

XDGAN: Multi-Modal 3D Shape Generation in 2D Space

Hassan Abu Alhajja¹
habualhajja@nvidia.com

Alara Dirik²
alara@pch-innovations.com

André Knörig²
andre@pch-innovations.com

Sanja Fidler^{1,3}
sfidler@nvidia.com

Maria Shugrina¹
mshugrina@nvidia.com

¹ NVIDIA

² PCH Innovations GmbH

³ University of Toronto

Abstract

Generative models for 2D images has recently seen tremendous progress in quality, resolution and speed as a result of the efficiency of 2D convolutional architectures. However it is difficult to extend this progress into the 3D domain since most current 3D representations rely on custom network components. This paper addresses a central question: *Is it possible to directly leverage 2D image generative models to generate 3D shapes instead?* To answer this, we propose XDGAN, an effective and fast method for applying 2D image GAN architectures to the generation of 3D object geometry combined with additional surface attributes, like color textures and normals. Specifically, we propose a novel method to convert 3D shapes into compact 1-channel geometry images and leverage StyleGAN3 and image-to-image translation networks to generate 3D objects in 2D space. The generated geometry images are quick to convert to 3D meshes, enabling real-time 3D object synthesis, visualization and interactive editing. Moreover, the use of standard 2D architectures can help bring more 2D advances into the 3D realm. We show both quantitatively and qualitatively that our method is highly effective at various tasks such as 3D shape generation, single view reconstruction and shape manipulation, while being significantly faster and more flexible compared to recent 3D generative models.

1 Introduction

Generative Adversarial Networks [1] have achieved remarkable progress in generating high-resolution realistic images, typically using convolutional architectures such as [2, 3]. Extending these advances to the 3D domain remains a challenge and an active area of research, with newest methods introducing implicit or explicit 3D awareness into the GAN generative process [4, 5, 6, 7, 8]. In this work, we show that it is possible to directly leverage a 2D GAN architecture designed for images to generate high-quality 3D shapes.

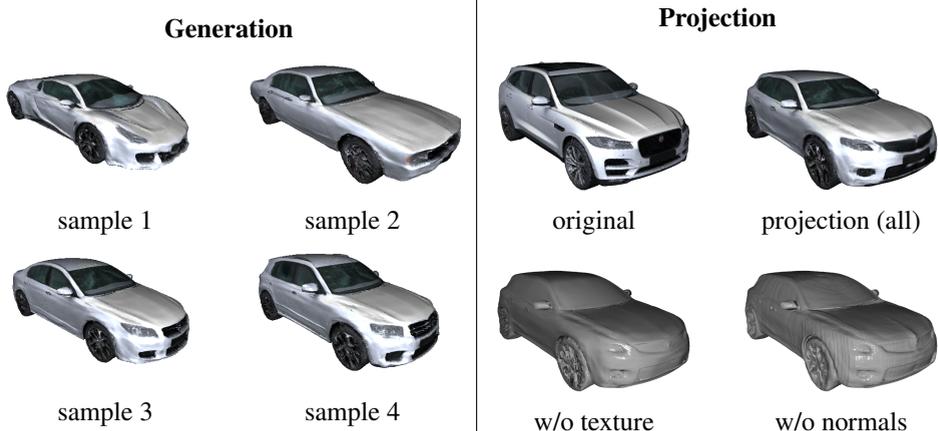


Figure 1: XDGAN allows generation of high-resolution textured 3D meshes, supports projection of 3D models into the latent space, where generation of multiple textures and semantic editing are possible. Results from the model trained with HQCars dataset.

The key to our approach is parameterization of 3D shapes as 2D planar *geometry images* [18], which we use as the training dataset for an image-based generator.

