

Adversarial Generation of Continuous Images

Ivan Skorokhodov¹, Savva Ignatyev², Mohamed Elhoseiny¹

¹King Abdullah University of Science and Technology (KAUST)

²Skolkovo Institute of Science and Technology

iskorokhodov@gmail.com

savvaignatiev@gmail.com

mohamed.elhoseiny@kaust.edu.sa

Abstract

In most existing learning systems, images are typically viewed as 2D pixel arrays. However, in another paradigm gaining popularity, a 2D image is represented as an implicit neural representation (INR) — an MLP that predicts an RGB pixel value given its (x, y) coordinate. In this paper, we propose two novel architectural techniques for building INR-based image decoders: factorized multiplicative modulation and multi-scale INRs, and use them to build a state-of-the-art continuous image GAN. Previous attempts to adapt INRs for image generation were limited to MNIST-like datasets and do not scale to complex real-world data. Our proposed INR-GAN architecture improves the performance of continuous image generators by several times, greatly reducing the gap between continuous image GANs and pixel-based ones. Apart from that, we explore several exciting properties of the INR-based decoders, like out-of-the-box superresolution, meaningful image-space interpolation, accelerated inference of low-resolution images, an ability to extrapolate outside of image boundaries, and strong geometric prior. The project page is located at <https://universome.github.io/inr-gan>.

1. Introduction

In deep learning, images are typically represented as 2D arrays of pixels. However, there is another paradigm which views an image as a continuous function $F(\mathbf{p}) = \mathbf{v}$ that maps a pixel's 2D coordinate $\mathbf{p} = (x, y) \in \mathbb{R}^2$ to its RGB value $\mathbf{v} = (r, g, b) \in \mathbb{R}^3$. The advantage of such a representation is that it gives a true continuous version of the underlying 2D signal instead of its cropped quantized counterpart like pixel-based representations do. In practice, we almost never know the underlying function $F(\mathbf{p})$ and thus have to work with its approximations. The most popular and expressive way to approximate $F(\mathbf{p})$ is through a neural network F_θ [81, 74]. It is called an *implicit neural representation (INR)* and is especially popular in 3D deep learning where working with voxels directly (i.e. pixels defined

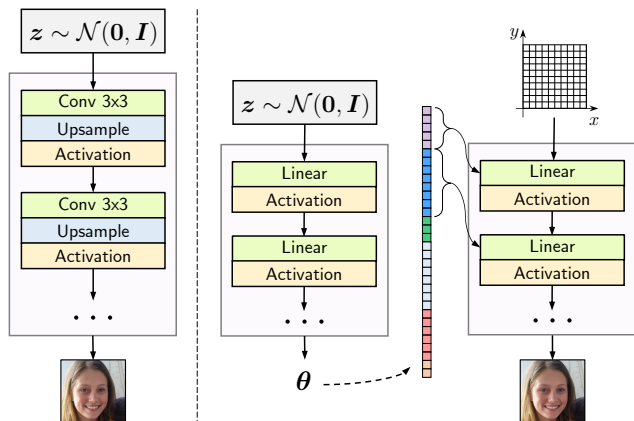


Figure 1: Comparison between a traditional convolutional generator (left) and an INR-based one (right). A traditional generator directly generates a pixel-based image representation given its latent code z . The INR-based one produces parameters of an MLP. The corresponding pixel-based representation is obtained by evaluating the INR at each coordinate location (x, y) of a specified grid.

on a 3D grid) is too expensive [51, 11, 67, 55].

Building such a decoder has two severe difficulties: 1) since it is a hypernetwork, i.e. a network that produces parameters for another network [26], it is unstable to train and requires too many parameters in general [9]; and 2) it is too costly to evaluate INRs for a dense high-resolution coordinates grid limiting their application to low-resolution images only. To alleviate these issues, we design two principled architectural techniques: *factorized multiplicative modulation (FMM)* layer for hypernetworks and *multi-scale INRs*. We use these techniques to build *INR-GAN*: a state-of-the-art continuous image generator that generates pictures in their INR representations. Previous attempts to build such a model were only conducted on small MNIST-like datasets [11, 3, 51] and do not scale to complex real-world data. In our case, we managed to achieve FID [28] scores of 5.09, 4.96 and 16.32 on LSUN Churches

256², LSUN Bedrooms 256² and FFHQ 1024², respectively, greatly reducing the gap between continuous image GANs and their pixel-based analogs. In their contemporary work, [1] achieved even better results by employing a large-scale INR-based decoder with learnable coordinate embeddings.

In our paper, we also shed light on many interesting properties of the INR-based decoders:

- *Extrapolating outside of image boundaries* (see Fig. 2): an ability to generate a “zoomed-out” version of an image without being trained explicitly to do this.
- *Geometric prior* (see Fig. 4): better encoding of geometric properties of a dataset in the latent space.
- *Accelerated low-resolution inference* (see Fig. 17): an ability to generate a low-resolution version of an image in shorter time than an image of the corresponding training-time resolution.
- Meaningful image space interpolation (see Fig. 5).
- *Out-of-the-box superresolution* (see Fig. 3): an ability to produce a higher-resolution version of an image without being trained for this task at all.

We emphasize that these features come naturally to INR-based decoders and do not require any additional training.

To summarize, our contributions are the following:

1. We propose a novel factorized multiplicative modulation (FMM) layer for hypernetworks. It makes it possible to generate INRs with a large number of parameters and stabilizes hypernetwork training.
2. We propose a novel *multi-scale* INR architecture. It makes it possible to represent high-resolution images in the INR-based form in a very efficient way.
3. Using the above two techniques we build INR-GAN: an INR-based GAN model that outperforms existing continuous image generators by several times on large real-world datasets.
4. We explore several exciting properties of INR-based decoders, like out-of-the-box superresolution, meaningful image-space interpolation, accelerated inference of low-resolution images, an ability to extrapolate outside of image boundaries, and geometric prior.

