

DECOR-GAN: 3D Shape Detailization by Conditional Refinement

Zhiqin Chen¹Vladimir G. Kim²Matthew Fisher²Noam Aigerman²Hao Zhang¹Siddhartha Chaudhuri^{2,3}¹ Simon Fraser University² Adobe Research³ IIT Bombay

Abstract

We introduce a deep generative network for 3D shape detailization, akin to stylization with the style being geometric details. We address the challenge of creating large varieties of high-resolution and detailed 3D geometry from a small set of exemplars by treating the problem as that of geometric detail transfer. Given a low-resolution coarse voxel shape, our network refines it, via voxel upsampling, into a higher-resolution shape enriched with geometric details. The output shape preserves the overall structure (or content) of the input, while its detail generation is conditioned on an input “style code” corresponding to a detailed exemplar. Our 3D detailization via conditional refinement is realized by a generative adversarial network, coined DECOR-GAN. The network utilizes a 3D CNN generator for upsampling coarse voxels and a 3D PatchGAN discriminator to enforce local patches of the generated model to be similar to those in the training detailed shapes. During testing, a style code is fed into the generator to condition the refinement. We demonstrate that our method can refine a coarse shape into a variety of detailed shapes with different styles. The generated results are evaluated in terms of content preservation, plausibility, and diversity. Comprehensive ablation studies are conducted to validate our network designs. Code is available at <https://github.com/czq142857/DECOR-GAN>.

1. Introduction

Creating high-quality detailed 3D shapes for visual design, AR/VR, gaming, and simulation is a laborious process that requires significant expertise. Recent advances in deep generative neural networks have mainly focused on learning low-dimensional [1, 8, 18, 46, 47] or structural representations [6, 26, 34, 43, 48] of 3D shapes from large collections of stock models, striving for plausibility and diversity of the generated shapes. While these techniques are effective at creating coarse geometry and enable the user to model rough objects, they lack the ability to represent, synthesize, and provide control over the finer geometric details.

In this work, we pose the novel problem of *3D shape detailization*, akin to stylization with the style defined by geo-




Figure 1: Our 3D detailization network, DECOR-GAN, refines a coarse shape (red, leftmost) into a variety of detailed shapes, each conditioned on a style code characterizing an exemplar detailed 3D shape (green, topmost).

metric details. We wish to address the challenge of creating large varieties of high-resolution and detailed 3D geometries from only a small set of detailed 3D exemplars by treating the problem as that of *geometric detail transfer*. Specifically, given a low-resolution coarse *content* shape and a *detailed style* shape, we would like to synthesize a novel shape that preserves the coarse structure of the content, while refining its geometric details to be similar to that of the style shape; see Figure 1. Importantly, the detail transfer should *not* rely on any annotations regarding shape details.

Our conditional detailization task cannot be accomplished by simply copying local patches from the style shape onto the content shape, since (a) it is unknown which patches represent details; (b) it may be difficult to integrate copied patches into the content shape to ensure consistency. To this end, we train a *generative* neural network that learns detailization priors over a collection of high-resolution detailed exemplar shapes, enabling it to refine a coarse instance using a *detail style code*; see top of Figure 1.

To date, there has been little work on generating high-resolution detailed 3D shapes. For example, surface-based representations that synthesize details on meshes [21, 28]

lack the ability to perform topological changes (as in several plant detailization results in Figure 1) and, due to complexities of mesh analysis, do not consider large shape contexts, limiting their notion of details to homogeneous geometric texture. To allow topological variations while leveraging a simpler domain for analyzing the content shape, we choose voxel representations for our task. However, 3D grids do not scale well with high resolution, which necessitates a network architecture that effectively leverages limited capacity. Also, regardless of the representation, careful choices of network losses must be made to balance the conflicting goals of detailization and (coarse) content preservation.

To tackle all these challenges, we design DECOR-GAN, a novel generative adversarial network that utilizes a 3D CNN generator to locally refine a coarse shape, via voxel upsampling, into a high-resolution detailed model, and a 3D PatchGAN discriminator to enforce local patches of the generated model to be similar to those in the training detailed shapes. Our generator learns local filters with limited receptive field, and thus effectively allocates its capacity towards generating local geometry. We condition our refinement method on a single detail style code that provides an intuitive way to control the model while also ensuring that generated details are consistent across the entire shape. We train our approach adversarially, using masked discriminators to simultaneously ensure that local details are stylistically plausible and coherent, while the global shape (i.e., downsampled version of the detailed shape) still respects the coarse source shape. Our method can be trained even with a small number of detailed training examples (up to 64 for all of our experiments), since our convolutional generator only relies on learning local filters.

We demonstrate that DECOR-GAN can be used to refine coarse shapes derived by downsampling a stock model or neurally generated by prior techniques. The user can control the style of the shape detailization either by providing a style code from an exemplar or by interpolating between existing styles. Thus, our technique offers a *complementary latent space for details* which could be used jointly with existing techniques that provide latent spaces for coarse shapes. We quantitatively evaluate our method using measures for content preservation, plausibility, and diversity. We also provide a comprehensive ablation study to validate our network designs, while demonstrating that simpler approaches (such as only using a reconstructive loss) do not provide the same quality of generated details.