Perhaps one of the greatest challenges to the development of generative models for 3D content is converging on the right representation. Direct extension of 2D pixel grids to 3D voxel grids [50] suffers from high memory demands of 3D convolutions, limiting resolution. Point cloud samples of surface geometry, while popular for generative tasks [0, 6, 28], are limited in their ability to model sharp features or high-resolution textures. Recently implicit 3D representations, such as NeRF [33] and DeepSDF [35], as well as hybrids, have shown great promise for generative tasks [0, 11, 52]. However, these approaches typically are too slow for interactive applications and their output does not interface with existing 3D tools, calling for costly or lossy conversion of models before use. To this date, 3D meshes, augmented with normals and textures, remain the most widely adopted 3D representation in 3D software, movies and games. Our method uses 2D generative architectures to produce fixed-topology textured meshes with a single forward pass, and is thus immediately practical.

We present XDGAN (X-Dimensional GAN), a method for using standard fast 2D GAN architectures for generating high-resolution 3D meshes with additional surface properties like textures and normals. XDGAN is trained on a collection of 3D shapes, which are first unwrapped into planar geometry images. Each pixel of a geometry image represents a vertex position of a fixed topology mesh, and is thus a direct representation of surface geometry. To our knowledge, ours is the first work to show that 2D convolutional generators can produce high-fidelity 3D shapes by operating on planar parameterization of 3D geometry, an idea that can help bring more advances in 2D architectures to the 3D realm. Specifically, we experiment with StyleGAN [23, 24] modified to train on higher-precision geometry images that represent 3D shapes. In order to augment generated geometry with additional per-vertex properties, we show that it is possible to directly apply an image-to-image translation framework [67] to predict these properties for each geometry image. We demonstrate that our approach can generate diverse, high-quality textured 3D meshes, beating existing generators in quality of surface detail, presence of color or practical usefulness of output representation. Further, we demonstrate that powerful properties of the StyleGAN latent space can also be exploited for 3D generation, reconstruction, as well as supervised and unsupervised

	Representations						
	Voxels	Points	Implicit	Mesh			
	[47, 54]	[50]	[11, 19, 35]	[45]	[48]	[16]	Ours
Real-time generation	✓	✓	✗	✓	✓	✓	✓
Real-time rendering	✓	✓	✗	✓	✓	✓	✓
High-quality surface	✗	✗	✓	✗	✗	✗	✓
Texture	✗	✗	✗	✗	✗	✓	✓
Variable topology	✓	✓	✓	✗	✗	✓	✗

Table 1: Comparison of 3D generative methods by the representation used. We include quantitative and qualitative comparisons with the starred methods in our experiments.

manipulation, just as it has been shown for 2D images.

2 Related Work

Generative Adversarial Networks (GANs) have become a popular technique for image generation since their introduction by Goodfellow et al. [15]. Striking progress in the quality and resolution of GAN-generated images has been achieved in recent years [6, 22, 23, 24], including in conditional settings [66]. While some degree of view control of such models has been exploited for downstream 3D tasks [66], these architectures have remained primarily confined to generation in the 2D domain. Most recently, approaches that leverage an intermediate 3D representation to improve 3D consistency and view control of GANs have also been proposed [4, 7, 8, 17, 34]. Our method is orthogonal to this line of work as we show that it is possible to leverage unmodified 2D GAN architectures, to learn *3D geometry* directly. Beyond unconditional generation, latent spaces of large-scale GANs have been successfully used for both coarse and fine-grained manipulations of images. While some techniques focus on editing real images [11, 9, 26, 50, 40], others devise meaningful exploration of the GAN latent space. Supervised approaches typically use pre-trained attribute classifiers to optimize edit directions [24, 43, 44]. Other works show that it is possible to find meaningful directions in latent space in an unsupervised way [20, 21, 49]. In our work, we use a modified version of InterFaceGAN [44] to demonstrate how latent space manipulation techniques, originally proposed for 2D image manipulation, can be leveraged for 3D generation as well.

While there have been remarkable advances in 3D shape generation in the recent years, deployment of 3D generative models poses challenges due to generation quality and speed, as well as compatibility of the output format with downstream tasks (see Tb.1 for summary).

The most direct 3D analog of the 2D pixel array is the 3D voxel grid, and early generative approaches adapted convolutional architectures to generate objects by operating on this 3D grid [50, 51]. Due to fixed resolution of the voxel grid and high memory demands of 3D convolutions these approaches are limited to low quality outputs. Later methods mitigate high memory requirements, for example, by using octrees to represent the 3D space more efficiently [47] or through local generation [54], but are still constrained in their ability to represent high-resolution smoothly varying geometry and texture (Tb.1, A).

