



Universiteit
Leiden

The Netherlands

Deep generative models for engineering design

Fan, J.

Citation

Fan, J. (2026, March 24). *Deep generative models for engineering design*.
Retrieved from <https://hdl.handle.net/1887/4298630>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4298630>

Note: To cite this publication please use the final published version (if applicable).

Chapter 7

Conclusion and Outlook

This thesis aims to bridge the gap between modern deep generative models and generative engineering design. This chapter summarizes all the content of the preceding chapters, responds to the main research questions raised in Chapter 1, and points out limitations and future directions.

7.1 Research Question Revisited

Chapter 1: In the introduction, this paper provides an overview for this study, including industrial design processes, generative engineering design, and the latest advances in deep generative modeling, and explains the motivation for using deep generative models to achieve generative engineering design. The main research question addressed in this paper is:

How to enable deep generative models to synthesize engineering designs?

To answer this question, we have broken it down into several sub-questions, which are discussed separately in different chapters. More precisely, this paper addresses the main question by tackling a series of challenges, namely data representation, model architecture, model training/sampling, and evaluation. A secondary objective is to apply the newly developed methods to industrial design processes and demonstrate their applicability.

Chapter 2: Although recent diffusion models have achieved impressive results on general images, their generated engineering blueprints are structurally unrecognizable,

i.e., poor plausibility. We assume that this can be caused by their prioritizing strategy for better visual quality. Herein, the challenge is:

How to prioritize plausibility in generation with deep generative models?

We observe that there is a range of noise levels associated with structure formation. Therefore, we design a new noise scheduling scheme for the diffusion model for training and sampling to focus on this selected noise level range. Our evaluation results show that this helps improve the rationality of the generated results. This proposed noise scheduling method is called plausibility-oriented diffusion model (PoDM).

Chapter 3: In the previous chapter, we discover the issue of poor plausibility in the generated designs, which are evaluated manually. This raise the following research question:

How to automatically evaluate the plausibility of designs generated by deep generative models?

Leveraging the advantage of denoising autoencoders in focusing on capturing underlying structures during training, we propose Fréchet Denoised Distance (FDD) as a new evaluation metric, which shows reliable performance in aligning with human judge on plausibility.

Chapter 4: High dependence on training autoencoders to enable deep generative models on high-dimensional data leads to significant computational costs and imperfect decoding results. Thus, the research question is:

How to use deep generative models to generate high-dimension designs more efficiently?

Several learning-free decomposition algorithms can achieve the same encoding-decoding function with adjustable information loss, such as singular value decomposition (SVD). Inspired by this, we design an SVD-based generation framework, SpoDify, which achieves performance comparable to state-of-the-art methods without training an autoencoder.

Chapter 5: Given that CAD is the most widely used data format in engineering design processes, enabling DGM for generative engineering design must address the

following research question:

How to enable deep generative models to directly synthesize CAD-native representation?

Existing work has directly generated B-Rep data (native CAD representation), but this progress relies on learning on structured point clouds sampled from geometries. Here, we address this limitation by proposing NeuroNURBS, a method that decomposes NURBS-based geometries into latent features that can be used for downstream tasks. This method demonstrates significant efficiency advantages in terms of storage, computation and speed.

Chapter 6: Industrial design encompasses many complex structures and knowledge. When using generative models to simulate engineer construction, our major concern is:

How can DGMs be beneficial for real-world industrial design applications?

Here, we use the car A-pillars and rims as examples to demonstrate how the methods we developed (SA-ALAE, PoDM, and SpoDify) can assist engineers in designing complex structures at certain stages of the design process. Through the attempts in Chapter 6, our research provides insights into the application of DGMs to real-world engineering design processes, thereby answering the core research question of this thesis.

7.2 Limitations and Future Work

While DGMs develop in a swift speed, applying the technologies for generative engineering design still faces several challenges. We list the limitations and future directions below, for more specific discussions, we refer to the conclusion sections of corresponding chapters.

Available dataset. Currently, despite the publication of numerous engineering design datasets, e.g., 2D design images (BIKED [137], Seeing3DChairs [7], etc.), sequential construction operations in sequence (DeepCAD [182], Fusion 360 Gallery [180], etc.), 3D B-Reps (ABC [76], DeepCAD in B-Rep [187], etc.), the existing open-source datasets for generative design are insufficient. More specifically, in industrial design

scenarios, design datasets often exhibit more complex structures (e.g., B-Rep with over 3 000 NURBS surfaces) but less in diversity. In this context, the challenge of achieving DGM shifts from generating diverse simple designs to effectively learning the complex data distribution of a single target. We note that the most advanced DGM currently used for B-Rep generation has been evaluated as having excellent performance, but the entities it generates have been deemed to be unreasonable designs. Here, we ask whether this problem is caused by poor DGM performance or excessive diversity in the training data set. Additionally, when we are developing the NeuroNURBS framework, these datasets cannot provide sufficient high-degree NURBS surfaces to challenge the model. Therefore, the potential of NeuroNURBS is underestimated. Considering that during the historical design phase, a significant amount of designs were created in a verifiable digital format, the release of new datasets containing these historical designs may be helpful in further advancing this field.

