Proceedings of the ASME 2024 19th International Manufacturing Science and Engineering Conference
MSEC2024
June 17 – June 21, 2024, Knoxville, Tennessee, USA

MSEC2024-125075

GENERATING OPTIMIZED 3D DESIGNS FOR MANUFACTURING USING A GUIDED VOXEL DIFFUSION MODEL

Premith Kumar Chilukuri[1], Binyang Song[1], SungKu Kang[2], Ran Jin[1,*]

[1]Virginia Tech, Blacksburg, Virginia
[2]Northeastern University, Boston, Massachusetts

ABSTRACT

In the age of digital manufacturing, efficient realization of high-performing designs into high-quality products through advanced manufacturing remains a challenge. One of the critical reasons for this challenge is the need for more design knowledge at the manufacturers' end, which severely hinders manufacturers' exploration and exploitation of innovative designs. Manufacturers often have no choice but to adopt designer-optimized designs without the ability to make enhancements. However, the advent of data-driven reverse engineering based on generative modeling has shown promise for addressing this issue. As industries increasingly adopt digital manufacturing strategies, integrating AI-assisted design generation and optimization tools could enable efficient production outcomes. By leveraging large datasets and powerful machine learning techniques, manufacturers can now explore the vast design space in an efficient, automated manner, enabling them to drive innovation by uncovering novel design opportunities. The expanded design knowledge gained through generative modeling allows manufacturers to move beyond just implementing others' designs and truly innovate themselves to create higher-quality and more optimized products. Nevertheless, while generative models have been widely explored for two-dimensional (2D) design and optimization, generative modeling of three-dimensional (3D) shapes still remains under-explored. Although a variety of 3D shape generative models have been developed, the resolutions and surface qualities of the generated designs can hardly meet manufacturing requirements. To bridge this gap, this paper proposes a generative denoising diffusion model (DDM) to generate new voxel-represented 3D designs, where the model is trained with historical 3D design data available to manufacturers. This research is important, not only because manufacturers' design knowledge can be improved, but also because design exploration can be expanded and expedited when more feasible designs can be generated faster to pursue designs with improved manufacturability, higher feasibility, and optimized performance. DDMs have been extensively explored to synthesize 2D pixel images, and also have achieved state-of-the-art generation quality. Motivated by the capabilities of DDMs in generating high-quality images, we adapt the application of DDMs from 2D space to 3D space and develop a voxel-DDM for high-quality 3D topology generation. Moreover, we develop a surrogate model for efficient and gradient-based evaluation of the manufacturability of 3D designs via additive manufacturing. We further integrate the surrogate model as a guidance module into the voxel-DDM to steer the denoising process toward generating new designs with optimized manufacturability. To showcase the proposed guided voxel-DDM and validate its effectiveness, we implement it through a real-world manufacturing case study to design a Microbial Fuel Cell (MFC) anode structure that enables an efficient, reliable, and high-quality manufacturing process. This research can help manufacturers overcome their lack of design knowledge in design space exploration and lead to high-quality designs for the manufacturing process in pursuit of optimized manufacturability.

Keywords: 3D Shape Generation, Design for Manufacturing, Generative AI for Design, Guided Denoising Diffusion

1. INTRODUCTION

Additive manufacturing (AM) has emerged as a highly disruptive technology that provides unprecedented design freedom and customizability for the production of end-use parts [1]. However, to realize the full potential of AM, it is necessary to overcome critical technical barriers, which currently limit the widespread adoption of AM across manufacturing industries. One major limitation that hampers the progress of AM is the lack of robust design tools and knowledge available to manufacturers to effectively utilize the capabilities of AM processes [2]. Currently, manufacturers are severely constrained in their ability to innovate AM designs and fully exploit the benefits of the technology. In most cases, manufacturers have no choice but to directly implement three-dimensional (3D) models provided by designers, without the ability to make further enhancements or refinements

*Corresponding author: jran5@vt.edu
Documentation for `asmeconf.cls`: Version 1.35, March 20, 2024.

themselves to improve manufacturability [3]. This overreliance on outside design expertise creates a bottleneck that restricts manufacturers' capacity to unlock the full innovation potential of AM through design optimization. There is a critical need for capabilities that empower manufacturers with expanded design knowledge and tools to take ownership of innovating optimized AM designs themselves.

The development of AI-powered generative design frameworks tailored for AM has the potential to overcome this barrier by providing manufacturers [4] with the expanded topology knowledge and automated exploration capabilities needed to take ownership of AM design innovation themselves. Democratizing robust design capabilities through AI assistants could greatly accelerate AM adoption across industrial sectors. By learning from large datasets of previous designs and simulations, generative models can rapidly explore extremely vast and complex design spaces to uncover novel high performance optimized designs for target manufacturing processes such as AM [4]. This advancement enables manufacturers not only to implement existing designs, but also to innovate and create higher quality and more optimized products. Specifically, generative modeling approaches leverage state-of-the-art image synthesis methodologies and have the potential to achieve high-fidelity 3D shape generation capabilities previously inaccessible to manufacturers.

Automated 3D shape synthesis [5] has immense potential to transform engineering design and manufacturing. 3D shape synthesis [6] can accelerate engineering design by automating high-fidelity 3D shape generation from abstract descriptions of representations such as text and sketches. This can expedite early design iterations through easier visualization and evaluation, as well as allow efficient design space exploration by significantly reducing the efforts needed for manual 3D modeling. Integrating AI tools to enhance manufacturability will be critical for manufacturers to keep pace with AM's rapid innovation and unlock its design freedom potential.

The key knowledge gap that drives this research is the lack of generative modeling for high-fidelity 3D shape generation and design optimization driven by manufacturability [7]. Although generative models have shown increasing success in 2D image synthesis [8], their extension to 3D design remains underexplored despite the massive potential impact. Existing 3D generative models [9] rely on voxel or point cloud representations, which impose limitations on resolution, surface quality, and manufacturability guidance. There is an unmet need for generative models producing manufacturing grade, printability-optimized 3D geometry. Most 3D design optimization [9] still relies on manual CAD modeling iterations requiring extensive expertise and trial and error. Automating intelligent 3D design space exploration with AI-guided generative models remains an open opportunity. Developing generative frameworks specifically designed for manufacturing could transform prototyping and innovation. Another gap is in incorporating manufacturability considerations into the design process. Generated designs are often theoretically sound but might be impractical due to manufacturing constraints. Integrating AI into 3D design also faces computational and data challenges. This research aims to develop rigorous generative models for high-fidelity, manufacturability-optimized 3D shape generation.

This paper aims to address knowledge gaps by introducing a generative modeling framework Figure 2 tailored for AM design. Innovation in this paper is the adaptation of state-of-the-art denoising diffusion models (DDMs) [10] for image synthesis from the 2D space to the 3D space for high-quality 3D shape generation. We develop a voxel-DDM (voxel-DDM) using a dataset of 3D designs. The model employs the 3D adaption of U-Net to handle voxel input and a factorized attention mechanism to learn both intra-layer and interlayer features for 3D generative modeling. The model is further enhanced by a regression guidance model [11] that is adapted from a pre-trained surrogate model. The surrogate model can provide efficient manufacturability evaluations to steer the denoising sampling process toward highly manufacturable designs by estimating a gradient to guide the sampling process. To validate the ability of the proposed framework to expand manufacturers' design knowledge and improve manufacturability, a case study is conducted to generate new 3D topologies of a Microbial Fuel Cell (MFC) anode [12] for an AM process. By learning from historical data, realistic 3D designs are generated. This strategy helps close the high-fidelity 3D design knowledge gap faced by manufacturers. The proposed guided voxel-DDM is tailored for design for AM and represents a significant advancement in the field of 3D generative modeling using voxel data.

