Generative Design for Manufacturing: Integrating Generation with Optimization Using a Guided Voxel Diffusion Model

Anonymous submission

Abstract

In digital manufacturing, converting advanced designs into quality products is hampered by manufacturers' limited design knowledge, restricting the adoption and enhancement of innovative solutions. This paper addresses this challenge through a novel generative denoising diffusion model (DDM) trained on historical 3D design data, enabling the creation of voxel-based designs that meet manufacturing standards. By integrating a surrogate model for evaluating the manufacturability of generated designs, the proposed DDM is able to optimize manufacturability during the generative process. This paper takes a leap forward from the predominant 2D focus of existing generative models towards 3D generative design, which not only broadens manufacturers' design capabilities but also accelerates the development of practical and optimized products. We demonstrate the efficacy of this approach via a case study on Microbial Fuel Cell (MFC) anode design, illustrating how this method can significantly enhance manufacturing workflows and outcomes. Our research offers a path for manufacturers to deepen their design expertise and foster innovation in digital manufacturing.

Introduction

Additive manufacturing (AM) revolutionizes design customizability (Tofail et al. 2018), yet its full potential is curtailed by the lack of sophisticated design tools for manufacturers to optimize and innovate independently (Mehrpouya et al. 2019). Current reliance on pre-optimized designs provided by designers limits manufacturers' innovation, necessitating AI-assisted tools that can endow them with the design knowledge required to exploit AM fully. AI-powered generative design, with its capacity to navigate complex design spaces through expansive datasets, has the potential to democratize design innovation in AM, enabling manufacturers to transition from mere implementers to innovators of optimized, high-quality products (Liu, Tian, and Kan 2022).

Addressing the shortfall in generative modeling for 3D shapes, this paper proposes adapting denoising diffusion models (DDMs) from 2D to 3D, aiming to overcome the limitations in resolution and manufacturability of current models. We introduce a voxel-based DDM (voxel-DDM), enriched with a regression guidance module for manufacturability evaluation, to guide the generative process toward practical, high-quality 3D designs. A case study on Microbial Fuel Cell (MFC) anode design exemplifies how this approach can equip manufacturers with a deeper understanding of design for AM, fostering a new era in high-fidelity, manufacturability-optimized 3D generative modeling.

This paper's significant contributions are as follows:

- A 3D topology dataset with labels, designed to enhance 3D surrogate and generative modeling efforts.
- A voxel denoising diffusion framework (Figure 2) extending advanced 2D image synthesis methodologies into the realm of 3D.
- A surrogate model that rapidly and precisely assesses manufacturability.
- The integration of guidance informed by manufacturability insights, derived from a pre-trained surrogate model, to refine the generative process.
- The augmentation of design knowledge and the enhancement of manufacturability for 3D MFC anode structures tailored for AM.

Literature Review

In engineering design and manufacturing, 3D representations are essential for accurately depicting designs and facilitating evaluation, optimization, and manufacturing in later design stages. Different forms of 3D representations, including structured voxels, unstructured meshes, point clouds, parametric, and implicit forms, cater to varied generative models for new design creation. Voxel representation, akin to 2D pixels extended to the 3D domain, leverages 3D convolutional neural network (CNN) architectures for shape synthesis (Lin et al. 2024; Mittal et al. 2022), while point clouds offer a compact alternative for efficient shape generation, despite lacking spatial order—a gap filled by models like PointNet through order-invariant operations(Tang et al. 2022).

Meshes offer a detailed unstructured representation, capturing not only vertices but also their topological connections, thus providing a high-fidelity depiction of 3D shapes, typically processed using graph neural networks (GNNs) (Hui et al. 2022). Implicit representations, predicting whether points are inside or outside shape boundaries, have integrated various embedding modules for latent feature learning (Ibing, Lim, and Kobbelt 2021). However, the complexity of these representations poses a challenge, often requiring complex architectures or multi-stage processes to

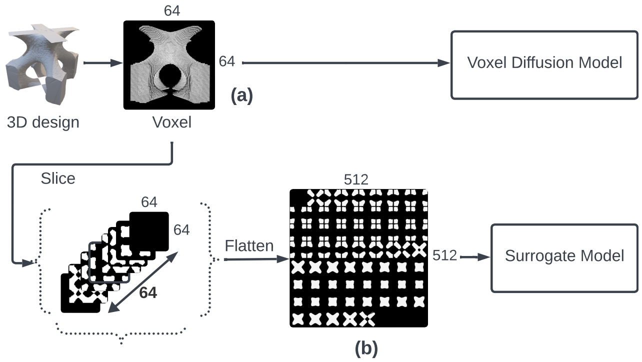

Figure 1: 3D Design data preprocessing to prepare training data for the voxel denoising diffusion model and surrogate model, respectively.

synthesize high-resolution outputs. The current literature reflects a burgeoning interest in generative models for 3D synthesis, extending traditional 2D techniques into the realm of 3D and exploring novel approaches for cross-modal synthesis to overcome the existing barriers in resolution and computational efficiency.