2. Related work

Deep learning models have explored many possible representations for understanding and generating 3D content. These include encoders and decoders for voxel [4, 10, 47], point cloud [14, 39], mesh [17, 35, 45], atlas chart [18, 50], and implicit occupancy function representations [9, 32, 49].

High-Resolution 3D Representations. While most representations can approximate the coarse shape of an object, specialized techniques have been developed for each representation to improve the quality of surface geometric detail, and hallucinate high-frequency detail from low-frequency input (“super-resolution”). For implicit functions, detail has been improved by training a coarse model, then increasing the accuracy and focusing on local details [12]. An alternative approach uses convolutional layers to generate input-point-specific features for the implicit network to decode [38]. Single-view 3D reconstruction quality has also been improved by using local image features to recover geometric detail in the implicit network [49]. For voxel representations, several hierarchical methods increase quality and resolution [19, 40, 46]. Smith et al. [42] take a different approach by reconstructing a voxel grid from high-resolution 2D depth maps. Taking this further, Mildenhall et al. [33] acquire high-resolution 3D geometry by merging several calibrated images into a neural implicit function. A PatchMatch approach has been used to reconstruct a partial 3D surface scan by directly copying patches from training models [11]. Wang et al. [52] proposed a patch-based up-scaling approach to super-resolve point clouds. Although these methods improve surface detail, they cannot be conditioned on an input style code and do not easily allow for controlled surface detail generation.

Mesh-based learning methods have been proposed to adjust the subdivision process of an input mesh to control surface details [28], synthesize surface textures using multi-scale neurally displaced vertices [21], or reconstruct a surface with reoccurring geometric repetitions from a point cloud [20]. While able to produce highly detailed surfaces, these methods cannot alter the topology of the input mesh or mix geometric details from a collection of style shapes. In a similar vein, our 3D detailization solution also differs from conventional image [25] or 3D upsampling [41] with the added controllability while requiring much fewer detailed high-resolution shapes during training.

Shape Detail Transfer. Mesh “quilting” is an early non-learning-based method for shape detail transfer; it tiles the surface of a coarse shape with copies of a geometric texture patch [53]. Takayama et al. [44] transfer a detailed patch from one shape to another by matching parametrizations. Chen et al. [7] extend 2D PatchMatch to surfaces in 3D. These methods are not automatic and/or require explicit factorization of the shape into content and detail layers. Ma et al. [30] solve the analogy – “shape : target :: exemplar : ?” – by automatically assembling exemplar patches. They require accurate surjective correspondences, local self-similarity and, most importantly, *three* input shapes to implicitly define the style-content factorization.

Among learning-based methods, Berkiten et al. [3] transfer displacement maps from source to target shapes in one-

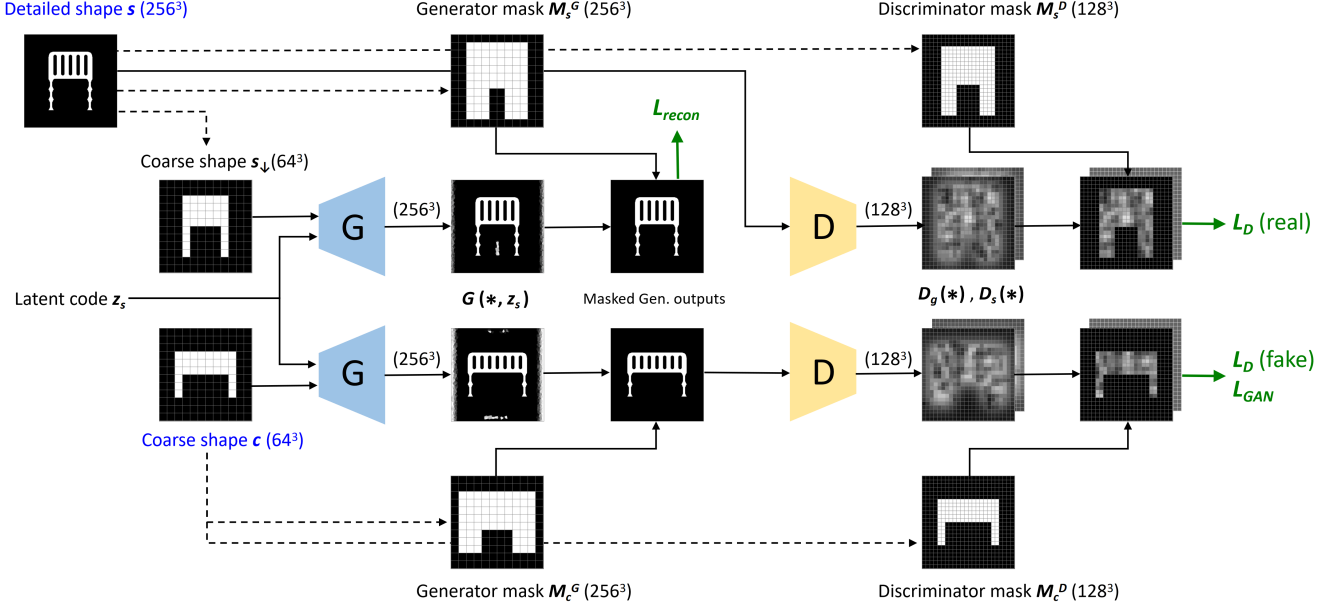


Figure 2: The network architecture. The training data is shown in blue and the loss functions are shown in green. Note that there is only one generator G and discriminator D – the blocks are duplicated for clarity.

