

Any-resolution Training for High-resolution Image Synthesis

Lucy Chai¹, Michaël Gharbi², Eli Shechtman²
Phillip Isola¹, and Richard Zhang²

¹MIT ²Adobe Research

Abstract. Generative models operate at fixed resolution, even though natural images come in a variety of sizes. As high-resolution details are downsampled away and low-resolution images are discarded altogether, precious supervision is lost. We argue that every pixel matters and create datasets with variable-size images, collected at their native resolutions. To take advantage of varied-size data, we introduce *continuous-scale* training, a process that samples *patches* at random scales to train a new generator with variable output resolutions. First, conditioning the generator on a target scale allows us to generate higher resolution images than previously possible, without adding layers to the model. Second, by conditioning on continuous coordinates, we can sample patches that still obey a consistent global layout, which also allows for scalable training at higher resolutions. Controlled FFHQ experiments show that our method can take advantage of multi-resolution training data better than discrete multi-scale approaches, achieving better FID scores and cleaner high-frequency details. We also train on other natural image domains including churches, mountains, and birds, and demonstrate arbitrary scale synthesis with both coherent global layouts and realistic local details, going beyond 2K resolution in our experiments. Our project page is available at: <https://chail.github.io/anyres-gan/>.

Keywords: Unconditional Image Synthesis, Generative Adversarial Networks, Continuous Coordinate Functions, Multi-scale Learning.

1 Introduction

The first step of typical generative modeling pipelines is to build a dataset with a fixed, target resolution. Images above the target resolution are downsampled, removing high-frequency details, and data of insufficient resolution is omitted, discarding structural information about low frequencies. Our insight is that this process wastes potentially learnable information. We propose to embrace the natural diversity of image sizes, processing them at their *native* resolution.

Relaxing the fixed-dataset assumption offers new, previously unexplored opportunities. One can potentially simultaneously learn global structure – for which large sets of readily-available low-resolution images suffice – and fine-scale details



Fig. 1: Trained on a dataset of varied-size images, our unconditional generator learns to synthesize patches at continuous scales to match the distribution of real patches. Here, we render crops of the image at different resolutions, indicating the target resolution for each. We indicate the region of each crop in the top-left panel, which is the image directly sampled without scale input.

– where even a handful of high-resolution images may be adequate, especially given their internal recurrence [48]. This enables generating images at higher resolutions than previously possible, by adding in higher-resolution images to currently available fixed-size datasets.

This problem setting offers unique challenges, both in regards to modeling and scaling. First, one must generate images across multiple scales in order to compare with the target distribution. Naïvely downsampling the full-resolution image is suboptimal, as it is common to even have images of $8\times$ difference in scale in the dataset. Secondly, generating and processing high-resolution images offers scaling challenges. Modern training architectures at 1024 resolution already push the current hardware to the limits in memory and training time, and are unable to fully make use of images above that resolution.

To bypass these issues, we design a generator to synthesize image crops at arbitrary scales, hence, performing *any-resolution* training. We modify the state-of-the-art StyleGAN3 [26] architecture to take a grid of continuous coordinates, defined on a bounded domain, as well as a target scale, inspired by recent work in coordinate-based conditioning [37,2,32,7]. By keeping the latent code constant and varying the crop coordinates and scale, our generator can output patches of the same image [31,44], but at various scales. This allows us to (1) efficiently generate at arbitrary scale, so that a discriminator critic can compare generations to a multi-resolution dataset, and (2) decouple high-resolution synthesis from increasing architecture size and prohibitive memory requirements.

We first experiment on FFHQ images [24] as a controlled setting, showing efficient data usage by comparing our training on image subsets to using the entire full resolution dataset. We find minimal degradation (FIDs varying by 0.3),

even at highly skewed distributions – 98% of the dataset at $4\times$ lower resolution, and just 2% at a higher, mixed resolutions. Practically, this means we can leverage large-scale ($>100k$) lower-resolution datasets, such as LSUN Churches [65], Flickr Mountains [41], and Flickr Birds collected by us, and add a relatively small amount of high-resolution images (~ 6000), for continuous resolution synthesis beyond the 1024 resolution limit of current generators. To summarize, we:

- propose to train on mixed-resolution datasets from images in-the-wild.
- modify the generator to be amenable to such data, sampling patches at arbitrary scales during our *any-resolution* training procedure.
- demonstrate successful generations beyond 1024×1024 , with fine details and coherent global structure, without a larger and more expensive generator.
- introduce a variant of the FID metric that captures image statistics at multiple scales, thus accounting for the details of high resolutions.

2 Related Works

Unconditional image synthesis. Recent generative models including GANs [16,27,26,3], Variational Autoencoders [29], diffusion models [38,17,50,51,11], and autoregressive models [30,54,40] such as transformers [55,6,14] are rapidly improving in quality. Of these, we focus on GANs, which offer state-of-the-art performance along with efficient inference and effective editing properties. A key innovation in GANs has been multi-resolution supervision during training. Works such as LapGAN [10], the Progressive/StyleGAN family [24,27,28,26], MSG-GAN [23], and AnyCost-GAN [33] have demonstrated stable training by growing the generator with additional layers that increase resolution by factors of two. Such a strategy even works for single-image GANs [44,47], based on the observation that images share statistics across scales. While several works [25,70,69] show data augmentations, such as small jitters in scale, can help stabilize training, they are processing the same, underlying fixed-resolution dataset. We draw upon the insights in these works for stable training, and seek to unlock training on an any-resolution dataset. Importantly, our generator does not use additional layers and can synthesize images at continuous scales, not only powers of two.

Coordinate-based functions. Coordinate-based encodings enable spatial conditioning and provide an inductive bias towards natural images [52]. Recent methods use point-based neural mappings to transform 2D or 3D coordinates to a color value for the purposes of unconditional generation [31,9,26,49,32,1], conditional generation [45], 3D view synthesis [37,43,5], or fitting arbitrary signals [7,36]. By oversampling the coordinate grid, one can generate a larger image at inference time. However, because these models keep the same fixed-scale dataset assumption during training, the outputs struggle to offer additional high-frequency details without a high-frequency training signal. We draw upon the innovations in coordinate-based functions to sample patches at different scales and locations, enabling us to efficiently train on multiple scales. MS-PIE [62] and MS-PE [9] add positional encodings for multi-scale synthesis, but retain a global image discriminator at smaller resolution. Concurrently, ScaleParty [39]

also samples patches, but their goal is to generate at multiple scales with cross-scale consistency while we focus on training with arbitrary size real images.

