, generating, and upsampling images at different resolutions.

encoding a given image up to a coarse scale, and generating the finer details conditioned on a new caption.

Currently, image generation is dominated by diffusion models such as Imagen 3 [23], DALL-E-3 [1] and Stable Diffusion 3 [16]. On the other hand, LLMs such as Gemini [46], GPT-4 [33], and Llama 3 [47] are based on autoregressive transformers. We believe autoregressive image generation is still worth pursuing because it might benefit from advances in LLMs, and multimodal applications might benefit of having a single architecture for all modalities. Recent work on image and video generation support this point [45, 48, 59]. Dieleman [12] recently interpreted denoising diffusion models as spectral autoregression, since, when looking at image spectra, the denoising procedure uncovers frequencies from low to high. In contrast, our method does literal spectral autoregression.

2. Related work

Image tokenization Several methods have been developed to map images to a small set of discrete tokens suitable for generative modeling. VQ-VAE [51] introduced vector-quantization in the latent space of a Variational Auto-Encoder to map images, audio and video to a set of discrete values. VQGAN [15] improved upon VQVAE by using perceptual and adversarial losses. We build on ViT-VQGAN [57], which improved upon VQGAN by using a Vision Transformer (ViT) [13] instead of convolutions, as well as codebook factorization and feature normalization.

In this paper, we are interested in multiscale image representations. VQ-VAE-2 [37] introduced multiscale latents

by keeping and quantizing intermediate downsampled convolutional features. RQ-VAE [26] quantized a set of residuals such that the latent vector is represented in a coarse-to-fine fashion and reconstructed by adding the code embeddings for each residual. Similarly, VAR [48] and STAR [28] split the latent space in multi-scale quantized residuals, so that it can be reconstructed by upsampling and summing. FAR [56] splits the latents into frequencies to train a next-frequency generative model. DQ-VAE [22] encoded smooth regions with fewer tokens that are generated first, while QG-VAE [14] and SEMANTICIST [53] learn non-local tokens that are ordered by importance.

The crucial difference between our approach and the aforementioned is that we operate on spectral coefficients of the input and not on latent features. This allows our tokenizer to be truly multiscale; it can take inputs at different scales and reconstruct up to some scale, and leverage that higher frequencies are more compressible in natural images.

Tangentially related to our work, Wave-ViT [55] modified the ViT self-attention by applying a DWT to the input and concatenating coefficients, effectively exchanging space for depth to reduce the sequence length. Zhu and Soricut [62] modified the ViT patchifier in a similar way, and introduced patch embeddings that leverage the sparsity of high frequency coefficients. Both methods are for discriminative tasks such as image classification and segmentation, while we focus on generation.

Autoregressive image generation Early approaches PixelRNN and PixelCNN [49, 50] generate images pixel by pixel by modeling the conditional distribution of the pixel given the previous pixels with recurrent layers of causal

convolutions. PixelSNAIL [5] improved on this model by introducing self-attention layer to better model long-range dependencies. VQ-VAE [51] introduced a two-stage approach with a tokenizer and a separate stage to model the distribution of tokens. VQGAN [15] greatly improved these results by both improving the tokenizer and using a transformer to model the distribution. Finally, ViT-VQGAN [57] proposed a transformer-based tokenizer, which Parti [58] used with a large autoregressive transformer capable of high-quality text-to-image generation.

MaskGIT [3] and Muse [4] highlighted the disadvantages of the typical autoregressive models raster-order conditioning, and proposed to generate all tokens in parallel iteratively, where each iteration keeps the highest confidence tokens. We address the same problem with a tokenizer whose sequence represents the image in a coarse-to-fine order instead of raster-order.

Multiscale image generation Multiscale image generation ideas have appeared in the context of VAEs [2], GANs [24, 39], and diffusion models [18, 25, 40], but have not been sufficiently explored with AR transformers.

Nash et al. [32] represented an image as a sequence of quantized and thresholded DCT coefficients, where the compression comes from the fact that many coefficients are zero and are omitted. By sorting the sequence by frequency, the method can model images from coarse-to-fine like we do. However, the compressed representation is handcrafted and results in long sequences. In a similar vein, Mattar et al. [30] handcrafted a tokenizer based on DWT coefficients, introducing tokens to represent large chunks of zeros. It also results in long sequences and is only applied to generations of small grayscale images. In contrast with these methods, instead of handcrafting a compressed input representation, our tokenizer learns to encode to, and decode from, a short compressed coarse-to-fine sequence.

3. Background

A discrete wavelet transform (DWT) is based on successive convolutions of the signal f with a pair of lowpass g and highpass h filters with the first step as follows

$$f_{\text{low}_1}[n] = f \star g = \sum_k f[k]g[n-k], \quad (1)$$

$$f_{\text{high}_1}[n] = f \star h = \sum_k f[k]h[n-k]. \quad (2)$$

The high and low outputs are subsampled by a factor of two, and the operation is repeated for the low channel, such that at level L we compute $f_{\text{low}_L} = f_{\text{low}_{L-1}} \star g$ and $f_{\text{high}_L} = f_{\text{low}_{L-1}} \star h$, subsample again and drop $f_{\text{low}_{L-1}}$. The output at level L comprises the approximation coefficients f_{low_L} and the detail coefficients $\{f_{\text{high}_k}\}$ for $1 \leq k \leq L$. This output has the same cardinality as the input and the

