3D GENERATION ON IMAGENET

Anonymous authors
Paper under double-blind review

ABSTRACT

All existing 3D-from-2D generators are designed for well-curated and alignable datasets: objects can be placed in the same position, similarly scaled and oriented, such that the camera always points to the center of the scene. This alignment procedure is infeasible for diverse, in-the-wild datasets: 1) it requires expensive annotation for each object category; and 2) most images are inherently "unalignable" (e.g., it is impossible to align a "cat face" with a "kitchen"). As a result, existing 3D generators are not scalable to large in-the-wild datasets. In this work, for the first time, we propose a 3D generator which works on non-aligned datasets. First, we develop a technique to use an off-the-shelf, imprecise depth estimator to incorporate the 3D inductive bias into a GAN-based generator. Then, we create a novel learnable camera parametrization which does not use any alignment assumptions and construct a camera gradient penalty regularization. Finally, we propose a simple distillation-based technique to transfer the knowledge from an off-the-shelf feature embedder, like ResNet50, into our discriminator. Our work is the first one which develops a 3D generator for non-aligned data. We conduct experiments on SDIP Dogs, SDIP Elephants, LSUN Horse, and ImageNet on the 256^2 resolution to demonstrate the effectiveness of our ideas. Visualizations: https://u2wjb9xxz9q.github.io/.

1 INTRODUCTION

Figure 1: Selected qualitative examples of the synthesized scenes by EG3D (Chan et al., 2022) and our method trained on ImageNet 256^2 (Deng et al., 2009). EG3D models the geometry in low resolution and renders either flat (when trained with the default camera distribution) or repetitive "layered" (when trained with a wide camera distribution) shapes. In contrast, our generator models the radiance field in full-resolution and learns high-fidelity details. Zoom-in for a better view.

We witness remarkable progress in the domain of 3D-aware image synthesis. The community is developing new methods to improve the image quality, 3D consistency and efficiency of the generators (e.g., (Chan et al., 2022; Deng et al., 2022; Skorokhodov et al., 2022; Zhao et al., 2022; Schwarz et al., 2022)). However, all the existing frameworks are designed for well-curated and aligned datasets consisting of objects of the same category, scale and global scene structure, like human or cat faces (Chan et al., 2021). Such curation requires a lot of specialized 3D knowledge about the object category at hand, since one needs to infer the underlying 3D keypoints to properly crop, rotate and scale the images (Deng et al., 2022; Chan et al., 2022). This makes it infeasible to perform a similar alignment procedure for diverse, in-the-wild datasets which contain numerous image categories and could even be inherently "non-alignable", i.e. there does not exist a single (ap-

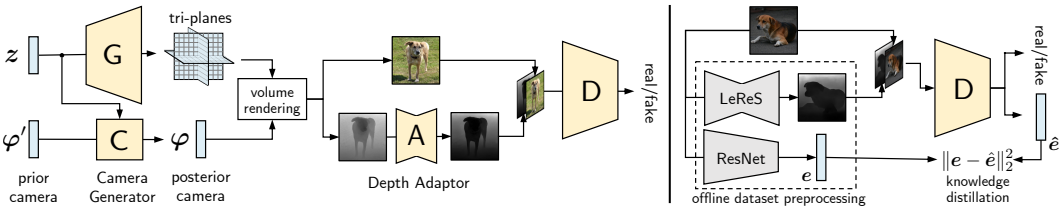

Figure 2: **3DGP overview**. Left: Our framework uses a tri-plane representation. To render an image, we first sample camera parameters from the camera prior and pass them to the camera generator, obtaining the camera posterior. We then render an image and the corresponding depth. The depth adapter reduces the gap between the rendered and the predicted depth. Right: our discriminator receives a 4-channel color-depth pair as an input. The generated sample includes the rendered image and the adapted depth. The real sample consists of a real image and its estimated depth. The discriminator has two outputs: the adversarial head and the knowledge distillation head.

proximate) canonical position which all the objects could be transformed into (e.g., it is impossible to align a landscape panorama with a spoon).

In this work, we argue that to extend 3D synthesis to in-the-wild datasets, one needs a framework which relies on more universal 3D priors. Thereafter, we make a step towards this direction and develop a generator which is guided only by (imperfect) depth predictions from an off-the-shelf monocular depth estimator. We show that these 3D cues are enough to make it learn to synthesize reasonable scenes from loosely curated, non-aligned datasets, such as ImageNet (Deng et al., 2009). Hence, we refer to our model as 3DGP: 3D generator with Generic Priors.

Training a 3D generator on non-aligned datasets comes with three main problems: 1) it is impossible to infer true camera parameters for the real images, which are needed to define a proper camera distribution for the generator; 2) objects appear in different shapes and scales, thus making it difficult to learn meaningful 3D geometry; and 3) the dataset typically contains a lot of variation in terms of texture and structure, which makes it challenging to fit even for 2D generators. As shown in Fig 1 (left), state-of-the-art 3D-aware generators, such as EG3D (Chan et al., 2022), struggle to learn the proper geometry in such a scenario. That is why we develop three novel techniques to address those problems and extend 3D synthesis to diverse non-aligned datasets.

Learnable "Ball-in-Sphere" camera distribution. Most existing methods utilize a restricted camera model (e.g., (Schwarz et al., 2020; Niemeyer & Geiger, 2021b; Chan et al., 2021)): the camera is positioned on a sphere with a constant radius, always points to the world center and has fixed intrinsics. But diverse, non-aligned datasets violate those assumptions: e.g., dogs datasets have images of both close-up snout photos and full-body dogs, which implies the variability in the focal length and look-at positions. Thus, we introduce a novel learnable camera model with 6 degrees of freedom, modeling the camera position on a fixed-radius sphere, field-of-view and the look-at position inside a fixed-radius ball. Besides, we observed that learning a camera distribution on complex datasets is prone to collapsing into delta distribution and thus introduce an efficient gradient penalty for the camera generator to prevent it.

Adversarial depth supervision (ADS). A generic image dataset features a wide diversity of objects with different shapes and poses. That is why learning a meaningful 3D geometry together with the camera distribution is an ill-posed problem, as the incorrect scale can be well compensated by an incorrect camera model (Hartley & Zisserman, 2003), or flat geometry (Zhao et al., 2022; Chan et al., 2022). To instill the 3D bias, we propose to provide the discriminator information about the scene geometry by concatenating the depth map of a scene as the 4-th channel of its RGB input. For the real images, we use their (*imperfect*) estimates from an off-the-shelf monocular depth predictor (Miangoleh et al., 2021). For the fake images, we render the depth from the synthesized radiance field, and process it with a shallow depth adaptor, bridging the distribution gap between the estimated and rendered depth maps. Exploiting monocular depth estimators for improving neural rendering was also explored in concurrent work (Yu et al., 2022), however their goal is just geometry reconstruction. Importantly, existing depth estimators generalize to a wide range of scenes, making 3DGP applicable to real-world datasets.