Due to their simplicity, point clouds, or unordered (x, y, z) samples of surface geometry, are a popular representation for 3D generation [2, 6, 28], reconstruction [29] and segmentation [13, 58, 59]. Despite continued progress in generative modeling of point clouds [53, 57], they cannot overcome the inherent limitations of the representation, which cannot model sharp surface details, or be rendered as high-fidelity textured meshes. We quanti-

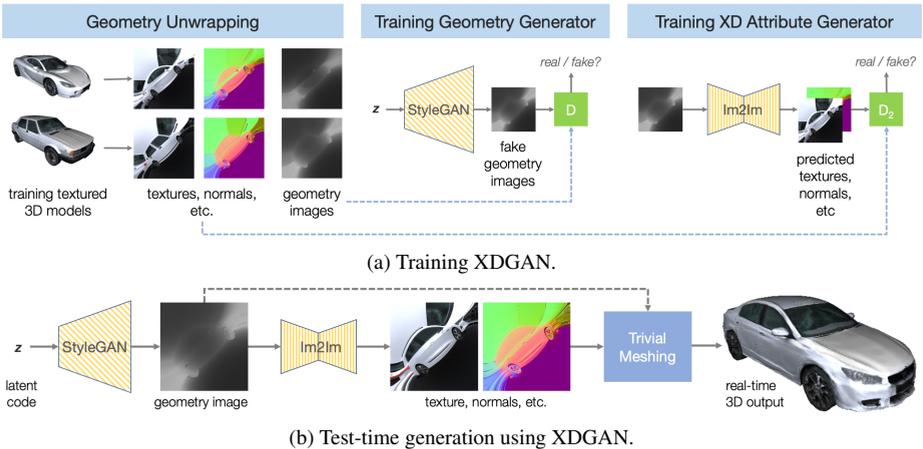


Figure 2: Overview of XDGAN: To train our model (a), we first convert a training dataset of textured 3D meshes into 2D geometry images with corresponding textures, normals or any other surface attributes. Next, we train a GAN model on geometry images to generate geometry, and an image-to-image translation network to generate X-channel attribute images for an input geometry image. At test-time (b), feed-forward evaluation of these networks followed by a trivial meshing step produces textured 3D output meshes in real-time.

tatively compare our method against the recently published Diffusion Point Cloud (DPC) [41] (See also Tb.1, B). DPC proposes a probabilistic diffusion model and uses point cloud representation of 3D shapes for shape generation, auto-encoding and shape completion.

Recently there has also been a surge in using implicit 3D representations, such as learned signed distance functions (SDFs) [49, 65] or occupancy [41, 62]. Like point clouds and voxel grids, these representations allow modeling varying object topology, but can also represent high-resolution surface detail. Most related to ours among implicit methods is IM-GAN [41], which proposes to use an implicit decoder IM-NET in conjunction with features learned by a latent-GAN model [2] to yield a general 3D generative model. Because the outputs of this and related methods are implicit in the network weights, high-resolution results typically cannot be visualized or exported in real-time, requiring several seconds on state-of-the-art GPU, and precluding interactive applications (See Tb.1, C). In addition, a number of hybrid approaches are emerging. For example, ShapeFormer [52] employs transformers in the space of discrete implicit shape elements to support shape completion, and DMTet [42] produces detailed tetrahedral meshes using sparse SDF samples. While DMTet generates a mesh-based representation quickly, like our method, it is only designed to work when conditioned on a rough input shape, not as a general generator.