Accordingly, we propose a new 3D shape generative model to synthesize 3D MFC anode designs for manufacturing. We receive inspiration from a few recent papers that extend DDMs from 2D space to 3D space for 3D spatiotemporal video generation (Ho et al. 2022; Xing et al. 2023).

Data and Method

In this study, we introduce voxel-DDM, a generative modeling framework optimized for AM, adapting DDMs for voxel-based 3D shape generation using a 3D U-Net (Çiçek et al. 2016). This framework is further refined with a manufacturing metric optimization guidance module, derived from a surrogate model, and its effectiveness is demonstrated through a 3D MFC design case study.

Data

The guided voxel-DDM developed in this study is trained and validated on a dataset of 2,735 3D MFC anode designs obtained from a prior study (Kang, Deng, and Jin 2021), originally in STL format and focused on structures with one or two unit cells. These designs are labeled with the "minimum feature size," a key manufacturability metric, and converted into voxel representations for model training. For effective surrogate modeling, these voxel representations are reorganized into 2D images by arranging the 64 layers into an 8×8 grid, forming a 512×512 image, each tagged with its corresponding manufacturability metric, as shown in 1. Our approach aims to enhance the manufacturability of MFC anode designs for additive manufacturing by maximizing the minimum feature size.

Voxel Denoising Diffusion Model

The voxel-DDM, aligning with standard DDMs, operates through two phases (Figure 2): forward diffusion, where Gaussian noise incrementally transforms input voxel data $z_0 \sim q(z_0)$ towards a Gaussian noise distribution over a number of steps (which is 400 in our study); and reverse denoising, which iteratively removes this noise to reconstruct the original data from its noised state. This Markovian forward process, governed by a controlled noise introduction at each step, ensures a smooth transition to the noise distribution. The reverse process, critical for the model's generative efficacy, is dictated by the model's ability to accurately predict the mean $\mu_\theta(z_t, t)$ and variance $\Sigma_\theta(z_t, t)$ at each reverse step, ultimately defining the fidelity and accuracy of the generated voxel data.

In our study, we leverage the U-Net framework adapted from 2D to 3D (Çiçek et al. 2016), specifically for voxel data processing, to predict the mean $\mu_\theta(z_t, t)$ and variance $\Sigma_\theta(z_t, t)$ in the reverse denoising phase. This architecture, enhanced with downsampling and upsampling paths connected by skip connections, effectively retains local and global information crucial for accurate 3D representation. The training of this model focuses on optimizing its parameters θ to minimize the reconstruction difference, employing a loss function $L(\theta) = E_{t,x,\epsilon}[||\epsilon - \epsilon_\theta(z_t, t)||^2]$ that improves the model's ability to predict and reverse the noise added during the forward diffusion process. Additionally, the model incorporates a novel factorized attention mechanism, comprising intra-layer and inter-layer attention blocks, to enhance learning of complex topological features in voxel data.

Once trained, the voxel-DDM's denoising module functions as a generator, iteratively refining random samples from a Gaussian distribution to create detailed, coherent designs, with each iteration bringing the sample closer to the training data's distribution. This process, enhanced by a factorized attention mechanism, demonstrates the model's potential in design for manufacturing, enabling the creation of complex designs previously difficult to materialize.

Surrogate Model

Another pivotal component of the proposed framework is the surrogate model, a modified ResNet-50 architecture (He et al. 2016), designed to predict the manufacturing-related metric "minimum feature size" for MFC anode designs. Adapted to accommodate our data, this model employs the pre-trained ResNet-50 as a feature extractor, freezing its parameters for knowledge retention, and replaces ResNet's final fully connected layers with a series of linear layers and ReLU activations. This structure sequentially transforms the feature space into progressively smaller dimensions — 512, 128, 64, and finally 8 — culminating in a linear layer that outputs the predicted metric. The surrogate model's architecture optimizes the ResNet-50 framework for the given prediction tasks in design for manufacturing.