Knowledge distillation into Discriminator. Prior works have shown benefits from transferring the knowledge from conventional 2D image encoders into a synthesis model (Sauer et al., 2022). The

state-of-the-art techniques utilize pretrained image classifiers as the discriminator backbone and develop regularization strategies on top of them (Sauer et al., 2021; Kumari et al., 2022). These techniques, however, are only applicable when the discriminator has a similar input distribution compared to what the encoder was trained on. This makes it difficult to use efficient patch-based discriminators (Skorokhodov et al., 2022) or passing depth maps as the 4-th channel to the discriminator. That is why, we develop a more general and more efficient knowledge transfer strategy based on knowledge distillation. It consists in forcing the discriminator to predict features of a pre-trained ResNet50 (He et al., 2016) model, effectively transferring the knowledge into our model. This technique has just 1% of computational overhead compared to standard training, but allows to improve FID for both 2D and 3D generators by at least 40%.

To highlight the advantages of our ideas, we first train on *non-aligned* single-category image datasets of SDIP Dogs (Mokady et al., 2022), SDIP Elephants (Mokady et al., 2022), and LSUN Horses (Yu et al., 2015) to show that our framework can generate images of different scale. It can zoom-in into a specific region of the scene, being able to generate animal faces as well as their bodies. We report extensive numerical and visual results supporting the importance of the proposed components. Finally, for the first time, we train our 3D generator on all the $1,000$ classes of ImageNet (Deng et al., 2009) and show that multi-categorical 3D synthesis is possible on non-aligned data. We show the selected samples with high-fidelity geometry, produced by our generator, in Fig. 1 (right).

2 Related works

3D-aware synthesis. The idea of generating 3D objects and manipulating the viewpoint was previously explored under the umbrella of learning disentangled image representations, with object pose being one of the factors of variation. Visual Object Networks (Zhu et al., 2018), similar to (Tatarchenko et al., 2017), require a volumetric dataset of object shapes to train a 3D-GAN and generate a voxel-grid for each object category, followed by viewpoint and texture sampling. This disentanglement is *explicit*, as explicit 3D supervision is available at training time. A further group of encoder-decoder-based generative methods uses multi-view images as a form of explicit supervision (Olszewski et al., 2019; Nguyen-Phuoc et al., 2020). Nguyen-Phuoc et al. (2019) developed a StyleGAN-based framework to generate view-dependent images of the same object. To do so, they sample an object pose from a uniform distribution and utilize a StyleGAN2 (Karras et al., 2020b) backbone to render the object. While showing impressive initial 3D-aware synthesis results, Nguyen-Phuoc et al. (2019) requires carefully curated single-category datasets.

Mildenhall et al. (2020) introduced Neural Radiance Fields (NeRF) for representing 3D scenes and objects. NeRFs offer a convenient way to integrate 3D bias and ignited a stream of 3D-aware generators (Schwarz et al., 2020; Niemeyer & Geiger, 2021b; Chan et al., 2021; Xue et al., 2022; Zhou et al., 2021; Wang et al., 2022; Gu et al., 2022; Or-El et al., 2021; Chan et al., 2022; Skorokhodov et al., 2022; Zhang et al., 2022; Xu et al., 2021; Bahmani et al., 2022). Some of them render low-resolution images directly (Schwarz et al., 2020; Niemeyer & Geiger, 2021b; Chan et al., 2021), the majority, however, resort to the use of a convolutional upsampler to boost image resolution (Gu et al., 2022; Chan et al., 2022; Xu et al., 2021; Or-El et al., 2021; Xue et al., 2022), while maintaining reasonable computational requirements. Adopting an upsampler comes at a cost of decreased fidelity of geometric details, as volumetric rendering is done at a lower resolution. To obtain higher resolution geometry, Skorokhodov et al. (2022) proposed EpiGRAF, a generative method that uses a patch-based strategy to efficiently train at the desired output resolution without the need of upsampler. In this work, we build on top of EpiGRAF as it offers an efficient solution for 3D synthesis at high resolution and with improved geometry.

Learning the camera poses. NeRFs require known camera poses, obtained from multi-view stereo (Schönberger et al., 2016) or structure from motion (Schönberger & Frahm, 2016). Alternatively, a group of works have been introduced to either automatically estimate camera poses (Wang et al., 2021) or finetune them during training (Lin et al., 2021; Kuang et al., 2022). The problem we are tackling in this work is fundamentally different as it requires to learn not the camera poses from multi-view observations, but a *distribution* of poses, while having access to sparse, single-view data of diverse object categories. In this respect, the work closest to ours is CAMPARI (Niemeyer & Geiger, 2021a) as it also learns a camera distribution. Beyond this the similarity ends, as our "Ball-in-Sphere" camera distribution allows for more flexibility and supports multiple object categories.

GANs with external knowledge. Several works observed improved convergence and fidelity of GANs when using existing, generic image-based models (Kumari et al., 2022; Sauer et al., 2021; 2022; Mo et al., 2020), the most notable being StyleGAN-XL (Sauer et al., 2022), which is built a pre-trained and fixed EfficientNet (Tan & Le, 2019) followed by a couple of discriminator layers to improve training on ImageNet (Deng et al., 2009). A similar technique is not suitable in our case as pre-training a generic RGB-D network on a large scale RGB-D dataset is problematic due to the lack of data. Another notable example is FreezeD Mo et al. (2020), which propose to distill discriminator features for GAN finetuning. Our work, on the other hand, rely on existing model for image classification.

3 METHOD

We chose EpiGRAF (Skorokhodov et al., 2022) as our backbone due to its fast training speed, image quality, and multi-view consistency. The key advantage of EpiGRAF (Skorokhodov et al., 2022) compared to other methods (Chan et al., 2022; Gu et al., 2022) is that it does not use a 2D upsampler. Instead, it relies on mutli-scale patch-wise training (Schwarz et al., 2020; Chai et al., 2022), rendering geometry at the target resolution.

Our generator architecture is similar to the one of EpiGRAF. Given a random latent code $\boldsymbol{z}$, our generator G produces a tri-plane representation for the scene. From this representation RGB color and density σ are obtained using a shallow 2-layer MLP decoder. Volumetric rendering (Mildenhall et al., 2020) is used to obtain pixel colors and depth values from an arbitrary viewpoint. Differently to prior works (Skorokhodov et al., 2022; Chan et al., 2022) that utilize fixed camera distribution, we sample camera from a learned camera generator C (see Sec. 3.1). We render depth, and process it by the depth adaptor, see Sec. 3.2, bridging the domains of rendered and estimated depth maps (Miangoleh et al., 2021). Our discriminator D follows the architecture of StyleGAN2 (Karras et al., 2020a), additionally accepting either adapted or estimated depth as fourth channel. To further improve image fidelity we propose a simple knowledge distillation technique, that enriches D with external knowledge obtained from ResNet (He et al., 2016), see Sec. 3.3. The overall model architecture is shown in Fig. 2.

3.1 LEARNABLE "BALL-IN-SPHERE" CAMERA DISTRIBUTION

Limitations of Existing Camera Parameterization. The camera parameterization of existing 3D generators follows an overly simplified distribution—its position is sampled on a fixed-radius sphere with fixed intrinsics, the camera always points to $(0, 0, 0)$. This parametrization has only two degrees of freedom: pitch and yaw ($\boldsymbol{\varphi}_{\text{pos}}$ in Fig. 3 (a)). This parametrization implicitly assumes that all objects could be centered, equally rotated and scaled with respect to some canonical alignment. However, 3D scenes are inherently non-alignable. For example, a scene could consist of multiple objects, such as "a cat in a kitchen". Furthermore, objects with highly articulated geometry assume significantly different shapes, rendering it impossible to establish a common camera convention for such data.

