

Layered Diffusion Model for One-Shot High Resolution Text-to-Image Synthesis

Emaad Khwaja^{*†‡§}
emaad@berkeley.edu

Abdullah Rashwan^{‡¶}
arashwan@google.com

Ting Chen^{||}
iamtingchen@google.com

Oliver Wang[¶]
olwang@google.com

Suraj Kothawade[¶]
skothawade@google.com

Yeqing Li[¶]
yeqing@google.com

Abstract

We present a one-shot text-to-image diffusion model that can generate high-resolution images from natural language descriptions. Our model employs a layered U-Net architecture that simultaneously synthesizes images at multiple resolution scales. We show that this method outperforms the baseline of synthesizing images only at the target resolution, while reducing the computational cost per step. We demonstrate that higher resolution synthesis can be achieved by layering convolutions at additional resolution scales, in contrast to other methods which require additional models for super-resolution synthesis.

1. Introduction

High-resolution synthesis of images in pixel space poses a challenge for diffusion models. Previous methods have addressed this by either learning in low-dimensional latent space [8, 13] or generating images via cascaded models focused on increasing resolutions [1, 7, 10]. However, these methods introduce additional complexity to the diffusion framework. Hooeboom et al. [5] proposed a simpler diffusion model that synthesizes high-resolution images by using shifted noise scheduling and model scaling, among other optimizations.

In this work, we extend the simple diffusion model framework to perform hierarchical image synthesis at multiple resolutions. Our goal is to design a light-weight model that can capture spatial image features at specific resolution scales within a singular model, rather than training with re-

spect to the highest resolution. We argue that this approach is preferable to cascaded models, as it allows the upscaling process to be guided by a multitude of learned features within the network, rather than from singular output image pixels for each successive super-resolution model.

Recently, Gu et al. [3] demonstrated a similar idea, termed "Matroyshka Diffusion." They report similar improvements when compared to multiple baseline approaches. Our model differs primarily in the use of scaled noise between resolutions (Section 2.2) and training optimizations employed for faster convergence (Section 3).

2. Architecture

2.1. Layered Model

We introduce a new image generation architecture based on a U-Net [9] structure that produces images conditioned on text prompts at multiple resolution scales (Fig. 1). Our architecture differs from a typical U-Net in a number of key ways:

- **Multiple Inputs:** At each resolution scale, we apply an input convolution to expand the input image to the hidden dimension size of the layer of the respective resolution scale.
- **Isolated Downsampling:** Within a standard U-Net architecture, downsampling convolutions are used to decrease tensor resolution. In our model, rather than downsample higher resolutions, the current tensor in the first part of the network, we utilize an input convolution of downsampled image using the technique described in Section 2.2. The information within these tensors is only retained only as skip connections concatenated to the upsampling layers. However, at the lowest resolution (128×128), the downsampling is allowed and the model operations proceed as normal (see Fig. 1). We believe

*UC Berkeley - UCSF Joint Bioengineering Graduate Program

† Computer Science Division, UC Berkeley

‡ These authors contributed equally to this work.

§ Work completed during an internship at Google.