Extrapolation. One method of generating “infinite” resolution is extrapolating an image. Early texture synthesis works [13,12,60] focus on stationary textures. Recent advances [72] explore non-stationary textures, with large-scale structures and inhomogeneous patterns. Similar approaches operate by outpainting images, extending images beyond their boundaries in a conditional setting [53,64,35,72]. Recent generative models synthesize large scenes [32,8], typically casting synthesis as an outpainting problem [64,35]. These methods are most effective for signals with a strong stationary component, such as landscapes, although extrapolation of structured scenes can be achieved in some domains [59]. Unlike textures, the images we wish to synthesize typically have a strong global structure. In a sense, we seek to extrapolate by “zooming in” or out, rather than “panning” beyond an image’s boundaries.

Super-resolution. An alternative approach to generating high-resolution imagery would be to start with an off-the-shelf generative model and feed its outputs to a super-resolution method [18,7,58,57], possibly exploiting the self-similarity properties of images [46,15,19,48]. Applying super-resolution models is challenging, in part because of the specific blur kernels super-resolution models are trained on [66]. Furthermore, though generations continue to improve, there remains a persistent domain shift between synthesized and real images [56,4]. Finally, super-resolution is a local, conditional problem where the global structure is dictated by the low-resolution input, and optionally an additional high-resolution reference image [71,63,34,22,61]. We synthesize plausible images unconditionally, leveraging a *set* of high-resolution images to produce both realistic global structure and coherent fine details.

3 Methods

In standard GAN training, all training images share a common fixed resolution, which matches the generator’s output size. We seek to exploit the variety of image resolutions available in the wild, learning from pixels that are usually discarded, to enable high- and continuously-variable resolution synthesis. We achieve this by switching from the common fixed-resolution thinking, to a novel ‘any-resolution’ approach, where the original size of each training image is preserved (Fig 2). We introduce a new class of GAN generators that learn from this multi-resolution signal to synthesize images at any resolution (§ 3.1), and show how to train them by sampling patches at multiple scales to jointly supervise the global-structure and fine image details (§ 3.2).

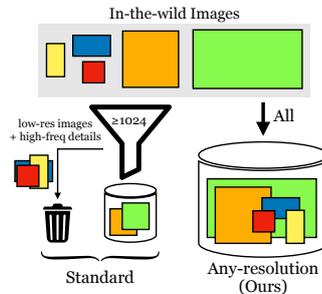


Fig. 2: Data comes at a variety of scales. Traditional dataset construction filters out low-resolution images and down-samples high-resolution images to a fixed, training resolution. We aim to keep images at their original resolution.

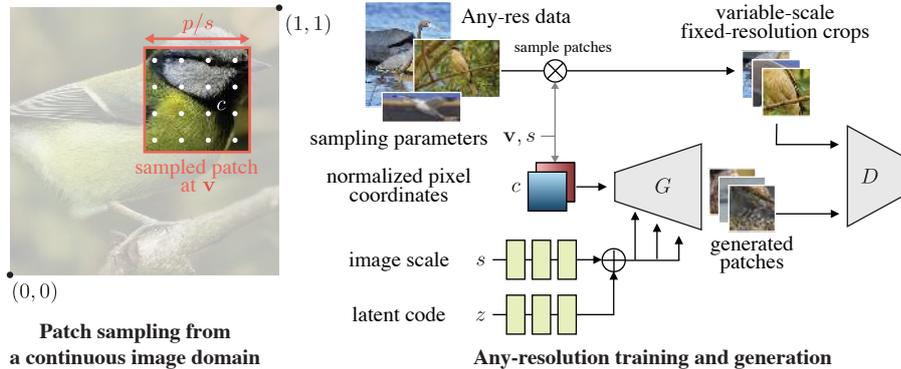


Fig. 3: Overview. (Left) We parameterize images (real or synthetic) as continuous functions over a normalized domain and extract random patches at various scales s , but constant resolution p . (Right) To train, we sample crops at random scales and offsets \mathbf{v} from full-size real images. The same crops are sampled from the generator, by passing it a grid of the desired coordinates $c_{\mathbf{v},s}$, and injecting the image scale s through modulation, in addition to a global latent code z .

3.1 Multi-resolution GAN

We design our approach to leverage state-of-the-art GANs. We keep the architecture of the discriminator unchanged. Since the discriminator operates at fixed resolution, we modify the generator to synthesize images at any resolution and receive the discriminator’s fixed-resolution supervision. Our implementation builds on the StyleGAN3 framework [26], which is conditioned on a fixed coordinate grid. We modify this grid for any-resolution and patch-based synthesis.

Continuous-resolution generator. We treat each image as a continuous function defined on a bounded normalized coordinate domain $[0, 1] \times [0, 1]$. The generator G always generates patches at a fixed pixel resolution $p \times p$, but each patch implicitly corresponds to a square sub-region, centered at $\mathbf{v} \in [0, 1]^2$, of the larger image. Denoting the resolution of the larger image as $s \times s$, we have that the patch size is p/s in normalized coordinates (see Figure 3, left). During training, we sample patches from images at multiple scales s , either from the generator or from the multi-resolution dataset, before passing them to the fixed-resolution discriminator D . Formally, our generator takes three inputs: a regular grid of normalized continuous pixel coordinates $c_{\mathbf{v},s} \in \mathbb{R}^{p \times p \times 2}$, the resolution $s \in \mathbb{N}$ of the (implicit) larger image the patch is extracted from, and z , the latent code representing this larger image. It synthesizes the patch’s pixel values at the sampled coordinates as:

$$G(z, c_{\mathbf{v},s}, s) = G(F(c_{\mathbf{v},s}); M(z, s)), \quad (1)$$

where F is a Fourier embedding of the continuous coordinates [26], and M is an auxiliary function that maps the latent code and sampling resolution into a set of modulation parameters for the StyleGAN3 generator (see § 3.3 for details). Our method therefore modifies two components from StyleGAN3. First,

we replace the fixed coordinate grid with *patch-dependent* coordinates to train on variable-resolution images; these coordinates are adjusted to account for up-sampling in StyleGAN3. Second, we append an auxiliary branch M to inject scale information throughout the generator.