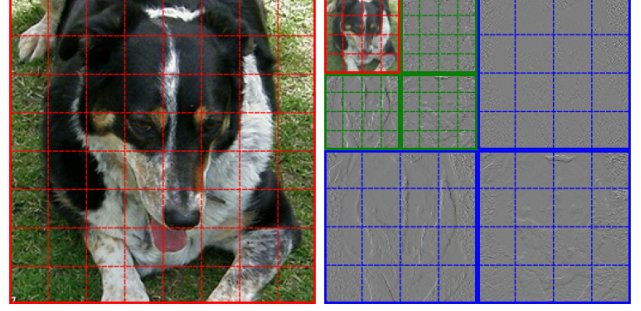


Figure 2. Input patchification. *Left*: typical patchification for Vision Transformers (ViT) [13], where the image is split in equal-sized patches. *Right*: we propose to patchify the coefficients of a discrete wavelet transform (DWT) instead. Each scale is shown in a different color. Scales other than the lowest contain three blocks representing horizontal, vertical and diagonal details; we concatenate the spatially correspondent patches of each block such that each scale is represented by the same sequence length. This results in larger patch sizes for higher scales, which are more compressible. The figure shows 3 scales and 16 tokens per scale; in our experiments we use 4 or 5 scales and 256 tokens per scale.

transformation is invertible, where the forward transform is typically called “analysis” and the backward “synthesis”. The simplest wavelet family is the Haar, where $g = [1, 1]^T$ and $h = [1, -1]^T$ (optionally scaled to unit norm).

For image analysis we use a 2D DWT, which is obtained by simply convolving the rows and columns with g and h . The approximation coefficients $f_{\text{low}_1} = f \star (gg^T)$, and the details are divided into horizontal $f_{H_1} = f \star (gh^T)$, vertical $f_{V_1} = f \star (hg^T)$ and diagonal $f_{D_1} = f \star (hh^T)$ details. Subsequent levels apply the same operations to the approximation coefficients. Fig. 2 shows a two-level transform, where we can see that the approximation coefficients correspond to a coarse version of the image.

We refer to the textbooks by Mallat [29] and Daubechies [9] for more information about wavelets.

4. Method

Our main contribution, SIT, is an image tokenizer that takes discrete wavelet transform (DWT) coefficients. The model follows ViT-VQGAN [57], with important changes that we describe in this section and visualize in Fig. 1. For image generation, we introduce AR-SIT which is based on Parti [58] with minor changes, using SIT as the tokenizer.

4.1. Tokenizer

Patchification and featurization The first step is to map the input to a sequence of patches. We apply the Haar DWT on the input image and patchify each scale separately. While Haar is the simplest wavelet and lacks properties found in other wavelets useful for compression, we found no benefits of using other wavelet families such as CDF [8].

Our design choice is to use the same number of patches for each scale. In a DWT, the higher scales correspond to high frequency details which are represented by more coefficients than the lower scales, but contribute less to the spatial pixel values. In other words, in natural images, most of the power spectrum is concentrated on lower frequencies. By representing each scale with the same number of tokens, we are compressing more the higher frequencies (since they are represented by more coefficients), similarly to what is done in image compression methods such as JPEG2000 [7].

The approximation (or lowpass) DWT coefficients correspond to a coarse version of the input image, which we consider the first scale. The following scales are divided in three blocks corresponding to horizontal, vertical, and diagonal details, where each coefficient relates to a specific spatial location. Thus, we can concatenate the three blocks such that each entry corresponds to the same spatial location. For example, the first scale will typically be split in patches that are $32 \times 32 \times 3$, with channels corresponding to RGB, the second scale will be $32 \times 32 \times 9$, where the channels correspond to the RGB of horizontal, vertical, and diagonal details, the third will be $64 \times 64 \times 9$ and so on. Fig. 2 shows an example of our patchification scheme.

Since patches have different resolutions (higher scales will have larger patches), the usual ViT linear embedding to map patches to features cannot be shared across all patches so we have different parameters per scale. Formally, given an image f , our patchifiers compute

$$f_{\text{low}_L}, \{f_{H_i}\}_{i \leq L}, \{f_{V_i}\}_{i \leq L}, \{f_{D_i}\}_{i \leq L} = \text{DWT}(f), \quad (3)$$

$$c_1 = P_0(f_{\text{low}_L}), \quad (4)$$

$$c_s = P_s(f_{H_{L-s+2}}, f_{V_{L-s+2}}, f_{D_{L-s+2}}), \quad 1 < s \leq L, \quad (5)$$

where $c_s = \{c_s^n\}_{1 \leq n \leq N}$ is the sequence of token embeddings at scale s , and $c_s^n \in \mathbb{R}^C$ is the embedding of the n -th token at the s -th scale, while L is the number of DWT levels, $S = L + 1$ is the number of scales and N the number of tokens per scale. For brevity, when there is no ambiguity, we may omit the set indexing and use $\{c_s\}$ to denote the set of tokens of all scales, for example.

The final projection after the decoder needs to map the features back to different sized patches, so it also has different parameters per scale. Those patches still represent DWT coefficients so an inverse DWT (IDWT) is finally applied to obtain an image output.