shot fashion. Wang et al. [51] propose neural cages to warp a detailed object to the coarse shape of another. Neither of these methods can mix details from multiple shapes or synthesize new topology. Chart-based methods [2, 18] map a common 2D domain to a collection of shape surfaces, which can be used to transfer details to and between shapes. However, they also cannot synthesize new topology, and can accurately represent only relatively simple base shapes.

Image Synthesis. Several methods control content and style in 2D image generation [16, 27, 37]. Recent approaches to generative imaging [15, 24, 37] employ a PatchGAN discriminator [23] that has also been used to construct a latent space for 3D shape categories [47]. One effective way to condition a generative image model is to inject a latent code into each level of the generator [13, 36]. We build upon these techniques in our generator design, which constructs a latent space over shape detail that is injected into the generator architecture to guide detail synthesis. We similarly employ a PatchGAN to ensure that the generated 3D shape resembles patches from training shapes.

3. Method

3.1. Network architecture and loss functions

The network structure for DECOR-GAN is shown in Figure 2. In the training stage, the network receives one of the M coarse shapes (64^3 voxels, referred to as “content shapes”), as well as a latent code describing the style of one of the N detailed shapes (256^3 voxels). The network is trained to upsample the content shape by 4 times according to the style of the detailed shape.

Network overview. We use a GAN approach and train a generative model that can generate a shape with local patches similar to those of the detailed shapes. For the generator we utilize a 3D CNN, and for the discriminator we use 3D CNN PatchGANs [23] with receptive fields of 18^3 . We use an embedding module to learn an 8-dimensional latent style code for each given detailed shape. Please refer to the supplementary material for the detailed architectures of the networks. Note that we only use one generator and one discriminator in our system, and Figure 2 shows duplicated networks solely for the sake of clarity.

Enforcing consistency between coarse input and fine output.

We guide the generator to generate upsampled voxels that are consistent with the coarse input voxels, in the sense that the downsampled version of the fine output should be identical to the coarse input content shape. To enforce this, we employ two masks: a generator mask and a discriminator mask. The generator mask ensures that empty voxels in the input content shape correspond with empty voxels in the output, and enables the generator to focus its capacity solely on producing plausible voxels within the valid region. We use a dilated mask to allow the generator to have some freedom to accommodate for different styles as well as allow for more topological variations. Figure 6 demonstrates the effect of using the generator mask. The discriminator mask ensures that each occupied voxel in the input content shape leads to creation of fine voxels in its corresponding area of the output. We apply discriminator masks on real and fake patches to only keep the signals of regions corresponding to occupied patches, so that lack of voxels in those patches will be punished. In our experi-

ments, we use discriminator masks with 1/2 of the resolution of the detailed shapes to fit the entire model into the GPU memory. More details of generating our masks are provided in the supplementary.

Preventing mode collapse. Employing one global discriminator may result in mode collapse: the generator may output the same style regardless of the input style code, and the discriminator may ignore the style when determining whether a patch is plausible or not. Therefore, our discriminator is split into $N + 1$ branches at the output layer, where N is the number of detailed shapes used in training, with the additional 1 standing for a global discriminator branch considering all styles. During training, the global branch computes a global GAN loss, and the appropriate style-specific branch computes a style-specific loss. The two losses are weighed in order to control the tradeoff between general plausibility and stylization.

Loss function. We now detail the different terms in our loss function. We denote the set of N detailed shapes by \mathcal{S} and the set of M coarse shapes by \mathcal{C} . The binary discriminator masks of shape $s \in \mathcal{S}$ and shape $c \in \mathcal{C}$ are denoted as M_s^D and M_c^D , respectively. The binary generator masks of shape $s \in \mathcal{S}$ and shape $c \in \mathcal{C}$ are denoted as M_s^G and M_c^G , respectively. We denote the latent style code of $s \in \mathcal{S}$ by z_s . We denote the generator and discriminator as G and D , respectively. The global branch of D is denoted as D_g and the branch for the detailed shape $s \in \mathcal{S}$ is denoted as D_s . \circ stands for element-wise product. We use the adversarial loss in LSGAN [31] as our GAN loss, therefore, our discriminator loss is composed of the global branch’s loss and the style branch’s loss:

$$L_D = L_D^{global} + L_D^{style}, \quad (1)$$

where

$$L_D^{global} = \mathbb{E}_{s \sim \mathcal{S}} \frac{\|(D_g(s) - 1) \circ M_s^D\|_2^2}{\|M_s^D\|_1} + \mathbb{E}_{s \sim \mathcal{S}, c \sim \mathcal{C}} \frac{\|D_g(G(c, z_s) \circ M_c^G) \circ M_c^D\|_2^2}{\|M_c^D\|_1}, \quad (2)$$

$$L_D^{style} = \mathbb{E}_{s \sim \mathcal{S}} \frac{\|(D_s(s) - 1) \circ M_s^D\|_2^2}{\|M_s^D\|_1} + \mathbb{E}_{s \sim \mathcal{S}, c \sim \mathcal{C}} \frac{\|D_s(G(c, z_s) \circ M_c^G) \circ M_c^D\|_2^2}{\|M_c^D\|_1}. \quad (3)$$

Our GAN loss for the generator is

$$L_{GAN} = L_{GAN}^{global} + \alpha \cdot L_{GAN}^{style}, \quad (4)$$

where

$$L_{GAN}^{global} = \mathbb{E}_{s \sim \mathcal{S}, c \sim \mathcal{C}} \frac{\|(D_g(G(c, z_s) \circ M_c^G) - 1) \circ M_c^D\|_2^2}{\|M_c^D\|_1}, \quad (5)$$

$$L_{GAN}^{style} = \mathbb{E}_{s \sim \mathcal{S}, c \sim \mathcal{C}} \frac{\|(D_s(G(c, z_s) \circ M_c^G) - 1) \circ M_c^D\|_2^2}{\|M_c^D\|_1}. \quad (6)$$

In addition, we add a reconstruction loss: if the input coarse shape and the style code stem from the same detailed shape, we expect the output of the generator to be the ground-truth detailed shape.

$$L_{recon} = \mathbb{E}_{s \sim \mathcal{S}} \frac{\|G(s_{\downarrow}, z_s) \circ M_s^G - s\|_2^2}{|s|} \quad (7)$$

where $|s|$ is the volume (height \times width \times depth) of the voxel model s , and s_{\downarrow} is the content shape downsampled from s . The reconstruction loss both speeds up convergence in the early stage of training, as well as helps avoid mode collapse.

The final loss function for the generator is a weighted sum of the GAN loss and the reconstruction loss.

$$L_{total} = L_{GAN} + \beta \cdot L_{recon} \quad (8)$$

3.2. Implementation details

In our implementation, we take several measures to handle the large memory footprint incurred when processing 3D high-resolution voxels. We use 256^3 voxels as our high-resolution detailed shapes and 64^3 voxels as low-resolution content shapes. In order to discard unused voxels, we crop each shape according to its dilated bounding box by cropping the content shape first, then use the upsampled crop region as the crop region of the high resolution models. Moreover, since most man-made shapes are symmetric, we assume all training shapes possess bilateral symmetry, and hence only generate half of the shape.

Another important consideration is that voxels are expected to hold binary values, as opposed to continuous intensities in images. As reported in the supplementary material of IM-NET [9], a naive GAN with CNN architectures performs poorly on binary images, as pixels with any value other than 0 or 1 will be considered as fake by the discriminator, thus preventing continuous optimization. We follow the approach in IM-NET and apply a Gaussian filter with $\sigma = 1.0$ to the training high-resolution voxel grids to make the boundary smoother, as a pre-processing step. We provide some analysis in Sec 4.4.

We set $\alpha = 0.2$ for car, $\alpha = 0.1$ for airplane, and $\alpha = 0.5$ for chair. We set $\beta = 10.0$. More discussion about the parameters can be found in Sec 4.4. We train individual models for different categories for 20 epochs on one Nvidia Tesla V100 GPU; training each model takes 6 to 24 hours depending on the category.

4. Results and evaluation

We conducted experiments on three categories from ShapeNet [5]: car, airplane, and chair. We use only the

bilaterally-symmetric models, apply an 80% / 20% train/test split of the coarse content shapes, and select 64 fine detailed shapes. We choose a 64^3 voxel resolution for coarsening airplanes and cars, and a coarser resolution of 32^3 for chairs to further remove topological details. Note that our network is designed to upsample any input by 4 times. We use marching cubes [29] to extract the surfaces visualized in the figures. More categories (table, motorbike, laptop, and plant) can be found in the supplementary, where we lift the bilateral symmetry assumption for some categories.

4.1. Style-content hybrids

In Figure 3, we show results obtained by upsampling a content shape with a latent code to guide the style of the output shape. While the *style-specific* discriminator encourages the generator to use style-appropriate patches, the *global* discriminator ensures that in case no such plausible patches exist, the generator will compromise on using patches from other styles and not generate implausible details. More results can be found in the supplementary.

4.2. Latent space

We can explore styles in a continuous way within the 8-dimensional latent space of styles, and have created a GUI app for it (see Sec 4.5). We visualize the chair style space in Figure 4, revealing grouping of similar styles.