2. Related work

Implicit Neural Representations. The original idea of augmenting neural networks with coordinates information




Figure 2: **Extrapolating outside of image boundaries.** After training our INR-based GAN on LSUN 256², we tried to evaluate it on a wider grid. During training, we used coordinates from $[0, 1]^2$ square, and here we evaluate it on coordinates from $[-0.3, 1.3]^2$ square. To our surprise, the model can produce meaningful extrapolation and generate the picture beyond the coordinates area it was trained on. Blue bounding box denotes $[0, 1]^2$ coordinates area — an image area the model was trained on. Very similar results were previously shown by [49].

was proposed in CPPN [78] which is a neuroevolution-based model that is trained to represent a 2D image. After that, there were several works that use coordinates as an additional source of information to neural networks [52, 87, 86, 68, 77], but the largest popularity of implicit neural representations (INRs) is observed in 3D deep learning, where it provides a cheap and continuous way to represent a 3D shape compared to mesh/voxel/pointcloud-based ones [55, 67, 51, 21, 20, 75]. They have also been used for other tasks, like representing textures [63], 3D shapes flow [60], scenes [57, 76], audios and differential equations [74], human grasps [41] and other information [56, 13]. Occupancy Networks [55, 69, 13, 36] model a probability function of a voxel being occupied by a 3D shape and typically employ a coordinate-based decoder that operates on top of single-view images. They use the multi-resolution surface extraction method, which is similar in nature to our proposed multi-scale INR. However, in our case, we share computation between neighboring pixels while they use surface extraction to find regions to refine the predictions on. In this way, they conduct the full inference for each low-resolution coordinate which is the opposite of what we try to achieve with multi-scale INRs. DeepSDF [67] models a signed distance function instead of the occupancy function and they additionally have an encoder, which transforms an image into a latent code. IM-NET [11] proposes to train a generative model on top of these latent codes and conduct experiments not only on ShapeNet objects [8], but on MNIST images as well. DeepMeta [51] models the occupancy function and predicts decoder parameters instead of the latent codes. A vital branch of research on INRs is concerned about the most efficient way to encode coordinates positions [84, 19]. Recent works show that using Fourier features greatly improves INR expressiveness [81, 74], which we observe in our experiments as well.

GANs. State-of-the-art results in image generation are



Figure 3: **Out-of-the-box superresolution properties of the INR-based decoder.** We train our INR-GAN on LSUN 128×128 and perform upsampling at test time to 256×256 resolution with either classical interpolation techniques (first 3 columns) or with “natural” INR-based upsampling by evaluating F_θ on a denser coordinates grid. INR-based superresolution is much crisper and obtained without any additional supervision. Numbers in parentheses denote the corresponding UpsampledFID scores.

held by generative adversarial networks (GANs) [23]. Two key challenges in GAN training are instability and mode collapse, that is why a big part of research in recent years was devoted to finding stronger objective formulation [2, 53, 48], regularizers [25, 54, 58] and architectural designs [38, 39, 40, 37] that would encourage stability, diversity and expressivity of the GAN-based models.

Generative models + coordinates. There were attempts to combine generative modeling and INR-based representations prior to our work. IM-NET [11] trains a generative model on top of latent codes of an occupancy auto-encoder. In our case, instead of feeding a latent code to a coordinate-based decoder, we produce its parameters with a hypernetwork-based generator. Besides, their image generation experiments were limited to small-scale MNIST-like datasets [46] only. CoordConv GAN [52] concatenates coordinates to each representation in the DCGAN model, which endows it with geometric translating behavior during latent space interpolation. [91] use CoordConv layers to perform superresolution. SBD [88] augments a VAE [43] decoder with coordinates information. SpatialVAE [3] additionally models rotation and translation separately from the latent codes COCO-GAN [49] generates images by patches given their spatial information and then assembles them into a single image. [93] uses coordinates-based convolutions for sky replacement and video harmonization. In their con-

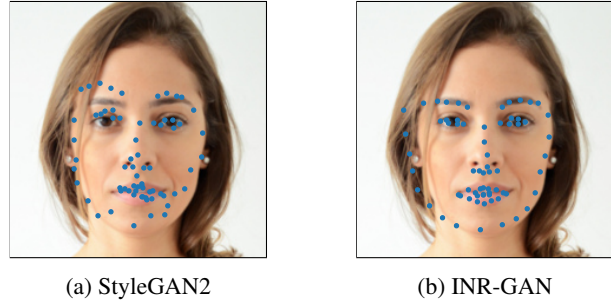


Figure 4: **Predicting keypoints from latent codes.** We train a linear model to predict face keypoints directly from the corresponding latent codes. Its performance shows how much geometric information is contained in a latent code and how accessible it is. For this benchmark, our model easily outperforms StyleGAN2 despite it being trained with the additional PPL loss [40] that forces better latent codes conditioning in the decoder. That demonstrates better *geometric prior* of our model. The corresponding scores are presented in Table 3.

temporary work, [1] builds an INR-based generator with learnable coordinate embeddings and achieves state-of-the-art results on several large-scale datasets.

Hypernetworks. Hypernetworks or meta-models are models that generate parameters for other models [26, 50]. Such parametrization provides higher expressivity [17, 18] and compression due to weight sharing through a meta-model [26]. Our factorized multiplicative modulation (FMM) is closely related to the squeeze-and-excitation mechanism [30]. But in contrast to it, we modulate weights instead of hidden representations, similar to [10], where authors condition kernel weights on an input via attention. Hypernetworks are known to be unstable to train [83] and [9] proposed a principled initialization scheme to remedy the issue. In our case, we found it to be unnecessary since our FMM module successfully reduces the influence of hypernetwork initialization on signal propagation inside an INR. Hypernetworks found many applications in other areas like few-shot learning [4], continual learning [85], architecture search [92] and others. In the case of generative modeling, [80] built a character-level language model, and [71] proposed a generative hypernetwork to produce parameters of neural classifiers. Similar to our FMM, [79] proposed a low-rank modulation by parametrizing a target model’s weight matrix as $W = W_s \odot (A \times B)$, where W_s is a shared component and A, B are rectangular matrices produced by a hypernetwork.