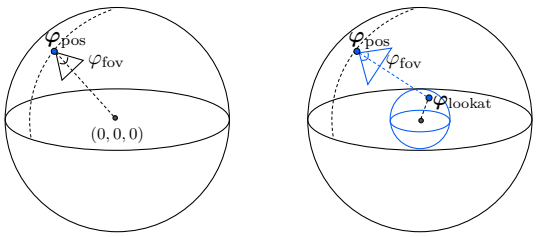

(a) Standard camera (b) Our camera

Figure 3: **Camera distribution.** The most commonly employed distribution (a) has 2 degrees of freedom and assumes the camera is on a sphere and looks at its center. (b) The proposed "Ball-in-Sphere" distribution has 4 additional degrees of freedom: field of view and the look at position. Learnable parameters are shown in blue.

Learnable "Ball-in-Sphere" Camera Distribution. We introduce a new camera parametrization which we call "Ball-in-Sphere". Contrary to the standard parametrization, it has four additional degrees of freedom: the field of view φ_{fov}, and pitch, yaw and radius of the inner sphere, specifying the look-at point within the outer sphere ($\boldsymbol{\varphi}_{\text{lookat}}$ in Fig. 3 (b)). Combining with the standard parameters on the outer sphere, our camera parametrization has six degrees of freedom $\boldsymbol{\varphi} = [\boldsymbol{\varphi}_{\text{pos}} \parallel \varphi_{\text{fov}} \parallel \boldsymbol{\varphi}_{\text{lookat}}]$, where $\parallel$ denotes concatenation. Note, that compared with the standard camera, ours allows us to learn the *scale* of objects and scenes, enabling zooming into different parts of the scene as shown in Fig. 3 (b).

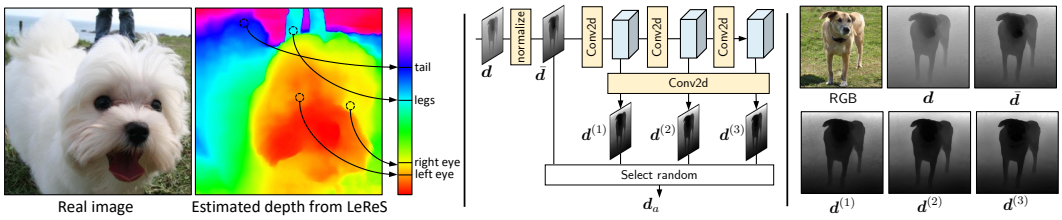

Figure 4: **Depth adapter.** Left: An example of a real image with its depth estimated by LeReS (Miangoleh et al., 2021). Note that the estimated depth has several artifacts. For example, the human legs are closer than the tail, eyes are spaced unrealistically, and far-away grass is predicted to be close. Middle: depth adapter meant to bridge the domains of predicted and NeRF-rendered depth. Right: a generated image with its adapted depth maps obtained from different layers of the adapter.

Instead of manually defining camera distributions, we learn the camera distribution during training for each dataset. In particular, we train the camera generator network C that takes camera parameters sampled from a sufficiently wide camera prior $\boldsymbol{\varphi}'$ and produces new camera parameters $\boldsymbol{\varphi}$. For a class conditional dataset, such as ImageNet where scenes have significantly different geometry, we additionally condition this network on the class label c and the latent code $\boldsymbol{z}$, i.e. $\boldsymbol{\varphi} = \mathsf{C}(\boldsymbol{\varphi}', \boldsymbol{z}, c)$. For a single category dataset we use $\boldsymbol{\varphi} = \mathsf{C}(\boldsymbol{\varphi}', \boldsymbol{z})$.

Camera Gradient Penalty. To the best of our knowledge, CAMPARI (Niemeyer & Geiger, 2021a) is the only work which also learns the camera distribution. They sample a set of camera parameters from a wide distribution and pass it to a shallow neural network, which produces a residual $\Delta\boldsymbol{\varphi} = \boldsymbol{\varphi} - \boldsymbol{\varphi}'$ for these parameters. However, since our camera model is richer, learning a residual for each parameter collapses the posterior distribution to a constant as we show in Sec. 4. A similar observation was made by Gu et al. (2022).

To prevent the camera generator from producing collapsed camera parameters, we seek a new regularization strategy. Ideally, it should prevent constant solutions, while at the same time, reducing the Lipschitz constant for C, which is shown to be crucial for stable training of generators (Odena et al., 2018). We can achieve both if we push the derivatives of the predicted camera parameters with respect to the prior camera parameters to either one or minus one, arriving at the following regularization term:

$$\mathcal{L}_{\varphi_i} = \left|\frac{\partial \varphi_i}{\partial \varphi_i'}\right| + \left|\frac{\partial \varphi_i}{\partial \varphi_i'}\right|^{-1}, \tag{1}$$

where $\varphi_i' \in \boldsymbol{\varphi}'$ is the camera sampled from the prior distribution and $\varphi_i \in \boldsymbol{\varphi}$ is produced by the camera generator. We refer to this loss as Camera Gradient Penalty. Note, that the first part of the loss prevents rapid changes in camera, facilitating stable optimization, while the second part avoids collapsed camera posteriors.

3.2 Adversarial depth supervision

To instill a 3D bias into our model, we develop a strategy of using depth maps predicted by an off-the-shelf estimator E (Miangoleh et al., 2021), for its advantages of being generic and readily applicable for many object categories.

The main idea is concatenating the depth map as a 4-th channel of the RGB as the input of the discriminator. The fake depth maps in this case is obtained with the help of neural rendering, while the real depth maps are estimated using monocular depth estimator E. However, naively utilizing the depth from E leads to training divergence. This happens because E could only produce relative depth, not metric depth. Moreover, monocular depth estimators are still not perfect, they produce noisy artifacts, ignore high-frequency details, and make prediction mistakes. Thus, we devise a mechanism that allows utilization of the imperfect depth maps. The central part of this mechanism is a learnable depth adaptor A, that should transform and augment the depth map obtained with neural rendering to look like a depth map from E.

More specifically, we first render raw depths $\boldsymbol{d}$ from NeRF via volumetric rendering:

$$\boldsymbol{d} = \int_{t_n}^{t_f} T(t)\sigma(r(t))t dt, \tag{2}$$

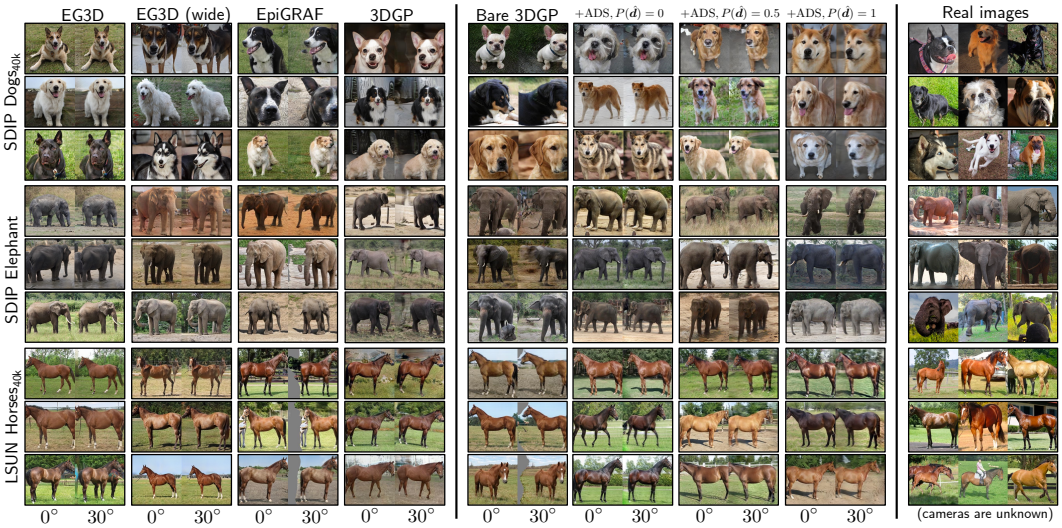

Figure 5: **Qualitative multi-view comparisons.** Left: samples from the models trained on single-category datasets with articulated geometry. Two views are shown for each sample. Middle: ablations following Tab. 1, where we change the probability of using the normalized rendered depth $P(\bar{\boldsymbol{d}})$. EG3D, EpiGRAF, and $P(\bar{\boldsymbol{d}}) = 0$ do not render realistic side views, due to the underlying flat geometry. Our full model instead generates realistic high-quality views on all the datasets. Right: randomly sampled real images. Zoom-in for greater detail.

where $t_n, t_f \in \mathbb{R}$ are near/far planes, $T(t)$ is accumulated transmittance, and $r(t)$ is a ray. Raw depth is shifted and scaled from the range of $[t_n, t_f]$ into $[-1, 1]$ to obtain normalized depth $\bar{\boldsymbol{d}}$:

$$\bar{\boldsymbol{d}} = 2 \cdot \frac{\boldsymbol{d} - (t_n + t_f + b)/2}{t_f - t_n - b}, \tag{3}$$

where $b \in [0, (t_n + t_f)/2]$ is an additional learnable shift needed to account for the empty space in the front of the camera. Note that depth values obtained using E span over the entire valid range of $[0, 2^{16}]$, thus empty space does not appear in them. Therefore, we map them into the $[-1, 1]$ range to get depth $\boldsymbol{d}_r$ for real input.

Learnable Depth Adaptor. Although $\bar{\boldsymbol{d}}$ is distributed in the same range as $\boldsymbol{d}_r$, $\bar{\boldsymbol{d}}$ is still not suitable for discriminator D. The reason is that $\bar{\boldsymbol{d}}$ forces G to learn the prediction artifacts from E and creates an additional confusion for G, given E provides only relative depth. Therefore, objects in the different images with the same metric depth may still have different relative depth. To overcome such issue, we introduce the depth adaptor A to produce an adapted depth map $\boldsymbol{d}_a = \mathsf{A}(\bar{\boldsymbol{d}}) \in \mathbb{R}^{h \times w}$, where $h \times w$ is a number of sampled pixels. The depth (fake $\boldsymbol{d}_a$ or real $\boldsymbol{d}_r$) is concatenated with the RGB input and passed to D.

The depth adaptor A models artifacts produced by E, so that the discriminator should focus only on the relevant high level geometry. However, if A is too powerful, it could fake the depth completely, and the generator G will not receive any meaningful signal. To this end, we utilize a depth adaptor that is based on 3-layer convolutional network (shown in Fig 4). After each layer, we produce a separated depth map with different levels of adaptation: $\boldsymbol{d}_a^{(1)}, \boldsymbol{d}_a^{(2)}$ and $\boldsymbol{d}_a^{(3)}$. The adapted depth $\boldsymbol{d}_a$ is randomly selected from the set of $\left\{\bar{\boldsymbol{d}}, \boldsymbol{d}_a^{(1)}, \boldsymbol{d}_a^{(2)}, \boldsymbol{d}_a^{(3)}\right\}$.

Such design can effectively learn good geometry while alleviating overfitting. For example, when the generator provides $\bar{\boldsymbol{d}}$ to the discriminator, it receives a strong signal for learning the geometry. On the other hand, if the discriminator sees highly adapted depth $\boldsymbol{d}_a^{(3)}$, it unlikely overfits to different unrelated estimation artifacts. Finally, $\boldsymbol{d}_a^{(1)}$ and $\boldsymbol{d}_a^{(2)}$ provide some trade-off between geometry learning and precise adaptation.

3.3 Knowledge distillation for Discriminator

Knowledge from pretrained classification networks was shown to improve training stability and generation quality in 2D GANs (Sauer et al., 2021; Kumari et al., 2022; Sauer et al., 2022; Casanova et al., 2021). A popular solution proposed by Sauer et al. (2021; 2022) is to use an off-the-shelf

model as a discriminator while freezing most of its weights. Unfortunately, this technique is not applicable in our scenario since we modify the architecture of the discriminator by adding an additional depth input (Sec. 3.2) and condition on the parameters of the patch similarly to EpiGRAF (Skorokhodov et al., 2022). Thus, we devise an alternative technique that can work with arbitrary architectures of the discriminator. Specifically, for each real sample, we obtain two feature representations: e from the pretrained ResNet (He et al., 2016) network and $\hat{e}$ extracted from the final representation of our discriminator D. Our loss simply pushes $\hat{e}$ to e as follows:

$$\mathcal{L}_{\text{dist}} = \|\boldsymbol{e} - \hat{\boldsymbol{e}}\|_2^2 \,. \tag{4}$$

$\mathcal{L}_{\text{dist}}$ can effectively distill knowledge from the pretrained ResNet (He et al., 2016) into our D.

3.4 Training

Our overall loss for generator G consists of two parts: adversarial loss and Camera Gradient Penalty for each camera parameter:

$$\mathcal{L}_{\mathsf{G}} = \mathcal{L}_{\text{adv}} + \sum_{\varphi_i \in \boldsymbol{\varphi}} \lambda_{\varphi_i} \mathcal{L}_{\varphi_i}, \tag{5}$$

where $\mathcal{L}_{\text{adv}}$ is the non-saturating loss (Goodfellow et al., 2014). We observe that a diverse distribution for camera origin is most important for leaning meaningful geometry, but it is also most prone to degrade to a constant solution. Therefore, we set $\lambda_{\varphi_i} = 0.3$ for $\boldsymbol{\varphi}_{\text{pos}}$, while set $\lambda_{\varphi_i} = 0.03$ for φ_{fov} and $\lambda_{\varphi_i} = 1e\text{-}3$ for $\boldsymbol{\varphi}_{\text{lookat}}$.

The loss for discriminator D, on the other hand, consists of three parts: adversarial loss, knowledge distillation, and $\mathcal{R}_1$ gradient penalty (Mescheder et al., 2018):

$$\mathcal{L}_{\mathsf{D}} = \mathcal{L}_{\text{adv}} + \lambda_{\text{dist}} \mathcal{L}_{\text{dist}} + \lambda_r \mathcal{R}_1. \tag{6}$$

We use the same optimizer and hyper-parameters as EpiGRAF. We observe that for depth adaptor, sampling adapted depth maps with equal probability is not always beneficial found that using $P(\bar{\boldsymbol{d}}) = 0.5$ leads to better geometry. For additional details, see Appx B.