4.3. Evaluation metrics

We now discuss the metrics used to evaluate our method. See the supplementary for the full details. **Strict-IOU** measures the Intersection over Union between the downsampled output voxels and the input voxels to evaluate how much the output respects the input. **Loose-IOU** is a relaxed version of Strict-IOU, which computes the percentage of occupied voxels in the input that are also occupied in the downsampled output. Local Plausibility (LP) is defined as the percentage of local patches in the output shape that are “similar” (according to IOU or F-score) to at least one local patch in the detailed shapes, which defines the metrics **LP-IOU** and **LP-F-score**. We evaluate the Diversity (Div) of the output shapes by computing the percentage of output shapes that are consistent with their input style code. An output shape with input style code z_s whose local patches are most similar to those of the detailed shape s is considered as a “consistent” output. We measure **Div-IOU** and **Div-F-score** according to similarity metrics for patches. We evaluate the plausibility of the generated shapes by training a classifier to distinguish generated and real shapes, and use the mean classification accuracy as **Cls-score**. Following Fréchet Inception Distance (FID) [22], we use FID to compare the output shapes with all available shapes in a category, denoted as **FID-all**; or only the detailed shapes providing styles, denoted as **FID-style**. For Cls-score and FID,

a lower score is better; for others, a higher score is better.

4.4. Ablation study

We now verify the necessity of the various parts of our network. We report the quantitative results for chairs in this section, and other categories in the supplementary.

Generator and Discriminator Masks. In Table 1, and Figure 5 (top) we compare our proposed method with several variations on it: **a. Recon. only**, in which we remove the discriminator and train the network with only L_{recon} , to validate the effectiveness of adversarial training. This results in the network essentially mode-collapsing (reflected by the low Div scores). **b. No Gen. mask**, in which we remove the generator mask. To still respect the input shape, we add a loss term to punish any voxels generated outside the generator mask. This results in comparable performance to our method locally, but since the generator needs to allocate a portion of its capacity to ensure no voxels are generated outside the generator mask, it is left with less capacity for generating the output shape. This is reflected in deterioration of global plausibility (reflected by the higher Cls-score). **c. Strict Gen. mask**, in which we use the un-dilated, “strict” generator mask, which results in a harsh drop in quality. **d. No Dis. mask**, in which we remove the discriminator mask, performs even worse, with some outputs completely empty, because patches with no occupied voxels are considered as real by the discriminator. **e. Conditional Dis. 1**, in which we remove both the generator mask and the discriminator mask, and condition the discriminator on the occupancy of the input coarse voxels, i.e., a conditional GAN. In this setting, each discriminator output will have a receptive field of 1^3 in the input coarse voxels. This leads to failure in generating diverse outputs, possibly due to the discriminator wasting capacity on parsing the input conditions. **f. Conditional Dis. 3** modifies the receptive field of *Conditional Dis. 1* to 3^3 , but runs into similar issues.

One may notice that there is a considerable difference between LP-IOU and LP-F-score, and they sometimes contradict each other. This is due to IOU being sensitive to small perturbations on the surface, especially when the structure is thin. F-score, on the other hand, is less sensitive. Due to the strictness of IOU, Div-IOU is usually higher than Div-F-score. In addition, we observe that Cls-score is consistent with apparent visual quality. Therefore, in the following, we only report Loose-IOU, LP-F-score, Div-IOU and Cls-score. The full tables can be found in supplementary.

We also show the effect of the generator mask in Figure 6. Evidently, the raw generator output (b) has various artifacts outside the masked region, which are removed upon applying the generator mask (a). However, new artifacts may be created in the process, such as the one shown in (c). Finding the connected components of (c) from (b) can remove such artifacts.




Figure 3: Results by upsampling coarse voxels with different style codes. In each table, we show on the top the detailed shapes that correspond to the input style codes. We show the input coarse voxel models on the left, where chairs are 32^3 and the others are 64^3 .

	Strict-IOU \uparrow	Loose-IOU \uparrow	LP-IOU \uparrow	LP-F-score \uparrow	Div-IOU \uparrow	Div-F-score \uparrow	Cls-score \downarrow	FID-all \downarrow	FID-style \downarrow
Recon. only	0.976	0.993	0.260	0.935	0.325	0.188	0.627	53.2	411.7
No Gen. mask	0.655	0.792	0.452	0.973	0.825	0.806	0.672	121.9	379.9
Strict Gen. mask	0.587	0.587	0.344	0.941	0.150	0.100	0.750	305.5	548.2
No Dis. mask	0.145	0.167	N/A	N/A	N/A	N/A	0.843	2408.9	2714.1
Conditional Dis. 1	0.947	0.981	0.259	0.949	0.291	0.194	0.593	51.3	402.7
Conditional Dis. 3	0.928	0.977	0.246	0.963	0.197	0.206	0.603	55.8	418.2
Proposed method	0.673	0.805	0.432	0.973	0.800	0.816	0.644	113.1	372.5

Table 1: Ablation on the generator and discriminator masks. “N/A” is due to empty outputs. Best results are marked in bold.