The key technical contributions of this paper include:

- A labeled 3D topology dataset supporting 3D surrogate modeling and generative modeling.
- A voxel denoising diffusion framework Figure 2 that adapts state-of-the-art image synthesis techniques from 2D space to 3D space for high-fidelity 3D shape generation.
- A surrogate model for fast and accurate prediction of manufacturability to guide the denoising process of the proposed Voxel-DDMs.
- Manufacturability-driven generative guidance based on a pre-trained surrogate model for manufacturability optimization during the generation process.
- Expanded design knowledge and improved manufacturability of 3D MFC anode structures for AM.

The rest of the paper is organized as follows: Related Work, Methodology, Results and Discussion, and Conclusions.

2. LITERATURE REVIEW

3D representations are critical in engineering design and manufacturing due to their faithful portrayal of shapes with high-level details, particularly in the middle to late design stages for design evaluation, optimization, and manufacturing. 3D shape synthesis, due to the complexity of 3D representations, currently stands as a significant challenge. Given that shapes can be represented by structured voxels [13], unstructured meshes [14] and point clouds [15], parametric representations [16], and implicit representations [17, 18], each representation enables the use of different generative models to generate new designs [19].

Voxel representation, the counterpart of 2D pixel data in the 3D domain, is commonly used for 3D shape generation. Leveraging the success of convolutional neural network (CNN) architectures in 2D image synthesis, various generative models exploit 3D CNN architectures to synthesize 3D shapes represented by voxels. Among them, 3D CNN-based multimodal auto-encoder (AE) models aim at 3D shape reconstruction by transforming one or more 2D views (with depth information) into 3D shapes [20, 21]. 3D CNN-based variational AEs (VAEs) [22] and generative adversarial networks (GANs) [23, 24] can generate new 3D shapes randomly or conditioned on textual and visual descriptions. Due to their high computational costs caused by dimension increase, these 3D CNN-based models often suffer from low generation resolution and low training speed. Accordingly, multiple-stage models have been explored to increase the resolution of the final output [25].

Compared to voxels, point clouds provide a more accurate 3D shape representation. Their high representation efficiency, compactness, and ease of acquisition attract intensive research interests for 3D shape generation tasks [26]. Since point clouds lack the characteristics of spatial order and arrangement, order-invariant operations such as max-pooling as in PointNet [27] or multiple transformation matrices before the convolution operator [28] are needed to compensate [26]. On this basis, various models founded on AE [29], VAE [30], and GAN [31] have been developed using fully connected layers, graph neural networks (GNNs) [32], or their variants with attention mechanisms [30] as decoders or generators for image or text guided point cloud generation. ShaperCrafter employs a different architecture, using a vector-quantized VAE as an encoder and a transformer-based autoregressive model as a decoder for text-guided 3D shape synthesis [33]. In recent years, denoising diffusion models (DDMs) have been drawing increasing attention for the generation of unconditional and text- or image-guided 3D shapes [34].

Meshes are another accurate, unstructured, 3D shape representation. In addition to vertex (i.e. point) coordinates, meshes further convey local face information of 3D shapes to record topological connections between mesh vertices, providing a higher quality 3D shape representation than a point cloud [26]. To effectively learn mesh features, various architectures such as GNNs [35] and their variants (e.g., graph convolutional networks [36]) have been developed by defining basic deep learning operators (e.g., convolution, pooling) based on specialized mesh structures. Similarly to 3D CNN to voxel representations, GNNs are the basic building blocks of a set of AE [37], VAE [38], and GAN [39] models for text or image-guided mesh generation. Another group of models aims to predict the displacement of each vertex to deform a given 3D mesh to generate new meshes [40].

Unlike the explicit representations reviewed above, implicit representations, formulated as implicit functions such as the occupancy function [41] or the signed distance function [17], can be used to infer the boundary of a 3D shape by predicting a value for each sampled point, which indicates that this point is inside or outside the boundary of the 3D shape. These implicit representations have been integrated with various embedding modules to learn latent representations of 3D shapes learned from images, text descriptions, or shapes [17], which are input to predict 3D shape boundaries. On this basis, different AE [42], VAE [43], and GAN [18] models have been developed for conditional shape generation. Additionally, CLIP-Forge marries an implicit representation with a flow-based model for the generation of 3D shapes guided by text or images [44].

Additionally, another strand of models maps 3D shapes to 2D parameter spaces, generates 2D samples using GANs, and then converts them into 3D meshes [45, 46]. Among them, Rank3DGAN [45] and XDGAN [47] extend this approach to conditional settings, allowing semantic manipulation during 3D shape generation.

Although a variety of 3D shape-generative models have been developed employing explicit and implicit 3D shape representations, these approaches share two common drawbacks. First, the resolution and quality of the 3D shapes generated by these models are generally low. This is caused by the characteristics of the representations. While the structural arrangement of voxel representations is favored by deep learning models, the increase in dimensionality from two to three significantly raises the required computational cost, which limits the representation resolution. Point-cloud and mesh representations consume less memory and computational cost than voxel representations of the same resolution because they only convey the boundary information rather than the entire volume. However, the irregularity of point clouds and meshes and the complexity of the mesh topology make it difficult to use for deep learning [26]. Continuous efforts are needed to explore ways of learning latent vector embeddings of point-cloud and mesh representations that can be used for effective 3D shape synthesis. Compared to all others, on the one hand, implicit representations allow for flexible shape topology and a continuous increase in representation resolution with reasonable memory consumption. On the other hand, an implicit representation, as the generator output of a generative model, does not convey the geometric features of the generated 3D shape in a straightforward manner. A preprocessing module is needed to convert the implicit representation into an explicit 3D shape for downstream tasks, which introduces additional loss and impairs the quality of the 3D shape.

Second, although conditional models can generate 3D shapes according to semantic or visual prompts [34], these models do not support the evaluation and optimization of engineering performance during the 3D shape generation process [48]. The performances of 3D shapes, such as manufacturability and strength, are critical in engineering applications. However, these requirements and constraints are not considered by the models developed in other domains. To mitigate these two issues, we propose a new 3D shape generative model to synthesize 3D MFC anode designs for manufacturing. We get inspiration from some recent papers that extend DDMs from 2D space to 3D space for the generation of 3D spatiotemporal video [49].