At test time, we can generate images at arbitrarily high resolutions by sampling the full continuous domain $[0, 1] \times [0, 1]$ at the desired sampling rate. Theoretically, the maximum resolution is infinite, but in practice the amount of detail that the model can generate is determined by generator resolution p and the resolutions of the training images.

3.2 Two-phase training

We train our generator in two phases. In the first, we want the generator to learn to generate globally-coherent images. For this, we disable the patch sampling mechanism and pretrain the generator at a fixed scale, corresponding to the full continuous image domain. That is, we fix $s = p$, and $\mathbf{v} = (0.5, 0.5)$, which is equivalent to standard fixed-resolution GAN training. Both the coordinate tensor $c_{\mathbf{v},s}$ and the scale conditioning variable s are constant in this phase, so we simply refer to the image generated as $G_{\text{fixed}}(z)$ and follow the training procedure of StyleGAN3. In the second phase, we enable patch sampling for both the real and synthetic images and continue training the generator using variable-scale patches, so it learns to synthesize fine details at any resolution. We found that using a copy of the pretrained fixed-scale generator G_{fixed} as a teacher model helps stabilize training in this phase.

Global fixed-resolution pretraining. During the pretraining phase, we effectively resample all the training images to a fixed resolution $p \times p$, as in standard GAN training. Let $x \sim \mathcal{D}_{\text{fixed}}$ denote an image sampled from this fixed size dataset. We optimize a standard GAN objective with non-saturating logistic loss and R_1 regularization on the discriminator:

$$\begin{aligned} V(D, G(z), x) &= D(x) - D(G(z)), & R_1(D, x) &= \|\nabla D(x)\|^2, \\ G_{\text{fixed}} &= \arg \min_G \max_D \mathbb{E}_{z, x \sim \mathcal{D}_{\text{fixed}}} V(D, G(z), x) - \frac{\lambda_{R_1}}{2} R_1(D, x). \end{aligned} \quad (2)$$

We follow the recommended values for λ_{R_1} , depending on the generator resolution p [26].

Mixed-resolution patch-based training. In the second phase, we enable multi-resolution sampling, alternating between extracting random crops from our any-resolution dataset and generating them with our continuous generator.

For synthetic patches, we sample a patch location \mathbf{v} uniformly in the continuous domain $[0, 1] \times [0, 1]$; and an arbitrary image resolution $s \geq p$, corresponding to the implicit full image around the square patch. From those, we derive the sampling coordinate grid $c_{\mathbf{v},s}$, and synthesize the patch image $G(z, c_{\mathbf{v},s}, s)$, as described earlier.

For ‘real’ patches, we sample an image from our dataset. Because this image can have any resolution $s_{\text{im}} \geq p$, we crop it to a random square matching its

smallest dimension, then Lanczos downsample this square to a random resolution $s \times s$ with $s_{\text{im}} \geq s \geq p$. Finally, we extract a random $p \times p$ crop from the downsampled image, recording its center \mathbf{v} . To preserve the generator’s global coherence and continuous generation ability, we sample at global scale $s = p$ and $\mathbf{v} = (0.5, 0.5)$ (similar to the pretraining step) with probability 50%. We found that image quality at global resolution $s = p$ degrades otherwise, and we refer to these generated full images of size $p \times p$ as “base images.” Our any-resolution GAN optimizes the following objective during this phase:

$$G^* = \arg \min_G \min_D \mathbb{E}_{z, \{x, s, \mathbf{v}\} \sim \mathcal{D}} V(D, G(z, c_{\mathbf{v}, s}, s), x) + \lambda_{\text{teacher}} \mathcal{L}_{\text{teacher}}(G, G_{\text{fixed}}, z) - \frac{\lambda_{R_1}}{2} R_1(D, x). \quad (3)$$

We use $\lambda_{\text{teacher}} = 5$; other values offer slight tradeoffs between similarity to the base teacher model G_{fixed} , and FID score (see supplemental). $\mathcal{L}_{\text{teacher}}$ is an auxiliary loss to encourage faithfulness to the pretrained fixed-resolution generator G_{fixed} . The architecture of D remains the same as in the pretraining step; we found that modifying the discriminator setup did not further improve results (see supplemental).

Teacher model. For the second training phase above, we initialize G with the pretrained weights of G_{fixed} . Weights for discriminator D are also kept for fine-tuning. We keep a separate copy of G_{fixed} with frozen weights, the teacher, for additional supervision. We design a loss function that encourages the generated patch (at any resolution), to match the teacher’s fixed-resolution output in the region corresponding to the patch, after downsampling and proper alignment[21]. Formally, this loss is given by:

$$\mathcal{L}_{\text{teacher}}(G, G_{\text{fixed}}, z) = d(m \odot w_{\mathbf{v}, s}(G(z, c_{\mathbf{v}, s}, s)), m \odot G_{\text{fixed}}(z)), \quad (4)$$

where d is the sum of a pixel-wise ℓ_1 loss, and the LPIPS perceptual distance [20,67]. The warp function $w_{\mathbf{v}, s}$ transforms and resamples the generated high-resolution patch using a band-limited Lanczos kernel, to project it in the coordinate frames of the low-resolution, global image $G_{\text{fixed}}(z)$. Because the warped patch does not cover the entire image domain, we multiply it with a binary mask m to indicate the valid pixels, prior to computing loss d .

3.3 Implementation details

Scale-conditioning. In addition to the pixel location $c_{\mathbf{v}, s}$, we also pass the global resolution information s to the generator. Knowledge of the global image scale is important to enable continuous scale variations and proper anti-aliasing [7,2]. We found it beneficial to explicitly inject this information into all intermediate layers of the generator. To achieve this, we use a dual modulation approach [68], embedding the latent code z and scale s separately using two independent sub-networks (we use the same mapping network architecture for each). The two outputs are summed to obtain a set of modulation parameters

$M(z, s)$, used to modulate the main generator features. Architectural details of the generator G and mapping network M , can be found in the supplemental.

Synthesizing large images. Our fully-convolutional generator can render image at arbitrary resolutions. But images larger than 1024×1024 require significant GPU memory. Equivalently, we can render non-overlapping tiles that we assemble into a larger image. Our patch-based multi-resolution training and the Fourier encoding of the spatial coordinates make the tile junctions seamless.