Generating feasible designs. Our work on PoDM is just a small step towards understanding the structure generated by DGM. We refer to the performance of generating reasonable designs as “plausibility”. However, in actual industrial design scenarios, engineering design needs to present feasible, manufacturable structures. In order for a design to be feasible and manufacturable, it must meet a large number of critical requirements, such as ease of assembly, sufficient strength, compliance with safety standards, etc. Currently, research in this field still struggles to explain why and how DGM generates realistic samples. To address this, we believe that the future direction is to iteratively generate designs, evaluate feasibility and design performance at each step, and feed the evaluation results back into the next iteration. On this, we have presented our insights in Section 2.4.

Conditional design generation. This thesis evaluates proposed models often under unconditional generation, as we aim to improve the generation efficiency and to enable the generation of special design representations for research purposes. However, in real-world industrial design processes, conditional design generation is required. Recent studies [88, 92] have achieved the generation of desired designs under specific condition (e.g., text descriptions, sketches, or mesh inputs), but industrial users prefer more intuitive inputs (e.g., dragging), which we have demonstrated in Section 2.4 but only for 2D images. To address this limitation, we believe that the future direction is to enable more intuitive input and evaluate the accuracy of the editing.

Design representation. This work begins with two-dimensional blueprints and delved into three-dimensional meshes and B-Reps, fully exploring the representational possibilities of design. However, to use DGM in real-world design cases, the current representation types are far from sufficient. More specifically, representing B-Reps with hierarchical structures leads to unnecessary large data volumes and poses challenges for the reversal process. Although recent work has achieved promising results [187], this framework suffers from high inefficiency during generation and poor performance when generating complex structures with thousands of surfaces. To address this issue, we believe that developing new representation methods or new concept of how B-Reps are formed, such as introducing interactions and combination functions between solids to form B-Rep volumes, may be helpful.

7.3 The Future of DGMs in Generative Engineering Design.

The limitations mentioned previously expose a significant gap but also reveal the enormous potential for growth in the field of DGM-based generative design. Our work is completed during a period when generative models began to shine, with researchers in the field continuously surpassing previous methods. It may take only a short time for the future directions discussed in this section to be surpassed, but we will still offer our insights.

Firstly, due to the development of large language models (LLMs), current LLMs can operate on a high diversity of data by treating it as “language”, such as code, tabular data, and even images, referred to as large multimodal models (LMMs). In fact, in historical industrial design processes, there is a rich repository of design data in various forms, e.g., code (JT and STEP files), engineering construction operations in sequence (from CADPART files), blueprint images (DWG files), or tabular data, which can be understood by LMMs according to recent research [185, 42, 9]. However, this concept remains in its early stages.

At the same time, there is great potential for research into using DGMs to directly learn from design data, such as B-Reps, as this enables a differentiable modeling process, which is useful for further applications such as numerical simulation. We have already provided this concept in Section 2.4. Most of the work in this thesis has been done in this direction, but it seems that there is still a long way to go. To enable more effective modeling of design data like B-Reps, DGMs need to access the natural pro-

cesses of construction, including but not limited to sketching, extruding, chamfering, drilling, and interaction. We believe that these construction operations are based on engineering expertise, and it is essential for DGMs to learn this knowledge in addition to structural information. It should be noted that representing these operations in text form does not necessarily enable the DGM to understand the construction knowledge and structural dependencies. DGMs also need to consider the changes these operations make to the structure and be penalized when constructing any infeasible entities. However, how to achieve this and how to represent it remains an open question.

Moreover, the current main trend in generative modeling applications is to deepen models exponentially. This trend has led to the emergence of giant, even mega models. Designers have observed that simply increasing the size of the model can effectively force it to learn hidden information better. However, we hope this does not become the future direction of research, as recent DGM frameworks have not yet been fully developed. Researchers in this field should strive forward to bring more advancements to DGM frameworks to address remaining challenges, such as low-sample learning, out-of-distribution learning or machine reasoning.