Guidance Module

To optimize the minimum feature size in MFC anode structures during generation, we integrate the pre-trained surrogate model into the voxel-DDM's denoising module as a guidance mechanism (Figure 2), creating what we term the

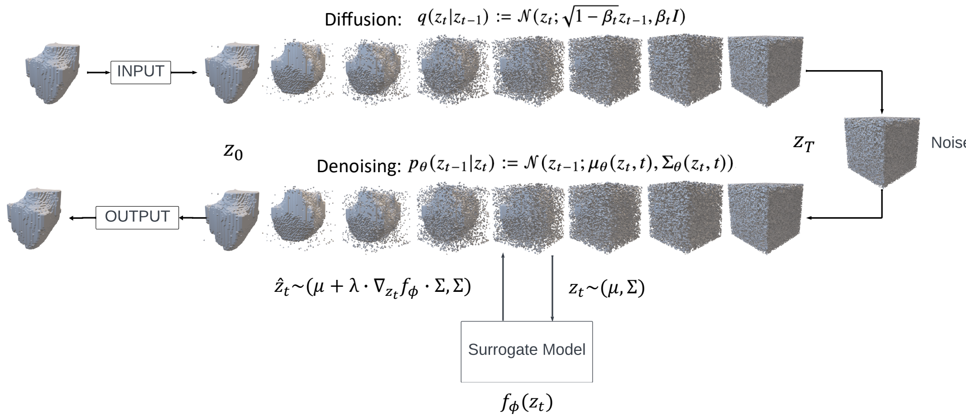

Figure 2: The proposed guided voxel-DDM that integrates the vanilla voxel-DDM with the surrogate model. In this study, the voxel-DDM is trained in 16K epochs with 400 sampling timesteps.

"guided voxel-DDM." This integration employs a regression guidance mechanism (Mazé and Ahmed 2023), adjusting the voxel-DDM's predicted mean based on the surrogate model's gradient, focusing on designs with larger minimum feature sizes. The guidance, activated when noise levels are below a certain threshold ($t < 200$), steers the generation process towards manufacturable designs by altering the trajectory in the latent space. This surrogate model, acting as both an evaluator and optimizer, provides real-time feedback during denoising, enabling the generation of designs that are not only novel and high-quality but also aligned with key manufacturing metrics, thereby enhancing manufacturability through strategic exploration of the latent space.

Results and Discussion

In this section, we validate the guided voxel-DDM for generating optimized 3D MFC anode designs, focusing on one critical topological attribute, minimum feature size. This experimental evaluation explores the model's generation capabilities, the performance of the surrogate model, and the effectiveness of the guided voxel-DDM in enhancing manufacturability, demonstrating its real-world application potential in complex geometry design for AM.

Quality of Generated Designs

The voxel-DDM, trained on 2,735 3D MFC anode designs, demonstrates its capability to generate diverse novel topologies, as assessed through key performance metrics like the Fréchet Inception Distance (FID) and the Inception Score (IS). With hyperparameters including a learning rate of 0.0001, input dimension of $64 \times 64 \times 64$, and 16,000 training epochs, the model's quality is evaluated using a test dataset comprising designs generated by both the vanilla and guided voxel-DDM. The vanilla voxel-DDM achieved an average FID score of 556.58 and an IS of 1.23, indicating a high degree of fidelity to real design data, as shown in Figure 3. Additionally, the average normalized minimum feature sizes (0.54 for vanilla and 0.63 for guided voxel-DDM, compared to 0.33 for the training set) further substantiate the model's efficacy in synthesizing manufacturable designs, underscoring its potential in real-world AM applications. This comprehensive evaluation, transcending mere visual quality to include manufacturability metrics, confirms the voxel-DDM's proficiency in producing functionally viable and diverse anode designs.

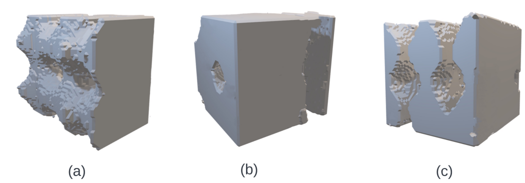

Figure 3: 3D designs generated by the proposed voxel DDM model

Optimization via Guidance Module

The surrogate model in our framework demonstrates exceptional predictive accuracy for the minimum feature size of 3D anode structures, achieving an impressive r^2 score of 0.96. It is integrated into the voxel-DDM as a guidance module, providing real-time feedback from the surrogate model during denoising. This significantly enhances the generative design process by optimizing manufacturability, resulting in designs with notably larger minimum feature sizes in the guided model compared to the vanilla model, as evidenced by statistical analysis and visual comparisons in Figure 4 and Figure 5. Additionally, the guided voxel-DDM's designs achieve an average FID score of 572.01 and an IS of 1.17, indicating a trade-off between optimized manufacturability and slightly lower diversity and quality compared to the vanilla model's broader range of designs. These results demonstrate the surrogate model's crucial role in enhancing manufacturability in generative design, marking a significant advancement in the field.