3. METHODOLOGY

In this paper, we develop a generative modeling framework, voxel-DDM, Figure 2 that allows the optimization of the manufacturing features tailored to the design of AM. This framework leverages the techniques of DDMs [10], a class of generative models renowned for their high-quality outputs, and adapts the

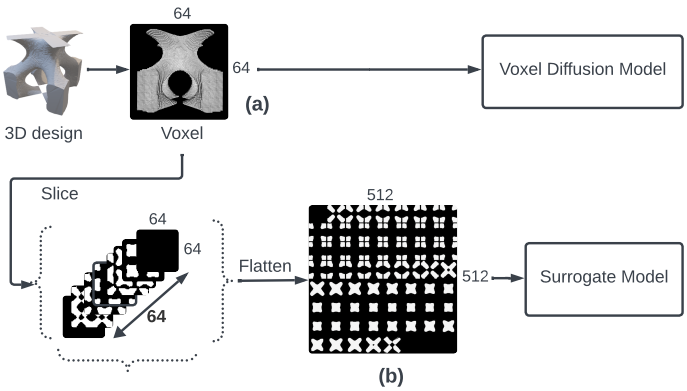

FIGURE 1: 3D Design data preprocessing to prepare data for training the voxel denoising diffusion model and surrogate model, respectively.

DDM architecture from image generation to voxel generation using a 3D U-Net [49], contributing a novel approach in the realm of 3D shape generation. Voxel-DDM is further enhanced with a manufacturability metric optimization guidance module, which is adapted from a surrogate model developed to predict manufacturability metrics and incorporated to optimize the metric during the generation process. We demonstrate and validate guided voxel-DDM using a case study of the 3D MFC design [12].

3.1 Dataset

The proposed guided voxel-DDM is trained and validated in a set of 2,735 3D MFC designs generated from a previous study [12]. As the proposed research is an exploratory research to propose AI-based design generation tools, the 2,735 3D MFC anode designs in our dataset are simplified, consisting only of designs with cell counts of 1 and 2. In the future, we will expand our dataset by including anode structures with larger cell counts. The original designs in the dataset are presented in STL format and labeled with a manufacturability metric, "minimum feature size" (MFS) (measured in millimeters, 'mm') of the MFC anode structure [12], which is obtained from simulations. The MFS is the smallest possible 3D printable feature of a 3D model, typically determined by the diameter of the nozzle and the 3D printer motors, which is 2 mm in our case study [12]. In this study, we wanted to maximize the MFS to improve the manufacturability of MFC anode designs via AM. We prepare the dataset for training and validating the guided voxel diffusion model by converting each input 3D design from STL format into a voxel representation with a dimension of $64 \times 64 \times 64$ (x, y, z; voxel-size of 0.2mm), as shown in Figure 1-(a).

To take advantage of advanced computer vision algorithms for the surrogate modeling of manufacturability metrics, we organize the voxel representation in a different way. The 3D design in STL format is converted into voxel representation, which is sliced 64 times along the Z-axis, resulting in 64 layers of 2D slices (64×64). Now, these 2D slices are arranged in a grid format with 8 rows and 8 columns (8×8, totaling 64)), forming a 512×512 image ($8 \times 64 (= 512)$) as shown in Figure 1-(b). Each image in our dataset is labeled with the corresponding manufacturability metric to train the surrogate model [50].

3.2 Voxel Denoising Diffusion Model

Similar to general DDMs [10, 49], the proposed voxel-DDM consists of two distinct phases: the forward diffusion process and the reverse denoising process, as shown in Figure 2. During the forward diffusion process, the model incrementally introduces Gaussian noise into the input voxel data $z_0 \sim q(z_0)$, transforming the data distribution into a Gaussian noise distribution $z_T \sim \mathcal{N}$ (T is the number of steps of the diffusion process, which is 400 in our study). This is a Markovian process in which each step introduces a controlled amount of noise to the data, as expressed in the equation $q(z_t|z_{t-1}) := \mathcal{N}(z_t; \sqrt{1-\beta_t} z_{t-1}, \beta_t I)$. Here, z_t represents the latent of the same dimension as z_0 at time t, and β_t is the variance of the noise added at the time step t, which controls the step size. The variance schedule $(\beta_1, ...\beta_t)$ is meticulously chosen to ensure a smooth transition from the data distribution to the noise distribution.

Conversely, the reverse denoising process [49], where the generative power of the model is manifested, involves iterative removal of Gaussian noise from the data to reconstruct the input data from its latent state, i.e., noised state. This reverse process is governed by $p_\theta(z_{t-1}|z_t) := \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t), \Sigma_\theta(z_t, t))$. In this equation, $\mu_\theta(z_t, t)$ and $\Sigma_\theta(z_t, t)$ are the predicted mean and variance at time step $t-1$by the model. The success of this reverse process is critical, as it directly impacts the fidelity and accuracy of the generated voxel data.

Similar to the adaptation of CNNs from 2D to 3D, the proposed diffusion model exploits a sophisticated adaptation of the U-Net framework [49] from 2D to 3D to handle voxel data [51]. The 3D U-Net model is trained to predict the mean $\mu_\theta(z_t, t)$ and variance $\Sigma_\theta(z_t, t)$ at each time step using the last sampled latent z_t as input in the reverse denoising process. This architecture, with its downsampling and upsampling paths connected by skip connections, excels in retaining both local and global information in voxel frames, a crucial feature for accurately representing 3D structures.

Training the model involves optimizing the parameters θ of the 3D U-Net [49] that define $\mu_\theta(z_t, t)$ and $\Sigma_\theta(z_t, t)$. The objective is to minimize the difference between the input data and the data reconstructed by the model through the reverse denoising process. The loss function used for training is formulated as $L(\theta) = \mathbb{E}_{t,x,\epsilon}[||\epsilon - \epsilon_\theta(z_t, t)||^2]$. Here, ϵ is the noise added during the forward process, and $\epsilon_\theta(z_t, t)$ is the noise removed during the reverse denoising process. This loss function guides the model to become better at predicting the noise added at each step of the forward process, thereby enabling it to effectively reverse this process. The proposed model's architecture is further refined with a factorized attention mechanism [49], including an intra-layer attention block and an inter-layer attention block. These attention blocks facilitate the learning of complex layer-wise and cross-layer topological features, respectively.

Once the voxel-DDM is well trained, we can apply the denoising module as a generator to sample detailed and coherent new designs. The denoising module takes a random sample drawn from the Gaussian distribution $\mathcal{N}$ as input and applies the learned reverse denoising process iteratively. With each iteration,