4 Experimental Results

Datasets. In our experiments, we use 4 *non-aligned* datasets: SDIP Dogs (Mokady et al., 2022), SDIP Elephants (Mokady et al., 2022), LSUN Horses (Yu et al., 2015), and ImageNet (Deng et al., 2009). The first three are single-category datasets that contain objects with complex articulated geometry, making them challenging for standard 3D generators. We found it useful to remove outlier images from SDIP Dogs and LSUN Horseswith the instance selection technique (DeVries et al., 2020), reducing their size to 40K samples each. We refer to the filtered versions of these datasets as SDIP Dogs$_{40\text{k}}$, and LSUN Horses$_{40\text{k}}$, respectively. We then validate our method on ImageNet, a real-world, multi-category dataset containing $1,000$ diverse object classes, with more than $1,000$ images per category. To train all our 3D generators on ImageNet (including the baselines), we use the same filtering strategy for it, keeping 1/3 of its images (note that we always measure the metrics for it on the full ImageNet).

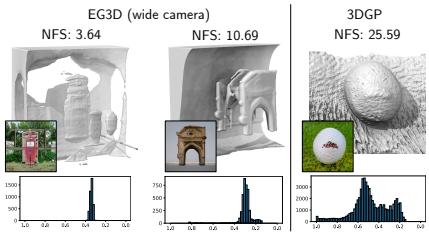

Figure 6: NFS for repetitive and diverse geometry with depth histograms.

Evaluation. We rely on FID (Heusel et al., 2017) to measure the image quality and FID$_{2\text{k}}$, which is computed on 2,048 images instead of 50k (as for FID) for efficiency. For ImageNet, we additionally compute Inception Score (IS) (Salimans et al., 2016). Note that while we train on the filtered ImageNet, we always compute the metrics on the full one. There is no established protocol to evaluate the geometry quality of 3D generators in general, but the state-of-the-art ones are tri-plane (Chan et al., 2022; Sun et al., 2022) or MPI-based Zhao et al. (2022), and we observed that it is possible to quantify their most popular geometry failure case: flatness of the shapes. For this, we propose Non-Flatness Score (NFS) which is computed as the average entropy of the normalized depth maps histograms. We depict its intuition in Fig 6 and provide the details in Appx F.

4.1 3D generation for single category datasets

We show quantitative results for single category datasets in Tab. 1a, where we compare with EG3D (Chan et al., 2022) and EpiGRAF (Skorokhodov et al., 2022). EG3D was mainly designed

Table 1: Comparisons on SDIP Dogs$_{40k}$, SDIP Elephants, and LSUN Horses$_{40k}$. (a) includes EG3D (with the standard and the wider camera range), EpiGRAF, and 3DGP. For completeness, we also provide 2D baseline StyleGAN2 (with KD). (b) includes the ablations of the proposed contributions. We report the total training cost for prior works, our model, and our proposed contributions.

(a) Comparison of EG3D, EpiGRAF and 3DGP.

Model	SDIP Dogs$_{40k}$		SDIP Elephants$_{40k}$		LSUN Horses$_{40k}$		Training cost
	FID$_{2k}$ ↓	NFS↑	FID$_{2k}$ ↓	NFS↑	FID$_{2k}$ ↓	NFS↑	(A100 days)
EG3D	16.2	11.91	4.78	2.59	3.12	13.34	3.7
+ wide camera	21.1	24.44	5.76	17.88	19.44	25.34	3.7
EpiGRAF	25.6	3.53	8.24	12.9	6.45	9.73	2.3
3DGP (ours)	8.74	34.35	5.79	32.8	4.86	30.4	2.6
StyleGAN2 (with KD)	6.24	0.0	3.94	0.0	2.57	0.0	1.5

(b) Impact of Adversarial Depth Supervision (ADS).

Bare 3DGP (w/o ADS, w/o C)	8.59	1.42	7.46	9.52	3.29	8.04	2.3
+ ADS, $P(\bar{\boldsymbol{d}}) = 0.0$	8.13	3.14	5.69	1.97	3.41	1.24	2.6
+ ADS, $P(\bar{\boldsymbol{d}}) = 0.25$	9.57	33.21	6.26	33.5	4.33	32.68	2.6
+ ADS, $P(\bar{\boldsymbol{d}}) = 0.5$	9.25	36.9	7.60	30.7	5.27	32.24	2.6
+ ADS, $P(\bar{\boldsymbol{d}}) = 1.0$	12.2	27.2	12.1	26.0	8.24	27.7	2.5

for FFHQ (Karras et al., 2019) and uses the true camera poses inferred from real images as the camera distribution for the generator. But in our scenario we do not have any knowledge about the true camera distribution, that is why to stay as close as possible to the setup EG3D was designed for, we use normal distribution for rotation and elevation angles with the same standard deviations as in FFHQ, which are equal to $\sigma_{\text{yaw}} = 0.3$ and $\sigma_{\text{pitch}} = 0.155$ respectively. Also, to make it learn better geometry, we trained it with a twice as wide camera distribution: $\sigma_{\text{yaw}} = 0.6$ and $\sigma_{\text{pitch}} = 0.3$. While it indeed help to improve the flatness of the shapes, it also considerably worsened the image quality, from 20% to to 500% as measured by FID$_{2k}$. For EpiGRAF, we could train it only with the default camera distribution of FFHQ and observed that it diverges with a wider one. Our model shows substantially lower (at least 2× better than EG3D, slightly higher than StyleGAN2) FID$_{2k}$and greater Non-Flatness Score on all the datasets. Low NFS indicates flat or repetitive geometry impairing the ability of the model to generate realistic side views. Indeed, as shown in Fig. 5 (left), both EG3D (with the default camera distribution) and EpiGRAF struggle to generate side views, while our method (3DGP) renders realistic side views on all three datasets.

Adversarial Depth Supervision. We evaluate the proposed Adversarial Depth Supervision (ADS) and the depth adaptor A. The only hyperparameter it has is the probability of using the non-adapted depth $P(\bar{\boldsymbol{d}})$ (see Sec. 3.2). We ablate this parameter in Tab. 1b. While FID scores are only slightly affected by varying $P(\bar{\boldsymbol{d}})$, we see substantial difference in the Non-Flatness Score. We first verify that NFS is the worst without ADS, indicating the lack of a 3D bias. When $P(\bar{\boldsymbol{d}}) = 0$, the discriminator is never presented with the rendered depth $\bar{\boldsymbol{d}}$, while the adaptor learns to fake the depth, leading to flat geometry. When $P(\bar{\boldsymbol{d}}) = 1$, the adaptor is never used, allowing the discriminator to easily determine the generated depth from the estimated depth, as there is a large domain gap (see Fig. 4), leading to reduced FID scores. We observe that the best results are obtained when $P(\bar{\boldsymbol{d}}) = 0.5$. The same observations can be visually verified in Fig. 5 (middle), where the Bare 3DGP, $P(\bar{\boldsymbol{d}}) = 0$ and $P(\bar{\boldsymbol{d}}) = 1$ are unable to render side views, while the model trained with $P(\bar{\boldsymbol{d}}) = 0.5$ features the overall best geometry and side views.

Knowledge Distillation. Here we further provide insights for discriminator D using our knowledge distillation strategy (see Sec. 3.3). We find that knowledge distillation provides an additional stability for adversarial training, along with the significant improvement in FID, which can be observed by comparing results of EpiGRAF in Tab. 1a and Bare 3DGP Tab. 1a. Additionally, we compare different knowledge distillation strategies in Appx C. However, as noted by Parmar et al. (2021), strategies which utilize additional classification networks, may provide a significant boost to FID, without corresponding improvement in visual quality.