Finally, incorporating industrial applications into consideration is an inevitable future direction for DGM research. In our work, we have observed that the development of DGMs for industrial purposes has not garnered much attention in the DGM community, while most recent work has focused on more general data representations such as point clouds, meshes, and implicit representations like neural radiative fields (NeRFs). This focus has led to DGM research often attracting widespread attention but struggling to find corresponding application cases in the real world. Meanwhile, generating engineering designs is a valuable application scenario where DGMs can make significant contributions. By bridging the gap between theoretical advancements and practical applications, we can unlock the full potential of DGMs, paving the way for innovative solutions that enhance efficiency and creativity in industrial applications.

Bibliography

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, pages 40–49. PMLR, 2018.
- [3] Md Ferdous Alam, Austin Lentsch, Nomi Yu, Sylvia Barmack, Suhin Kim, Daron Acemoglu, John Hart, Simon Johnson, and Faez Ahmed. From Automation to Augmentation: Redefining Engineering Design and Manufacturing in the Age of NextGen-AI. *An MIT Exploration of Generative AI*, mar 27 2024. <https://mit-genai.pubpub.org/pub/9s6690gd>.
- [4] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [5] Silvia Ansaldi, Leila De Floriani, and Bianca Falcidieno. Geometric modeling of solid objects by using a face adjacency graph representation. *SIGGRAPH Comput. Graph.*, 19(3):131–139, July 1985.
- [6] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [7] Mathieu Aubry, Daniel Maturana, Alexei A. Efros, Bryan C. Russell, and Josef Sivic. Seeing 3d chairs: Exemplar part-based 2d-3d alignment using a large dataset of cad models. In *2014 IEEE CVPR*, pages 3762–3769, 2014.
- [8] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multi-MAE: Multi-modal multi-task masked autoencoders. *European Conference on Computer Vision*, 2022.
- [9] Akshay Badagabettu, Sai Sravan Yarlagadda, and Amir Barati Farimani. Query2cad: Generating cad models using natural language queries, 2024.

-
- [10] Nicholas Baker, Hongjing Lu, Gennady Erlikhman, and Philip J. Kellman. Deep convolutional networks do not classify based on global object shape. *PLoS Computational Biology*, 14, 2018.
 - [11] Shane T. Barratt and Rishi Sharma. A note on the inception score. *ArXiv*, abs/1801.01973, 2018.
 - [12] Yoshua Bengio, Eric Laufer, Guillaume Alain, and Jason Yosinski. Deep generative stochastic networks trainable by backprop. In *International Conference on Machine Learning*, pages 226–234. PMLR, 2014.
 - [13] Henrik Berg and Adam Brandt. Investigation of aerodynamic wheel design. 2018.
 - [14] Eyal Betzalel, Coby Penso, Aviv Navon, and Ethan Fetaya. A study on the evaluation of generative models. *CoRR*, abs/2206.10935, 2022.
 - [15] Gregory Beylkin, Ronald Coifman, and Vladimir Rokhlin. Fast wavelet transforms and numerical algorithms i. *Communications on pure and applied mathematics*, 44(2):141–183, 1991.
 - [16] Mikolaj Binkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD gans. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
 - [17] Andrii Blyndaruk and Olena Shapovalova. End-to-end differentiable nurbs-geom-gnn-tgn pipeline for spatiotemporal identification of moving objects. In *Scientific and Practical Journal” Materials of Scientific Conferences of the Petro Mohyla Black Sea National University”*, number 1, pages 16–19, 2025.
 - [18] Ali Borji. Pros and cons of gan evaluation measures: New developments. *Computer Vision and Image Understanding*, 215:103329, 2022.
 - [19] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
 - [20] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
 - [21] Lucas F Buzuti and Carlos E Thomaz. Fréchet autoencoder distance: A new approach for evaluation of generative adversarial networks. *Computer Vision and Image Understanding*, 235:103768, 2023.
 - [22] Daniel R Canelhas, Erik Schaffernicht, Todor Stoyanov, Achim J Lilienthal, and Andrew J Davison. An eigenshapes approach to compressed signed distance fields and their utility in robot mapping. *arXiv preprint arXiv:1609.02462*, 2016.