Flexible sequence length The tokenizer encoder and decoder transformers operate on the sequence of patch features of length SN . The sequence length is a major factor in resource utilization so we want to keep it constrained. Our method is flexible since we can choose the number of scales and the patch size per scale, while most ViT-based models such as ViT-VQGAN [57] are more restricted. They use the same patch size for the whole image; thus, keeping the

same patch size and doubling each image dimension would increase the sequence length by a factor of 4, where our method is capable of including additional scales which only increases the sequence length by multiples of N .

For example, for a 256×256 input, ViT-VQGAN uses 8×8 patches to obtain a sequence of 1024 elements. Our SIT can use 4 scales and 256 tokens per scale for the same sequence length. When increasing the resolution to 512×512 , the baseline can either increase the patch size to 16×16 , resulting in the same sequence length, or keep it 8×8 , resulting in a $4\times$ longer sequence. SIT, for example, can vary the number of scales from 4 to 6, resulting in sequence lengths of 1024, 1280, 1536 which are more manageable.

Transformers After featurization, the single sequence containing all scales passes through a transformer encoder, followed by quantization and a transformer decoder. We propose two optional modifications to the transformers, which are otherwise identical to the ones used in ViT [13].

First, we propose a scale-causal attention mask, where an element at some scale attends to all elements of its own scale and lower scales, represented by a block-triangular mask pattern. With dense attention, we write the application of the encoding transformer T_{enc} as $\{z_s\} = T_{\text{enc}}(\{c_s\})$, and for scale-causal we write $z_s = T_{\text{enc}}(\{c_k\}_{1 \leq k \leq s})$ for each s . The scale-causal attention can be applied independently to the encoder and decoder, enabling different applications. For the multiscale reconstruction experiments in Sec. 5.1, we need to both encode and decode multiple resolutions, so both encoder and decoder use scale-causal masks. For coarse-to-fine image generation in Sec. 5.2, only the decoder needs to be scale-causal in order to decode the partially generated sequence. For the text-guided image upsampling in Sec. 5.3, only the encoder needs to be scale-causal to encode the lower resolution inputs. For the image editing experiments in Sec. 5.4, the scale-causal encoder prevents information leaking from high to low scale.

Second, we propose to use different transformer parameters for the approximation (first scale) and details coefficients (other scales). In our model, the subsequences corresponding to coefficients of each type of coefficient come from quite distinct distributions, so it makes sense to treat them differently. This contrasts with the spatial representation of images where each patch can be considered as coming from the same distribution. Thus, the parameters of the key, query, and value embeddings, the layer norms, and the MLP on each transformer layer are not shared between the approximation and details coefficients. This Approximation-Details Transformer (ADTransformer) still takes a single sequence composed of all coefficients. We experimented with different transformers per sequence with cross-attention for information sharing, but it performed worse. This change leads to more memory utilization to store the extra parameters, but the training/inference speed

	LPIPS ↓	PSNR ↑	L1 ↓	FID ↓	IS ↑	images/s ↑
<i>Resolution: 512 × 512</i>						
ViT-VQGAN	0.320	22.4	0.042	6.92	151.5	593
SIT-5 (Ours)	0.260	22.0	0.046	2.65	192.0	410
SIT-6 (Ours)	0.239	23.1	0.040	1.74	203.7	320
<i>Resolution: 256 × 256</i>						
ViT-VQGAN (reported)	-	24.8	0.032	1.99	184.4	-
ViT-VQGAN (reproduced)	0.167	25.0	0.031	2.33	184.0	-
ViT-VQGAN (no LL)	0.163	23.8	0.038	1.20	194.6	626
SIT-4 (Ours)	0.144	24.0	0.037	1.20	199.5	596
SIT-5 (Ours)	0.135	24.5	0.035	0.97	202.3	411
SIT-SC-5 (Ours)	0.161	24.1	0.037	1.33	193.7	411
<i>Resolution: 128 × 128</i>						
ViT-VQGAN	0.185	26.3	0.030	3.77	117.3	626
SIT-SC-5 (ours)	0.159	27.1	0.027	2.13	129.3	582
<i>Resolution: 64 × 64</i>						
ViT-VQGAN	0.129	28.8	0.023	3.53	21.0	627
SIT-SC-5 (ours)	0.111	31.3	0.017	1.39	30.1	847
<i>Resolution: 32 × 32</i>						
ViT-VQGAN	0.214	23.3	0.045	-	3.7	627
SIT-SC-5 (ours)	0.029	36.8	0.010	0.31	3.5	825
<i>Resolution: 16 × 16</i>						
ViT-VQGAN	0.127	24.9	0.039	-	1.7	627
SIT-SC-5 (ours)	0.013	41.3	0.006	0.09	1.8	2620

Table 1. Multiscale reconstruction on ImageNet. “SC” denotes scale-causal attention, which slightly reduces performance at the highest resolution but enables multiscale reconstruction without downsampling/upsampling or retraining. The ViT-VQGAN values from Yu et al. [57] used a logit-laplace loss (LL) which was later considered harmful [58], so we retrain without it. Our SIT improves reconstruction metrics, is significantly faster at lower resolutions, and robust when increasing resolution. The ViT-VQGAN baseline suffered from instability during training on 512×512 inputs so we selected the best values before divergence. We report test time throughput of an encoding/decoding cycle for the max batch size that fits on a TPU v5e.

is similar because the number of operations is unchanged.