"Ball-in-Sphere" Camera Distribution. We investigate the impact of our learnable camera distribution model. Note, that if the camera distribution is sampled from a narrow distribution, the discriminator will not be presented with side views during training, leading to flat geometry. On the other hand, if the camera distribution is very wide and covers more views than it is available in the dataset, the discriminator will easily win, leading to collapsed training. The proposed "Ball-in-

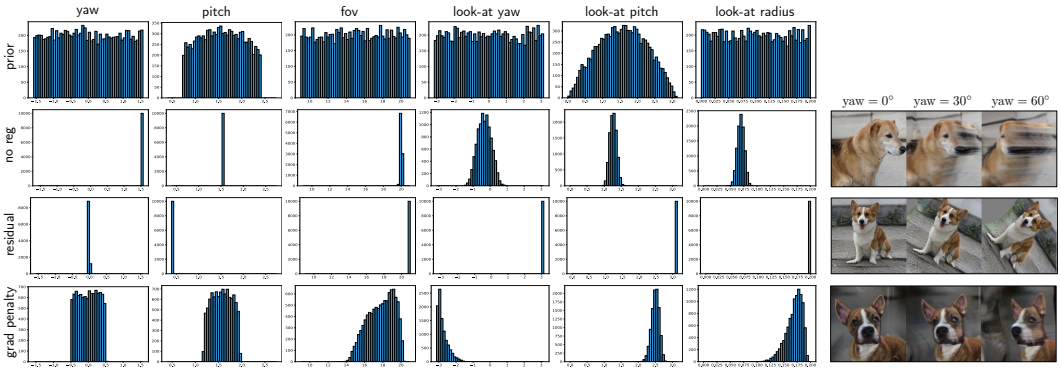

Figure 7: **Comparisons of regularization strategies.** Left: selected generic and wide prior for each of the 6 DoFs of our camera model (top row). Without any regularization (*no reg*) or with the residual-based model (*residual*) as in Niemeyer & Geiger (2021a), the camera generator collapses to highly concentrated distributions. In contrast, the proposed regularization (*grad penalty*) leads to a wider posterior. Right: random samples for each strategy. For (*no reg*) and (*residual*), no meaningful geometry is learned; only our method (*grad penalty*) leads to good geometry.

Sphere" model is able to strike a balance between the two extremes. In Fig. 7, we analyze different strategies for learning the camera distribution: sampling the camera φ from the prior φ' without learning, predicting residuals $\varphi - \varphi'$, and using our proposed camera generator C with Camera Gradient Penalty. For the first two cases, the learned distributions are highly concentrated around constant values. Not surprisingly, no meaningful side views can be generated as geometry becomes flat. In contrast, our Camera Gradient Penalty provides sufficient regularization to C, enabling it to converge to a sufficiently wide posterior, resulting into valid geometry and realistic side views.

4.2 3D SYNTHESIS ON IMAGENET

ImageNet (Deng et al., 2009) is significantly more difficult than single category datasets. Following prior work, we use class labels c as an additional guidance for the generator (Sauer et al., 2022; Brock et al., 2018). We modify the baselines accordingly. The quantitative results of ImageNet experiments are presented in Tab. 2. For the reference, we also report results of 2D generators, such as GAN-based methods of BigGAN (Brock et al., 2018) and StyleGAN-XL (Sauer et al., 2022) and a diffusion-based approach of ADM (Dhariwal & Nichol, 2021). As expected, 2D generators show superior FID and IS, compared to the 3D counterparts as the latter use their capacity to learn geometric details. As discussed before, the "wideness" of the camera distribution is an important factor to consider, hence, for fair comparisons, we report the standard camera ranges as well as wide ones. Despite our best efforts to find a reasonable camera distribution, both EG3D and EpiGRAF produce flat or repetitive geometry, confirming similar experiments on single category articulated datasets. 3DGP, on the other hand, produces geometry with rich details (see Fig. 1).

Table 2: Comparison on ImageNet 256^2. First three rows: 2D methods. Next four rows: 3D baselines with narrow and wide camera distribution. Last row: scores for our full method. Last column: training cost in A100 days.

Method	FID ↓	IS ↑	NFS↑	A100 days ↓
BigGAN	8.7	142.3	0.0	60
StyleGAN-XL	2.30	265.1	0.0	163+
ADM	4.59	186.7	0.0	458
EG3D	26.7	61.4	3.70	18.7
+ wide camera	25.6	57.3	9.83	18.7
EpiGRAF	47.56	26.68	3.93	15.9
3DGP (ours)	26.47	73.09	27.52	28.2

5 LIMITATIONS AND CONCLUSION

In this work, we present the first 3D GAN that is scalable to in-the-wild multi-category data such as ImageNet. We demonstrate how to utilize generic priors in the form of monocular depth and latent feature representation to improve the visual quality and learn geometry. Moreover, we propose the new "Ball-in-Sphere" camera parameterization and a novel regularization term that enable learning meaningful camera distribution. On the other hand, our work still has several limitations, such as the gluing of foreground and background, lower visual quality compared to 2D generators, and no reliable quantitative measure of generated geometry. More discussion are provided in Appx A.

Besides further improving the discussed limitations, other future directions could be scaling 3D GANs to even larger datasets, which may further improve the discovery of meaningful geometry. Such scaling could be achieved by conditioning 3D generation on natural language.

REFERENCES

Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Hao Tang, Gordon Wetzstein, Leonidas Guibas, Luc Van Gool, and Radu Timofte. 3d-aware video generation. *arXiv preprint arXiv:2206.14797*, 2022.

Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.

Arantxa Casanova, Marlène Careil, Jakob Verbeek, Michal Drozdzal, and Adriana Romero-Soriano. Instance-conditioned gan. In *Proceedings of the Neural Information Processing Systems Conference*, 2021.

Lucy Chai, Michael Gharbi, Eli Shechtman, Phillip Isola, and Richard Zhang. Any-resolution training for high-resolution image synthesis. *arXiv preprint arXiv:2204.07156*, 2022.

Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 2009.

Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. Gram: Generative radiance manifolds for 3d-aware image generation. In *IEEE Computer Vision and Pattern Recognition*, 2022.

Terrance DeVries, Michal Drozdzal, and Graham W Taylor. Instance selection for gans. *Advances in Neural Information Processing Systems*, 2020.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.

Ainaz Eftekhar, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *ICCV*, 2021.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the Neural Information Processing Systems Conference*, 2014.

Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Machine Learning*, 2022.

Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

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 *Proceedings of the Neural Information Processing Systems Conference*, 2017.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676*, 2020a.

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 Conference on Computer Vision and Pattern Recognition*, 2020b.

Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021.

Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *CVPR*, 2022.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Zhengfei Kuang, Kyle Olszewski, Menglei Chai, Zeng Huang, Panos Achlioptas, and Sergey Tulyakov. Neroic: Neural rendering of objects from online image collections. In *Special Interest Group on Computer Graphics and Interactive Techniques*, 2022.

Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for gan training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021.

Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, 2018.

S Mahdi H Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yagiz Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9685–9694, 2021.

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. In *European conference on computer vision*, 2020.

Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze the discriminator: a simple baseline for fine-tuning gans. *arXiv preprint arXiv:2002.10964*, 2020.

Ron Mokady, Michal Yarom, Omer Tov, Oran Lang, Daniel Cohen-Or, Tali Dekel, Michal Irani, and Inbar Mosseri. Self-distilled stylegan: Towards generation from internet photos. *arXiv preprint arXiv:2202.12211*, 2022.

Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019.