3 Methodology

Our approach leverages the power of 2D convolutional generator networks in order to generate 3D geometry. Given a training set of (optionally textured) 3D objects, we first convert them into geometry images and corresponding normal maps, textures or other attribute maps (§3.1). The one-channel geometry images represent vertex positions of a fixed topology mesh, and are used to train an image-based GAN model to generate such plausible geometry images (§3.2). In order to augment output with per-vertex attributes like textures and normals, we similarly train an image-to-image translation network on these aligned attribute

maps (§3.2). After training (Fig.2a), these two image-based generators produce detailed textured 3D meshes in real-time (§2b). Fast generation makes interactive exploration of the GAN latent space especially attractive. Although our method generates 3D output, we reap the benefits of latent space exploration techniques developed for the 2D domain to explore and edit 3D shapes (§3.3).

3.1 Conversion to Geometry Images

Geometry images [18] are a representation that captures 3D surface as a 2D array of (x, y, z) values via an implicit surface parameterization, which is a bi-directional mapping from an object’s surface to a 2D plane. Such parameterizations are often used in computer graphics to map 2D textures and other attributes onto a 3D surface mesh. Once the parameterization is determined, the plane is sampled to create a 3-channel $n \times n$ image where each pixel represents an (x, y, z) vertex location of a new geometry image mesh, which is independent from the original object’s mesh topology or representation. The order and connectivity of each vertex within this geometry image mesh is encoded by its pixel index, where faces connect neighboring pixel-vertices using a regular structure. Geometry images can contain additional channels encoding surface properties like color texture, normals, object segmentation and similar. Because the same surface parameterization is used when sampling these additional properties, all channels of the geometry image are aligned, with each pixel representing location and other properties of the same 3D vertex. In our work, we generate normals and texture geometry images and propose applying 2D convolutional generators to these aligned X-channel images.

In order to map a closed 3D surface onto a 2D plane, we need to determine seams, make cuts across the surface to conform to a chosen topology (e.g. disk, sphere, cylinder), and unwrap it onto a flat square surface. Creating appropriate parameterizations that are well-suited to individual 3D object topology and shape is an active research area [18, 27, 57, 44, 46]. Since our goal is to train a generative model, we want to reduce the representation variability as much as possible to allow the model to learn the true data distribution instead of the variability in parameterizations.

For this reason, we choose to work with a specific family of spherical parameterizations, and a fixed seam structure to convert 3D shapes to geometry images and create consistent input for training the 2D generator. We start with two realistic assumptions about the geometry of our shapes. First, we assume that most of our 3D shapes are symmetric, which means we only need to represent half of the shape in the geometry image. Cutting the closed surface of the shape in half also defines an edge of the surface that is needed for the parameterization. Second, we assume our shapes are mostly genus-zero or close to it. This means that the surface can be mapped onto a surface of a sphere. While the spherical projection method imposes limitations on output geometry, our aim is to explore generation in the most consistent setting first.

Given above assumptions, we define a three step process to obtain a geometry image from a collection of 3D shapes (Refer to Fig.3). We first scale and align all input shapes

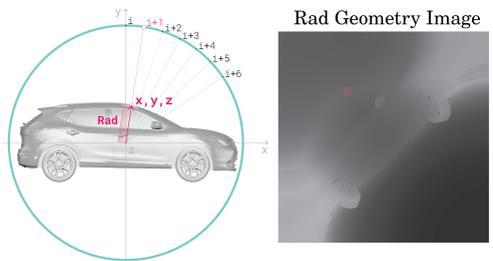


Figure 3: Generation of geometry images from 3D data. See §3.1.

to ensure consistent orientation and scale and place them within the projection sphere. We then project the surface of the shape onto a spherical domain using the normal direction of each point on the sphere. This creates a unique mapping between each point on the sphere’s surface and the shape’s surface. Finally, we cut the sphere in half along the symmetry axis of the shape and map the surface of the hemisphere onto a 2D plane using a fixed, predefined area-preserving parameterization. Because projection from the sphere onto a 3D surface has only one degree of freedom (distance along the normal direction of the sphere, see Fig.3), we represent the 3D location of a surface point as a single value measuring the distance (which we call “Rad”) from the origin along the projection ray, effectively reducing the degrees of freedom of the representation from 3 channels per pixel to only 1. To compute the original (x, y, z) coordinates of all the vertices, one only needs to multiply the Rad image by the image representing the normals of the projection sphere, fixed for the whole dataset. In addition to positions, we produce aligned images of normals and color textures.