¶ Google

|| Google Deepmind

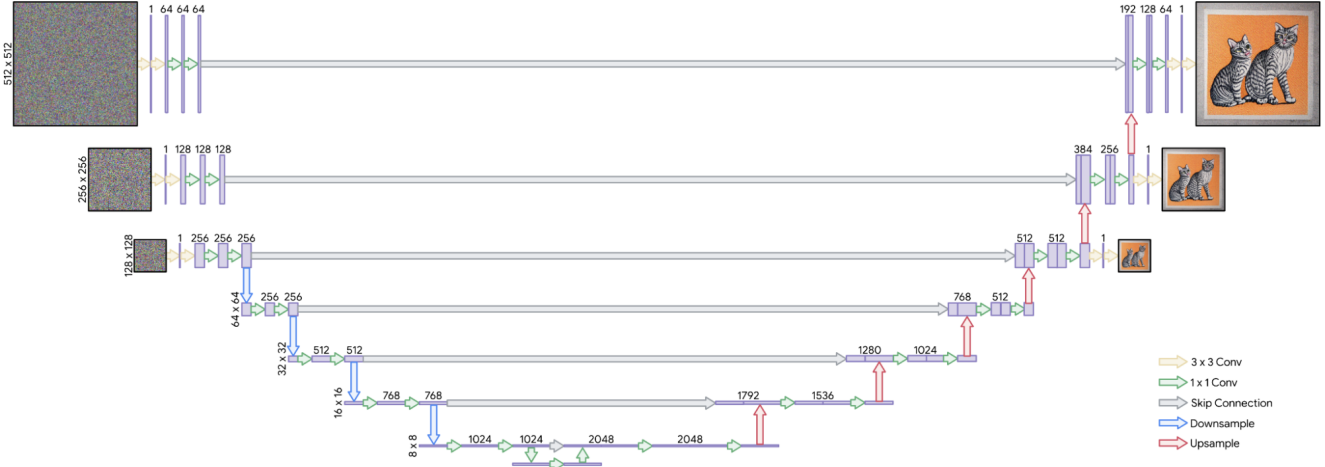


Figure 1. Layered diffusion model architecture. We provide noisy inputs at 3 resolutions and generate images of corresponding size. Note that information is only shared between layers during upsampling via skip connection.

this encourages robust learning at every level rather than optimizing for high resolution features, which is important at the lowest resolution where image content is decided [10]. Additionally, eliminating high resolution convolutions decreases the total number of floating point operations (FLOPs) within the model.

- **Multiple Outputs:** We apply out convolutions at relevant resolution scales in the upsampling part of the network which are used to evaluate loss (mean squared error) during training. While higher resolution information is isolated during downsampling, it is concatenated with the output of the lower model layers via skip connection once upsampled. We scale loss with respect to resolution scale. During inference, only the output convolution of the highest resolution is necessary.

Our model aims to learn spatial frequencies at different levels of the network. We adopt a hierarchical design philosophy for this model. We begin with a base model initially trained at a resolution of 128×128 with an embedding dimension of 256. To include a higher resolution, we include largely parallel convolutional layers atop this structure with (e.g. 256×256 with hidden dimension of 128). We show that our multi-scale model can generate higher quality images than a single resolution-scale model that directly generates images of the same resolution conditioned on text (where hidden dimension sizes are identical).

We compare models using Fréchet inception distance (FID) [4] and Inception Score (IS) [11] on the MS-COCO [6] validation set.

2.2. Noise Scaling

Rather than employing independently sampled Gaussian noise for each resolution scale, denoted as ϵ , we explore a strategy which uses the sinc (also known as Whittaker-Shannon) interpolation formula whereby we randomly sample noise at the highest resolution, and downsample to the lower resolutions. Sinc interpolation acts as a lowpass filter to ensure the preservation of the noise signal across layers, allowing effective sharing of pixel information across different resolutions while maintaining the Gaussian nature of the noise [12]. We formulate our downsampling operation, \mathcal{S} , as:

$$\mathcal{S}(\epsilon_0(\hat{m}, \hat{n}), T) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1}$$

$$\epsilon_0[m, n] \left(\text{sinc} \left(\frac{\hat{m} - mT}{T} \right) \cdot \text{sinc} \left(\frac{\hat{n} - nT}{T} \right) \right) \quad (1)$$

where ϵ_0 represents noise randomly sampled at the highest resolution via $\epsilon \sim \mathcal{N}(0, I)$, M and N are the dimensions of the target downsampled resolution, $\epsilon_0[m, n]$ are the pixel values of the high resolution noise at location (m, n) , (\hat{m}, \hat{n}) denotes the corresponding pixel positions in the scaled noise, and T is the downsampling scale. This formula represents a 2D version of the sinc interpolation formula applied to each noise term.