-
- [23] Weijuan Cao, Trevor T Robinson, Yang Hua, Andrew Colligan, and Wanbin Pan. Graph representation of 3d cad models for machining feature recognition with deep learning. In *ASME 2020 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference: Proceedings*, November 2020. International Design Engineering Technical Conferences & Computers and Information in Engineering Conference ; Conference date: 17-08-2020 Through 19-08-2020.
- [24] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS’20, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [25] Amaresh Chakrabarti, Kristina Shea, Robert Stone, Jonathan Cagan, Matthew Campbell, Noe Vargas Hernandez, and Kristin L. Wood. Computer-Based Design Synthesis Research: An Overview. *Journal of Computing and Information Science in Engineering*, 11(2):021003, June 2011. eprint: https://asmedigitalcollection.asme.org/computingengineering/article-pdf/11/2/021003/5566343/021003_1.pdf.
- [26] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository, 2015.
- [27] Wei Chen and Faez Ahmed. Padgan: Learning to generate high-quality novel designs. *Journal of Mechanical Design*, 143:031703, 11 2020.
- [28] Wei Chen, Kevin Chiu, and Mark D Fuge. Airfoil design parameterization and optimization using bézier generative adversarial networks. *AIAA Journal*, 58(11):4723–4735, 2020.
- [29] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5939–5948, 2019.
- [30] Shiyang Cheng, Michael Bronstein, Yuxiang Zhou, Irene Kotsia, Maja Pantic, and Stefanos Zafeiriou. Meshgan: Non-linear 3d morphable models of faces. *arXiv preprint arXiv:1903.10384*, 2019.
- [31] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8185–8194, 2020.
- [32] Gene Chou, Yuval Bahat, and Felix Heide. Diffusion-sdf: Conditional generative modeling of signed distance functions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2262–2272, 2023.

- [33] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 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*, pages 424–432. Springer, 2016.
- [34] CIMdata. Generative design: What’s that? <https://www.cimdata.com/en/news/item/8402-generative-design-what-s-that>, April 2023. Accessed: 2023-04-24.
- [35] Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the national academy of sciences*, 102(21):7426–7431, 2005.
- [36] Andrew R. Colligan, Trevor T. Robinson, Declan C. Nolan, Yang Hua, and Weijuan Cao. Hierarchical cadnet: Learning from b-reps for machining feature recognition. *Computer-Aided Design*, 147:103226, 2022.
- [37] Corinna Cortes and Vladimir Vapnik. *Support-vector networks*. Number 3. Springer, 1995.
- [38] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [39] Anjana Deva Prasad, Aditya Balu, Harshil Shah, Soumik Sarkar, Chinmay Hegde, and Adarsh Krishnamurthy. Nurbs-diff: A differentiable programming module for nurbs. *Comput. Aided Des.*, 146(C), May 2022.
- [40] Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [41] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- [42] Yuhao Du, Shunian Chen, Wenbo Zan, Peizhao Li, Mingxuan Wang, Dingjie Song, Bo Li, Yan Hu, and Benyou Wang. Blenderllm: Training large language models for computer-aided design with self-improvement, 2024.
- [43] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pages 2650–2658, 2015.
- [44] G Farin. Handbook of computer aided geometric design. *Elsevier science google schola*, 2:577–580, 2002.
- [45] Robert L Foote. Regularity of the distance function. *Proceedings of the American Mathematical Society*, 92(1):153–155, 1984.

-
- [46] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *ArXiv*, abs/1811.12231, 2018.
 - [47] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
 - [48] Ian Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
 - [49] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
 - [50] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
 - [51] Jana Gulanová, Ladislav Gulan, Michal Forrai, and Mario Hirz. Generative engineering design methodology used for the development of surface-based components. *Computer-Aided Design and Applications*, 14(5):642–649, 2017.
 - [52] Florentin Guth, Simon Coste, Valentin De Bortoli, and Stephane Mallat. Wavelet score-based generative modeling, 2022.
 - [53] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. Meshcnn: A network with an edge. *ACM Transactions on Graphics (TOG)*, 38(4):90:1–90:12, 2019.
 - [54] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
 - [55] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
 - [56] Negar Heidari and Alexandros Iosifidis. Geometric deep learning for computer-aided design: A survey. *arXiv preprint arXiv:2402.17695*, 2024.
 - [57] Katherine L. Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. *arXiv: Computer Vision and Pattern Recognition*, 2019.