3.2 Generating Geometry Images and Other Attributes

Once our training dataset is converted to Rad geometry images (§3.1), we train standard 2D GAN architectures to produce "realistic" Rad images. The only modification required is to adapt them to high-precision geometry images instead of the traditional 8-bit RGB images. Specifically, we experiment using StyleGAN3 [25] for geometry generation, but nothing suggests that our approach will not generalize to other GAN architectures (see Supplemental for experiments with StyleGAN2 [24]). We found training over single-channel Rad images to differ little in its stability from standard image-based GAN training.

During the generation of Rad geometry images (§3.1), we also produce aligned normal and texture maps using the same surface parameterization. Normal maps describe the local curvature of a 3D surface and are often used in computer graphics to add fine details to a 3D shape without adding actual geometry. To generate smooth normal maps even for shapes with low resolution, we use the smooth shading technique which interpolates the normals along a flat mesh triangle to give the illusion of smooth shading. The computed smooth normals are then mapped onto a 2D image using the same parameterization. The texture maps are generated by using an environment map around the object which can generate both realistic lighting and reflection for the object. The result of the rendering is then baked as a texture and mapped into a 2D image using the same procedure as with normal maps.

We use these aligned image channels to train an image-to-image translation network to produce X attribute channel images given an input Rad image of the same resolution. Specifically, we experiment with SPADE image-to-image translation network [66], but nothing in our setup is specific to this particular network choice. We note that SPADE is designed to be able to produce diverse outputs given a single input.

3.3 Exploiting the Latent Space

A major advantage of our 3D generative model is that it is based on standard GAN architectures, which allows us to directly apply various techniques developed for image-based GANs. To demonstrate this, we integrate the existing technique of InterFaceGAN [44] in order to find editing directions within our trained latent space. To this end, we first project our training shapes into the latent space W to create a set of projected latent vectors w with labels (e.g. *sports-car*, *SUV*, *etc.*). We then train one linear SVMs for each label using these labeled examples, and use the unit normal vector to the separation hyperplane of the SVM \mathbf{n}^T as an editing direction in latent space. The target attribute of a real or generated 3D shape can

then be enhanced or negated by simply steering its latent vector p along this editing direction \mathbf{n}^T by t steps with a step size of α to produce a new latent vector $p' = p + (\alpha \times t \times \mathbf{n}^T)$

4 Experiments

We evaluate our model’s performance on shape projection, editing and single-view reconstruction tasks. We compare our method to three state-of-the-art 3D shape generation methods that use triangular meshes (AtlasNet [14]), point-clouds (DPC [5]) and implicit representation (IM-Net [15]). The experiments show that our method achieves significantly better results on some object classes and competitive results on others, all while performing inference several orders of magnitude faster than the closest competitor IM-NET [15].

4.1 Datasets and Experimental Setup

For the quantitative comparisons, we use the ShapeNetV2 [4] *car* and *airplane* categories consisting of 3509 and 4045 shapes respectively and use 10% of each set for testing. We choose these two categories as they are sufficiently distinct structurally and demonstrate the range of our method. Additionally, we use a proprietary new dataset consisting of 3050 high-quality car models which we call *HQCars*. In addition to the higher quality of meshes, this dataset provides high quality texture maps and semantic labels, like model, maker, year and type, which are not available in ShapeNet.

To create our training set, we first converted the set of 3D meshes to 1-channel Rad geometry images of size 512 x 512 and the corresponding normals and texture maps. (See Supplementary Material for comparisons with geometry images using 3 (x, y, z) channels) We use the Rad geometry images to train a modified version of the official StyleGAN3 implementation [2] to process high-precision values instead of the 8-bit RGB images. We train this StyleGAN3-r variant without any data augmentation with batch size 8 and gamma value of 1 for 500 epochs. To generate the additional texture and normal maps, we train an image-to-image translation model separately for each using the official implementation of SPADE [6], similarly modified to deal with the range and precision of geometry images. Each model is trained on aligned pairs of ground-truth geometry images and texture or normals maps, for 300 epochs with batch size of 32.