Figure 4: Visualization of 64 latent codes for chairs via T-SNE embedding. For each latent code, the corresponding style shape is displayed in its location.

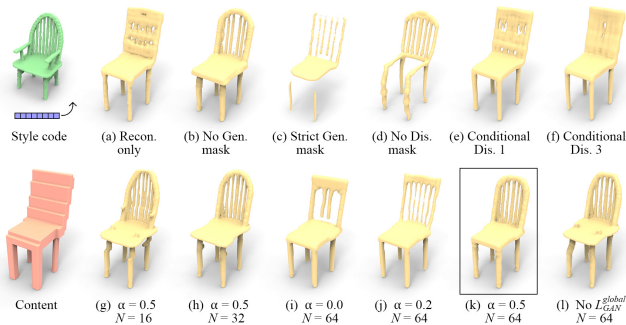


Figure 5: Ablation study. The input content shape and style code are shown on the left. The result with our default parameters is placed inside a box.

Parameter α and the number of styles. In Tables 2 we show results of setting the parameter α to values 0.0, 0.2, 0.5, as well as completely omitting L_{GAN}^{global} , which can be seen as using a very large α that would dominate the other terms. α controls the trade-off between global plausibility (and respecting the coarse content), and adherence to the style code. As shown by the increase in Div-IOU, the higher α is, the more stylized the output will be. However,

16 detailed shapes as styles				
	Loose-IOU \uparrow	LP-F-score \uparrow	Div-IOU \uparrow	Cls-score \downarrow
$\alpha = 0.0$	0.840	0.956	0.147	0.695
$\alpha = 0.2$	0.750	0.971	0.875	0.667
$\alpha = 0.5$	0.738	0.970	0.997	0.690
No L_{GAN}^{global}	0.735	0.963	1.000	0.692
32 detailed shapes as styles				
$\alpha = 0.0$	0.864	0.962	0.184	0.598
$\alpha = 0.2$	0.812	0.974	0.838	0.636
$\alpha = 0.5$	0.757	0.974	0.934	0.662
No L_{GAN}^{global}	0.728	0.969	0.997	0.690
64 detailed shapes as styles				
$\alpha = 0.0$	0.868	0.983	0.163	0.589
$\alpha = 0.2$	0.864	0.985	0.353	0.619
$\alpha = 0.5$	0.805	0.973	0.800	0.644
No L_{GAN}^{global}	0.741	0.965	0.950	0.669

Table 2: Ablation on parameter α and the number of styles.

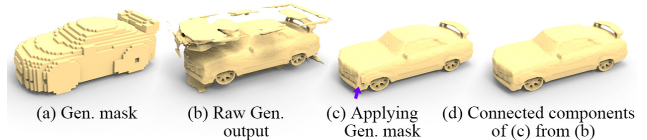


Figure 6: The effectiveness of the generator mask. This example is taken from row 2, column 5 of Figure 3.

the decrease in Loose-IOU and increase in Cls-score hint that a higher α also makes the output less considerate of the content and less globally plausible. This can also be seen in the qualitative results in Figure 5 (bottom).

Simultaneously, we also test the effect of varying the size of the detailed-shape dataset used in training, between 16, 32, and 64. As expected, both Cls-score and Loose-IOU monotonically improve with the increase in dataset size, showing plausibility improves, because the network has a larger collection of patches to choose from. This can be seen in Figure 5 (g), where the network was trained with 16 detailed shapes, and generated small bumps on the seat and the back, clearly stemming from a vestige of the arm. The same phenomenon can be found in (l) where there is no global discriminator. By increasing the size of the detailed-shape dataset, the vestige disappears and the back becomes more square-like, as shown in (g) (h) (k).

Gaussian filter. The effect of the preprocessing with the Gaussian filter blurring of the detailed shapes is shown in Table 3 and Figure 7. The larger σ is, the more blurry the training high-resolution voxels are, and the better the optimization goes. Without the Gaussian filter ($\sigma = 0.0$), the output looks like (a) *Recon. only* in Figure 5, indicating that

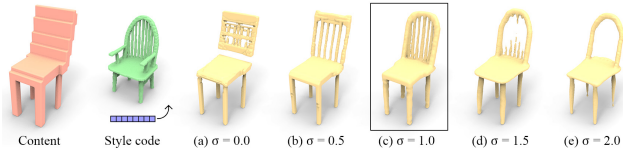


Figure 7: Ablation study on σ of the Gaussian filter. The input content shape and style code are shown on the left. The result with default parameters is placed inside a box.