- [58] 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. *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, 2018.
- [59] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.
- [60] Mario Hirz, Patrick Rossbacher, and Jana Gulánová. Future trends in cad—from the perspective of automotive industry. *Computer-Aided Design and Applications*, 14(6):734–741, 2017.
- [61] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020.
- [62] Danijela Horak, Simiao Yu, and Gholamreza Salimi Khorshidi. Topology distance: A topology-based approach for evaluating generative adversarial networks. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 7721–7728. AAAI Press, 2021.
- [63] J. Hoschek and D. Lasser. *Fundamentals of Computer Aided Geometric Design*. Ak Peters Series. Taylor & Francis, 1996.
- [64] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [65] Sergey Ioffe. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [66] Pradeep Kumar Jayaraman, Joseph George Lambourne, Nishkrit Desai, Karl Willis, Aditya Sanghi, and Nigel J. W. Morris. Solidgen: An autoregressive model for direct b-rep synthesis. *Transactions on Machine Learning Research*, 2023. Featured Certification.
- [67] Pradeep Kumar Jayaraman, Aditya Sanghi, Joseph G. Lambourne, Karl D.D. Willis, Thomas Davies, Hooman Shayani, and Nigel Morris. Uv-net: Learning from boundary representations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [68] Mark W Jones. Distance field compression. *The Journal of WSCG*, 12(2):199–204, 2004.
- [69] Nikolai Kalischek, Torben Peters, Jan D Wegner, and Konrad Schindler. Tetradiffusion: Tetrahedral diffusion models for 3d shape generation. *arXiv e-prints*, pages arXiv–2211, 2022.

-
- [70] Tero Karras. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [71] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022.
- [72] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [73] Mohamed Kas, Abderrazak Chahi, Ibrahim Kajo, and Yassine Ruichek. Eigen-gan: An svd subspace-based learning for image generation using conditional gan. *Knowledge-Based Systems*, 293:111691, 2024.
- [74] Katherine Keegan, Tanvi Vishwanath, and Yihua Xu. A tensor svd-based classification algorithm applied to fmri data. *arXiv preprint arXiv:2111.00587*, 2021.
- [75] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [76] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [77] B. Kolarevic. *Architecture in the Digital Age: Design and Manufacturing*. Taylor & Francis, 2004.
- [78] Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, Thomas Unterthiner, and Xiaohua Zhai. An image is worth 16x16 words: Transformers for image recognition at scale. 2021.
- [79] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021.
- [80] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/kriz/cifar.html>, 5(4):1, 2010.
- [81] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [82] Sarah C. Kucker, Larissa K. Samuelson, Lynn K. Perry, Hanako Yoshida, Eliana Colunga, Megan G Lorenz, and Linda B. Smith. Reproducibility and a unifying explanation: Lessons from the shape bias. *Infant behavior & development*, 54:156–165, 2019.

- [83] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fréchet inception distance. In *The Eleventh International Conference on Learning Representations*, 2023.
- [84] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Neural Information Processing Systems*, 2019.
- [85] Joseph George Lambourne, Karl Willis, Pradeep Kumar Jayaraman, Longfei Zhang, Aditya Sanghi, and Kamal Rahimi Malekshan. Reconstructing editable prismatic cad from rounded voxel models. In *SIGGRAPH Asia 2022 Conference Papers*, SA '22, page 1–9. ACM, November 2022.
- [86] Barbara Landau, Linda B. Smith, and Susan Scanlon Jones. The importance of shape in early lexical learning. *Cognitive Development*, 3:299–321, 1988.
- [87] Sang Hun Lee and Kunwoo Lee. Partial entity structure: a compact non-manifold boundary representation based on partial topological entities. In *International Conference on Smart Media and Applications*, 2001.
- [88] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J. Mitra. Sketch2cad: Sequential cad modeling by sketching in context. *ACM Trans. Graph. (Proceedings of SIGGRAPH Asia 2020)*, 39(6):164:1–164:14, 2020.
- [89] Jiahao Li, Yunpeng Bai, Yongkang Dai, Hao Guo, Hongping Gan, and Yilei Shi. Autoregressive generation with b-rep holistic token sequence representation. *arXiv preprint arXiv:2601.16771*, 2026.
- [90] Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12642–12651, 2023.
- [91] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyang Wu, Bir Bhanu, Richard J Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8642–8651, 2020.
- [92] Yilin Liu, Duoteng Xu, Xingyao Yu, Xiang Xu, Daniel Cohen-Or, Hao Zhang, and Hui Huang. Hola: B-rep generation using a holistic latent representation, 2025.
- [93] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. *Advances in neural information processing systems*, 32, 2019.
- [94] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

-
- [95] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015.
 - [96] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998.
 - [97] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022.
 - [98] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2837–2845, 2021.
 - [99] Antoine Maiorca, Youngwoo Yoon, and Thierry Dutoit. Evaluating the quality of a synthesized motion with the fréchet motion distance. In *ACM SIGGRAPH 2022 Posters, SIGGRAPH ’22*, New York, NY, USA, 2022. Association for Computing Machinery.
 - [100] Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting particle solutions of fokker-planck equations through gradient-log-density estimation. *Entropy*, 22(8):802, 2020.
 - [101] Peter Meltzer, Hooman Shayani, Amir Khasahmadi, Pradeep Kumar Jayaraman, Aditya Sanghi, and Joseph Lambourne. Uvstyle-net: Unsupervised few-shot learning of 3d style similarity measure for b-reps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9690–9699, 2021.
 - [102] Victoria Miles, Stefano Giani, and Oliver Vogt. Approaching step file analysis as a language processing task: A robust and scale-invariant solution for machining feature recognition. *Journal of Computational and Applied Mathematics*, 427:115166, 2023.
 - [103] Leon Mirsky. Symmetric gauge functions and unitarily invariant norms. *The quarterly journal of mathematics*, 11(1):50–59, 1960.
 - [104] Mehdi Mirza. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
 - [105] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 306–315, 2022.