Quantizer The encoder outputs a sequence of features that are quantized to a fixed-sized codebook, similarly to ViT-VQGAN [57], VQVAE [51], and VQGAN [15]. We modify the quantizer such that approximation and details have different codebooks for the same reasons discussed previously. Thus, codebook sizes and feature dimensions are the same as the ViT-VQGAN baseline, but we have different features for the same code at different positions. Formally, we apply $q_s^n = Q_{\text{approx}}(z_s^n)$ for $s = 1$ and $q_s^n = Q_{\text{details}}(z_s^n)$ for $s \geq 2$ for each pair (s, n) , where q_s^n is chosen from the codebook for s and can be associated with its discrete position in the codebook, denoted $\lfloor q_s^n \rfloor$.

Training We follow the ViT-VQGAN training protocol and use the same weighting for the L2, perceptual, adversarial, and quantization losses. We remove the logit-laplace loss that was shown detrimental in follow-up work [58].

All losses are applied to the spatial domain images reconstructed by the inverse discrete wavelet transform on the decoder output: $\hat{f} = \text{IDWT}(T_{\text{dec}}(\{q_s\}))$. We noticed insta-

	FID ↓	IS ↑	images/s ↑	images/Gb ↑
<i>Resolution: 256 × 256</i>				
Parti350M (reported)	14.1	-	-	-
Parti350M	12.4	36.5	7.8	12.0
AR-SIT-SCD-4 (Ours)	12.6	37.3	6.5	8.0
<i>Resolution: 128 × 128</i>				
Parti350M	11.2	33.5	7.6	12.0
AR-SIT-SCD-4 (Ours)	11.4	33.2	12.6	12.0
<i>Resolution: 64 × 64</i>				
Parti350M	10.5	16.9	7.6	12.0
AR-SIT-SCD-4 (Ours)	11.4	18.6	24.5	16.0
<i>Resolution: 32 × 32</i>				
Parti350M	5.8	2.9	7.7	7.7
AR-SIT-SCD-4 (Ours)	7.6	3.2	74.7	28.0

Table 2. Coarse-to-fine image generation on MS-COCO [27]. With SIT, the autoregressive generation is stopped early for a coarse version of the image. The Parti350M baseline does not have this property, so we generate at full resolution and downsample for comparison. AR-SIT matches the baseline performance at the training resolution but is several times faster and more memory efficient at lower resolutions, even when trained only on higher resolution data. We report throughput and memory utilization during generation given the max batch size that fits on a TPU v6e.

bility during training due to the adversarial loss, which was fixed by applying spectral normalization following Miyato et al. [31], which simply divides the discriminator weight matrices by their largest singular value.

4.2. Autoregressive image generation

We use our tokenizer for autoregressive image generation, by training a second stage transformer model similar to Parti [58], with some modifications. Formally, the autoregressive transformer T models categorical distributions over the discrete codes

$$P(\lfloor q_s^n \rfloor \mid \{\lfloor q_i \rfloor\}_{1 \leq i < s} \cup \{\lfloor q_s^i \rfloor\}^{1 \leq i < n}) = T(c) \quad (6)$$

which can be sampled one code at a time for generation. T can be conditioned on a textual description processed by a transformer encoder for text-to-image generation. For training, we model the distribution of input codes as $P(\{\lfloor q_s \rfloor\}) = \prod_{1 \leq s \leq N}^{1 \leq n \leq N} P(\lfloor q_s^n \rfloor)$ and minimize the negative log-likelihood $-\log P(\{\lfloor q_s \rfloor\})$ over the training set.

With different codebooks for approximation and details, the same token id might have a different meaning depending on its position. Thus, the AR model has different token embeddings for approximation and details tokens. The last layer for logit prediction is also different.

For the generative applications in Secs. 5.2 and 5.3, we introduce mild changes in order to interrupt the generation after all tokens up to a certain scale are generated, and to start the generation with given tokens up to a certain scale.