We train our model and all baselines from scratch on the *car* and *airplane* training data subsets *separately* and evaluate on the respective test subsets. We use the official Pytorch implementations of AtlasNet, DPC and IM-Net with the default parameters.

For all experiments, we use reconstruction-based tasks and compare the generated shapes to ground truth shapes using Chamfer distance (CD) and Earth Mover’s Distance (EMD) metrics. CD and EMD metrics are computed by sampling 5000 points from the generated and ground truth shapes, matching the nearest points across the sets of points and computing their sum of Euclidean and Manhattan distances respectively. In addition to CD and EMD, we use the Light Field Distance (LFD) [16] as a visual similarity metric. To compute LFD, the generated and ground truth shapes are rendered from various camera angles. The rendered images are then encoded using Zernike moments and Fourier descriptors to compute similarity. Finally, we report the average inference time of ours and competing methods as real-time generation is critical for interactive applications.

	Reconstruction error						Infr. Time
	Chamfer ↓		EM Dist. ↓		LFD Dist. ↓		Seconds ↓
	Car	Airplane	Car	Airplane	Car	Airplane	Mean
AtlasNet-Sph	2.089	6.241	2.43	4.85	2610	5006	0.078
AtlasNet-25	1.786	5.216	2.22	3.94	2565	4730	0.112
IM-Net	1.439	0.406	1.49	1.06	2243	4759	13.280
DPC	6.146	4.450	4.58	4.11	-	-	0.061
Ours	0.977	2.721	1.34	1.96	2407	7435	0.020

Table 2: Quantitative results on ShapeNet reconstruction. Note that the CD values are multiplied by 10^3 and EMD are multiplied by 10^2 . Since DPC produces only point clouds, it is not possible to compute the LFD distance for its outputs.

4.2 Quantitative Evaluation

For evaluating on the task of 3D shape reconstruction, we use an iterative latent projection approach to obtain a reconstruction of ground-truth test model. We start from a randomly sampled latent vector and iteratively compute the loss between the generated geometry image from this vector and the ground-truth geometry image, then back-propagate the loss through the generator to update the latent vector. We repeat this procedure for 500 iterations for each test model. For the single-view reconstruction task, we train a ResNet-based encoder on the rendered RGB images [12] of a ShapeNet model from a random viewpoint, and predicts a latent vector which is then used to generate a geometry image. For both tasks, the final evaluation metrics are computed between the triangle mesh created from the generated geometry image and the original ShapeNet test mesh.

Results in Table 2 show that our method achieves by far the best CD and EMD on the car category and ranks second on the airplane category after IM-Net, while being faster by factor of 500. This finding is not surprising, as we have made some simplifying assumptions during conversion to geometry images (§3.1), and it is harder to map airplanes with many sharp and spiky features onto the spherical domain than it is for relatively compact cars. These results also demonstrate that when the 2D parameterization is appropriate (as it is for cars), 2D generative models using our technique outperform even implicit 3D representation models, which is an unexpected outcome. We further note that our method is the fastest by at least a factor of 3 over any baseline. For Single-View reconstruction experiment (Table 3), the results are similar where our methods achieves very good results on the car category while falling short from IM-Net on the airplane category.

	SVR error			
	Chamfer ↓		EM Dist. ↓	
	Car	Airplane	Car	Airplane
AtlasNet-Sph	2.335	5.514	2.558	4.917
AtlasNet-25	2.107	4.605	2.264	3.737
IM-Net	1.437	0.907	1.492	1.282
Ours	1.325	3.055	1.511	2.030

Table 3: Quantitative results on Single View Reconstruction. Note that the CD values are multiplied by 10^3 and EMD are multiplied by 10^2 .

Input Model					
AtlasNet-Sphere					
AtlasNet-25					
IM-Net					
Ours (no normals)					
Ours					

Table 4: Qualitative Comparison of Reconstruction Results (Zoom for more details). Note that our method can generate highest-fidelity 3D meshes including detailed surface normals.