- [106] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [107] Volodymyr Mnih. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [108] Sergey Moiseev, Andrey Kuznetsov, Viktor Gamayunov, Anton Ezhov, Sergey Serebryakov, and Sergey Kolesnikov. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021.
- [109] Jens-Dominik Mueller, Xingchen Zhang, Siamak Akbarzadeh, and Yang Wang. Geometric continuity constraints of automatically derived parametrisations in cad-based shape optimisation. *International Journal of Computational Fluid Dynamics*, 33:1–17, 11 2019.
- [110] Phillip Mueller, Talip Uenlue, Sebastian Schmidt, Marcel Kollovieh, **Jiajie Fan**, Stephan Günnemann, and Lars Mikelsons. Geodiffusion: A training-free framework for accurate 3d geometric conditioning in image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6374–6384, October 2025.
- [111] Orest Mykhaskiv, Mladen Banovic, Salvatore Auriemma, Pavanakumar Mohanamuraly, Andrea Walther, Herve Legrand, and Jens-Dominik Mueller. Nurbs-based and parametric-based shape optimization with differentiated cad kernel. *Computer-Aided Design and Applications*, 15:1–11, 04 2018.
- [112] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunje Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7176–7185. PMLR, 13–18 Jul 2020.
- [113] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011.
- [114] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [115] Sangeun Oh, Yongsu Jung, Seongsin Kim, Ikjin Lee, and Namwoo Kang. Deep Generative Design: Integration of Topology Optimization and Generative Models. *Journal of Mechanical Design*, 141(11):111405, September 2019. eprint: <https://asmedigitalcollection.asme.org/mechanicaldesign/article-pdf/141/11/111405/6578473/md.141.11.111405.pdf>.
- [116] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

-
- [117] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- [118] OpenAI. Chatgpt. <https://www.openai.com/chatgpt/>, 2022. Accessed: [date accessed].
- [119] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024.
- [120] Edouard Oyallon, Stéphane Mallat, and Laurent Sifre. Generic deep networks with wavelet scattering, 2014.
- [121] Emanuele Palumbo, Imant Daunhawer, and Julia E Vogt. MMVAE+: Enhancing the generative quality of multimodal VAEs without compromises. In *The Eleventh International Conference on Learning Representations*, 2023.
- [122] Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. Drag your GAN: interactive point-based manipulation on the generative image manifold. In Erik Brunvand, Alla Sheffer, and Michael Wimmer, editors, *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH 2023, Los Angeles, CA, USA, August 6-10, 2023*, pages 78:1–78:11. ACM, 2023.
- [123] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019.
- [124] T. Paviot. pythonocc, 2022.
- [125] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.
- [126] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [to appear].
- [127] L. Piegl. On nurbs: a survey. *IEEE Computer Graphics and Applications*, 11(1):55–71, 1991.
- [128] L. Piegl and W. Tiller. *The NURBS Book*. Monographs in Visual Communication. Springer Berlin Heidelberg, 1996.

- [129] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*, 2016.
- [130] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [131] Yixuan Qiu, Qingyi Gao, and Xiao Wang. Adaptive learning of the latent space of wasserstein generative adversarial networks. *Journal of the American Statistical Association*, pages 1–13, 2024.
- [132] J. Ross Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.
- [133] Alec Radford. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [134] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021.
- [135] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European conference on computer vision (ECCV)*, pages 704–720, 2018.
- [136] J Redmon. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [137] Lyle Regenwetter, Brent Curry, and Faez Ahmed. BIKED: A Dataset for Computational Bicycle Design With Machine Learning Benchmarks. *Journal of Mechanical Design*, 144(3), 10 2021. 031706.
- [138] Lyle Regenwetter, Amin Heyrani Nobari, and Faez Ahmed. Deep generative models in engineering design: A review. *Journal of Mechanical Design*, 144(7):071704, 2022.
- [139] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022.
- [140] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. *Advances in neural information processing systems*, 30, 2017.
- [141] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lučić, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