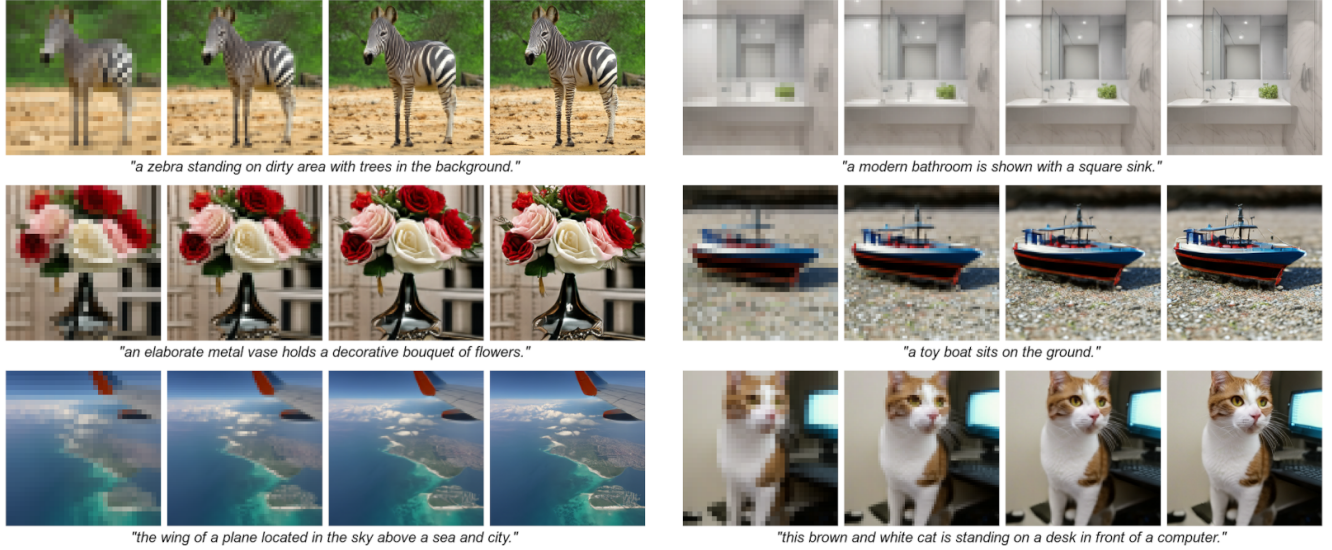


Figure 3. Coarse-to-fine text-to-image on MS-COCO [27] prompts. Each quadruple shows generations from AR-SIT-SCD for the given prompt, with only the first 25%, 50%, 75% and 100% of tokens generated, corresponding to resolutions of 32×32 to 256×256 . Training uses only 256×256 images. Our model enables quick generation of coarse image candidates that can be further improved if needed.

5. Experiments

We focus our experiments on demonstrating the properties and applications enumerated in P.1-P.6 (Sec. 1). Our ideas build on ViT-based tokenizers, so the most fair comparisons are against ViT tokenizers and AR generative models; namely ViT-VQGAN [57] and Parti [58]. In these comparisons we are able to match the architectures and training protocol exactly. Nevertheless we also compare with a broader set of methods in Sec. 5.5.

5.1. Multiscale reconstruction

We train our tokenizer on ImageNet [10] and evaluate its reconstruction performance. We follow the ViT-VQGAN [57] “Small” architecture and training protocol.

The model “SIT-5” is the Spectral Image Tokenizer with 5 scales from 16×16 to 128×128 resolutions, where each scale is represented by 256 tokens. The variation “SIT-SC” uses scale-causal attention on both the encoder and decoder, which enables handling inputs and outputs of different resolutions, even though it is only trained at 256×256 . For example, when the input image is 64×64 , only the first three scales are used, resulting in shorter sequence lengths which reduces memory utilization and processing time. Both models employ the ADTransformer described in Sec. 4.1.

The ViT-VQGAN baseline only works for 256×256 inputs, so to evaluate against it fairly, we upsampled the low-resolution inputs to that resolution. This brings the input patches away from the training distribution, which might explain the drop in reconstruction quality.

We also evaluate at 512×512 . The ViT-VQGAN baseline increases its patch size from 8×8 to 16×16 , keep-

ing the sequence length constant. Interestingly, the baseline suffered heavy instability during training, which was not resolved by reducing the learning rate, using spectral normalization, or the logit-laplace loss. In contrast, our SIT variations trained successfully with no hyperparameter changes.

Tab. 1 shows the reconstruction metrics; Fig. 7 shows reconstruction samples at multiple resolutions. Results provide evidence for P.1 and demonstrate P.2 (Sec. 1).

5.2. Coarse-to-fine text-to-image generation

We tackle text-to-image generation by using an autoregressive transformer to model the distribution of discrete tokens output by our tokenizer. Since the SIT sequence of tokens represents an image in a coarse-fine fashion, the autoregressive model generation has the same property, which means that we can interrupt the generation after a certain number of tokens and decode a coarse version of the generation.

For this to work, the SIT decoder must be scale-causal, while there is no constraint for the encoder so we use dense attention (we denote this variant Scale-Causal Decoder, or “SIT-SCD”). It uses a “small” encoder and “large” decoder as described in ViT-VQGAN [57], and 4 scales.

We follow the Parti350M [58] architecture and training protocol and name our generative model auto-regressive SIT, or “AR-SIT”. SIT and AR-SIT are trained on a subset of WebLI [6] of around 128M images, where each image is seen at most once during training. Evaluation is on 30k examples of MS-COCO [27]. We show metrics and generations at different resolutions in Fig. 3 and Tab. 2. AR-SIT matches the baseline performance but is several times faster and more memory efficient at lower resolutions. These results demonstrate the property P.4 (Sec. 1).