Thu H Nguyen-Phuoc, Christian Richardt, Long Mai, Yongliang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. In *Proceedings of the Neural Information Processing Systems Conference*, 2020.

Michael Niemeyer and Andreas Geiger. Campari: Camera-aware decomposed generative neural radiance fields. In *2021 International Conference on 3D Vision (3DV)*, 2021a.

Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021b.

Augustus Odena, Jacob Buckman, Catherine Olsson, Tom Brown, Christopher Olah, Colin Raffel, and Ian Goodfellow. Is generator conditioning causally related to gan performance? In *International conference on machine learning*, 2018.

Kyle Olszewski, Sergey Tulyakov, Oliver Woodford, Hao Li, and Linjie Luo. Transformable bottleneck networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.

Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-Resolution 3D-Consistent Image and Geometry Generation. *arXiv preprint arXiv:2112.11427*, 2021.

Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On buggy resizing libraries and surprising subtleties in fid calculation. *arXiv preprint arXiv:2104.11222*, 2021.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Proceedings of the Neural Information Processing Systems Conference*, 2016.

Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. *Advances in Neural Information Processing Systems*, 34:17480–17492, 2021.

Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022.

Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision*, 2016.

Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Proceedings of the Neural Information Processing Systems Conference*, 2020.

Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. *ARXIV*, 2022.

Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. In *Proceedings of the Neural Information Processing Systems Conference*, 2022.

Jingxiang Sun, Xuan Wang, Yichun Shi, Lizhen Wang, Jue Wang, and Yebin Liu. Ide-3d: Interactive disentangled editing for high-resolution 3d-aware portrait synthesis. *arXiv preprint arXiv:2205.15517*, 2022.

Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 2019.

Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.

Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. NeRF$--$: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.

Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and Bolei Zhou. 3d-aware image synthesis via learning structural and textural representations. *arXiv preprint arXiv:2112.10759*, 2021.

Yang Xue, Yuheng Li, Krishna Kumar Singh, and Yong Jae Lee. Giraffe hd: A high-resolution 3d-aware generative model. *arXiv preprint arXiv:2203.14954*, 2022.

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. *arXiv preprint arXiv:1506.03365*, 2015.

Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *arXiv preprint arXiv:2206.00665*, 2022.

Xuanmeng Zhang, Zhedong Zheng, Daiheng Gao, Bang Zhang, Pan Pan, and Yi Yang. Multi-view consistent generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G. Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. In *Proc. ECCV*, 2022.

Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. Cips-3d: A 3d-aware generator of gans based on conditionally-independent pixel synthesis. *arXiv preprint arXiv:2110.09788*, 2021.

Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: Image generation with disentangled 3d representations. In *Proceedings of the Neural Information Processing Systems Conference*, 2018.

A LIMITATIONS AND FUTURE WORKS

Lower texture quality compared to 2D generators. Despite providing a more reasonable representation of the underlining scene, 3D generators still have a lower visual quality than 2D generators. Closing this gap is essential for a wide adaptation of 3D generators.

Background sticking. One common 3D artifact of 3DGP is gluing of the foreground and the background. In other words our model predicts a single shape for both and there is no clear separation between the two. One potential cause of this artifact is the dataset bias, where most of the photos are frontal, thus it is not beneficial for the model to explore backward views. Another reason is the bias of tri-planes toward flat geometry. However, all our attempts to replace tri-planes with an MLP-based NeRF led to much worse results (see Appx D). Inventing a different efficient parameterization may be an important direction.

No reliable quantitative measure of generated geometry. In this work we introduce Non-Flatness Score as a proxy metric for evaluating the quality of the underlying geometry. However it can capture only a single failure case, specific for generators based on tri-planes: flatness of geometry. Devising a reasonable metric applicable to a variety of scenarios could significantly speed up the progress in this area.

B IMPLEMENTATION DETAILS

This section provides additional architectural and training details. Also note that we will release the source code. Our generator G consists of M, S and T. Mapping network M takes noise $\boldsymbol{z} \in \mathbb{R}^{512}$ and class label $c \in 0, ..., K-1$, where K is the number of classes, and produces the style code $\boldsymbol{w} \in \mathbb{R}^{512}$. Similar to StyleGAN2-ADA Karras et al. (2020a) and EpiGRAF Skorokhodov et al. (2022), M is 2-layer MLP network with LeakyReLU activations and 512 neurons in each layer. Synthesis network is a decoder network same as StyleGAN2 Karras et al. (2020b), except that it produce tri-plane feature features $\boldsymbol{p} = (\boldsymbol{p}^{xy}, \boldsymbol{p}^{yz}, \boldsymbol{p}^{xz}) \in \mathbb{R}^{3\times(512\times512\times32)}$. A feature vector $\boldsymbol{f}_{xyz} \in \mathbb{R}^{32}$ located $(x, y, z) \in \mathbb{R}^3$ is computed by projecting the coordinate back to the tri-plane representation, followed by bi-linearly interpolating the nearby features and averaging the features from different planes. Following EpiGRAF (Skorokhodov et al., 2022), tri-plane decoder is a two-layer MLP with Leaky-ReLU activations and 64 neurons in the hidden layer, that takes a tri-plane feature $\boldsymbol{f}_{xyz}$ as input and produced the color and density (RGB, σ) in that point. We use same volume rendering procedure as EpiGRAF (Skorokhodov et al., 2022).

Table 3: Names of the networks used in our method:

G	Generator
M	Mapping network
S	Synthesis network
T	Tri-plane decoder
C	Camera generator
A	Depth adaptor
D	Discriminator
E	Depth Estimator
F	Feature Extractor

Camera generator consists of linear layers with SoftPlus activations, it architecture is depicted in Fig. 8. We found it crucial to use SoftPlus activation and not LeakyReLU, since optimization of Camera Gradient Penalty for non-smooth functions is unstable for small learning rates (smaller 0.02). For depth adaptor, we use a three layer convolutional neural network with 5×5 kernel sizes with LeakyReLU activations and 64 filters in each layer. Additionally we use one shared convolutional layer that converts $64 \times h \times w$ features to the depth maps. We use the same architecture for D as EpiGRAF (Skorokhodov et al., 2022), but additionally concatenate a 1-channel depth to the 3-channel RGB input. Finally E and F is pretrainined LeReS (Miangoleh et al., 2021) and ResNet50 (He et al., 2016) networks without any modifications. We used the `timm` library to extract the features for real images.

Similarly to EpiGRAF (Skorokhodov et al., 2022) we set $\lambda_r = 0.1$ and $\lambda_{\text{dist}} = 1$. We train all the models with Adam optimizer (Kingma & Ba, 2014) using the learning rate of $2e$-3 and $\beta_1 = 0.0, \beta_2 = 0.99$. Following EpiGRAF (Skorokhodov et al., 2022), our model uses patch-wise training with 64×64-resolution patches and uses their proposed β scale sampling strategy without any modifications. We use the batch size of 64 in all the experiments, since in early experiments we didn't find any improvements from using a large batch size neither for our model nor for StyleGAN2, as observed by Brock et al. (2018) and Sauer et al. (2022).