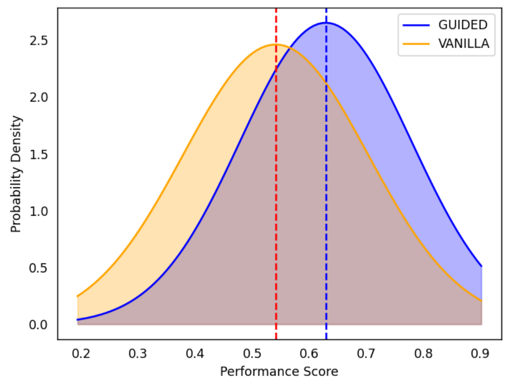

Figure 4: Minimum feature size distributions of 100 random designs respectively generated by the guided and vanilla (unguided) voxel-DDMs

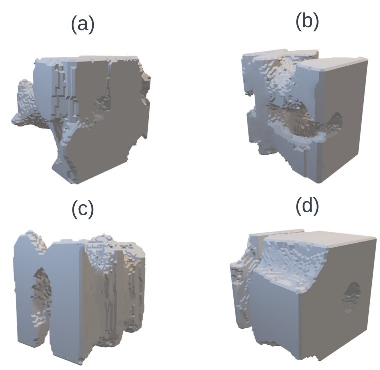

Figure 5: Two pairs of designs respectively generated by the guided and vanilla (unguided) voxel-DDMs using two same noise inputs. Top row: generated by the Vanilla Voxel-DDM, and bottom row: generated by the guided Voxel-DDM

Conclusion

We introduce a guided voxel denoising diffusion model (voxel-DDM) that extends denoising diffusion techniques from 2D to 3D, enabling the generation of complex 3D topologies optimized for additive manufacturing (AM). This model integrates a pre-trained surrogate model to guide the denoising process, focusing on generating designs with enhanced manufacturability, addressing a significant knowledge gap in the field. Our experimental results, particularly in designing Microbial Fuel Cell (MFC) anode structures, showcase the model's capability to broaden design possibilities, improve manufacturability, and foster innovation in data-driven 3D generative design for AM.

References

Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S. S.; Brox, T.; and Ronneberger, O. 2016. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19*, 424–432. Springer.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video diffusion models. *arXiv:2204.03458*.

Hui, K.-H.; Li, R.; Hu, J.; and Fu, C.-W. 2022. Neural Template: Topology-Aware Reconstruction and Disentangled Generation of 3D Meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18572–18582.

Ibing, M.; Lim, I.; and Kobbelt, L. 2021. 3D Shape Generation With Grid-Based Implicit Functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13559–13568.

Kang, S.; Deng, X.; and Jin, R. 2021. 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*, 21(6): 061008.

Lin, H.; Xu, Q.; Xu, H.; Xu, Y.; Zheng, Y.; Zhong, Y.; and Nie, Z. 2024. 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*, 24(1): 11005–11006.

Liu, C.; Tian, W.; and Kan, C. 2022. When AI meets additive manufacturing: Challenges and emerging opportunities for human-centered products development. *Journal of Manufacturing Systems*, 64: 648–656.

Mazé, F.; and Ahmed, F. 2023. Diffusion models beat gans on topology optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), Washington, DC*.

Mehrpouya, M.; Dehghanghadikolaei, A.; Fotovvati, B.; Vosooghnia, A.; Emamian, S. S.; and Gisario, A. 2019. The potential of additive manufacturing in the smart factory industrial 4.0: A review. *Applied Sciences*, 9(18): 3865.

Mittal, P.; Cheng, Y.-C.; Singh, M.; and Tulsiani, S. 2022. AutoSDF: Shape Priors for 3D Completion, Reconstruction and Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 306–315.

Tang, Y.; Qian, Y.; Zhang, Q.; Zeng, Y.; Hou, J.; and Zhe, X. 2022. WarpingGAN: Warping Multiple Uniform Priors for Adversarial 3D Point Cloud Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6397–6405.

Tofail, S. A.; Koumoulos, E. P.; Bandyopadhyay, A.; Bose, S.; O'Donoghue, L.; and Charitidis, C. 2018. Additive manufacturing: scientific and technological challenges, market uptake and opportunities. *Materials today*, 21(1): 22–37.

Xing, Z.; Feng, Q.; Chen, H.; Dai, Q.; Hu, H.; Xu, H.; Wu, Z.; and Jiang, Y.-G. 2023. A Survey on Video Diffusion Models. *arXiv preprint arXiv:2310.10647*.