4 Experiments

We introduce a modified image quality metric that computes FID over multi-scale image patches without downsampling, which is largely correlated to the standard FID metric when ground truth high-resolution images are available (FFHQ), yet more sensitive to the quality of larger resolutions. We then compare our model to alternative approaches for variable scale synthesis and super-resolution on other natural image domains (§ 4.1). Finally, we investigate variations of our model and training procedure to validate our design decisions (§ 4.2).

Data. Our method is general and can work on collections of any-resolution data. As such, when targeting high-resolution generation, rather than starting over, we can add additional high-resolution (HR) images to existing, fixed-size, low-resolution (LR) datasets. Our datasets and their statistics are listed in Tab. 1. Figure 4 shows the resolution distribution in each dataset.

We begin initial experiments with a controlled setting of FFHQ, which contains 70K images at 1024 resolution. From these, we construct a varied-size dataset by (1) using 256 resolution for all images (2) downsampling a 5K subset between 512-1024 (uniformly distributed) and (3) add 1k subset at full 1024. The last step enables us to judiciously compare to methods that are limited to synthesizing images at strict powers of 2. We refer to this mixture as FFHQ6k. We use the full 1024 dataset as ground-truth for evaluation metrics.

In the remaining domains, we push current generation results to higher resolution by scraping HR images from Flickr. In cases where a standard fixed-size dataset is available (LSUN Churches [65] and Mountains [41]), we select the additional HR images to approximately match the LR domain. Our final generators synthesize realistic details despite the majority of the training set being LR. For Birds and Churches, $> 92\%$ of the training set is at 256 resolution but our model maintains quality beyond 1024; for Mountains $> 98\%$ of training set at 1024 but our model can generate beyond 2048. These categories cover a range between objects (but without the strong alignment of FFHQ) and outdoor scenes.

4.1 Continuous Multi-Scale Image Synthesis.

Qualitative examples. We show qualitative examples in Fig. 5. Our generated images preserve the fine details of HR structures, such as bricks, rocky slopes, feathers, or hair. Pushing the inference resolution towards and beyond the higher

Table 1: Any-resolution datasets and generator settings. We build upon low-resolution (LR) datasets, and use it for fixed-size dataset pre-training. We add additional high-resolution (HR) images, of mixed resolutions. Note that the number of HR images is small ($\sim 2\text{-}8\%$ of LR size). Patches of size p are sampled from both subsets during training, with average sampled scale $\mathbb{E}[s]$.

Domain	Dataset						Generator			
	Source	# Imgs		Resolutions			Config	Resolutions		
		LR	HR	LR	HR _{min}	HR _{med}		HR _{max}	p	$\mathbb{E}[s]$
Faces	FFHQ	70,000	6000	256	512	819	1024	R	256	458
Churches	LSUN & Flickr	126,227	6253	256	1024	2836	18,000	T	256	1061
Birds	Flickr	112,266	7625	256	512	1365	2048	T	256	585
Mountains	Flickr	507,495	9361	1024	2049	3168	12,431	T	1024	1823

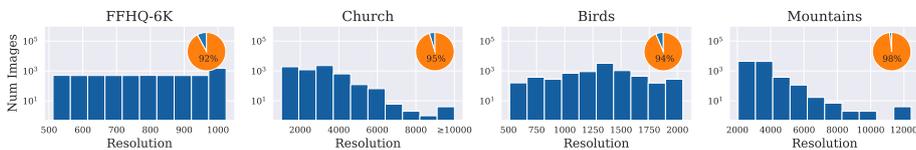


Fig. 4: Training set size distributions. Histogram shows the size distribution of the HR images (y-axis in log scale); pie chart indicates the proportion of LR to HR images.

resolutions of training images, we find that textured surfaces typically deteriorate first before edge boundaries deteriorate eventually (Fig. 6).

Patch-FID metric. Standard FID evaluates *global structure* by first downsampling all images to a common size of 299. By design, this ignores high-resolution details (and itself can cause artifacts [42]). Therefore, we propose a modification, which we dub ‘patch-FID’, to specifically evaluate *texture* synthesized at higher resolutions. Our patch-FID randomly resizes and crops patches from the HR dataset, and computes FID on real and generated patches, sampled at corresponding scales and locations. We use 50k patches, matching standard FID. By avoiding downsampling, our patch-FID is more sensitive to blurriness or artifacts at higher resolutions, resulting in larger absolute difference compared to standard FID at 1024 resolution. As a sanity check, when a full HR ground-truth is available, we find it is largely correlated to standard FID (see Table 2).

To summarize, to evaluate structure, we compute standard FID on images generated at specified resolutions, e.g., FID (256). To evaluate texture, we sample patches at random scales and locations and measure our patch-FID, which we denote as pFID (random). Lower numbers are better in both cases.

Alternative methods of multi-size training and generation. Our generator is encouraged to synthesize realistic high-resolution textures at training, even when the discriminator does not get to see the full image. While MS-PE [9] also enables continuous resolution synthesis, the discriminator learns only at a single resolution and the generator is not trained patch-wise. We find that this

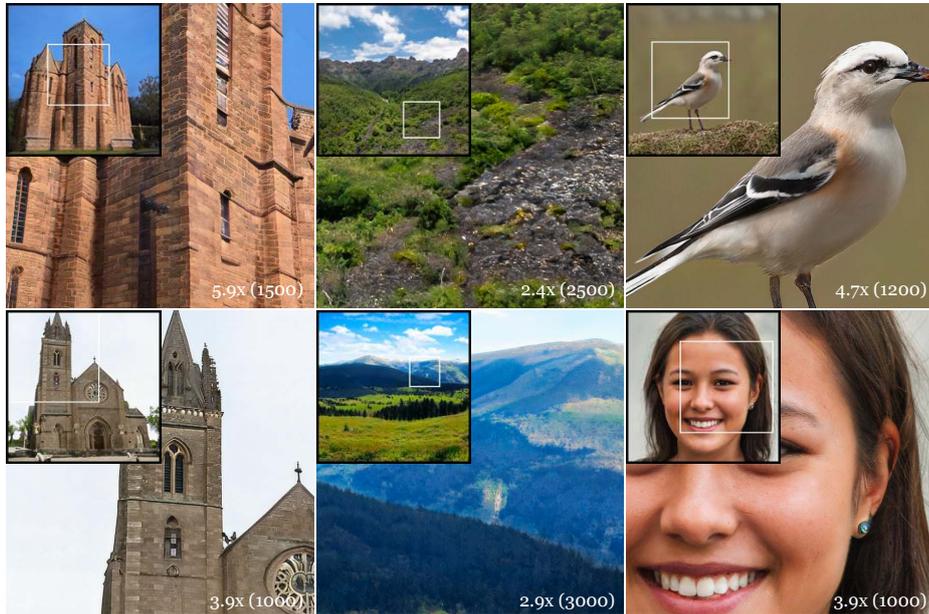


Fig. 5: The inset shows the entire generated, high-resolution images (between 1000-3000 resolution), with enlarged regions of interest outlined in the white box. Note that our model can render the image (or any sub-region) at any resolution.