	Loose-IOU \uparrow	LP-F-score \uparrow	Div-IOU \uparrow	Cls-score \downarrow
$\sigma = 0.0$	0.952	0.943	0.153	0.544
$\sigma = 0.5$	0.919	0.952	0.172	0.580
$\sigma = 1.0$	0.805	0.973	0.800	0.644
$\sigma = 1.5$	0.719	0.985	0.944	0.667
$\sigma = 2.0$	0.614	0.982	0.575	0.711

Table 3: Ablation on σ of the Gaussian filter.

the network may have reached a state where it is not easy to optimize a local patch to other styles, because the patches of different styles are not continuous for the optimization. As σ is increased, the output shape becomes more stylized, however above some threshold, the shape becomes too blurry to be reconstructed, especially thin structures. Progressive training, where the network is trained with larger σ then switches to smaller and smaller σ s during multiple steps, may work better.

4.5. GUI application

As an application of our method, we have created a GUI application where a user can explore the style space, to facilitate interactive design. Please refer to the supplementary.

4.6. GAN application

Since state-of-the-art 3D GANs are unable to generate detailed outputs, our method can be used directly to upsample a GAN-generated shape into a detailed shape, as long as the GAN output can be converted into a voxel grid. In Figures 8, we show results on upsampling shapes from IM-GAN [9]. See more results in the supplementary.

5. Conclusion, limitation, and future work

This paper introduces the first method to perform high-resolution conditional detail-generation for coarse 3D models represented as voxels. The coarse input enables control over the general shape of the output, while the input style code enables control over the type of details generated, which in tandem yield a tractable method for generating plausible variations of various objects, either driven by humans or via automatic randomized algorithms.

One main limitation is memory, similar to other voxel methods: a high-resolution voxel model, e.g., 256^3 , overflows GPU memory when upsampled to 1024^3 . We would like to explore more local networks to upsample patch-by-patch, or a recursive network. A second limitation is that



Figure 8: Upsampling GAN-generated shapes into detailed shapes. In the first row we show a sequence of generated shapes from IM-GAN via linearly interpolating two random latent codes. The last two rows show our upsampled results.

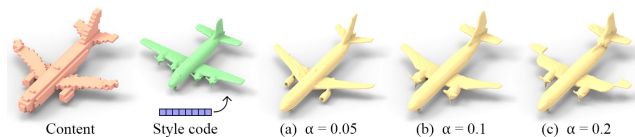


Figure 9: Sensitivity on parameter α . The input content shape and style code are shown on the left.

we mainly transfer local patches from the training shapes to the target. Therefore, we cannot generate unseen patterns, e.g., a group of slats on a chair back with a different frequency from those on training shapes, see last row of Figure 8. The network also lacks awareness to global structures and in some cases the generated shapes may present topological inconsistencies, see, e.g., second-to-last row of Figure 3. Lastly, as is often the case with GAN training, parameter tuning may be elaborate and fragile, see Figure 9.

Many immediate follow-ups suggest themselves. One would be to learn a meaningful, smooth, latent space, so that all latent codes will produce valid styles and latent space interpolation produce smooth transitions. Likewise, exploring hierarchies of details could enable more elaborate and consistent output. A complete decoupling of fine details, the coarse shape, and the semantic shape category is also interesting as it would enable training across larger collections of shapes, with the same styles employed across different shape categories. Lastly, we of course eye various advancements in voxel representation to reach higher resolutions.

We are excited about the future prospects of this work for detailization of 3D content. Immediate applications include amplification of stock data, image-guided 3D style generation, and enabling CAD-like edits that preserve fine detail.