Hypernetworks + generative modeling. Combining hypernetworks and generative models is not new. In [71] and [27], authors built a GAN model to generate parameters of a neural network that solves a regression or clas-

sification task and demonstrate its favorable performance for uncertainty estimation. HyperVAE [59] is designated to encode any target distribution by producing generative model parameters given distribution samples. HCNAF [64] is a hypernetwork that produces parameters for a conditional autoregressive flow model [42, 66, 31]. [7] trains a generator to produce a 3D object in the INR form and uses mFiLM[70, 14] to compress the output space.

Hypernetworks + INRs. There are also works that combine hypernetworks and INRs. [76] represents a 3D scene as an INR, rendered by differentiable ray-marching, and trains a hypernetwork to learn the space of such 3D scenes. [45] proposes to represent an image dataset using a hypernetwork and perform super-resolution by passing denser coordinate grid into it. DeepMeta [51] builds an encoder that takes a single-view 3D shape image as input and outputs parameters of an INR. Authors also trained a model to encode a MNIST image into a temporal sequence of digits.

Computationally efficient models. Among other things, we demonstrate that our INR-based decoder enjoys faster inference speed compared to classical convolutional ones. Each INR layer can be seen as a 1×1 convolution [29, 73], i.e. a convolution with the kernel size of 1. The core difference is that layer’s weights are produced with a hypernetwork G. Since INR may have a lot of parameters, we factorize them with low-rank factorization [90]. There is a vast literature on using low-rank matrix approximations to compress deep models or accelerate them [61, 12, 34, 15]. In our case, we found it sufficient just to output a weight matrix as a product of two low-rank matrices $\mathbf{W} = \mathbf{A} \times \mathbf{B}$ without using any specialized techniques. SENet [30] proposed to apply squeeze-and-excitation mechanism on the hidden representations, and we apply them on the INR weights to make the model more stable to train and faster to converge.

3. Image Meta Generation

3.1. Model overview

We adopt the standard GAN training setup and replace a convolutional generator with our INR-based one, illustrated in Figure 1. We build upon the StyleGAN2 framework and keep every other component except for the generator untouched, including the discriminator, losses, optimizers, and the hyperparameters. The details are in Appendix A.

Our generator G is hypernetwork-based: it takes latent code $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ as input and generates parameters θ for an INR model F_θ . To produce an actual image, we evaluate F_θ at all the locations of a predefined coordinates grid, which size is determined by the dataset resolution. For example, for LSUN bedroom 256^2 we compute pixel values at $256^2 = 65536$ grid locations of $[0, 1]^2$ square, that are positioned uniformly. We use recently proposed Fourier fea-



(a) Image interpolation in the pixel-based form.



(b) Image interpolation in the INR-based form.

Figure 5: Images have meaningful interpolation when represented in the INR-based form. To interpolate between F_{θ_1} and F_{θ_2} , we compute interpolation parameters $\theta = \alpha\theta_1 + (1 - \alpha)\theta_2$ and evaluate F_θ for the provided coordinates grid.

tures to embed the (x, y) coordinates [81, 74]. From the implementation perspective, these Fourier features is just a simple linear layer $\mathbf{u} = \sin(\gamma \mathbf{U} \mathbf{p})$ with sine (or cosine) activation that maps a raw coordinate vector $\mathbf{p} = (x, y)$ into a feature vector $\mathbf{u} \in \mathbb{R}^{d_f}$. Note that our embedding matrix \mathbf{U} is not kept fixed and shared but predicted by the generator from z . This gives the model flexibility to select feature frequencies that are the most appropriate for a given image.

3.2. Factorized Multiplicative Modulation (FMM)

A hypernetwork is a model that generates parameters for another model. In our case, we want to generate INR’s parameters given the noise vector z . Imagine, that we need to produce the weights $\mathbf{W}^\ell \in \mathbb{R}^{n_{in} \times n_{out}}$, $\mathbf{b}^\ell \in \mathbb{R}^{n_{out}}$ of the ℓ -th linear layer of F_θ . A naive implementation would output them directly, but this is extremely inefficient: if the hypernetwork has the hidden dimensionality of size h , then its output projection matrix will have the size $h \times (n_{in} \cdot n_{out} + n_{out})$. Even for small n_{in}, n_{out} this is prohibitively expensive.

The main problem lies in generating \mathbf{W}^ℓ since it contains most of the weights, thus factorizing it via low-rank matrix decomposition $\mathbf{W}^\ell = \mathbf{A}^\ell \times \mathbf{B}^\ell$ might seem like a reasonable idea. However, our preliminary experiments showed that it severely decreases the performance because having low-rank weight matrices is equivalent to having a lot of zero singular values, leading to severe instabili-