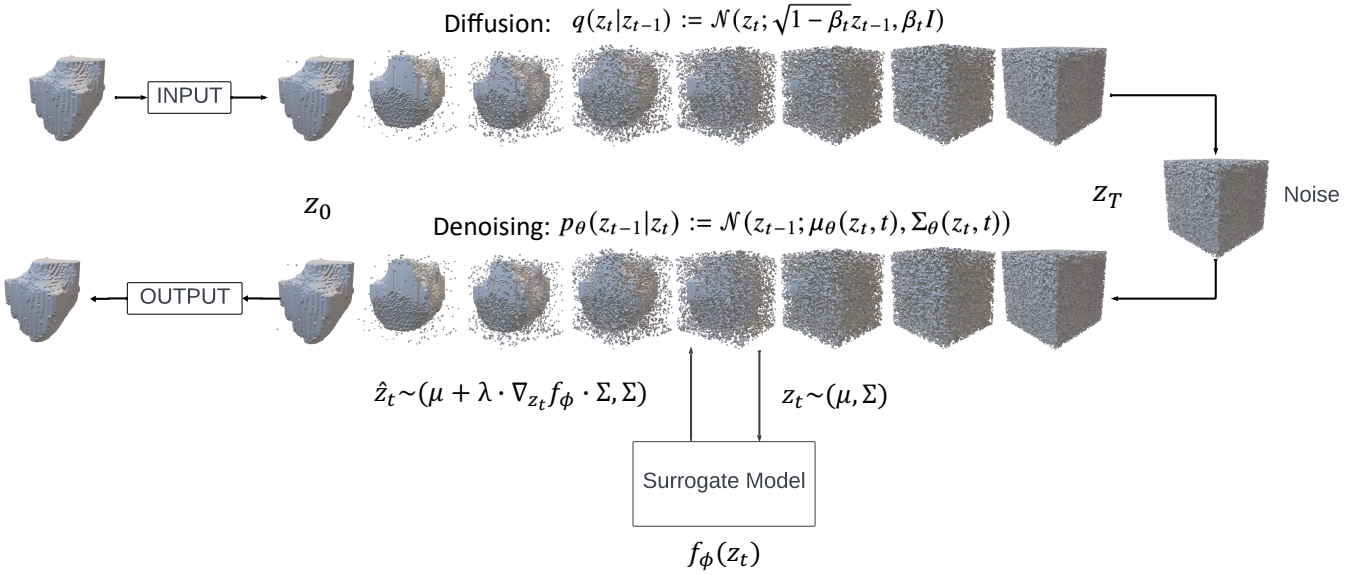

FIGURE 2: The proposed guided voxel-DDM that integrate the vanilla voxel-DDM with the surrogate model. In this study, the voxel-DDM is trained in 16K epochs with 400 sampling timesteps.

the model refines the sample, reducing its noise and bringing it closer to the distribution of the training data. This iterative refinement is the crux of the sampling process. The factorized attention mechanism employed in this model can further enhance the quality of the sample. This model not only showcases the potential of AI in manufacturing design but also offers a path for creating complex designs that were previously challenging to conceptualize and materialize.

3.3 Surrogate Model

In the proposed research, the surrogate model plays an integral role in the generative framework (Figure 2) for the design of the manufacturing. The surrogate model, a modified version of the ResNet-50 [50] architecture, is specifically tailored to predict the manufacturability metric, MFS, of the MFC anode designs [12] represented as 521×512 images. We modify the architecture of ResNet-50 to suit our data. The model initially uses the pre-trained ResNet-50 as a feature extractor, where the parameters of these layers are frozen to retain the knowledge they have already acquired. The final fully connected layers of ResNet are replaced with additional linear layers. This sequence consists of several blocks of linear layers interspersed with Rectified Linear Unit (ReLU) activation functions. Specifically, it includes a transformation from the original feature space to a 512-dimensional space, followed by subsequent reductions to 128, 64, and finally to an 8-dimensional space, each step accompanied by ReLU activation for non-linearity. The final layer is a linear layer that maps the 8-dimensional space to a 1D output value (the predicted metric). Table 1 summarizes the architecture of the surrogate model.

The MFS values of the MFC anode designs in our dataset range from 1.51 mm to 15.97 mm. To facilitate the training process, we normalized the values of the MFS to the range [0,1]. The mean squared error (MSE) is used as a loss function, which

TABLE 1: Architecture of the modified ResNet-50 surrogate model

Layer Type	Parameters	Output Dim
Input from ResNet-50	-	2048 features
Linear + ReLU	$2048 \rightarrow 512$	512 nodes
Linear + ReLU	$512 \rightarrow 128$	128 nodes
Linear + ReLU	$128 \rightarrow 64$	64 nodes
Linear + ReLU	$64 \rightarrow 8$	8 nodes
Linear + ReLU (Output)	$8 \rightarrow 1$	1 node

is a standard choice for regression tasks. The MSE loss [52] is defined as:

$$\text{MSE Loss} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

. Herein, y_i is the true MFS and $\hat{y}_i$ is the predicted value. An Adam optimizer is used for training, with a learning rate (lr) initially set to 0.001. The lr is adjusted during training using an lr scheduler: reduces the lr by 10% for every 7 epochs. This approach helps to refine the model weights more precisely as training progresses. The model is trained and validated over multiple folds (3 in this case) to ensure the generalizability and robustness of the trained model. Since the trained surrogate will be used to guide the denoising process, the intermediate noisy output from the denoising phase of the voxel-DDMs will be fed to the surrogate model to calculate the gradient for guidance based on the predicted manufacturability metric.

3.4 Guidance Module

In order to evaluate and optimize the MFS values of the MFC anode structures during the generation process, we integrate the trained surrogate model into the denoising module (Figure 2) as a guidance module. We refer to the voxel-DDM with the

guidance module as the "guided voxel-DDM". This is achieved through a regression guidance mechanism [53], which shifts the mean predicted by the voxel-DDM by $\Sigma \cdot \nabla_{x_t} f_\phi(x_t)$. In this paper, Σ is the variance of the Gaussian distribution representing $p_\theta(z_t|z_{t+1})$, $f_\phi(x_t)$ is the surrogate neural network predicting the normalized MFS values of the MFC anode designs, and ∇_{x_t} is the gradient with respect to x_t. Consequently, the updated mean is formulated as $\hat{\mu}_\theta(z_t, t) = \mu_\theta(z_t, t) + \lambda \cdot \Sigma \cdot \nabla_{x_t} f_\phi(x_t)$, where λ is the gradient scale hyperparameter defining the weight of the guidance module. This method modifies the predicted distribution from which we sample at each step by favoring designs with large MFS values. In the early phase of the denoising process, x_t is too noisy to be fed as input to the surrogate model Figure 7 to predict the metric accurately, leading to meaningless guidance to the denoising module. Consequently, we determine a maximum level of noise beyond which the regression guidance should not be included. In our case, the guidance module is only applied when t < 200, which is selected through a few pilot studies.

Integrating the surrogate model as the guidance module of voxel-DDM is crucial (Figure 2). This integration enables the voxel-DDM to generate designs that are not only novel and high-quality but also align with desired manufacturability metrics. During the denoising process of the voxel-DDM, the surrogate model provides real-time feedback on the predicted metric of the generated design. This feedback is then used to guide the denoising module to generate optimized designs, steering it towards designs with larger MFS values. The surrogate model effectively acts as an evaluator and optimizer, influencing the trajectory of the generation process in the latent space. By favoring generated designs with good features through gradient manipulations, the surrogate model biases the voxel-DDM to explore regions in the latent space that correspond to more manufacturable designs.