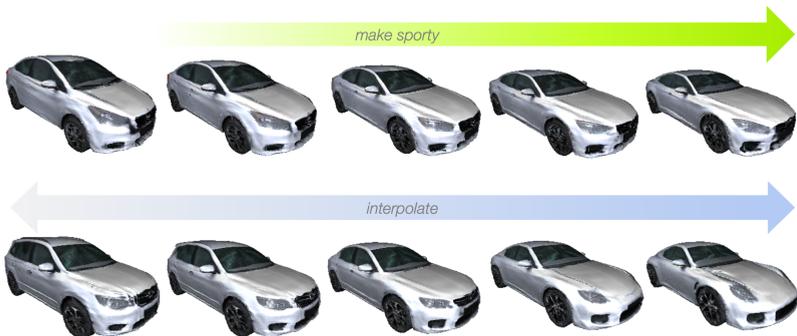


Figure 4: The latent space of XDGAN allows meaningful latent space exploration such as semantic editing (top, see §3.3) and interpolation between two input models (bottom).

4.3 Qualitative Results

Table 4 show 3D meshes projected with XDGAN using the method in §4.2, visually compared to other competitors. Our methods generates water-tight meshes with high-fidelity that preserve even small details like car mirrors and wheel patterns. The inclusion of generated surface normals also add another level of detail and smoothness where needed, making the results from XDGAN visually outstanding compared to other methods.

Furthermore, Fig.4 shows the the powerful properties of XDGAN’s latent space. Editing directions optimized using InterFaceGAN [14] (§3.3) allow semantic editing of the 3D object *in real time*, for example by making a car appear sportier (Fig.4, top). The latent space additionally allows for generating intermediates between two input models by linearly interpolating their projected latent space vectors.

5 Conclusion

We presented XDGAN, the first method to show that 2D convolutional architectures can be directly used for generative modeling of 3D objects with high fidelity. Our approach is

based on a novel combination of planar image parameterization with an image-based GAN. We show that our method generates high-quality textured 3D meshes, allows unconditional sampling, projection, and can be used in several conditional settings. In addition, XDGAN's speed and convenient mesh output allow for a number of applications that can be directly employed in practice. The shortcoming of our approach is its reliance on a fixed topology mesh, associated with the geometry image. We believe that there are many ways future work could mitigate this limitation, for example through part-based generation. Despite this limitation, we have shown impressive generation results from our method, and hope that this work may help close the gap between advances in 2D generative architectures and comparable approaches in 3D.

6 Acknowledgments

We thank Stefan Liske and Timotheus Gmeiner from PCH Innovations GmbH for the insightful discussions and making this collaboration possible. Part of this research work has been done while Hassan Abu Alhaija was at PCH Innovations GmbH.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8296–8305, 2020.
- [2] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas J. Guibas. Learning representations and generative models for 3d point clouds. In *ICML*, 2018.
- [3] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6711–6720, 2021.
- [4] Hassan Abu Alhaija, Siva Karthik Mustikovela, Justus Thies, Varun Jampani, Matthias Nießner, Andreas Geiger, and Carsten Rother. Intrinsic autoencoders for joint deferred neural rendering and intrinsic image decomposition. In *2020 International Conference on 3D Vision (3DV)*, pages 1176–1185. IEEE, 2020.
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *CoRR*, abs/1809.11096, 2018. URL <http://arxiv.org/abs/1809.11096>.
- [6] Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snavely, and Bharath Hariharan. Learning gradient fields for shape generation. In *European Conference on Computer Vision*, pages 364–381. Springer, 2020.
- [7] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. *arXiv preprint arXiv:2112.07945*, 2021.

- [8] 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/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021.
- [9] 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.
- [10] Ding-Yun Chen, Xiao-Pei Tian, Edward Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3d model retrieval. *Computer Graphics Forum*, 22, 2003.
- [11] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5932–5941, 2019.
- [12] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [13] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017.
- [14] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. Ganalyze: Toward visual definitions of cognitive image properties. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5744–5753, 2019.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- [16] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [17] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis, 2021.
- [18] Xianfeng Gu, Steven J. Gortler, and Hugues Hoppe. Geometry images. *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, 2002.
- [19] Zekun Hao, Hadar Averbuch-Elor, Noah Snavely, and Serge J. Belongie. Dualsdf: Semantic shape manipulation using a two-level representation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7628–7638, 2020.
- [20] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *arXiv preprint arXiv:2004.02546*, 2020.