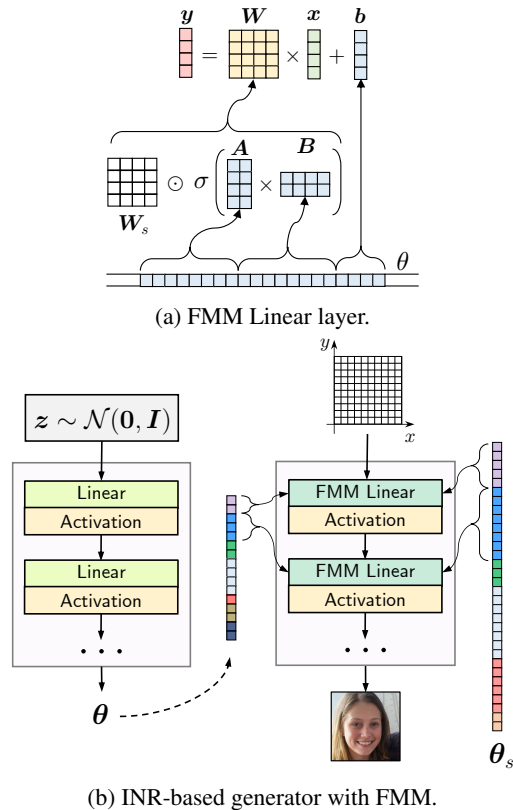


Figure 6: (a) FMM linear layer for inputs x , shared matrix W_s and output y ; (b) INR-based generator with the FMM mechanism: its parameters are split into θ_s (shared) and θ (predicted by the hypernetwork). This mechanism makes our architecture be similar to StyleGAN2 [40]: the hypernetwork becomes a *mapping network* and the INR F_θ becomes the *synthesis network* (decoder).

ties in GAN training [5, 58]. That is why we change the parametrization so that the full rank is preserved while the hypernetwork output is still factorized.

Namely, we first define a shared parameters matrix $W_s^\ell \in \mathbb{R}^{n_{out}^\ell \times n_{in}^\ell}$, which is learnable and shared across all samples. Our hypernetwork produces two rectangular matrices $A^\ell \in \mathbb{R}^{n_{out}^\ell \times r}$ and $B^\ell \in \mathbb{R}^{r \times n_{in}^\ell}$, which are multiplied together to obtain a low-rank *modulating* matrix $W_h^\ell = A^\ell \times B^\ell$. We compute the final weight matrix W^ℓ as $W^\ell = W_s^\ell \odot \sigma(W_h^\ell)$ where σ is sigmoid function. Bias vector b^ℓ is produced directly since it is small. In all the experiments, we set $r = 10$ for all the layers of F_θ except the first and the last ones which are small enough to be generated directly. We illustrate FMM in Figure 6.

The above modulation is very similar to the one proposed in [79] with the difference that [79] does not use the sigmoid activation. However, as our experiments in 4 demonstrate, using the activation is crucial since it bounds

the activations and makes the training more stable. Note that the FMM-based INR-GAN becomes very close architecturally to StyleGAN2 [40], which uses a mapping network to predict the style vector used to modulate the convolutional weights of its decoder via multiplication.

3.3. Multi-scale INRs

Scaling a traditional INR-based decoder to large image resolutions is too expensive: to produce a 1024^2 image, we have to input $\approx 10^6$ coordinates into F_θ . Such a huge batch size makes it impossible to use large hidden layers' sizes due to excessive computation and memory consumption. To circumvent the issue, we propose a *multi-scale INR* architecture: we split F_θ into K blocks, where each block operates on its own resolution and only the final block operates on the target one. Earlier blocks compute low-res features that are then replicated and passed to the next level. This process is illustrated on Fig. 7 and is equivalent to using different grid sizes depending on the resolution and using the nearest neighbor interpolation for the hidden representations. For the multi-scale INRs, we use more neurons for lower resolutions and fewer ones for the more expensive high-resolution blocks. The use of multi-scale INRs makes our architecture very similar to classical convolutional decoders, which grow the resolution progressively with depth [62]. In our case, one of the additional benefits of having the multi-scale architecture is that neighboring pixels get conditioned on the common context computed at a previous resolution.

We start with the resolution of 64^2 for the first block and increase it by 2 for each next one until we reach the target resolution. Each block contains 2-4 layers, and Fourier coordinates features are concatenated to a hidden representation at the beginning of each block. Replicating low-resolution features for each high-dimensional pixel is equivalent to upsampling the inner grid with nearest neighbors interpolation, and this is the way we implement it in practice. Further implementation details can be found in the supplementary material or the attached source code.

4. Experiments

4.1. Standard GAN training

Datasets. We conduct experiments on four datasets: LSUN Bedrooms 128^2 , LSUN Bedrooms 256^2 , LSUN Churches 256^2 and FFHQ 1024^2 . LSUN Bedrooms and LSUN Churches consist on 3M and 125k images of in-the-wild bedrooms and churches [89], respectively. FFHQ is a high-resolution dataset of 70k human faces [39]. For all the datasets, we apply random horizontal flip for data augmentation.

Evaluation metrics. We evaluate the model using Frechet Inception Distance (FID) [28] metric using 50k im-

Table 1: We start with the INR-based decoder conditioned on the latent code [11, 3] and progressively improve it. O/M denotes “out-of-memory” error: we couldn’t train the model even for a batch size of 1 on a 32GB NVidia V100 GPU.

Decoder type	LSUN 128 ²			LSUN 256 ²			FFHQ 1024 ²		
	GMACs	#params	FID	GMACs	#params	FID	GMACs	#params	FID
Latent-code conditioned INR decoder [44, 3]	30.09	7.1M	229.9	120.33	7.1M	253.3	1925.22	7.1M	O/M
+ Hypernetwork-based decoder [51]	23.54	2055.5M	28.83	88.01	2055.5M	O/M	1377.52	2055.5M	O/M
+ Fourier embeddings from [74, 81] (ours)	28.02	2312.02M	23.07	105.34	2312.02M	O/M	1651.52	2312.02M	O/M
+ Factorized Multiplicative Modulation (ours)	25.87	108.2M	11.51	103.19	108.2M	15.68	1649.37	108.2M	O/M
+ Multi-Scale INR (ours)	21.58	107.03M	5.69	38.76	107.03M	6.27	47.35	117.3M	16.32
StyleGAN2 generator [39]	-	-	-	84.36	30.03M	2.65	143.18	30.37M	4.41