-
- [142] Ruslan Salakhutdinov and Geoffrey Hinton. Deep boltzmann machines. In *Artificial intelligence and statistics*, pages 448–455. PMLR, 2009.
- [143] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *ArXiv*, abs/1606.03498, 2016.
- [144] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [145] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation, 2023.
- [146] Thomas Schütz. *Hucho-Aerodynamik des Automobils: Strömungsmechanik, Wärmetechnik, Fahrdynamik, Komfort*. Springer-Verlag, 2013.
- [147] Adam Sebestyen, Albert Wiltche, Milena Stavric, and Ozan Özdenizci. From nurbs to neural networks: Efficient geometry encoding for generative ai in architectural design. *International Journal of Architectural Computing*, 23(3):720–741, 2025.
- [148] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.
- [149] S. S. SHAPIRO and M. B. WILK. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3-4):591–611, December 1965. eprint: <https://academic.oup.com/biomet/article-pdf/52/3-4/591/962907/52-3-4-591.pdf>.
- [150] Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*, 2023.
- [151] Jaehyeok Shim, Changwoo Kang, and Kyungdon Joo. Diffusion-based signed distance fields for 3d shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20887–20897, June 2023.
- [152] Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-only transformers. *arXiv preprint arXiv:2311.15475*, 2023.
- [153] Wojciech Skarka. Application of moka methodology in generative model creation using catia. *Engineering Applications of Artificial Intelligence*, 20(5):677–690, 2007.

- [154] Colin Smith. On vertex-vertex meshes and their use in geometric and biological modeling. *Retrieved March*, 13:2017, 2006.
- [155] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [156] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020.
- [157] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.
- [158] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.
- [159] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [160] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [161] George Stein, Jesse C. Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Leigh Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L. Caterini, J. Eric T. Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. *CoRR*, abs/2306.04675, 2023.
- [162] Gilbert W Stewart. On the early history of the singular value decomposition. *SIAM review*, 35(4):551–566, 1993.
- [163] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2015.
- [164] **Jiajie Fan**, Amal Trigui, Thomas Bäck, and Hao Wang. Enhancing plausibility evaluation for generated designs with denoising autoencoder. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision – ECCV 2024*, pages 88–105, Cham, 2025. Springer Nature Switzerland.
- [165] **Jiajie Fan**, Laure Vuaille, Thomas Bäck, and Hao Wang. Adversarial latent autoencoder with self-attention for structural image synthesis. In *2024 IEEE Conference on Artificial Intelligence (CAI)*, pages 119–124, 2024.

-
- [166] **Jiajie Fan**, Babak Gholami, Thomas Bäck, and Hao Wang. NeuroNURBS: Learning efficient surface representations for 3D solids. *CoRR*, abs/2411.10848, 2024.
 - [167] **Jiajie Fan**, Laure Vuaille, Thomas Bäck, and Hao Wang. On the noise scheduling for generating plausible designs with diffusion models. *CoRR*, abs/2411.10848, 2023.
 - [168] **Jiajie Fan**, Amal Trigui, Andrea Bonfanti, Felix Dietrich, Thomas Bäck, and Hao Wang. A mesh is worth 512 numbers: Spectral-domain diffusion modeling for high-dimension shape generation. *arXiv preprint arXiv/2503.06485*, 2025.
 - [169] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018.
 - [170] University of Illinois at Urbana-Champaign. UIUC Airfoil Database. http://m-selig.ae.illinois.edu/ads/coord_database.html, 2022. Accessed on 18.10.2022.
 - [171] U.S. Department of Transportation National Highway Traffic Safety Administration. Laboratory Test Procedure For FMVSS No.201U. Technical Report TP201U-01, U.S. Department of Transportation National Highway Traffic Safety Administration, 1998.
 - [172] Muhammad Usama, Mohammad Sadil Khan, Didier Stricker, and Muhammad Zeshan Afzal. Nurbgen: High-fidelity text-to-cad generation through llm-driven nurbs modeling, 2025.
 - [173] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
 - [174] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
 - [175] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
 - [176] Pascal Vincent, H. Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *International Conference on Machine Learning*, 2008.
 - [177] Kevin J Weiler. Topological structures for geometric modeling (boundary representation, manifold, radial edge structure). In *Rensselaer Polytechnic Institute*, 1986.
 - [178] Francis Williams. Point cloud utils, 2022. <https://www.github.com/fwilliams/point-cloud-utils>.