4. RESULTS AND DISCUSSION

This section presents the experimental results validating the proposed guided voxel-DDM for generating optimized 3D MFC anode designs [12] for AM. MFC anodes are critical components that require precise topological attributes to ensure optimal performance. These attributes include MFS, minimum cavity size, connectivity of features/cavities, high surface area to volume ratio, and electrical resistance [12]. The design of these anodes is often constrained by traditional manufacturing methods, which limit the ability to experiment with complex geometries that may enhance manufacturability metrics. In this paper, we only consider one manufacturability metric, which is the MFS, which can be used to evaluate the manufacturability of a design through AM. This case study offers us an exploration of the applicability of the proposed voxel-DDM in a real-world manufacturing context. In the following, we will present the generation capability of the vanilla voxel-DDM, the performance of the surrogate model [50], and the effectiveness of the guided voxel-DDM (Figure 2).

4.1 Quality of Designs Generated by Voxel Denoising Diffusion Model

The proposed voxel-DDM (Figure 2) is trained on a dataset of 2,735 3D MFC anode designs. The hyperparameters of the voxel-DDM are as follows: lr - 0.0001, input dimension - $64 \times 64 \times 64$, loss function - MSE, sampling time steps - 400, training epochs - 16,000, batch size - 2 and ema decay - 0.995. Figure 3 shows some designs generated by the vanilla voxel-DDM. By visually inspecting the generated designs, we can see that the model is able to generate a wide diversity of novel 3D MFC anode topologies by sampling the learned latent space. The quality of the generated designs is assessed through the Fréchet Inception Distance (FID) score and the Inception Score (IS) computed against the test data distribution. Generally, diffusion models suffer from unstable performances, so, to evaluate the performance of the proposed Voxel-DDM, we performed a reproduction test by generating a test dataset which consists of 100 designs generated by the guided voxel-DDM and another 100 designs generated by vanilla voxel-DDM using the same noise seeds. The normalized MFS values of the generated designs were evaluated following the data preprocessing pipeline introduced in Section 3.1). FID measures [54] the similarity of the feature distributions between generated and real samples, with lower values indicating greater fidelity. IS evaluates [54] both image quality and diversity, with the higher being better. Table 2 reports the average FID score, the average IS, and the average normalized MFS values of 100 random designs generated by the vanilla (unguided) voxel-DDM.

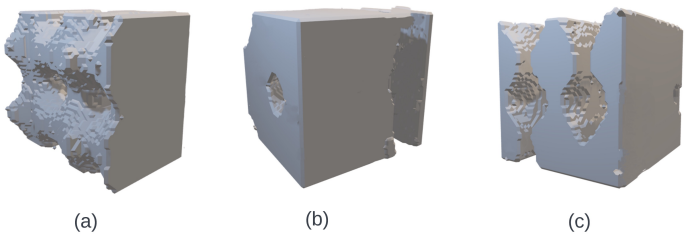

FIGURE 3: 3D designs generated by the proposed voxel-DDM model

TABLE 2: Performance comparison between the guided and vanilla (unguided) voxel-DDMs in terms of average Fréchet inception distance (FID) score, Inception Scores (IS), and average normalized minimum feature size (MFS)

	FID	IS	Normalized MFS
Vanilla	556.58	1.23 ± 0.0	0.54
Guided	572.01	1.17 ± 0.0	0.63

The FID score [54] provides an assessment of the similarity between the generated designs and the real distributions of the design data. The inception score (IS) [55] evaluates both image quality and diversity, with the highest being better. A lower FID score indicates a closer resemblance to real data, which is often associated with higher-quality designs. In our study, the vanilla voxel-DDM achieves an average FID score of 556.58 and an average IS of 1.23. This FID score is competitive compared to the FID scores achieved by other existing 3D shape generative models [49], suggesting that the proposed voxel-DDM model is capable of synthesizing designs that are statistically aligned with real data. The Inception Score (IS) complements the FID by evaluating the clarity and diversity of the designs.

A deeper analysis of these metrics reveals their inherent lim-

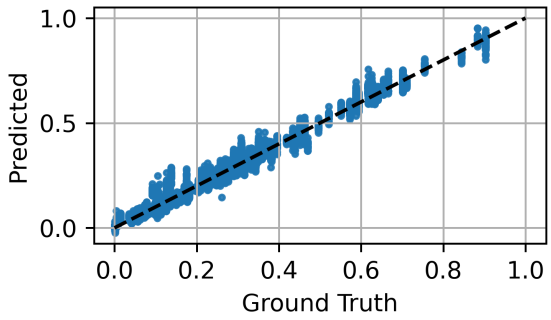

FIGURE 4: The comparison between the ground truth and the predicted normalized MFS values of the 3D MFC anode designs in our validation dataset.

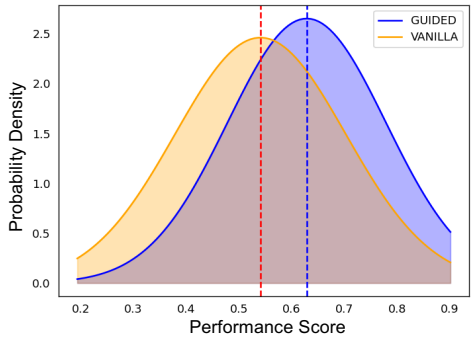

FIGURE 5: Normalized minimum feature size distributions of 100 random designs respectively generated by the guided and vanilla (unguided) voxel-DDMs

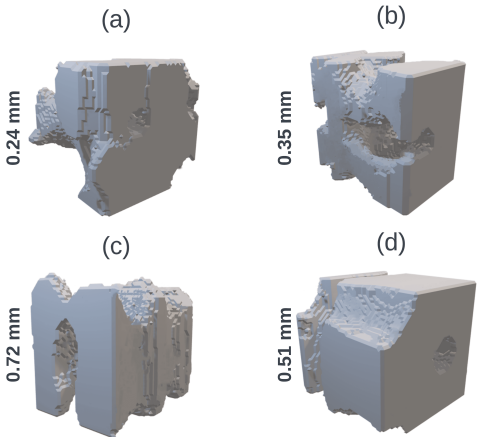

FIGURE 6: Two pairs of designs respectively generated by the guided and vanilla (unguided) voxel-DDMs using two same noise inputs along with their corresponding normalized MFS values. Top - Vanilla voxel-DDM and bottom - guided voxel-DDM.

itations and the complexities involved in assessing the quality of the design. The FID score [54], while insightful, does not account for the functional aspects of the designs. Therefore, our evaluation extends beyond this score to include one manufacturability metric predicted by the surrogate model, which provides a direct indicator of each design's manufacturability via AM. In general, the generated design exhibits an average normalized MFS of 0.54 for vanilla voxel-DDM and 0.63 for the guided voxel-DDM. Quantitative metrics validate the ability of the proposed voxel-DDM to produce novel high-fidelity anode designs on par with existing ones (ground truth (ie, non-optimized designs w.r.t. minimum anode feature size) average normalized MFS = 0.328) [12].

4.2 Performance of the Surrogate Model

The surrogate model, leveraging a modified ResNet-50 architecture [50], exhibits impressive predictive performance in evaluating the manufacturability metric, normalized MFS. Through rigorous training and validation, the model achieves a training loss of 0.002 and a validation loss of 0.0015, with a high r^2 score of 0.96, denoting the variance explained by the model. These metrics suggest that the surrogate model can make highly accurate predictions, as evidenced by the strong linear correlation between the predicted and ground truth values in Figure 4.