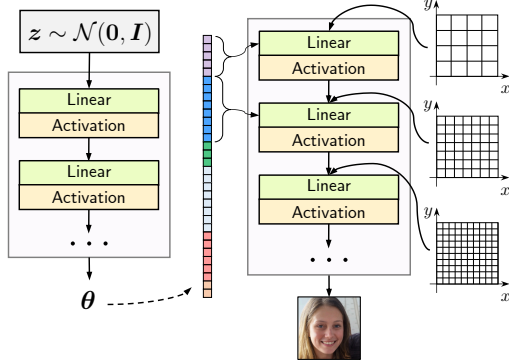


Figure 7: **Multi-Scale INR-based GAN (without FMM).** Each block operates on a different resolution, determined by the granularity of an input grid. We increase the granularity with depth: this allows to share computation between neighbouring pixels *and* condition them on a common context. We depict the multi-scale mechanism without FMM not to clutter the illustration. In practice, we use both FMM and the multi-scale architecture for our INR-GAN.

ages to compute the statistics. We also compute the number of parameters for a given model and the amount of multiply-accumulate operations (MACs) — a standard measure for accessing the model’s computational efficiency [29, 16].

Models and training details. All our models have equivalent training settings and differ only in the generator architecture. We use two existing architectures as our baseline:

1. A non-hypernetwork-based generator, which has shared parameters for all INRs and each INR is conditioned on latent code $w = G(z)$ [11, 3, 67]. I.e., instead of producing parameters for F_θ , we pass w into it as an additional input $v = F_\theta(x, y, w)$.
2. A hypernetwork-based generator, which produces parameters θ for F_θ but does not use any factorization techniques to reduce the output matrix size [51].

We build upon the above baselines by incorporating Fourier positional embeddings of the coordinates [81, 74], incorporating our FMM layer and incorporating our multi-scale

INR architecture. In all the experiments, G is a 4-layer MLP with residual connections that takes $z \in \mathbb{R}^{512}$ as input and produces INR parameters θ as output. For the first baseline, it produces the transformed latent codes instead. For hypernetwork-based models, we additionally apply a learned linear layer to w to obtain INR parameters θ . In all the experiments, a ResNet-based discriminator from StyleGAN2 [40] is used. However, since beating the scores is not the goal of the paper, we used its “small” version (corresponding to config-e). R_1 -regularization [54] with the penalty weight of 10 is used. All the models are trained for 800k iterations on 4 NVidia V100 32GB GPUs.

Results. The results are reported in Table 1. As we can see, the latent code conditioned baseline cannot fit training data at all since it lacks expressivity and cannot capture the whole image representation in the latent code. Its hypernetworks-based counterpart achieves much higher performance but is still not competitive. An attempt to fit the baselines on FFHQ 1024² resulted in out-of-memory errors even for a batch size of 1 for a 32GB GPU. The non-factorized hypernetwork-based decoder [51] was too expensive even for 256² resolution.

Our INR-based decoder with FMM layer and multi-scale INR architecture achieves competitive FID scores and greatly reduces the gap between continuous and pixel-based image generation. It is still inferior to StyleGAN2 [40] in terms of performance, but is three times more efficient. One should also note that we didn’t use any of the numerous training tricks employed by StyleGAN2 and that our discriminator architecture is smaller (it corresponds to config-e [40]). Besides, as we show in Section 4.2, our model naturally gives rise to a lot of additional properties that convolutional decoders lack. Our model also has three times more parameters than its convolution-based counterpart. The reason for it is the huge size of the output projection matrix of G , which has the dimensionality of $\approx 10^3 \times 10^5$ occupying 90-99% of parameters of the entire model. Compressing this matrix is an ongoing hypernetworks research topic [26, 15] and we leave this for future work. We emphasize that despite its enormous size, the output layer’s inference time is almost negligible due to highly optimized matrix-matrix multiplications on modern GPUs.

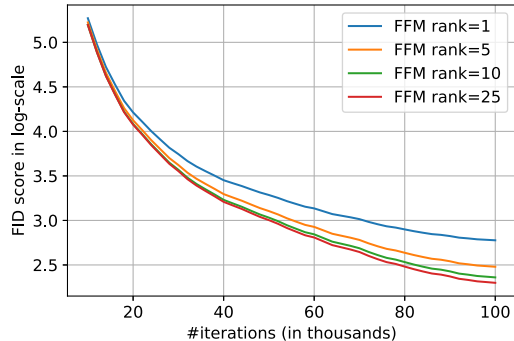


Figure 8: FID scores (in log scale) for different rank values of the proposed FMM modulation. Increasing the rank beyond 1 gives clear advantage, but the returns are diminishing for the subsequent rank increase.

Table 2: FID scores for additional ablations of our multi-scale INR-based GAN (INR-GAN). Removing $\sigma(x)$ from FMM *worsens* the scores. Incorporating StyleGAN2’s architecture and bilinear upsampling allows the INR-based generator to rival its convolution-based counterpart. Img/sec of G is measured on 1 NVidia V100 32GB.

Decoder type	Churches ↓	Bedrooms ↓	img/sec ↑
INR-GAN w/o FMM activation	10.35	11.73	267.3
INR-GAN	7.12	6.27	267.1
+ StyleGAN2 architecture	5.09	4.96	265.1
+ bilinear upsampling	3.12	3.41	203.4
StyleGAN2	3.86	2.65	85.5

Ablating FMM rank. In all the experiments, we use an FMM rank of 10. We ablate over its importance for LSUN bedroom 256^2 and report the resulting convergence plots on Fig. 8. As one can see, the model benefits from the increased rank, but the advantage is diminishing for further rank increase. These results show that vanilla or “full-rank” hypernetworks are heavily overparameterized and there is no need to predict each parameter separately. On the other that is the evidence that there is much potential for “going further” than squeeze-and-excitation [30] and AdaIN [33, 39] approaches, i.e. predicting more than a single modulating value for each neuron.

