

Generating 3D House Wireframes with Semantics

Xueqi Ma^①, Yilin Liu^②, Wenjun Zhou^③, Ruowei Wang^④, and Hui Huang^{*⑤}

Visual Computing Research Center, Shenzhen University
hhzhiyan@gmail.com

Abstract. We present a new approach for generating 3D house wireframes with semantic enrichment using an autoregressive model. Unlike conventional generative models that independently process vertices, edges, and faces, our approach employs a unified wire-based representation for improved coherence in learning 3D wireframe structures. By re-ordering wire sequences based on semantic meanings, we facilitate seamless semantic integration during sequence generation. Our two-phase technique merges a graph-based autoencoder with a transformer-based decoder to learn latent geometric tokens and generate semantic-aware wireframes. Through iterative prediction and decoding during inference, our model produces detailed wireframes that can be easily segmented into distinct components, such as walls, roofs, and rooms, reflecting the semantic essence of the shape. Empirical results on a comprehensive house dataset validate the superior accuracy, novelty, and semantic fidelity of our model compared to existing generative models. More results and details can be found on <https://vcc.tech/research/2024/3DWire>.

Keywords: Wireframe · 3D Generation · Autoregressive Model

1 Introduction

The 3D wireframe serves as a crucial data structure in the fields of computer vision and graphics, offering a clean and compact abstraction of an object’s shape through a composition of points and lines. This form of representation is pivotal for the creation of complex and detailed models, with a particular advantage in representing polyhedral objects such as mechanical parts, furniture, and buildings, due to its simplicity and efficacy in preserving sharp features. Despite its wide use and benefits, the generation of 3D wireframes remains a sophisticated and demanding process. It involves the meticulous abstraction of the object’s geometry into a series of precise line segments, a task that necessitates both accuracy and creativity. Consequently, there is a significant need for an automated approach to efficiently generate accurate and innovative 3D wireframes.

Early studies in wireframe generation have mainly focused on reconstructing wireframes from various data sources, like images [32, 54, 55] or point clouds [4, 21]. These approaches typically employ learning-based techniques for detecting corners and edges, followed by a handcraft assembling process to extract the

* Corresponding author

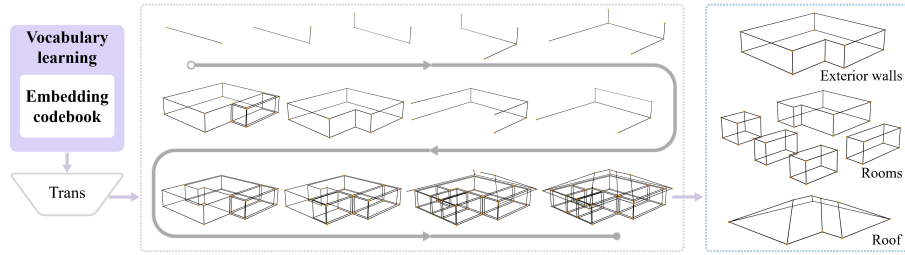


Fig. 1: Our method creates 3D wireframes through autoregressive sampling from a trained transformer model, generating tokens from a learned geometric vocabulary. These tokens are decoded into line segments to form the final wireframe. Based on the nodes’ connectivity, the resulting wireframes can be easily split into multiple parts, such as walls, roofs, and rooms, reflecting the underlying semantic meaning of shapes.

wireframe structures from the input data. Despite their effectiveness in achieving high levels of accuracy, these methods do not possess the functionality to create novel wireframe data. In contrast, the latest research has shown an increasing interest in the learning-based 3D generation [15, 27, 36]. These approaches aim to learn the distribution patterns of primitive sequences (vertices, edges, faces) during the training phase. Consequently, they can generate new data by sampling from these distributions and predicting primitives in an autoregressive manner. However, these methods treat different types of primitives (vertices, edges, faces) as distinct sequences and model them separately. The error in the generation of one type of primitive can propagate to the subsequent generation of other primitives, leading to a more complex learning process. Moreover, unlike text generation tasks, which inherently follow a sequential order, geometric primitives do not have a natural order. Current methods predominantly organize these primitive sequences based on their spatial coordinates. The absence of high-level correlations among primitives might introduce ambiguity during the distribution modeling, which can adversely impact the quality of the generated models.

In this paper, we introduce a new method for the generation of semantically grouped 3D house wireframes. Our method adopts a wire-based representation, which, unlike traditional methods that model vertices and edges *separately*, focuses on constructing *pure wire* sequences based on their semantic correlations. A wireframe under our model is conceptualized as a graph where nodes correspond to wires (line segments), and edges denote the connectivity among them. Notably, a wireframe may comprise disconnected components—for example, the roof, or the exterior walls of a house are usually not connected to the interior rooms. We model each disconnected component as a distinct sub-graph, organizing wires based on their topological connectivity to mirror the semantic relationships among them. The sequence of wires is established through a breadth-first search (BFS) traversal of the graph, ensuring the generation of a wireframe that is both coherent and semantically structured, as shown in Fig. 1.