4.3 Manufacturability Metrics Optimization through Guidance Module

The guidance module integrated into the voxel-DDM plays a key role in aligning the generative design process with MFS optimization. It utilizes real-time feedback from the surrogate model during the denoising process, guiding the generation of new designs towards optimal solutions. Based on the evaluation test dataset (see section 4.1), the MFS distributions of the 3D anode designs, respectively, generated by the guided and vanilla voxel-DDMs are illustrated in Figure 5. In the evaluation test dataset, the 100 random designs generated by the guided model have a larger MFS mean value than the MFS mean value of 100 random designs generated by the vanilla model ($mean_{guided}$ = 0.63, $mean_{vanilla}$ = 0.54). Furthermore, during the pairwise Student's t test, guided voxel-DDM achieved: $p < 0.000106$, and T-statistics = 3.954. Figure 6 shows two pairs of 3D anode designs, respectively, generated by the vanilla (top) and guided (bottom) voxel-DDMs using the same noise input. The normalized MFS values of the generated designs were estimated using the trained surrogate model. The top designs (a) and (b) generated by the vanilla voxel-DDM have normalized MFS values of 0.24 and 0.35, respectively; while that of the designs (c) and (d) generated by the guided Voxel-DDM are 0.72 and 0.51. This comparison illustrates that the designs on the bottom have larger normalized MFS values compared to those on the top, suggesting the effectiveness of the guided voxel-DDM in consistently generating better designs achieved compared to the vanilla model. This comparison highlights the transformative impact of the module.

We further compare the denoising process of the vanilla and guided voxel-DDMs. As illustrated in Figure 7, the design patterns of the generated designs emerge at earlier stages during the guided denoising process, while the patterns appear later during the vanilla denoising process. This suggests that the surrogate model is manipulating the gradient to generate the designs at an early stage during denoising and then further optimize it to maximize the minimum feature size.

Furthermore, the designs generated by the guided voxel-DDM achieve an average FID score [54] of 572.01 and an average IS of 1.17 [54], while the vanilla model scores 556.58 and

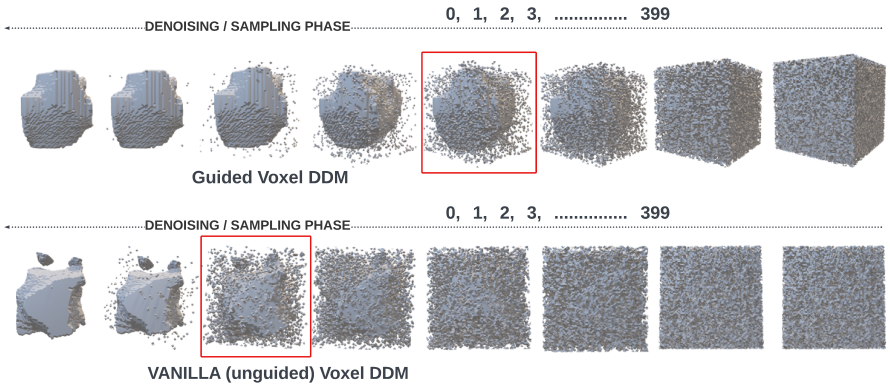

FIGURE 7: Impact of gradient manipulation of the surrogate model on Denoising process of the guided voxel-DDMs.

an average IS of 1.23, as reported in Table 2. These scores reflect a nuanced trade-off between the models: the guided model, driven by metric optimization, produces designs with slightly lower quality and diversity, while the unguided model, free from such constraints, generates a slightly broader range of designs with slightly higher quality. These findings in this section also underscore the effectiveness of the surrogate model in evaluating design features critical to manufacturability and its potential to significantly impact the field of generative design for manufacturing.

The integration of the guidance module makes voxel-DDM acutely responsive to the predictive insights of the surrogate model, ensuring that the generated designs meet not only the complexity and detail requirements, but also the predefined manufacturing metrics or constraints. As a navigational tool, the guidance module directs the design generation towards promising areas within the vast design space, specifically targeting MFC anode structures with higher manufacturability. This guided voxel-DDM empowered by the surrogate model allows for a more concentrated exploration of the design landscape, especially beneficial in the realm of MFC anode design where multiple manufacturability metrics need to be carefully balanced. The improved manufacturability metric achieved by the guided model underscores the practicality and advantages of incorporating data-driven guidance into the design process. The case study also highlights the framework's ability to expand the knowledge of viable anode design spaces.

4.4 Limitations and Future Directions

While the proposed guided voxel-DDM framework achieves promising results, this study has a few limitations and leaves room for further improvement. First, the dataset used in this study is relatively small and needs greater variety, which may limit the ability of the guided voxel-DDM to explore the solution space in a more broad way. Future studies should aim to collect larger datasets and retrain the generative model accordingly. Second, in this study, only one manufacturability metric, the minimum feature size (measured in millimeters 'mm'), is considered to demonstrate the proposed exploratory framework, leaving other metrics unexplored, such as the minimum cavity size, anode connectivity, anode resistance, and surface area. We will incorporate more metrics in the future to assess the manufacturability, feasibility, and performance of the generated designs more comprehensively. Third, the proposed framework has not been compared with other benchmark models in a quantitative way. Various experiments will be conducted to compare the effectiveness of the proposed model with other 3D shape generative models in terms of generation quality and performance optimization in the future. Additionally, we will seek to adapt the developed framework for various conditional generation scenarios, such as text-, image-, or 3D-guided generation. These new features will enable more controlled design refinements in the future. Finally, diffusion models, which rely on Gaussian noise for a predictable progression towards a noise distribution, face challenges when introduced to non-Gaussian noise, affecting model training, inference, and the quality of generated outputs. In future work, adapting to non-Gaussian noise involves modifications to the model's architecture and loss function, requiring technical adjustments and empirical evaluation to maintain or improve performance.

5. CONCLUSION

In this paper, we develop a guided voxel denoising diffusion model (voxel-DDM) that adapts state-of-the-art generative modeling techniques to achieve high-fidelity 3D shape generation optimized for additive manufacturing (AM). The key innovation is to adapt a denoising diffusion model (DDM) from the 2D space to the 3D space to generate intricate 3D topologies. Manufacturability guidance is incorporated through a pretrained surrogate model to steer the denoising sampling process toward designs with higher manufacturability. Our work fills the knowledge gap on generative design manufacturability by incorporating a pre-trained surrogate model to guide the denoising process to generate better manufacturable designs. Experiments on the design of Microbial Fuel Cell anode structures demonstrate the ability of the framework to expand design knowledge, improve manufacturability, and unlock greater innovation potential. By combining data-driven generative modeling with manufacturing design, this methodology enables the discovery of novel, high-quality 3D designs tailored for target manufacturing processes.

ACKNOWLEDGMENTS

Premith Kumar Chilukuri's and Ran Jin's efforts are supported by NSF CMMI-2208864. The authors thank Dr. Giorgio Giannone and Mayuranath Sureshkumar for discussing and sharing research ideas with the second author.

REFERENCES

[1] Tofail, Syed AM, Koumoulos, Elias P, Bandyopadhyay, Amit, Bose, Susmita, O'Donoghue, Lisa and Charitidis, Costas. "Additive manufacturing: scientific and technological challenges, market uptake and opportunities." *Materials today* Vol. 21 No. 1 (2018): pp. 22–37.