Additional ablations. All the above experiments were conducted on top of a vanilla INR-based decoder which uses nearest neighbour interpolation (to make neighboring pixels be computed independently as the vanilla INR does). It also lacks important StyleGAN2’s techniques which improve the performance, like equalized learning rate [38], style mixing, pixel normalization [38] and noise injection [39]. In all our experiments, we also used a small-size version of StyleGAN2’s discriminator to make the training run faster. Incorporating the above design choices and giving up the nearest neighbour upsampling for the bilinear one (which might be reasonable for some applications) sig-

nificantly improves our generator’s performance. We conduct those experiments on LSUN Churches 256^2 and LSUN Bedrooms 256^2 and report the results in Table 2. For the INR-based decoder with bilinear upsampling, we even surpass StyleGAN2’s performance on Churches.

Another critical question is how important the activation in FMM is. As discussed in Section 3.2, it allows to stabilize training by bounding the magnitudes of the weights. We ablate its influence by replacing σ with the identity mapping and report the resulted FID in Table 2. The drop in performance demonstrates that it carries a crucial influence on the image quality.

Performance on multi-class datasets. To test if the proposed architectural techniques improve the performance on diverse multi-class datasets, we also conduct experiments on LSUN-10 256^2 and MiniImageNet 128^2 . We provide the details of these experiments in Appendix F.

4.2. Exploring the properties

Extrapolating outside of image boundaries. At test time, we sample pixels beyond the coordinates grid that the model was trained on and present the results on Fig. 2. It shows that an INR-based decoder is capable of generating meaningful content outside of the grid boundaries, which is equivalent to a zooming-out operation. It is a surprising quality indicating that the coordinates features produced by the generator are exploited by an INR in a generalizable manner instead of just being simply memorized.

Meaningful interpolation. Image space interpolation is known for its poor behavior [22]. However, when images are represented in the INR-based form and not the pixel-based one, the interpolation becomes reasonable. We illustrate the difference on Fig. 5.

Keypoints prediction. Direct access to coordinates provides more fine-grained control over the geometrical properties of images during the generation process. Thus one can expect the INR-based decoder to better embed the geometric structure of an image into the latent space [52]. To test this hypothesis, we perform the following experiment. We take 10k samples from our model trained on FFHQ 256^2 and extract face keypoints with a pre-trained model [6]. After that, we fit a simple linear regression model to predict these keypoints coordinates given the corresponding latent code. We compute this model’s test loss on the real-world images from FFHQ 256^2 and coin this metric Keypoints Prediction Loss (KPL). The above procedure is formally described in Algorithm 1 in Appendix B. The corresponding latent codes for the real-world images are obtained by projecting an image into the corresponding latent space through gradient descent optimization. We use the standard protocol from [40] to do this for both models. KPL is computed for both \mathcal{Z} -space and \mathcal{W} -space. For our model, this corresponds to G input noise vectors z and penultimate hidden

Table 3: Keypoints Prediction Loss (KPL) metric for our INR-based decoder and StyleGAN2 generator computed on FFHQ 256×256 dataset for different latent spaces. The qualitative difference is illustrated on Fig. 4.

Generator	KPL (\mathcal{Z})	KPL (\mathcal{W})	KPL (random)
StyleGAN2	$2.1 \cdot 10^{-4}$	$7.6 \cdot 10^{-5}$	$4.6 \cdot 10^{-4}$
INR-GAN	$6.0 \cdot 10^{-5}$	$5.2 \cdot 10^{-5}$	$4.7 \cdot 10^{-4}$

representations (i.e. hidden representations before the output projection into INR parameter space). For StyleGAN2, this corresponds to the mapping network’s input and output space. We also compute KPL on top of random vectors to check that a better performance in keypoints prediction is not due to their reduced variability. The results of this experiment are provided in Table 3.

Accelerated inference of lower-resolution images. Our INR-GAN has a natural capability of generating a low-resolution sample faster because it can be directly evaluated at a low-resolution coordinates grid without performing the full generation. We explore this in Appendix G.

Out-of-the-box superresolution. Our INR-based decoder is able to produce images of higher resolution than it was trained on. For this, we evaluate our model on a denser coordinates grid. To measure this quantitatively, we propose UpsampledFID: a variant of FID score where fake data statistics are computed based on upsampled low-resolution images. We compare our UpsampledFID score to three standard upsampling techniques: nearest neighbor, bilinear, and bicubic interpolation. Fig. 3 demonstrates that INR-upsampling improves the score by up to 50%.

5. Additional potential of INR-based decoders

In Section 4.2, we explored several exciting properties of INR-based decoders and tested them in practice. Here, we highlight their additional potential and leave its exploration for future work.

An ability to backpropagate through pixels positions. Since coordinate positions (x, y) are transformed via well-differentiable operations, it gives a possibility to back-propagate through coordinate positions. This opens a large range of possible work like using spatial transformer layers [35] in the generator and not only discriminator. Or producing a coordinates grid with the discriminator for zooming-in into specific parts of an image.

Faster inference speed. Since INRs do not use local context information during inference, it does not spend computation on aggregating it as convolutions do. This is equivalent to using convolutions with a kernel size of 1 and thus works much faster [29]. Table 1 demonstrates that our model uses three times fewer MACs for its inference process, which is due to ignoring context information.

Parallel pixel computation Non-autoregressive models like Parallel WaveNet [65, 24] are valued by their ability to generate an arbitrarily long sequence in parallel, so their inference speed decreases linearly with more compute being added. INR-based decoders also have this property due to their independent pixel generation nature. Convolutional decoders, in contrast, are forced to use local context during the inference process, which limits their parallel inference.