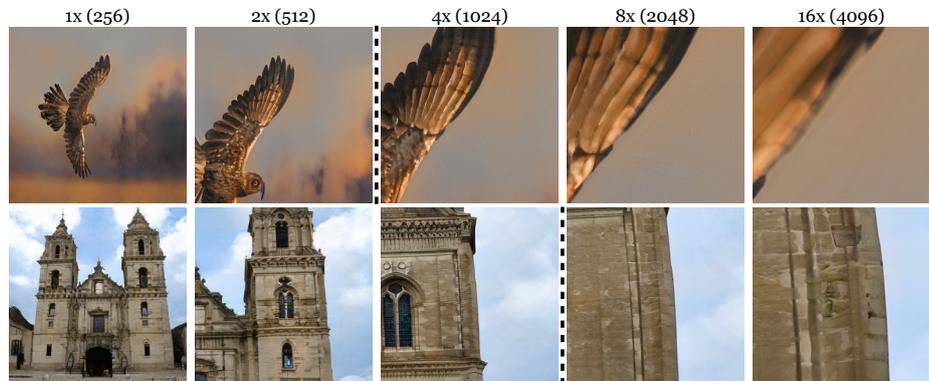


Fig. 6: Extrapolation limits. We test the extrapolation capabilities of our model by specifying the inference scale s . Typically, textures such as bricks and feathers deteriorate first before edges degrade. The dotted line indicates when generation starts to exceed the average scale sampled in training $\mathbb{E}[s]$ (which is 585 for birds and 1061 for churches).

Table 2: Varied-size training and inference. Random-resize MS-PE [9] performs varied-size synthesis, but assumes a fixed-size dataset. AnyCost-GAN handles varied training at powers of 2. Our method directly utilizes training images at any size, achieving better results by FID. († = copied from paper)

FFHQ6K	FID		pFID	
	256	512	1024	random
MS-PE [9]†	6.75	30.41	–	–
Anycost [33]	4.24	5.94	6.47	18.39
Ours	3.34	3.71	4.06	2.96

Table 3: Comparison to super-resolution using patch-FID (pFID). For each domain, we compare our model to continuous-scale (LIIF) and fixed-scale super-resolution (Real-ESRGAN) models. Lower pFID suggests that our model can generate realistic details at high resolutions, not achievable with super-resolution alone.

	pFID (random)			
	FFHQ6K	Church	Bird	Mountain
LIIF [7]	22.93	83.88	30.19	23.10
Real-ESRGAN [57]	16.92	23.04	16.10	19.05
Ours	2.96	9.89	6.52	7.99

downsampling for the discriminator is detrimental to image quality at higher resolutions. Anycost-GAN[33] performs image synthesis at powers-of-two resolutions by adding additional synthesis blocks. For comparison, we modify it to handle a multi-size dataset by downsampling images to the nearest power of two and training each layer only on the valid image subset. Compared to Anycost-GAN, our model is more data-efficient, due to weight sharing for generation at multiple scales. Anycost-GAN learns a separate module for each increase in resolution, creating artifacts at higher resolutions when fewer HR training images are available and higher FID scores (Tab. 2). Additionally, Anycost-GAN increases the generator and discriminator size for synthesis at higher resolutions, whereas our model incurs a constant training cost, regardless of the inference scale.

Comparison to super-resolution. Most super-resolution methods require LR/HR image pairs, whereas there is no ground-truth HR counterpart to a LR image synthesized by our fixed-scale generator G_{fixed} . The teacher regularization encourages similarity between G_{fixed} and G 's outputs, but unlike super-resolution, this supervision occurs at low-resolution, allowing variations in fine details. Figure 7 compares our model to super-resolution methods applied to the output of G_{fixed} . Our method produces much sharper details than LIIF [7], a recent continuous-scale super-resolution technique, and cleaner images than the state-of-the-art Real-ESRGAN [57]. The latter is a fixed-resolution model, so we run it iteratively until exceeding a target resolution, and then Lanczos down-sample the result to the target size. Real-ESRGAN's outputs are either overly smooth, or exhibit grid-like artifacts. Our method generates realistic textures based on the low-resolution output of G_{fixed} and reaches a lower pFID (Tab. 3).

4.2 Model Variations

Using the full high-resolution FFHQ dataset as a benchmark, we investigate individual components of our architecture and training process. We train each model variation for 5M images and record metrics from the best FID@1024 checkpoint. We only report quantitative metrics in the main paper and refer to the supplemental for further evaluations and visual comparisons.

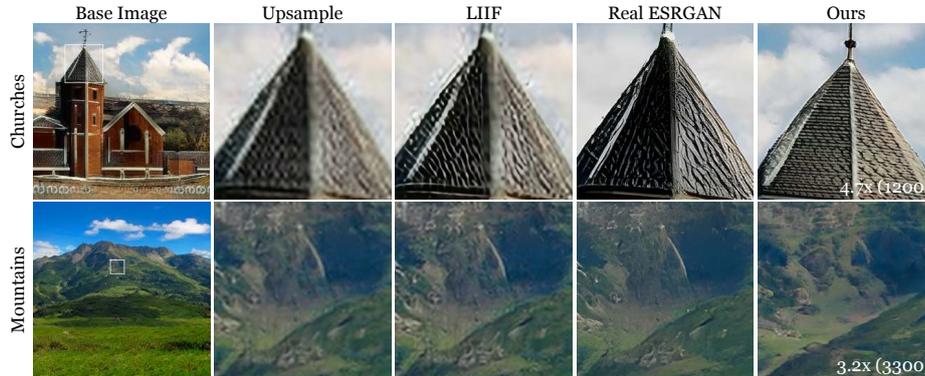


Fig. 7: Super-resolution Comparisons. Qualitative comparisons of Lanczos upsampling a patch from the base image (upsample), continuous (LIIF [7]) and fixed-factor (Real ESRGAN [57]) super-resolution models, and our trained model. LIIF tends to amplify artifacts from the base image (e.g. the JPEG artifacts around the church). While Real-ESRGAN is better at suppressing artifacts, it tends to overly smooth surfaces or synthesize grid-like textures (mountain). Our model is not a super-resolution model; it can add additional details to the low-resolution image but tolerates slight distortions in structure which are regularized with the teacher weight.