[2] Mehrpouya, Mehrshad, Dehghanghadikolaei, Amir, Fotovvati, Behzad, Vosooghnia, Alireza, Emamian, Sattar S and Gisario, Annamaria. "The potential of additive manufacturing in the smart factory industrial 4.0: A review." *Applied Sciences* Vol. 9 No. 18 (2019): p. 3865.

[3] Patterson, Albert E. and Allison, James T. "Mapping and Enforcement of Minimally Restrictive Manufacturability Constraints in Mechanical Design." *ASME Open Journal of Engineering* Vol. 1 (2022): p. 014502.

[4] Liu, Chenang, Tian, Wenmeng and Kan, Chen. "When AI meets additive manufacturing: Challenges and emerging opportunities for human-centered products development." *Journal of Manufacturing Systems* Vol. 64 (2022): pp. 648–656.

[5] Luo, Andrew, Li, Tianqin, Zhang, Wen-Hao and Lee, Tai Sing. "SurfGen: Adversarial 3D Shape Synthesis with Explicit Surface Discriminators." *CoRR* Vol. abs/2201.00112 (2022).

[6] Zhao, Zibo, Liu, Wen, Chen, Xin, Zeng, Xianfang, Wang, Rui, Cheng, Pei, Fu, Bin, Chen, Tao, Yu, Gang and Gao, Shenghua. "Michelangelo: Conditional 3D Shape Generation based on Shape-Image-Text Aligned Latent Representation." *arXiv preprint arXiv:2306.17115* (2023).

[7] Regenwetter, Lyle, Nobari, Amin Heyrani and Ahmed, Faez. "Deep Generative Models in Engineering Design: A Review." *CoRR* Vol. abs/2110.10863 (2021).

[8] Wu, Xian, Xu, Kun and Hall, Peter. "A survey of image synthesis and editing with generative adversarial networks." *Tsinghua Science and Technology* Vol. 22 (2017): pp. 660–674.

[9] Shi, Zifan, Peng, Sida, Xu, Yinghao, Geiger, Andreas, Liao, Yiyi and Shen, Yujun. "Deep Generative Models on 3D Representations: A Survey." (2023).

[10] Kulikov, Vladimir, Yadin, Shahar, Kleiner, Matan and Michaeli, Tomer. "SinDDM: A Single Image Denoising Diffusion Model." *Proceedings of the 40th International Conference on Machine Learning*. 2023. JMLR.org.

[11] Arechiga, Nikos, Permenter, Frank, Song, Binyang and Yuan, Chenyang. "Drag-guided diffusion models for vehicle image generation." (2023).

[12] Kang, SungKu, Deng, Xinwei and Jin, Ran. "A Cost-Efficient Data-Driven Approach to Design Space Exploration for Personalized Geometric Design in Additive Manufacturing." *Journal of Computing and Information Science in Engineering* Vol. 21 No. 6 (2021): p. 061008.

[13] Ma, Shuai, Tang, Qian, Liu, Ying and Feng, Qixiang. "Prediction of mechanical properties of three-dimensional printed lattice structures through machine learning." *Journal of Computing and Information Science in Engineering* Vol. 22 No. 3 (2022): p. 031008.

[14] Nguyen, Cong Hong Phong and Choi, Young. "Triangular mesh and boundary representation combined approach for 3D CAD lightweight representation for collaborative product development." *Journal of Computing and Information Science in Engineering* Vol. 19 No. 1 (2019): p. 011009.

[15] Tucker, Thomas M. and Kurfess, Thomas R. "Point Cloud to CAD Model Registration Methods in Manufacturing Inspection." *Journal of Computing and Information Science in Engineering* Vol. 6 No. 4 (2006): pp. 418–421.

[16] Mata, Marta Perez, Ahmed-Kristensen, Saeema and Shea, Kristina. "Implementation of design rules for perception into a tool for three-dimensional shape generation using a shape grammar and a parametric model." *Journal of Mechanical Design* Vol. 141 No. 1 (2019).

[17] Park, Jeong Joon, Florence, Peter, Straub, Julian, Newcombe, Richard and Lovegrove, Steven. "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.

[18] Chen, Zhiqin and Zhang, Hao. "Learning Implicit Fields for Generative Shape Modeling." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.

[19] Toscano, Juan Diego, Zuniga-Navarrete, Christian, Siu, Wilson David Jo, Segura, Luis Javier and Sun, Hongyue. "Teeth Mold Point Cloud Completion Via Data Augmentation and Hybrid RL-GAN." *Journal of Computing and Information Science in Engineering* Vol. 23 No. 4 (2023).

[20] Xie, Haozhe, Yao, Hongxun, Sun, Xiaoshuai, Zhou, Shangchen and Zhang, Shengping. "Pix2Vox: Context-Aware 3D Reconstruction From Single and Multi-View Images." *IEEE/CVF International Conference on Computer Vision*. 2019.

[21] Yang, Shuo, Xu, Min, Xie, Haozhe, Perry, Stuart and Xia, Jiahao. "Single-View 3D Object Reconstruction From Shape Priors in Memory." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 3152–3161. 2021.

[22] Mittal, Paritosh, Cheng, Yen-Chi, Singh, Maneesh and Tulsiani, Shubham. "AutoSDF: Shape Priors for 3D Completion, Reconstruction and Generation." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 306–315. 2022.

[23] Lin, Hongbin, Xu, Qingfeng, Xu, Handing, Xu, Yanjie, Zheng, Yiming, Zhong, Yubin and Nie, Zhenguo. "Three-Dimensional-Slice-Super-Resolution-Net: A Fast Few Shooting Learning Model for 3D Super-Resolution Using Slice-Up and Slice-Reconstruction." *Journal of Computing and Information Science in Engineering* Vol. 24 No. 1 (2024): pp. 11005–11006.

[24] Chen, Zhiqin, Kim, Vladimir G., Fisher, Matthew, Aigerman, Noam, Zhang, Hao and Chaudhuri, Siddhartha. "DECOR-GAN: 3D Shape Detailization by Conditional Re-

finement." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 15740–15749. 2021.

[25] Gkioxari, Georgia, Malik, Jitendra and Johnson, Justin. "Mesh R-CNN." *IEEE/CVF International Conference on Computer Vision*. 2019.

[26] Xu, Qun-Ce, Mu, Tai-Jiang and Yang, Yong-Liang. "A survey of deep learning-based 3D shape generation." *Computational Visual Media* Vol. 9 No. 3 (2023): pp. 407–442.

[27] Garcia-Garcia, A., Gomez-Donoso, F., Garcia-Rodriguez, J., Orts-Escolano, S., Cazorla, M. and Azorin-Lopez, J. "PointNet: A 3D Convolutional Neural Network for real-time object class recognition." *2016 International Joint Conference on Neural Networks (IJCNN)*: pp. 1578–1584. 2016.

[28] Li, Yangyan, Bu, Rui, Sun, Mingchao, Wu, Wei, Di, Xinhan and Chen, Baoquan. "PointCNN: Convolution On X-Transformed Points." *Advances in Neural Information Processing Systems*, Vol. 31. 2018. Curran Associates, Inc.