Universal decoder architecture One can employ the same decoder architecture for training a decoder model in any domain: 2D images, 3D shapes, video, audio, etc. They would only differ at what coordinates are being passed as input to the corresponding INR. For 2D images, these would be $(x, y) \in \mathbb{R}^2$ coordinates, for 2D video, this would be $(x, y, t) \in \mathbb{R}^2 \times \mathbb{R}_+$ with the additional timestep $t \in \mathbb{R}_+$ input, for 3D shapes — $(x, y, z) \in \mathbb{R}^3$ coordinates, etc. That has already been partially explored in [51, 74].

Biological plausibility. While convolutional *encoders* have a very intimate connection to how a human eye works [47], convolutional *decoders* do not have much resemblance to any human brain mechanism. In contrast, [72] argue that hypernetworks have a very close relation to how a prefrontal cortex influences other brain parts. It does so by modulating the activity in several different areas at once, precisely how our hypernetwork-based generator applies modulation to different INR layers.

6. Limitations

The core limitation of INR-based decoders comes from their strengths: not using spatially common context between neighboring pixels. Multi-Scale INR architecture partially alleviates this issue by grounding them on the same low-resolution representation, but as our experiment with adding bilinear interpolation demonstrates (see Table 2), it is not enough. Also, we noticed that the INR-based decoders might become too sensitive to high-frequency coordinates features. This may produce “texture” artifacts: a generated image having a random transparent texture spanned on it which becomes more noticeable for higher resolutions.

7. Conclusion

This paper explored the adversarial generation of continuous images represented in the implicit neural representations (INRs) form. We proposed two principled architectural techniques: factorized multiplicative modulation and multi-scale INRs, which allowed us to obtain solid state-of-the-art results in continuous image generation. We explored several attractive properties of INR-based decoders and discussed their future potential and limitations.

References

- [1] Ivan Anokhin, Kirill Demochkin, Taras Khakhulin, Gleb Sterkin, Victor Lempitsky, and Denis Korzhnikov. Image generators with conditionally-independent pixel synthesis, 2020. [2](#), [3](#)
- [2] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017. [3](#)
- [3] Tristan Bepler, Ellen Zhong, Kotaro Kelley, Edward Brignole, and Bonnie Berger. Explicitly disentangling image content from translation and rotation with spatial-vae. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alche-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 15409–15419. Curran Associates, Inc., 2019. [1](#), [3](#), [6](#)
- [4] Luca Bertinetto, João F. Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 523–531. Curran Associates, Inc., 2016. [3](#)
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018. [5](#)
- [6] Adrian Bulat and Georgios Tzimiropoulos. Super-fan: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with gans. *arXiv preprint arXiv:1712.02765*, 2017. [7](#), [13](#)
- [7] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. *arXiv preprint arXiv:2012.00926*, 2020. [4](#)
- [8] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [2](#)
- [9] Oscar Chang, Lampros Flokas, and Hod Lipson. Principled weight initialization for hypernetworks. In *International Conference on Learning Representations*, 2020. [1](#), [3](#)
- [10] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [3](#)
- [11] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#), [2](#), [3](#), [6](#)
- [12] Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, and David P Woodruff. Algorithms for ℓ_p low rank approximation. *arXiv preprint arXiv:1705.06730*, 2017. [4](#)
- [13] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016. [2](#)
- [14] Vincent Dumoulin, Ethan Perez, Nathan Schucher, Florian Strub, Harm de Vries, Aaron Courville, and Yoshua Bengio. Feature-wise transformations. *Distill*, 2018. <https://distill.pub/2018/feature-wise-transformations>. [4](#), [6](#)
- [15] Elad Eban, Yair Movshovitz-Attias, Hao Wu, Mark Sandler, Andrew Poon, Yerlan Idelbayev, and Miguel A. Carreira-Perpinan. Structured multi-hashing for model compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [4](#), [6](#)
- [16] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 6202–6211, 2019. [6](#)
- [17] Tomer Galanti and Lior Wolf. Comparing the parameter complexity of hypernetworks and the embedding-based alternative. *arXiv preprint arXiv:2002.10006*, 2020. [3](#)
- [18] Tomer Galanti and Lior Wolf. On the modularity of hypernetworks. *Advances in Neural Information Processing Systems*, 33, 2020. [3](#)
- [19] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252, 2017. [2](#)
- [20] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Deep structured implicit functions. *arXiv preprint arXiv:1912.06126*, 2019. [2](#)
- [21] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7154–7164, 2019. [2](#)
- [22] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>. [7](#)
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [3](#), [13](#)
- [24] Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*, 2017. [8](#)
- [25] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017. [3](#)
- [26] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. [1](#), [3](#), [6](#)
- [27] Christian Henning, Johannes von Oswald, João Sacramento, Simone Carlo Surace, Jean-Pascal Pfister, and Benjamin F Grewe. Approximating the predictive distribution via adversarially-trained hypernetworks. *2018 NeurIPS Workshop on Bayesian Deep Learning*, 2018. [3](#)