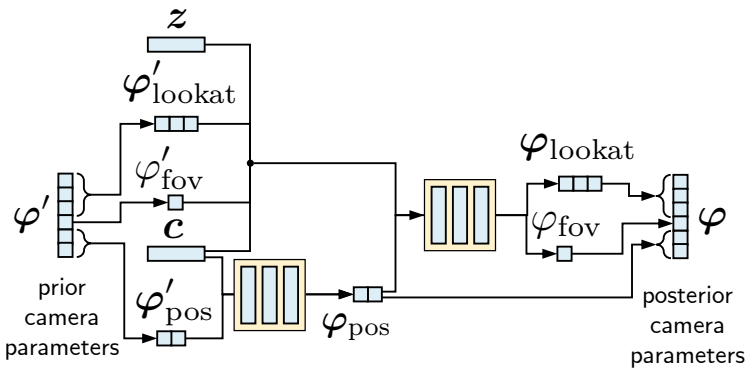

Figure 8: **Camera generator** C**.** We condition C on class labels $\boldsymbol{c}$ when generating the camera position $\hat{\boldsymbol{\varphi}}$ since it might be different for different classes. And we condition it on $\boldsymbol{z}$ when generating the look-at position and field-of-view since it might depend on the object shape (e.g., there is a higher probability to synthesize a close-up view of a dog's snout rather than its tail). Each MLP consists on 3 layers with Softplus non-linearities.

As being said in Section 4, we use the instance selection technique by DeVries et al. (2020) to remove image outliers from the datasets since they might negatively affect the geometry. This procedure works by first extracting a 2,048-dimensional feature vector for each image, then fitting a multivariate gaussian distribution on the obtained dataset and removing the images, which features have low probability density. For SDIP Dogs and LSUN Horses, we fit a multi-variate gaussian distribution for the whole dataset. For ImageNet, we fit a separate model for each class with additional diagonal regularization for covariance, which is needed due to its singularity: feature vector has more dimensions than the number of images in a class. We refer to (DeVries et al., 2020) for additional details.

C KNOWLEDGE DISTILLATION COMPARISON

In this section we compare our knowledge distillation strategy with projected D (Sauer et al., 2021; 2022), another popular technique of utilizing existing classification models for GAN training. Note that projected D relies on pretrained EfficientNet (Tan & Le, 2019), it freezes most of the weights and adds some random convolutions layers that project intermediate outputs of EfficientNet to the features, that will be later processed by discriminator. Note that this strategy relies on specific architecture of Discriminator D, thus in order to adapt it for 3DGP we have to rely on two discriminators. First one is our disciminator D with additional depth input and conditioned on patch parameters and second one is projected D. In Tab 4 we show comparison between projected D and our knowledge distillation strategy on ImageNet (Deng et al., 2009) and SDIP $\text{Dogs}_{40\text{k}}$ Mokady et al. (2022) datasets. First we compare these strategies for 2D generator architectures: StyleGAN3-t (large) (Karras et al., 2021) and StyleGAN2 (Karras et al., 2020a). It can be observed that projected D is not generic and it results heavily varies depending on generator architecture. Moreover training cost of projective D is higher. Finally we test this strategy for our 3D generator, and again we can observe that this technique lead to inferior results in both visual quality and training cost.

D FAILED EXPERIMENTS

In this section we describe several ideas that have been tried in this project, which however did not work for us.

Gradient Penalty for Depth Adaptor. To prevent depth adaptor A from completely faking the depth and forcing it to still provide a useful signal to the generator, one could employ a gradient penalty that will force gradient of A close to one. We tried it in two setups: window-to-pixel gradient and pixel-to-pixel gradient. Even with the large weight for this loss and shallow adaptor with $k = 3$ kernel size, it didn't enforce good geometry

Table 4: Comparing our knowledge transfer strategy with the one from StyleGAN-XL.

Method	ImageNet 128^2 @ 10M		SDIP Dogs$_{40k}$ 256^2 @ 5M		Restricts D's
	FID↓	Training cost ↓	FID ↓	Training cost ↓	architecture?
StyleGAN3-t (large)	28.1	11.9	6.42	7.3	No
- with Projected D	22.8	11.6	22.6	8.1	Yes
- with Knowledge Distillation	16.3	11.9	2.47	7.3	No
StyleGAN2	33.7	1.81	9.79	1.3	No
- with Projected D	160.5	3.83	4.10	2.1	Yes
- with Knowledge Distillation	20.75	1.83	2.08	1.4	No
3DGP	53.6	6.33	21.3	2.6	No
- with Projected D	105.9	8.0	10.5	4.2	Yes
- with Knowledge Distillation	27.8	6.83	4.51	2.6	No

3DGP with Projected D. Before devoting to knowledge distillation, we spend significant resources trying to incorporate projected D in our setup. For double discriminators setting, see Appx C, we tested: different weighting strategies, less frequent updates for projected D, enabling projected D after some geometry was already learned. We also tested single projected D, where we learned additional depth encoder that can consume our depth inputs. In all these experiments learned geometry was significantly inferior to the setting without projected D. Most of the time geometry becomes completely flat.

MLPs instead of tri-planes. We noticed that tri-planes are extremely biased towards flat generation. Therefore one natural idea would be to utilize different representations, one possible candidate is MLPs Chan et al. (2021). However our experiments with MLPs always lead to significantly inferior visual quality. Our hypothesis is that they just don't have enough capacity to model large scale datasets.

Few-shot NeRF regularization. We also try to improve learned geometry by incorporating ray entropy minimization loss from InfoNeRF Kim et al. (2022) with different weights. But the model always diverges. We hypothesize that this is because it is significantly more complicated to find proper geometry, if the model is stuck in local minimum where the object densities are very sharp.

Normal supervision. Another prominent idea was utilizing normal supervision, since generic networks for normals Eftekhar et al. (2021) are also widely available and provide good results on the arbitrary data. Since computing normals requires gradient of density σ with respect to (x, y, z) and pytorch does not provide second order derivatives for F.grid_sample [1], we develop a custom cuda kernel that computes second derivative for F.grid_sample. Unfortunately we found that normals supervision is inferior to depth and it is much easier for the generator G to fake normals.

E Additional samples

Since it is much easier to visualize the samples from NeRF-based generators as videos rather than RGB images, we provide all the additional visualizations on https://u2wjb9xxz9q.github.io/.

F Non-Flatness Score details

To detect and quantify the flatness of 3D generators, we propose a Non-Flatness Score (NFS). To compute NFS, we sample N latent codes with their corresponding radiance fields. For each radiance field, we perform integration following Eq. 2 to obtain its depth, however, we first set the 50% of lowest density to zero. This is necessary to cull spurious density artifacts.

We then normalize each depth map according to the corresponding near and far planes and compute a histogram with B bins, showing the distribution of the depth values. To analyze how much the depth values are concentrated versus spread across the volume, we compute its entropy for each

[1] https://github.com/pytorch/pytorch/issues/34704

distribution. Averaging over N depth maps gives the sought score:

$$\text{NFS} = \frac{1}{N}\sum_{i=1}^{N}\exp\left[-\frac{1}{h\cdot w}\sum_{j=1}^{B} b(\boldsymbol{d}^{(i)})_j\right], \tag{7}$$

where $b(\boldsymbol{d}^{(i)})_j$ is the normalized number of depth values in the j-th bin of the i-th depth map $\boldsymbol{d}^{(i)}$, with $N = 256$ and $B = 64$. NFS does not directly evaluate the quality of the geometry. Instead, it helps to detect and quantify how flat the generated geometry is. In Fig. 6, we show examples of repetitive geometry generated by EG3D and a more diverse generation produced by 3DGP along with their depth histograms and NFS.