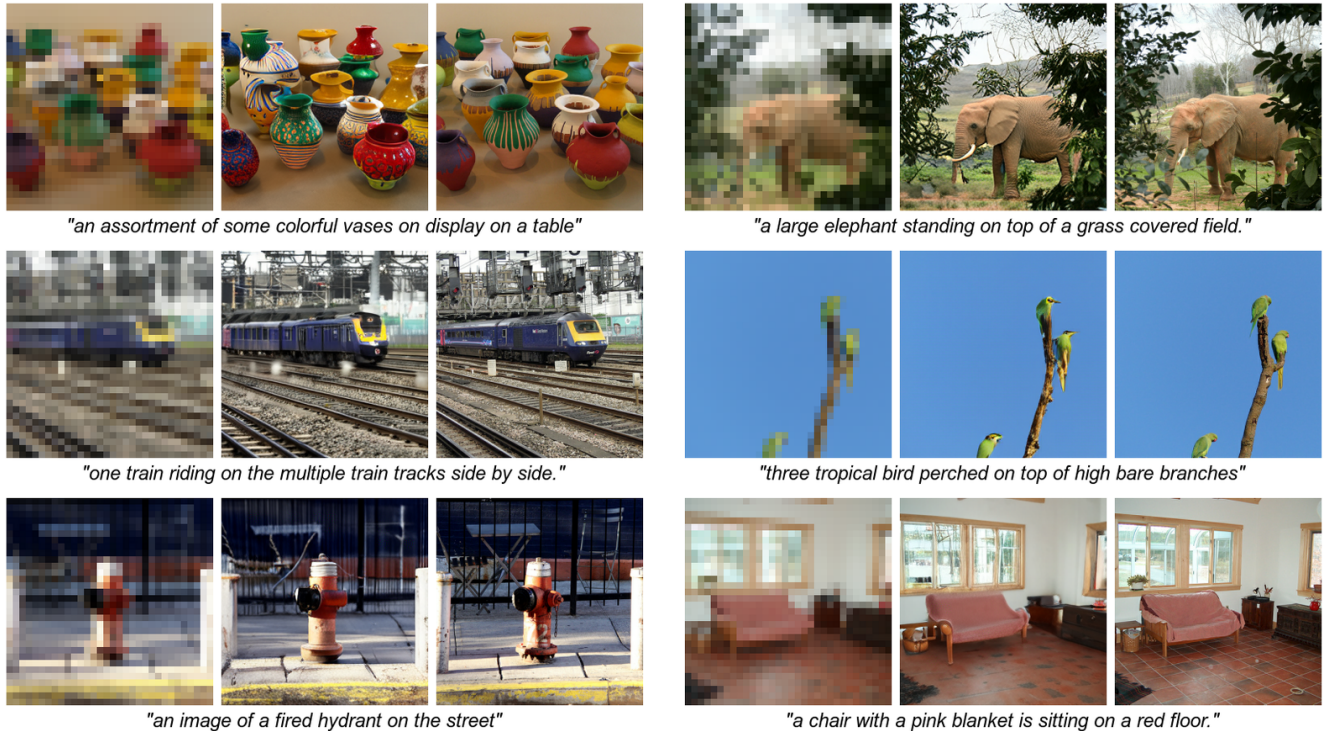


Figure 4. Text-guided image upsampling on MS-COCO [27]. Our coarse-to-fine generative models can take a low-resolution image, encode it as the first few tokens of a sequence, and generate the rest of sequence, which, when decoded, effectively upsamples the input. Each triplet shows the given 32×32 image, our 256×256 reconstruction and the ground truth.

5.3. Text-guided image upsampling

We leverage the coarse-to-fine nature of SIT to apply a pre-trained AR-SIT for text-guided image upsampling, where we are given a low resolution image and a caption.

The idea is to encode the low-resolution image, which will give the first tokens of our high resolution output. AR-SIT then takes these tokens and generates the rest of the sequence. For this to work, the SIT encoder must be scale-causal to properly tokenize low-resolution inputs, while there is no constraint for the decoder so we use dense attention. We denote this variant Scale-Causal Encoder, or “SIT-SCE”. The model is otherwise as the one used in Sec. 5.2.

Here the inputs are 32×32 . Fig. 4 shows text-guided upsampling examples on MS-COCO [27]. We obtain an FID of 6.2 when evaluating these generations, compared to 12.6 when given only the prompts. Fig. 8 and Tab. 4 show additional results upsampling from 16×16 .

5.4. Text-guided image editing

Our coarse-to-fine AR-SIT enables a text-guided image editing application, where we want to change image details while keeping the same overall appearance, which corresponds to freezing lower scales while generating higher.

We apply an AR-SIT trained only on the maximum likelihood objective as follows. The given image is tokenized only up to the first scale, corresponding to a coarse repre-

sentation. The tokenizer encoder must be scale-causal in order to avoid leakage from high to low scales, so we use SIT-SCE here. Now we use AR-SIT to generate the rest of the sequence, conditioning on the textual caption.

In this experiment, we use a 5-scale model such that the lowest resolution is 16×16 ; we found that starting with higher resolutions limits the changes the model can generate. Fig. 5 shows some examples of text-guided editing on MS-COCO [27], where we lightly modify the original captions, for example by swapping “elephants” with “cows”. We show additional results in Fig. 6 in the supplemental.

5.5. Class-conditional image generation

Here we evaluate class-conditional generation on 256×256 ImageNet. We use the SIT-4 tokenizer from Tab. 1 and follow ViT-VQGAN’s [57] VIM-Base AR architecture and training protocol. Tab. 3 shows we outperform the baseline in this fair comparison which is evidence for the improved conditioning described in P.3. Fig. 9 shows samples.