- [28] 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 *Advances in neural information processing systems*, pages 6626–6637, 2017. 1, 5
- [29] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 4, 6, 8
- [30] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3, 4, 7
- [31] Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron Courville. Neural autoregressive flows. In *International Conference on Machine Learning*, pages 2078–2087, 2018. 4
- [32] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*, 2017. 20
- [33] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 7
- [34] Yerlan Idelbayev and Miguel A. Carreira-Perpinan. Low-rank compression of neural nets: Learning the rank of each layer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4
- [35] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015. 8
- [36] Timothy Jeruzalski, Boyang Deng, Mohammad Norouzi, John P Lewis, Geoffrey Hinton, and Andrea Tagliasacchi. Nasa: Neural articulated shape approximation. *arXiv preprint arXiv:1912.03207*, 2019. 2
- [37] Animesh Karnewar and Oliver Wang. Msg-gan: Multi-scale gradients for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7799–7808, 2020. 3, 13
- [38] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 3, 7, 13
- [39] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3, 5, 6, 7, 13
- [40] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3, 5, 6, 7, 13, 14, 16
- [41] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps, 2020. 2
- [42] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016. 4
- [43] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [44] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial generation of continuous implicit shape representations. *arXiv preprint arXiv:2002.00349*, 2020. 6
- [45] Sylwester Kloczek, Łukasz Maziarka, Maciej Wołczyk, Jacek Tabor, Jakub Nowak, and Marek Śmieja. Hypernetwork functional image representation. In *International Conference on Artificial Neural Networks*, pages 496–510. Springer, 2019. 4
- [46] Y. LECUN. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 3
- [47] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 1995. 8
- [48] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017. 3
- [49] Chieh Hubert Lin, Chia-Che Chang, Yu-Sheng Chen, Da-Cheng Juan, Wei Wei, and Hwann-Tzong Chen. Cogan: Generation by parts via conditional coordinating. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2, 3
- [50] Etai Littwin, Tomer Galanti, and Lior Wolf. On the optimization dynamics of wide hypernetworks. *arXiv preprint arXiv:2003.12193*, 2020. 3
- [51] Gidi Littwin and Lior Wolf. Deep meta functionals for shape representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 1, 2, 4, 6, 8
- [52] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9605–9616. Curran Associates, Inc., 2018. 2, 3, 7, 17
- [53] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017. 3
- [54] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Which training methods for gans do actually converge? In *International Conference on Machine Learning (ICML)*, 2018. 3, 6, 13
- [55] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 1, 2
- [56] Mateusz Michalkiewicz, Jhony K. Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit sur-

- face representations as layers in neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [57] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *arXiv preprint arXiv:2003.08934*, 2020. 2
- [58] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018. 3, 5
- [59] Phuoc Nguyen, Truyen Tran, Sunil Gupta, Santu Rana, Hieu-Chi Dam, and Svetha Venkatesh. Hypervae: A minimum description length variational hyper-encoding network, 2020. 4
- [60] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5379–5389, 2019. 2
- [61] Alexander Novikov, Dmitrii Podoprikin, Anton Osokin, and Dmitry P Vetrov. Tensorizing neural networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 28, pages 442–450. Curran Associates, Inc., 2015. 4
- [62] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 2016. 5
- [63] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4531–4540, 2019. 2
- [64] Geunseob Oh and Jean-Sebastien Valois. Hcnaf: Hyperconditioned neural autoregressive flow and its application for probabilistic occupancy map forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4
- [65] Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International conference on machine learning*, pages 3918–3926. PMLR, 2018. 8
- [66] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017. 4
- [67] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 1, 2, 6
- [68] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [69] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [70] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 4
- [71] Neale Ratzlaff and Li Fuxin. HyperGAN: A generative model for diverse, performant neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 5361–5369, Long Beach, California, USA, 09–15 Jun 2019. PMLR. 3
- [72] Jacob Russin, Randall O’Reilly, and Yoshua Bengio. Deep learning needs a prefrontal cortex. *Bridging AI and Cognitive Science ICLR 2020 Workshop*, 2020. 8
- [73] Laurent Sifre and Stéphane Mallat. Rigid-motion scattering for image classification. *P.h.D. thesis*, 2014. 4
- [74] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 1, 2, 4, 6, 8, 15
- [75] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. 2
- [76] Vincent Sitzmann, Michael Zollhofer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 1121–1132. Curran Associates, Inc., 2019. 2, 4
- [77] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6d object pose estimation under hybrid representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [78] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007. 2
- [79] Joseph Suarez. Character-level language modeling with recurrent highway hypernetworks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3269–3278, 2017. 3, 5
- [80] Joseph Suarez. Language modeling with recurrent highway hypernetworks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 30, pages 3267–3276. Curran Associates, Inc., 2017. 3
- [81] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020. 1, 2, 4, 6, 15
- [82] Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Con-*

- ference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE, 2016. [20](#), [21](#)
- [83] Kenya Ukai, Takashi Matsubara, and Kuniaki Uehara. Hypernetwork-based implicit posterior estimation and model averaging of cnn. In *Asian Conference on Machine Learning*, pages 176–191, 2018. [3](#)
- [84] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#), [15](#)
- [85] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F. Grewe. Continual learning with hypernetworks. In *International Conference on Learning Representations*, 2020. [3](#)
- [86] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic, faster and stronger. *arXiv preprint arXiv:2003.10152*, 2020. [2](#)
- [87] Zhenyi Wang and Olga Veksler. Location augmentation for cnn. *arXiv preprint arXiv:1807.07044*, 2018. [2](#)
- [88] Nicholas Watters, Loic Matthey, Christopher P Burgess, and Alexander Lerchner. Spatial broadcast decoder: A simple architecture for learning disentangled representations in vaes. *arXiv preprint arXiv:1901.07017*, 2019. [3](#)
- [89] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop, 2016. [5](#)
- [90] Xiyu Yu, Tongliang Liu, Xinchao Wang, and Dacheng Tao. On compressing deep models by low rank and sparse decomposition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7370–7379, 2017. [4](#)
- [91] K. Zafeirouli, A. Dimou, A. Axenopoulos, and P. Daras. Efficient, lightweight, coordinate-based network for image super resolution. In *2019 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, pages 1–9, 2019. [3](#)
- [92] Chris Zhang, Mengye Ren, and Raquel Urtasun. Graph hypernetworks for neural architecture search. In *International Conference on Learning Representations*, 2019. [3](#)
- [93] Zhengxia Zou. Castle in the sky: Dynamic sky replacement and harmonization in videos, 2020. [3](#)