- [179] Karl DD Willis, Pradeep Kumar Jayaraman, Hang Chu, Yunsheng Tian, Yifei Li, Daniele Grandi, Aditya Sanghi, Linh Tran, Joseph G Lambourne, Armando Solar-Lezama, and Wojciech Matusik. Joinable: Learning bottom-up assembly of parametric cad joints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [180] Karl DD Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G Lambourne, Armando Solar-Lezama, and Wojciech Matusik. Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. *ACM Transactions on Graphics (TOG)*, 40(4):1–24, 2021.
- [181] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *Advances in neural information processing systems*, 29, 2016.
- [182] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6772–6782, October 2021.
- [183] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *ArXiv*, abs/1708.07747, 2017.
- [184] Bojun Xiong, Si-Tong Wei, Xin-Yang Zheng, Yan-Pei Cao, Zhouhui Lian, and Peng-Shuai Wang. OctFusion: Octree-based diffusion models for 3d shape generation. *arXiv*, 2024.
- [185] Jingwei Xu, Zibo Zhao, Chenyu Wang, Wen Liu, Yi Ma, and Shenghua Gao. Cad-mllm: Unifying multimodality-conditioned cad generation with mllm, 2024.
- [186] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaodong He, and Jianfeng Gao. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1316–1324, 2018.
- [187] Xiang Xu, Joseph G Lambourne, Pradeep Kumar Jayaraman, Zhengqing Wang, Karl DD Willis, and Yasutaka Furukawa. Brepgen: A b-rep generative diffusion model with structured latent geometry. *arXiv preprint arXiv:2401.15563*, 2024.
- [188] Xiang Xu, Karl DD Willis, Joseph G Lambourne, Chin-Yi Cheng, Pradeep Kumar Jayaraman, and Yasutaka Furukawa. Skexgen: Autoregressive generation of cad construction sequences with disentangled codebooks. In *International Conference on Machine Learning*, pages 24698–24724. PMLR, 2022.
- [189] Shaoliang Yang and Jun Wang. Triple-parametric autoencoder for 2d reparameterization via bézier, b-spline, and nurbs representations. *Computer-Aided Design*, 189:103936, 2025.

-
- [190] Isabela M. Yepes and Manasvi Goyal. Image classification using singular value decomposition and optimization, 2024.
 - [191] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, 2024.
 - [192] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
 - [193] Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales, and Chen Change Loy. Sketch me that shoe. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 799–807, 2016.
 - [194] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
 - [195] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023.
 - [196] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023.
 - [197] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019.
 - [198] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 586–595. Computer Vision Foundation / IEEE Computer Society, 2018.
 - [199] Zheyang Zhang, Yongxing Wang, Peter K Jimack, and He Wang. Meshingnet: A new mesh generation method based on deep learning. In *International conference on computational science*, pages 186–198. Springer, 2020.
 - [200] Junsheng Zhou, Weiqi Zhang, Baorui Ma, Kanle Shi, Yu-Shen Liu, and Zhizhong Han. Udiff: Generating conditional unsigned distance fields with optimal wavelet diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
 - [201] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5826–5835, 2021.

- [202] Wang Zhou. Image quality assessment: from error measurement to structural similarity. *IEEE transactions on image processing*, 13:600–613, 2004.

List of Abbreviations

ALAE	Adversarial Latent Autoencoder
B-Rep	Boundary Representation
CAD	Computer-Aided Design
CAE	Computational-Aided Engineering
CD	Chamfer Distance
DAE	Denoising Autoencoder
DDIM	Denoising Diffusion Implicit Mode
DDPM	Denoising Diffusion Probabilistic Model
DGM	Deep Generative Model
DPS	Design Plausibility Score
FDD	Fréchet Denoised Distance
FID	Fréchet Inception Distance
FMH	Free Motion Headform
FPS	Frames per Second
GAN	Generative Adversarial Network
GED	Generative Engineering Design
GenAI	Generative Artificial Intelligence
HIC	Head Injury Criterion
KDD	Kernel Denoised Distance
LDM	Latent Diffusion Model
LLM	Large Language Model
LMM	Large Multimodality Model
MMD	Maximum Mean Discrepancy
MSE	Mean Squared Error
NeRF	Neural Radiance Fields

NURBS	Non-Uniform Rational B-Splines
NWD	Neural Wavelet-domain Diffusion
ODE	Ordinary Differential Equation
PDR	Plausible Design Rat
PoDM	Plausible-oriented Diffusion Mode
ReLU	Rectified Linear Unit
SA-ALAE	Self-Attention Adversarial Latent Autoencoder
SDE	Stochastic Differential Equation
SDF	Signed Distance Field
SVD	Singular Value Decomposition
SVM	Support Vector Machine
VAE	Variational Autoencoder