Impact of teacher regularization. Our full model uses an “inverse” teacher regularizer to encourage a downsampled HR patch to match the low-resolution teacher as described in 3.2. We also explored a variant with a “forward” teacher loss, in which the generated patch is encouraged to match the *upsampled* teacher output. This variation is qualitatively inferior and blurs details; it has worse FID at higher resolutions (see supplemental for details and visuals). Removing the teacher altogether improves pFID but degrades FID. Qualitatively, the generated patches diverge significantly from the fixed-size global image. We hypothesize that the global change in structure negatively impacts overall image quality, causing global FID metrics to increase, but this cannot be captured from evaluating patches alone. We found $\lambda_{\text{teacher}} = 5$ to provide the best balance between global and local image quality, but we observe minimal differences in FID and pFID for other values, evidence that the model can tolerate a range of values for this parameter. See supplemental for a parameter sweep with full scores.

Removing scale conditioning degrades quality. We inject the scale information to intermediate layers of the generator through scale-conditioning. Adding this improves FID@1024 from 4.88 to 4.47, and pFID from 4.67 to 4.28.

Multi-size training improves fidelity at all scales. Our multi-size data pipeline lets our model learn to synthesize at continuous scales, which is a strictly more challenging than learning at a fixed scale. In Table 4, we investigate to what extent learning from images of varied sizes offers benefits over fixed-scale training on a smaller dataset. Visual comparisons can be found in supplemental. In a first alternative, using the same FFHQ-6K dataset, we resize all images

Table 4: Multisize Training. Downsampling or upsampling all images to a common size, or using only the subset of the largest images, worsens FID compared to our training strategy. (*) indicates our default setting.

	FID			pFID
	256	512	1024	random
Resize down to 512	3.31	4.11	19.18	26.83
Resize up to 1024	3.46	13.43	4.86	6.65
Train 1024 subset	3.46	12.41	4.67	5.43
Multisize training (*)	3.37	4.41	4.47	4.28

Table 5: Number of HR images. Our method is robust to a wide range of HR images, even when only 1K images at HR are available (<2% of the full ground-truth dataset).

	FID			pFID
	256	512	1024	random
1k (1.4%)	3.43	5.13	4.38	3.73
5k (7.1%) (*)	3.36	4.97	4.54	3.48
10k (14.2%)	3.46	4.96	4.65	3.54
70k (100%)	3.42	4.88	4.52	3.42

down to 512 and train the model to generate patches at 512×512 . In this case, the model performs well up to 512 scale, but does not generalize beyond (e.g., 1024) since it cannot exploit the information lost in downsampling. In two other variants, we train models for 1024 resolution in the first case by upsampling all images up to 1024, in the second by keeping only the 1K subset of images at 1024 resolution. Both variants are trained to output images specifically at 1024 resolution. The former approach (upsampling) increases blurriness. In the latter, FID@1024 remains worse than that of our multi-size training, which can take advantage of more data despite most of it being *smaller* than 1024.

Impact of number of HR images Due to the patch-based training procedure, we find that our model can be trained with a small fraction of HR images, compared to the 70k LR images in the dataset. In Table 5, we use progressively larger subsets of HR images: 1k, 5k, 10k. We found that the FID scores are largely similar (within 0.3) to using the entire 70k HR images. However, training with 1k or fewer HR images shows evidence of divergence during training (see supplemental), but stabilizes by 5k HR images. For the remaining domains, we collect roughly 5K-10K images to construct the HR dataset.

4.3 Properties of Multi-Scale Generation

Correcting artifacts from low resolution. Because our model is not directly trained with corresponding LR and HR image pairs, we find that there can be small distortions between the upsampled base image and the HR generation from the same latent code. In some cases, this can be a desirable property (Fig. 8). For instance, the base generator on the birds dataset can struggle in synthesizing the eye of the bird, which is less apparent at low resolutions, but more salient at high resolution. Consequently, our HR generation will add the missing eye, and also synthesizes additional feather and beak details. In the churches domain, because the LR and HR datasets are collected separately, we find that the synthesized watermarks and JPEG artifacts at the base resolution disappear at higher resolution, because the HR dataset we used is of higher quality and does not have any watermark. The similarity between the LR and HR generations can be tuned using λ_{teacher} during training.

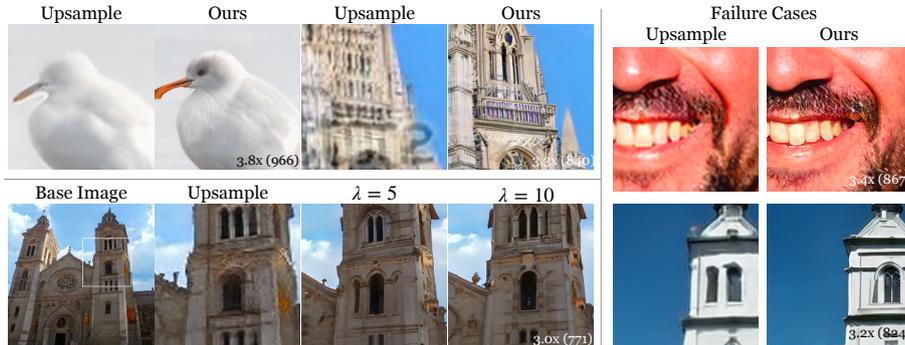


Fig. 8: Model Properties and Failure Cases. As fine details can be more difficult to learn at low resolution, our model is capable of adding corrections when generating at higher resolutions. In the case of inconsistencies between the LR and HR data sources, the model deletes patterns that are not present in the HR dataset (*e.g.* watermarks and compression artifacts), influenced by the teacher regularization weight. Failure cases include biases towards circular or ring-like structures.