[29] Komarichev, Artem, Hua, Jing and Zhong, Zichun. "Learning geometry-aware joint latent space for simultaneous multimodal shape generation." *Computer Aided Geometric Design* Vol. 93 (2022): p. 102076.

[30] Kim, Jinwoo, Yoo, Jaehoon, Lee, Juho and Hong, Seunghoon. "SetVAE: Learning Hierarchical Composition for Generative Modeling of Set-Structured Data." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 15059–15068. 2021.

[31] Yang, Ximing, Wu, Yuan, Zhang, Kaiyi and Jin, Cheng. "CPCGAN: A controllable 3D point cloud generative adversarial network with semantic label generating." *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4: pp. 3154–3162. 2021.

[32] Shu, Dong Wook, Park, Sung Woo and Kwon, Junseok. "3D Point Cloud Generative Adversarial Network Based on Tree Structured Graph Convolutions." *IEEE/CVF International Conference on Computer Vision*. 2019.

[33] Fu, Rao, Zhan, Xiao, CHEN, YIWEN, Ritchie, Daniel and Sridhar, Srinath. "ShapeCrafter: A Recursive Text-Conditioned 3D Shape Generation Model." *Advances in Neural Information Processing Systems*, Vol. 35: pp. 8882–8895. 2022. Curran Associates, Inc.

[34] Nichol, Alex, Jun, Heewoo, Dhariwal, Prafulla, Mishkin, Pamela and Chen, Mark. "Point-E: A System for Generating 3D Point Clouds from Complex Prompts." (2022).

[35] Scarselli, Franco, Gori, Marco, Tsoi, Ah Chung, Hagenbuchner, Markus and Monfardini, Gabriele. "The Graph Neural Network Model." *IEEE Transactions on Neural Networks* Vol. 20 No. 1 (2009): pp. 61–80.

[36] Kipf, Thomas N. and Welling, Max. "Semi-Supervised Classification with Graph Convolutional Networks." *International Conference on Learning Representations*. 2017.

[37] Hui, Ka-Hei, Li, Ruihui, Hu, Jingyu and Fu, Chi-Wing. "Neural Template: Topology-Aware Reconstruction and Disentangled Generation of 3D Meshes." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 18572–18582. 2022.

[38] Li, Xingang, Xie, Charles and Sha, Zhenghui. "A Predictive and Generative Design Approach for Three-Dimensional Mesh Shapes Using Target-Embedding Variational Autoencoder." *Journal of Mechanical Design* Vol. 144 No. 11 (2022).

[39] Chen, Wenzheng, Ling, Huan, Gao, Jun, Smith, Edward, Lehtinen, Jaakko, Jacobson, Alec and Fidler, Sanja. "Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer." *Advances in Neural Information Processing Systems*, Vol. 32. 2019. Curran Associates, Inc.

[40] Michel, Oscar, Bar-On, Roi, Liu, Richard, Benaim, Sagie and Hanocka, Rana. "Text2Mesh: Text-Driven Neural Stylization for Meshes." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 13492–13502. 2022.

[41] Mescheder, Lars, Oechsle, Michael, Niemeyer, Michael, Nowozin, Sebastian and Geiger, Andreas. "Occupancy Networks: Learning 3D Reconstruction in Function Space." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019.

[42] Liu, Shi-Lin, Guo, Hao-Xiang, Pan, Hao, Wang, Peng-Shuai, Tong, Xin and Liu, Yang. "Deep Implicit Moving Least-Squares Functions for 3D Reconstruction." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 1788–1797. 2021.

[43] Tang, Jia-Heng, Chen, Weikai, Yang, jie, Wang, Bo, Liu, Songrun, Yang, Bo and Gao, Lin. "OctField: Hierarchical Implicit Functions for 3D Modeling." *Advances in Neural Information Processing Systems*, Vol. 34: pp. 12648–12660. 2021. Curran Associates, Inc.

[44] Sanghi, Aditya, Chu, Hang, Lambourne, Joseph G., Wang, Ye, Cheng, Chin-Yi, Fumero, Marco and Malekshan, Kamal Rahimi. "CLIP-Forge: Towards Zero-Shot Text-To-Shape Generation." *IEEE/CVF Conference on Computer Vision and Pattern Recognition*: pp. 18603–18613. 2022.

[45] Saquil, Yassir, Xu, Qun-Ce, Yang, Yong-Liang and Hall, Peter. "Rank3DGAN: Semantic mesh generation using relative attributes." *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 04: pp. 5586–5594. 2020.

[46] Song, Binyang, Yuan, Chenyang, Permenter, Frank, Arechiga, Nikos and Ahmed, Faez. "Data-Driven Car Drag Prediction with Depth and Normal Renderings." *Journal of Mechanical Design* .

[47] Alhaija, Hassan Abu, Dirik, Alara, Knorig, Andr'e, Fidler, Sanja and Shugrina, Maria. "XDGAN: Multi-Modal 3D Shape Generation in 2D Space." *British Machine Vision Conference*. 2022.

[48] Song, Binyang, Zhou, Rui and Ahmed, Faez. "Multi-modal Machine Learning in Engineering Design: A Review and Future Directions." *Journal of Computing and Information Science in Engineering* (2023): pp. 1–32.

[49] Ho, Jonathan, Salimans, Tim, Gritsenko, Alexey, Chan, William, Norouzi, Mohammad and Fleet, David J. "Video diffusion models." *arXiv:2204.03458* (2022).

[50] He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing and Sun, Jian. "Deep residual learning for image recognition." *Proceedings of the IEEE conference on computer vision and pattern recognition*: pp. 770–778. 2016.

[51] Çiçek, Özgün, Abdulkadir, Ahmed, Lienkamp, Soeren S, Brox, Thomas and Ronneberger, Olaf. "3D U-Net: learning dense volumetric segmentation from sparse annotation." *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19*: pp. 424–432. 2016. Springer.

[52] Ciampiconi, Lorenzo, Elwood, Adam, Leonardi, Marco, Mohamed, Ashraf and Rozza, Alessandro. "A survey and taxonomy of loss functions in machine learning." (2023).

[53] Mazé, François and Ahmed, Faez. "Diffusion models beat gans on topology optimization." *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), Washington, DC*. 2023.

[54] Chong, Min Jin and Forsyth, David A. "Effectively Unbiased FID and Inception Score and where to find them." *CoRR* Vol. abs/1911.07023 (2019).

[55] Betzalel, Eyal, Penso, Coby, Navon, Aviv and Fetaya, Ethan. "A Study on the Evaluation of Generative Models." (2022).

[56] Choy, Christopher B., Xu, Danfei, Gwak, Jun Young, Chen, Kevin and Savarese, Silvio. "3D-R2N2: A unified approach for single and multi-view 3D object reconstruction." *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* Vol. 9912 LNCS (2016): pp. 628–644.

[57] Yang, Yaoqing, Feng, Chen, Shen, Yiru and Tian, Dong. "FoldingNet: Point Cloud Auto-Encoder via Deep Grid Deformation." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.