Acknowledgements. We thank the anonymous reviewers for their valuable comments. This work was completed while the first author was carrying out an internship at Adobe Research; it is supported in part by an NSERC grant (No. 611370) and a gift fund from Adobe.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas J Guibas. Learning representations and generative models for 3D point clouds. In *ICML*, 2018. 1
- [2] Heli Ben-Hamu, Haggai Maron, Itay Kezurer, Gal Avineri, and Yaron Lipman. Multi-chart generative surface modeling. *ACM Trans. Graphics*, 2018. 3
- [3] Sema Berkiten, Maciej Halber, Justin Solomon, Chongyang Ma, Hao Li, and Szymon Rusinkiewicz. Learning detail transfer based on geometric features. In *Computer Graphics Forum*, 2017. 2
- [4] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016. 2
- [5] 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. 4
- [6] Siddhartha Chaudhuri, Daniel Ritchie, Jiajun Wu, Kai Xu, and Hao Zhang. Learning generative models of 3D structures. *Computer Graphics Forum (Eurographics STAR)*, 2020. 1
- [7] Xiaobai Chen, Thomas Funkhouser, Dan B Goldman, and Eli Shechtman. Non-parametric texture transfer using Mesh-Match. *Adobe Technical Report 2012-2*, 2012. 2
- [8] Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. BSP-NET: Generating compact meshes via binary space partitioning. In *CVPR*, 2020. 1
- [9] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 2, 4, 8
- [10] 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 *ECCV*, 2016. 2
- [11] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3D-encoder-predictor cnns and shape synthesis. In *CVPR*, 2017. 2
- [12] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. Curriculum DeepSDF. *ECCV*, 2020. 2
- [13] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *ICLR*, 2017. 3
- [14] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3D object reconstruction from a single image. In *CVPR*, 2017. 2
- [15] Noa Fish, Lilach Perry, Amit Bermano, and Daniel Cohen-Or. SketchPatch: Sketch stylization via seamless patch-level synthesis. In *SIGGRAPH Asia*, 2020. 3
- [16] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 3
- [17] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV*, 2019. 2
- [18] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A papier-mâché approach to learning 3D surface generation. In *CVPR*, 2018. 1, 2, 3
- [19] Christian Häne, Shubham Tulsiani, and Jitendra Malik. Hierarchical surface prediction for 3D object reconstruction. In *3DV*, 2017. 2
- [20] Rana Hanocka, Gal Metzger, Raja Giryes, and Daniel Cohen-Or. Point2mesh: A self-prior for deformable meshes. *ACM Trans. Graph.*, 2020. 2
- [21] Amir Hertz, Rana Hanocka, Raja Giryes, and Daniel Cohen-Or. Deep geometric texture synthesis. *ACM Trans. Graphics*, 2020. 1, 2
- [22] 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 *NeurIPS*, 2017. 5
- [23] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 3
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 3
- [25] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. 2
- [26] Jun Li, Kai Xu, Siddhartha Chaudhuri, Ersin Yumer, Hao Zhang, and Leonidas Guibas. GRASS: Generative recursive autoencoders for shape structures. *ACM Trans. Graphics*, 2017. 1
- [27] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NeurIPS*, 2017. 3
- [28] Hsueh-Ti Derek Liu, Vladimir G. Kim, Siddhartha Chaudhuri, Noam Aigerman, and Alec Jacobson. Neural subdivision. *ACM Trans. Graphics*, 2020. 1, 2
- [29] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH*, 1987. 5
- [30] Chongyang Ma, Haibin Huang, Alla Sheffer, Evangelos Kalogerakis, and Rui Wang. Analogy-driven 3D style transfer. *Computer Graphics Forum (Proc. Eurographics)*, 2014. 2
- [31] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. 4
- [32] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, 2019. 2
- [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. 2

- [34] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas Guibas. StructureNet: Hierarchical graph networks for 3D shape generation. *ACM Trans. Graphics*, 2019. 1
- [35] Charlie Nash, Yaroslav Ganin, SM Eslami, and Peter W Battaglia. Polygen: An autoregressive generative model of 3D meshes. *arXiv preprint arXiv:2002.10880*, 2020. 2
- [36] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 3
- [37] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei Efros, and Richard Zhang. Swapping auto-encoder for deep image manipulation. *NeurIPS*, 2020. 3
- [38] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. *ECCV*, 2020. 2
- [39] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017. 2
- [40] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3D representations at high resolutions. In *CVPR*, 2017. 2
- [41] Irina Sánchez and Verónica Vilaplana. Brain MRI super-resolution using 3d generative adversarial networks. In *First International Conference on Medical Imaging with Deep Learning, Amsterdam*, 2018. 2
- [42] Edward Smith, Scott Fujimoto, and David Meger. Multi-view silhouette and depth decomposition for high resolution 3D object representation. In *NeurIPS*, 2018. 2
- [43] Minhyuk Sung, Hao Su, Vladimir G. Kim, Siddhartha Chaudhuri, and Leonidas Guibas. ComplementMe: Weakly-supervised component suggestions for 3d modeling. *SIGGRAPH Asia*, 2017. 1
- [44] Kenshi Takayama, Ryan Schmidt, Karan Singh, Takeo Igarashi, Tamy Boubekeur, and Olga Sorkine-Hornung. GeoBrush: Interactive mesh geometry cloning. *Computer Graphics Forum (Proc. Eurographics)*, 2011. 2
- [45] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D mesh models from single RGB images. In *ECCV*, 2018. 2
- [46] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-CNN: Octree-based convolutional neural networks for 3D shape analysis. In *SIGGRAPH*, 2017. 1, 2
- [47] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *NeurIPS*, 2016. 1, 2, 3
- [48] Ruidi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen. PQ-NET: A generative part Seq2Seq network for 3D shapes. In *CVPR*, 2020. 1
- [49] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. DISN: Deep implicit surface network for high-quality single-view 3D reconstruction. In *NeurIPS*, 2019. 2
- [50] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. FoldingNet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, 2018. 2
- [51] Wang Yifan, Noam Aigerman, Vladimir G Kim, Siddhartha Chaudhuri, and Olga Sorkine-Hornung. Neural cages for detail-preserving 3D deformations. In *CVPR*, 2020. 3
- [52] Wang Yifan, Shihao Wu, Hui Huang, Daniel Cohen-Or, and Olga Sorkine-Hornung. Patch-based progressive 3D point set upsampling. In *CVPR*, 2019. 2
- [53] Kun Zhou, Xin Huang, Xi Wang, Yiyang Tong, Mathieu Desbrun, Baining Guo, and Heung-Yeung Shum. Mesh quilting for geometric texture synthesis. *ACM Trans. Graphics*, 2006. 2