AR-SIT-4* departs from ViT-VQGAN and Parti hyperparameters to compare with more recent works, by increasing the perceptual and adversarial loss weights, doubling the tokenizer codebook, reducing its latent dimension, training with a constant learning rate, including 2D RoPE [44] and GeGLU [42], in both the tokenizer and generator. We train the AR model longer and since there is no text encoder, we



Figure 5. Text-guided image editing on MS-COCO [27]. Our coarse-to-fine generative models can do text-guided editing by encoding a given image but keeping only the lower scales, and using a pre-trained AR-SIT to re-generate the higher scales conditioned on the textual prompt. Each triplet shows the given image, its 16×16 reconstruction using only the coefficients used to start the generation, and the edited image after generating the whole sequence.

	params	seq len	CFG	FID ↓	IS ↑
AR-ViT-VQGAN (reported)	650M	1024	-	8.81	110.8
AR-ViT-VQGAN (reproduced)	650M	1024	-	8.37	111.8
AR-SIT-4 (Ours)	650M	1024	-	6.95	138.3
RQ-VAE [26]	480M	256	-	15.72	86.8
DQ-VAE [22]	355M	640	-	7.34	152.8
LlamaGen-L [45]	343M	256	1.5	4.08	198.5
VAR [48]	310M	680	2.0	3.30	274.4
AR-SIT-4* (Ours)	350M	1024	1.5	4.06	190.9

Table 3. ImageNet 256×256 class conditional generation. Top half shows a fair comparison against ViT-VQGAN, which we clearly outperform. AR-SIT-4* uses different hyperparameters for comparison against a broader class of methods. While we do not outperform all of them, there are differences in the architectures and training schedules that influence the results and are orthogonal to our contributions. Moreover, none of the baselines has the multiscale capabilities described in P.2 and P.4 (Sec. 1).

double the number of decoder layers. The bottom half of Tab. 3 compares with similarly-sized AR models including some with multiscale/residual latents (although none exhibit properties P.2 and P.4). AR-SIT-4* performance matches LlamaGen but not to VAR, but the models are not exactly comparable since they use a convolutional tokenizers and other improvements such as AdaLN [34] and attention normalization; VAR also trains the tokenizer on larger

data than ImageNet. These are orthogonal to our contributions and can be incorporated in future work.

We also train an AR-SIT-5* on 512×512 ImageNet and obtain an FID of 5.97 and IS of 215.22, showing robustness at higher resolutions, while most prior work requires separate stages for upsampling. Fig. 10 shows samples.

6. Conclusion and limitations

We presented a spectral image tokenizer (SIT), and an autoregressive generative transformer trained with it (AR-SIT). SIT is naturally multiscale and leverages spectral properties of natural images for improved reconstruction quality. AR-SIT enables applications such as rapid generation of coarse images that can be refined later, and text-guided image upsampling and editing.

While our methods improved tokenizer reconstruction accuracy and class-conditional generation, the text-to-image metrics were similar to the fair baseline. Chang et al. [4] similarly observed that a better tokenizer does not necessarily lead to a better generative model. Nevertheless, our method has multiscale properties and enables new applications not possible with prior work. We only experiment with relatively small AR-SITs of 350-650M parameters, while the Parti [58] baseline goes up to 22B parameters.

7. Acknowledgments

We thank Leonardo Zepeda-Núñez for reviewing this manuscript and offering interesting discussions and suggestions, and Jon Barron for sharing useful code we relied on.

References

- [1] James Betker, Gabriel Goh, Li Jing, † TimBrooks, Jianfeng Wang, Linjie Li, † LongOuyang, † JuntangZhuang, † JoyceLee, † YufeiGuo, † WesamManassra, † PrafullaDhariwal, † CaseyChu, † YunxinJiao, and Aditya Ramesh. Improving image generation with better captions. [2](#)
- [2] Lei Cai, Hongyang Gao, and Shuiwang Ji. *Multi-Stage Variational Auto-Encoders for Coarse-to-Fine Image Generation*, page 630–638. Society for Industrial and Applied Mathematics, 2019. [3](#)
- [3] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative image transformer. In *CVPR*, 2022. [1, 3](#)
- [4] Huiwen Chang, Han Zhang, Jarred Barber, Aaron Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Patrick Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li, and Dilip Krishnan. Muse: Text-to-image generation via masked generative transformers. In *ICML*, 2023. [1, 3, 8](#)
- [5] XI Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. PixelSNAIL: An improved autoregressive generative model. In *ICML*, 2018. [3](#)
- [6] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Nan Ding, Keran Rong, Hassan Akbari, Gaurav Mishra, Linting Xue, Ashish Thapliyal, James Bradbury, Weicheng Kuo, Mojtaba Seyedhosseini, Chao Jia, Burcu Karagol Ayan, Carlos Riquelme, Andreas Steiner, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. Pali: A jointly-scaled multilingual language-image model, 2022. [6](#)
- [7] C. A. Christopoulos, T. Ebrahimi, and A. N. Skodras. Jpeg2000: the new still picture compression standard. In *Proceedings of the 2000 ACM Workshops on Multimedia*, page 45–49, New York, NY, USA, 2000. Association for Computing Machinery. [4](#)
- [8] A. Cohen, Ingrid Daubechies, and J.-C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 45(5):485–560, 1992. [3](#)
- [9] Ingrid Daubechies. *Ten lectures on wavelets*. Society for Industrial and Applied Mathematics, USA, 1992. [3](#)
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [6](#)
- [11] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. [1](#)
- [12] Sander Dieleman. Diffusion is spectral autoregression, 2024. [2](#)
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. [2, 3, 4](#)
- [14] Tim Elsner, Paula Usinger, Victor Czech, Gregor Kobsik, Yanjiang He, Isaak Lim, and Leif Kobbelt. Quantised global autoencoder: A holistic approach to representing visual data, 2024. [2](#)
- [15] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. [1, 2, 3, 5](#)
- [16] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yan-nik Marek, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis, 2024. [2](#)
- [17] Rinon Gal, Dana Cohen Hochberg, Amit Bermanto, and Daniel Cohen-Or. Swagan: a style-based wavelet-driven generative model. *ACM TOG*, 2021. [1](#)
- [18] Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Joshua M. Susskind, and Navdeep Jaitly. Matryoshka diffusion models. In *ICLR*, 2024. [3](#)
- [19] Philippe Hansen-Estruch, David Yan, Ching-Yao Chung, Orr Zohar, Jialiang Wang, Tingbo Hou, Tao Xu, Sriram Vishwanath, Peter Vajda, and Xinlei Chen. Learnings from scaling visual tokenizers for reconstruction and generation, 2025. [1](#)
- [20] 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. In *NeurIPS*, 2017. [2](#)
- [21] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022. [1](#)
- [22] Mengqi Huang, Zhendong Mao, Zhuowei Chen, and Yongdong Zhang. Towards accurate image coding: Improved autoregressive image generation with dynamic vector quantization. In *CVPR*, 2023. [2, 8](#)
- [23] Imagen-Team-Google. Imagen 3, 2024. [2](#)
- [24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *ICLR*, 2018. [3](#)
- [25] Vladimir Kulikov, Shahar Yadin, Matan Kleiner, and Tomer Michaeli. Sinddm: A single image denoising diffusion model. In *ICML*, pages 17920–17930, 2023. [3](#)
- [26] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *CVPR*, 2022. [2, 8](#)
- [27] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*. [5, 6, 7, 8, 2, 3](#)
- [28] Xiaoxiao Ma, Mohan Zhou, Tao Liang, Yalong Bai, Tiejun Zhao, Huaian Chen, and Yi Jin. Star: Scale-wise text-to-image generation via auto-regressive representations, 2024. [2](#)