Failure Cases. Our model tends to inherit the artifacts from StyleGAN3, such as a centered front tooth in FFHQ. In instances in which the base resolution image contains uneven surfaces, the model may fail to fully mitigate them at higher resolutions. These artifacts are often subtle at the low resolution, but become more apparent when upsampling the base image or generating at a larger target scale. In some cases, our model also has a tendency to generate “watery” circular or ring-like artifacts (Fig. 8).

5 Conclusion

We propose an image synthesis approach that can train on images of varied resolution and perform inference at continuous resolutions. This lifts the fixed-resolution requirement of prior generative models, which discard higher-resolution details. To do this, we train a generator jointly on a low-resolution dataset to learn global structure, and on patches from the varied-size dataset to learn details. At inference time, we can synthesize an image at any resolution by supplying the appropriate coordinate grid and scale factor to the generator. By using training images at their native resolutions and a single model for continuous-resolution synthesis, our method can efficiently leverage information present in only a handful of high-resolution images to complement a large set of low-resolution images. This approach enables high-resolution synthesis without a larger generator or large dataset of fixed-size, high-resolution images.

Acknowledgements. We thank Assaf Shocher for feedback and Taesung Park for dataset collection advice. LC is supported by the NSF Graduate Research Fellowship under Grant No. 1745302 and Adobe Research Fellowship. This work was started while LC was an intern at Adobe Research.

References

1. Anokhin, I., Demochkin, K., Khakhulin, T., Sterkin, G., Lempitsky, V., Korzhnikov, D.: Image generators with conditionally-independent pixel synthesis. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 14278–14287 (2021) [3](#)
2. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Int. Conf. Comput. Vis. pp. 5855–5864 (2021) [2](#), [7](#)
3. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: Int. Conf. Learn. Represent. (2018) [3](#)
4. Chai, L., Bau, D., Lim, S.N., Isola, P.: What makes fake images detectable? understanding properties that generalize. In: Eur. Conf. Comput. Vis. (2020) [4](#)
5. Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5799–5809 (2021) [3](#)
6. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: Int. Conf. Machine Learning. pp. 1691–1703. PMLR (2020) [3](#)
7. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 8628–8638 (2021) [2](#), [3](#), [4](#), [7](#), [11](#), [12](#)
8. Cheng, Y.C., Lin, C.H., Lee, H.Y., Ren, J., Tulyakov, S., Yang, M.H.: In&out: Diverse image outpainting via gan inversion. arXiv preprint arXiv:2104.00675 (2021) [4](#)
9. Choi, J., Lee, J., Jeong, Y., Yoon, S.: Toward spatially unbiased generative models. In: Int. Conf. Comput. Vis. (2021) [3](#), [9](#), [11](#)
10. Denton, E., Chintala, S., Szlam, A., Fergus, R.: Deep generative image models using a laplacian pyramid of adversarial networks. In: Adv. Neural Inform. Process. Syst. (2015) [3](#)
11. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. vol. 34 (2021) [3](#)
12. Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 341–346 (2001) [4](#)
13. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: Int. Conf. Comput. Vis. vol. 2, pp. 1033–1038. IEEE (1999) [4](#)
14. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 12873–12883 (2021) [3](#)
15. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: Int. Conf. Comput. Vis. pp. 349–356. IEEE (2009) [4](#)
16. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Int. Conf. Learn. Represent. (2015) [3](#)
17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Adv. Neural Inform. Process. Syst. vol. 33, pp. 6840–6851 (2020) [3](#)
18. Hu, X., Mu, H., Zhang, X., Wang, Z., Tan, T., Sun, J.: Meta-sr: A magnification-arbitrary network for super-resolution. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1575–1584 (2019) [4](#)
19. Huang, J.B., Singh, A., Ahuja, N.: Single image super-resolution from transformed self-exemplars. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5197–5206 (2015) [4](#)

20. Huh, M., Zhang, R., Zhu, J.Y., Paris, S., Hertzmann, A.: Transforming and projecting images into class-conditional generative networks. In: *Eur. Conf. Comput. Vis.* (2020) [7](#)
21. Irani, M., Peleg, S.: Improving resolution by image registration. *CVGIP: Graphical models and image processing* **53**(3), 231–239 (1991) [7](#)
22. Jiang, Y., Chan, K.C., Wang, X., Loy, C.C., Liu, Z.: Robust reference-based super-resolution via c2-matching. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2103–2112 (2021) [4](#)
23. Karnewar, A., Wang, O.: Msg-gan: Multi-scale gradients for generative adversarial networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 7799–7808 (2020) [3](#)
24. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: *Int. Conf. Learn. Represent.* (2018) [2](#), [3](#)
25. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. In: *Adv. Neural Inform. Process. Syst.* (2020) [3](#)
26. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. In: *Adv. Neural Inform. Process. Syst.* vol. 34 (2021) [2](#), [3](#), [5](#), [6](#)
27. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2019) [3](#)
28. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2020) [3](#)
29. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013) [3](#)
30. Larochelle, H., Murray, I.: The neural autoregressive distribution estimator. In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. pp. 29–37. *JMLR Workshop and Conference Proceedings* (2011) [3](#)
31. Lin, C.H., Chang, C.C., Chen, Y.S., Juan, D.C., Wei, W., Chen, H.T.: Cocom-gan: Generation by parts via conditional coordinating. In: *Int. Conf. Comput. Vis.* pp. 4512–4521 (2019) [2](#), [3](#)
32. Lin, C.H., Lee, H.Y., Cheng, Y.C., Tulyakov, S., Yang, M.H.: Infinitygan: Towards infinite-resolution image synthesis. *Int. Conf. Learn. Represent.* (2021) [2](#), [3](#), [4](#)
33. Lin, J., Zhang, R., Ganz, F., Han, S., Zhu, J.Y.: Anycost gans for interactive image synthesis and editing. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 14986–14996 (2021) [3](#), [11](#)
34. Lu, L., Li, W., Tao, X., Lu, J., Jia, J.: Masa-sr: Matching acceleration and spatial adaptation for reference-based image super-resolution. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6368–6377 (2021) [4](#)
35. Ma, Y., Ma, J., Zhou, M., Chen, Q., Ge, T., Jiang, Y., Lin, T.: Boosting image outpainting with semantic layout prediction. *arXiv preprint arXiv:2110.09267* (2021) [4](#)
36. Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: *Int. Conf. Comput. Vis.* pp. 14214–14223 (2021) [3](#)
37. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: *Eur. Conf. Comput. Vis.* pp. 405–421. Springer (2020) [2](#), [3](#)
38. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *Int. Conf. Machine Learning*. pp. 8162–8171. PMLR (2021) [3](#)