We observed superior image synthesis when using sinc interpolation (Fig. A.2), with markedly lower FID and higher IS (Table A.1). We believe this approach is beneficial when compared to randomly sampled noise because pixel-level spatial frequencies for input noise are preserved at each resolution. While we use bilinear scaling for downsampling images (Equation 3), applying the same operation

to ϵ_0 is problematic because bilinear interpolation nonlinearly modifies the signal profile and does not ensure that the output noise maintains a Gaussian distribution [2].

2.3. Cosine Schedule Shifting

Hooeboom et al. [5] demonstrated that shifting the cosine schedules used to noise images helped overcome the obstacle of global image structure being defined for a long duration during image synthesis (see Fig. A.3). We similarly implement shifted cosine schedules for the higher resolution layers of the model (Fig. A.4). We first found the optimal delay for a layered model with highest resolution synthesis at 256×256 to be $\log(\frac{1}{8})$. We then maintain this offset at the 256×256 resolution and swept to find the optimal delay for a model trained to synthesize at 512×512 , which was $\log(\frac{1}{32})$ (Fig. A.5). By doing so, our model is encouraged to synthesize fundamental image features earlier and incorporate textured information at higher levels later on in the diffusion process.

2.4. Mathematical Formulation

A text-to-image diffusion model \hat{x}_θ is trained on a denoising objective of given by:

$$\mathbb{E}_{x,c,\epsilon,t} \left[w_t \|\hat{x}_\theta(\alpha_t x + \sigma_t \epsilon, c) - x\|_2^2 \right] \quad (2)$$

where (x, c) are data-conditioning pairs, $t \sim \mathcal{U}([0, 1])$, and α_t , σ_t , and w_t are functions of t that influence sample quality.

To extend this objective function to multiple resolutions as discussed in Section 2.1, we introduce a summation over different resolutions as:

$$\mathbb{E}_{x_0,c,\epsilon_0,t} \left[w_t \sum_{i=0}^K \|\hat{x}_{\theta,i}(\alpha_{t,i} \mathcal{D}(x_0, i) + \sigma_{t,i} \mathcal{S}(\epsilon_0, i), c) - \mathcal{D}(x_0, i)\|_2^2 \right] \quad (3)$$

where K is the number of resolutions, $\hat{x}_{\theta,i}$ is the predicted image at resolution i , $\alpha_{t,i}$ and $\sigma_{t,i}$ are obtained via shifted cosine schedules described in Section 2.3, $\mathcal{D}(x_0, i)$ represents the bilinear downsampling operation applied directly to the highest resolution image x_0 , and $\mathcal{S}(\epsilon_0, i)$ represents the sinc interpolation from Equation 1.

3. Training Optimizations

3.1. Strategic Cropping

Saharia et al. [10] demonstrated accelerated training convergence with respect to TPU time by training the 1024×1024 super-resolution model on $64 \times 64 \rightarrow 256 \times 256$ crops. We

similarly found that we could utilize cropping to increase training speed by cropping all input images to 128×128 . However, to accommodate our layered model architecture, a crop must also be applied within the a tensor immediately before upscaling (Fig. A.6). This was done with the following procedure:

1. Randomly select a 64×64 square within the base 128×128 image.
2. Identify the corresponding 128×128 positions within the 256×256 image via simple multiplication.
3. Crop 256×256 image to 128×128 .
4. Repeat for higher resolutions by utilizing the previous resolution.
5. Apply the corresponding 64×64 crops in the upsample stack of the U-Net prior to the upsampling convolution to bring the output image to 128×128 .

3.2. Model Stacking

A limitation of this implementation is that it requires images of at least the largest resolution of the model, which fundamentally limits the training set size based on depending on the highest resolution present. For example, a model trained on images of $[512 \times 512, 256 \times 256, 128 \times 128]$ would require a minimum base image resolution of 512×512

To address this issue, we leveraged the hierarchical nature of our model by iterative rounds of training and checkpointing. Each resolution scale is constructed upon its predecessor, allowing us to initialize a higher resolution model with weights from a model trained with a smaller maximum resolution.