- [29] Stephane Mallat. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, Inc., USA, 3rd edition, 2008. 3
- [30] Wael Mattar, Idan Levy, Nir Sharon, and Shai Dekel. Wavelets are all you need for autoregressive image generation, 2024. 3
- [31] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 5
- [32] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W. Battaglia. Generating images with sparse representations. In *ICML*, 2021. 3
- [33] OpenAI. Gpt-4 technical report, 2023. 1, 2
- [34] William Peebles and Saining Xie. Scalable diffusion models with transformers. 2023. 8
- [35] Hao Phung, Quan Dao, and Anh Tran. Wavelet diffusion models are fast and scalable image generators. In *CVPR*, 2023. 1
- [36] Vivek Ramanujan, Kushal Tirumala, Armen Aghajanyan, Luke Zettlemoyer, and Ali Farhadi. When worse is better: Navigating the compression-generation tradeoff in visual tokenization, 2024. 1
- [37] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *NeurIPS*, 2019. 1, 2
- [38] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2021. 1
- [39] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *ICCV*, 2019. 3
- [40] Dohoon Ryu and Jong Chul Ye. Pyramidal denoising diffusion probabilistic models, 2022. 3
- [41] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016. 2
- [42] Noam Shazeer. Glu variants improve transformer, 2020. 7
- [43] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 1
- [44] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding, 2021. 7
- [45] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 2, 8
- [46] Gemini Team. Gemini: A family of highly capable multi-modal models, 2023. 1, 2
- [47] Llama team. The llama 3 herd of models, 2024. 1, 2
- [48] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. In *NeurIPS*, 2024. 2, 8
- [49] Aaron van den Oord, Nal Kalchbrenner, Lasse Espeholt, koray kavukcuoglu, Oriol Vinyals, and Alex Graves. Conditional image generation with pixelcnn decoders. In *NeurIPS*, 2016. 2
- [50] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016. 2
- [51] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 1, 2, 3, 5
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1
- [53] Xin Wen, Bingchen Zhao, Ismail Elezi, Jiankang Deng, and Xiaojuan Qi. "principal components" enable a new language of images, 2025. 2
- [54] Jingfeng Yao, Bin Yang, and Xinggang Wang. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *CVPR*, 2025. 1
- [55] Ting Yao, Yingwei Pan, Yehao Li, Chong-Wah Ngo, and Tao Mei. Wave-vit: Unifying wavelet and transformers for visual representation learning. In *ECCV*, 2022. 2
- [56] Hu Yu, Hao Luo, Hangjie Yuan, Yu Rong, and Feng Zhao. Frequency autoregressive image generation with continuous tokens, 2025. 2
- [57] 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. 1, 2, 3, 4, 5, 6, 7
- [58] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *Transactions on Machine Learning Research*, 2022. Featured Certification. 1, 3, 5, 6, 8
- [59] Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A Ross, and Lu Jiang. Language model beats diffusion - tokenizer is key to visual generation. In *ICLR*, 2024. 2
- [60] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. In *NeurIPS*, 2024. 1
- [61] 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. 2
- [62] Zhenhai Zhu and Radu Soricut. Wavelet-based image tokenizer for vision transformers, 2024. 2