39. Ntavelis, E., Shahbazi, M., Kastanis, I., Timofte, R., Danelljan, M., Van Gool, L.: Arbitrary-scale image synthesis. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022) [3](#)
40. Van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al.: Conditional image generation with pixelcnn decoders. In: Adv. Neural Inform. Process. Syst. vol. 29 (2016) [3](#)
41. Park, T., Zhu, J.Y., Wang, O., Lu, J., Shechtman, E., Efros, A.A., Zhang, R.: Swapping autoencoder for deep image manipulation. In: Adv. Neural Inform. Process. Syst. (2020) [3](#), [8](#)
42. Parmar, G., Zhang, R., Zhu, J.Y.: On aliased resizing and surprising subtleties in gan evaluation. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022) [9](#)
43. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3d-aware image synthesis. In: Adv. Neural Inform. Process. Syst. vol. 33, pp. 20154–20166 (2020) [3](#)
44. Shaham, T.R., Dekel, T., Michaeli, T.: Singan: Learning a generative model from a single natural image. In: Int. Conf. Comput. Vis. pp. 4570–4580 (2019) [2](#), [3](#)
45. Shaham, T.R., Gharbi, M., Zhang, R., Shechtman, E., Michaeli, T.: Spatially-adaptive pixelwise networks for fast image translation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 14882–14891 (2021) [3](#)
46. Shechtman, E., Irani, M.: Matching local self-similarities across images and videos. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1–8. IEEE (2007) [4](#)
47. Shocher, A., Bagon, S., Isola, P., Irani, M.: Ingan: Capturing and retargeting the “dna” of a natural image. In: Int. Conf. Comput. Vis. pp. 4492–4501 (2019) [3](#)
48. Shocher, A., Cohen, N., Irani, M.: “zero-shot” super-resolution using deep internal learning. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 3118–3126 (2018) [2](#), [4](#)
49. Skorokhodov, I., Ignatyev, S., Elhoseiny, M.: Adversarial generation of continuous images. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10753–10764 (2021) [3](#)
50. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: Int. Conf. Learn. Represent. (2021) [3](#)
51. Song, Y., Ermon, S.: Improved techniques for training score-based generative models. In: Adv. Neural Inform. Process. Syst. vol. 33, pp. 12438–12448 (2020) [3](#)
52. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: Adv. Neural Inform. Process. Syst. vol. 33, pp. 7537–7547 (2020) [3](#)
53. Teterwak, P., Sarna, A., Krishnan, D., Maschinot, A., Belanger, D., Liu, C., Freeman, W.T.: Boundless: Generative adversarial networks for image extension. In: Int. Conf. Comput. Vis. pp. 10521–10530 (2019) [4](#)
54. Van Oord, A., Kalchbrenner, N., Kavukcuoglu, K.: Pixel recurrent neural networks. In: Int. Conf. Machine Learning. pp. 1747–1756. PMLR (2016) [3](#)
55. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Adv. Neural Inform. Process. Syst. vol. 30 (2017) [3](#)
56. Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A.: Cnn-generated images are surprisingly easy to spot... for now. In: IEEE Conf. Comput. Vis. Pattern Recog. (2020) [4](#)
57. Wang, X., Xie, L., Dong, C., Shan, Y.: Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: Int. Conf. Comput. Vis. pp. 1905–1914 (2021) [4](#), [11](#), [12](#)

58. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European conference on computer vision (ECCV) workshops (2018) 4
59. Wang, Y., Tao, X., Shen, X., Jia, J.: Wide-context semantic image extrapolation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1399–1408 (2019) 4
60. Wexler, Y., Shechtman, E., Irani, M.: Space-time completion of video. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(3), 463–476 (2007) 4
61. Xia, B., Tian, Y., Hang, Y., Yang, W., Liao, Q., Zhou, J.: Coarse-to-fine embedded patchmatch and multi-scale dynamic aggregation for reference-based super-resolution. *arXiv preprint arXiv:2201.04358* (2022) 4
62. Xu, R., Wang, X., Chen, K., Zhou, B., Loy, C.C.: Positional encoding as spatial inductive bias in gans. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 13569–13578 (2021) 3
63. Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5791–5800 (2020) 4
64. Yang, Z., Dong, J., Liu, P., Yang, Y., Yan, S.: Very long natural scenery image prediction by outpainting. In: *Int. Conf. Comput. Vis.* pp. 10561–10570 (2019) 4
65. Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015) 3, 8
66. Zhang, K., Danelljan, M., Li, Y., Timofte, R., Liu, J., Tang, J., Wu, G., Zhu, Y., He, X., Xu, W., et al.: Aim 2020 challenge on efficient super-resolution: Methods and results. In: *Eur. Conf. Comput. Vis.* pp. 5–40. Springer (2020) 4
67. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2018) 7
68. Zhao, S., Cui, J., Sheng, Y., Dong, Y., Liang, X., Chang, E.I., Xu, Y.: Large scale image completion via co-modulated generative adversarial networks. In: *Int. Conf. Learn. Represent.* (2021) 7
69. Zhao, S., Liu, Z., Lin, J., Zhu, J.Y., Han, S.: Differentiable augmentation for data-efficient gan training. In: *Adv. Neural Inform. Process. Syst.* vol. 33, pp. 7559–7570 (2020) 3
70. Zhao, Z., Zhang, Z., Chen, T., Singh, S., Zhang, H.: Image augmentations for gan training. *arXiv preprint arXiv:2006.02595* (2020) 3
71. Zheng, H., Ji, M., Wang, H., Liu, Y., Fang, L.: Crossnet: An end-to-end reference-based super resolution network using cross-scale warping. In: *Eur. Conf. Comput. Vis.* pp. 88–104 (2018) 4
72. Zhou, Y., Zhu, Z., Bai, X., Lischinski, D., Cohen-Or, D., Huang, H.: Non-stationary texture synthesis by adversarial expansion. *ACM Trans. Graph.* (2018) 4