We visualize our results in Table A.2. At 256×256 , we compare a randomly initialized model (a) and a model where the lower weights are initialized with a 128×128 model trained to 500k steps. Similarly at 512×512 , we compare a randomly initialized model with ones initialized with the (a) and (b) 256×256 models trained to 500k steps.

Interestingly, we observed an increase in FID when using randomly initialized models. Despite this, our images demonstrated enhanced text-image alignment when the models were pre-loaded (Fig. A.7). More exploration is required to balance the image fidelity-prompt alignment trade off.

4. Results

In Table 1, we evaluate our layered diffusion model against baseline models trained at singular target resolution from the perspective of image generation quality (FID and IS) as well as total floating point operations (FLOPs). We observe that our multi-scale synthesis outperforms single-scale models when the maximum output resolution was both 256×256 and 512×512 Note that the 256×256 comparison is evaluated on a model trained with on two input images $[128 \times 128,$

Model	Target Resolution	FLOPs	FID	IS
Single Resolution	128×128	1.24×10^{12}	11.53	$24.48 \pm .41$
Single Resolution	256×256	2.20×10^{12}	14.61	$28.61 \pm .44$
Layered	256×256	2.04×10^{12}	13.38	29.62 \pm .57
Single Resolution	512×512	3.24×10^{12}	53.89	$23.21 \pm .29$
Layered	512×512	2.79×10^{12}	40.05	28.74 \pm .47

Table 1. “Target Resolution“ denotes the highest resolution of image used within the model. Models denoted as “Single Resolution“ utilize a single resolution image at that scale in a typical U-Net architecture. “Layered“ models utilize our multi-resolution model. Possible image inputs are 128×128 , 256×256 , and 512×512 . Layered models are provided with multiple input images. For example, if the target resolution is 256×256 , the layered model is provided with 128×128 and 256×256 input images.

256×256], and not the individual outputs of the 512×512 model.)

At 512×512 resolution, our model reduces the FID score by 25% and increases the IS score by 24%, while also decreasing the FLOPs by 14%. This indicates that our model can generate more realistic and diverse images with less computation. Our multi-scale model also improves the FID and IS scores at 256×256 resolution by 8% and 4%, respectively, compared to the single-scale model. We observe image features in accordance with these findings (Fig. A.1), where higher resolution features can be seen in outputs from the layered model.

5. Conclusion

In this study, we propose a novel one-shot text-to-image diffusion model that can generate high-resolution images from natural language descriptions in a single forward pass. Our model employs a layered U-Net architecture that predicts images at multiple resolution scales simultaneously, using a sinc interpolation formula to scale noise and a shifted cosine schedule for the diffusion steps. We show that this method outperforms the baseline of synthesizing images only at the target resolution, while reducing the computational cost per step. Moreover, our model does not need any additional super-resolution models to enhance the image quality. We also explore various techniques to improve the training efficiency and the image quality, such as cropping and stacking. We believe that our model is an interesting path to explore further which streamlines text-to-image synthesis to a one-shot method.

References

- [1] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Qinsheng Zhang, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. eDiff-i: Text-to-image diffusion models with an ensemble of expert denoisers. **1**
- [2] Paolo Banelli. Non-linear transformations of gaussians and gaussian-mixtures with implications on estimation and information theory. **3**
- [3] Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Josh Susskind, and Navdeep Jaitly. Matryoshka diffusion models. **1**
- [4] 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. **2**
- [5] Emiel Hooeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. **1, 3**
- [6] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common objects in context. **2**
- [7] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. **1**
- [8] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. **1**
- [9] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. **1**
- [10] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. **1, 2, 3**
- [11] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. **2**
- [12] J. M. Whittaker. On the cardinal function of interpolation theory. *1(1)*:41–46. **2**
- [13] Zeyue Xue, Guanglu Song, Qiushan Guo, Boxiao Liu, Zhuofan Zong, Yu Liu, and Ping Luo. RAPHAEL: Text-to-image generation via large mixture of diffusion paths. **1**