- [21] Ali Jahanian, Lucy Chai, and Phillip Isola. On the "steerability" of generative adversarial networks. *arXiv preprint arXiv:1907.07171*, 2019.
- [22] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation, 2018.
- [23] 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*, pages 4401–4410, 2019.
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8107–8116, 2020.
- [25] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021.
- [26] Hyunsu Kim, Yunjey Choi, Junho Kim, Sungjoo Yoo, and Youngjung Uh. Exploiting spatial dimensions of latent in gan for real-time image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 852–861, 2021.
- [27] Minchen Li, Danny M. Kaufman, Vladimir G. Kim, Justin M. Solomon, and Alla Sheffer. Optcuts: joint optimization of surface cuts and parameterization. *ACM Trans. Graph.*, 37:247:1–247:13, 2018.
- [28] Ruihui Li, Xianzhi Li, Ka-Hei Hui, and Chi-Wing Fu. Sp-gan: Sphere-guided 3d shape generation and manipulation. *ArXiv*, abs/2108.04476, 2021.
- [29] Chen-Hsuan Lin, Chen Kong, and Simon Lucey. Learning efficient point cloud generation for dense 3d object reconstruction. In *AAAI*, 2018.
- [30] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [31] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2836–2844, 2021.
- [32] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4455–4465, 2019.
- [33] 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 *ECCV*, 2020.
- [34] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021.

- [35] Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2019.
- [36] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [37] Roi Poranne, Marco Tarini, Sandro Huber, Daniele Panozzo, and Olga Sorkine-Hornung. Autocuts: simultaneous distortion and cut optimization for uv mapping. *ACM Trans. Graph.*, 36:215:1–215:11, 2017.
- [38] C. Qi, L. Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017.
- [39] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [40] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021.
- [41] Alla Sheffer and John C. Hart. Seamster: inconspicuous low-distortion texture seam layout. *IEEE Visualization, 2002. VIS 2002.*, pages 291–298, 2002.
- [42] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34, 2021.
- [43] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020.
- [44] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *TPAMI*, 2020.
- [45] Ayan Sinha, Jing Bai, and Karthik Ramani. Deep learning 3d shape surfaces using geometry images. In *European conference on computer vision*, pages 223–240. Springer, 2016.
- [46] Olga Sorkine-Hornung, Daniel Cohen-Or, Rony Goldenthal, and Dani Lischinski. Bounded-distortion piecewise mesh parameterization. *IEEE Visualization, 2002. VIS 2002.*, pages 355–362, 2002.
- [47] 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*, pages 2088–2096, 2017.
- [48] Nobuyuki Umetani. Exploring generative 3d shapes using autoencoder networks. In *SIGGRAPH Asia 2017 technical briefs*, pages 1–4. 2017.

- [49] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *International Conference on Machine Learning*, pages 9786–9796. PMLR, 2020.
- [50] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Joshua B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NIPS*, 2016.
- [51] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, 2015.
- [52] Xingguang Yan, Liqiang Lin, Niloy J Mitra, Dani Lischinski, Danny Cohen-Or, and Hui Huang. Shapeformer: Transformer-based shape completion via sparse representation. *arXiv preprint arXiv:2201.10326*, 2022.
- [53] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4541–4550, 2019.
- [54] Dongsu Zhang, Changwoon Choi, Jeonghwan Kim, and Young Min Kim. Learning to generate 3d shapes with generative cellular automata. *arXiv preprint arXiv:2103.04130*, 2021.
- [55] Jason Y. Zhang, Gengshan Yang, Shubham Tulsiani, and Deva Ramanan. NeRS: Neural reflectance surfaces for sparse-view 3d reconstruction in the wild. In *Conference on Neural Information Processing Systems*, 2021.
- [56] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yanan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. In *International Conference on Learning Representations*, 2020.
- [57] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021.