The wireframe generation is split into two stages. In the initial stage, we learn a vocabulary of latent geometric tokens representing the wireframe’s wires. This involves employing a Graph Convolutional Network (GCN) [52] to encode the local geometric and topological features of the wires, complemented by an attention-based encoder designed to extract the global information of the wireframe. The subsequent stage utilizes a transformer-based [41] decoder to autoregressively produce a sequence of tokens from the trained vocabulary. These tokens are then decoded into spatial coordinates to construct the wires, which are used in the assembly of the final wireframe. Extensive experiments show that our method can faithfully generate 3D house wireframes with semantics, which outperforms state-of-the-art generative models with higher accuracy and novelty, while enhancing the semantic integrity of the generated 3D wireframes.

In summary, our approach introduces two significant contributions:

- We propose a wire-based representation model for learning wireframe distributions, significantly enhancing the accuracy of 3D wireframe generation.
- Our method introduces a semantic-aware sequence construction technique that reduces ambiguity in the learning phase. Moreover, it allows for the wireframe to be segmented into distinct parts, each reflecting the semantic essence of the underlying 3D house model.

2 Related Work

The concept of wireframes as a fundamental data structure has been extensively explored over several decades. This section covers three primary areas of research that are closely related to our study: wireframe reconstruction, floorplan generation, and the application of autoregression in 3D model generation.

2.1 Wireframe reconstruction

Reconstructing Wireframes from Images: A significant area of research focuses on deriving wireframe representations from images of man-made environments [13, 24, 45, 47, 54]. This involves the use of deep convolutional neural networks [19] to accurately detect vertices and line segments, facilitating the creation of wireframe models. Various studies have introduced different geometric priors to enhance this process, including a trainable Hough transform block [20], the utilization of global structural regularities [55], and the exploitation of deep image gradients [32]. Additionally, recent efforts [23, 46] have extended these techniques to generate wireframes from sequences of images, demonstrating advancements in temporal data processing for 3D reconstruction.

Reconstructing Wireframes from Point Clouds: Recent research focuses on deriving 3D wireframes from point clouds. The PC2WF model [21] encodes point clouds to feature points and connects them to establish 3D wireframes. WireframeNet [4] refines the point cloud using a medial axis transform (MAT) before

predicting edge points and constructing the wireframe based on their connectivity. Matveev et al. [25] use a scalar distance field for creating a wireframe’s topological graph from detected corners and curves. Furthermore, Tan et al. [39] introduce a method that integrates coarse and fine pruning modules for wireframe optimization, employing particle swarm optimization to maintain the wireframe’s correct topology.

However, current research predominantly utilizes images and point clouds as prerequisites for generating 3D wireframes, highlighting a gap in the capability to produce novel 3D wireframes without pre-specified conditions. This underscores the ongoing challenge of unconditional 3D wireframe generation.

2.2 Floorplan Generation

Wireframes serve a pivotal role in floor plan generation, where building layouts are depicted as 2D nodes and edges. Wu et al. [44] pioneered house layout generation by simulating the human design process, initially determining room locations followed by wall placement. HouseGAN [28] and House-GAN++ [29] utilize graph-constrained relational and conditional GANs [31] to create realistic and diverse house layouts. HouseDiffusion [35] adopts bubble diagrams to generate vector floor plans via a dual diffusion strategy. FLNet [40] merges graph convolutional networks with spatial layout networks for crafting floor layouts within user-defined constraints. Graph2Plan [12] integrates user feedback to automate floor plan generation from building boundaries and layout diagrams. WallPlan [37] models wall structures as graphs, using intersections as nodes and sections as edges, to incrementally develop floor plans through graph traversal. RoomFormer [50] utilizes a Transformer-based approach for generating multiple room polygons in floorplan reconstruction.

Predominantly, these methodologies concentrate on generating 2D floor plans, *relying* on constraints such as graphs [12, 28, 29, 35, 40], building boundaries [12, 37, 40, 43, 44], textual inputs [5, 18], or point clouds [50]. Converting these 2D layouts into 3D architectural models typically necessitates additional post-processing.

2.3 Autoregression for 3D Generation

Recently, auto-regressive models have made significant strides in 2D image generation [8, 30, 33] and have also extended to 3D tasks [6, 27, 36, 38, 48]. Autosdf [26] tackles multimodal 3D tasks using an autoregressive prior, grounded in a discretized, low-dimensional latent representation. Ibing et al. [14] use octrees for hierarchical 3D shape representation and introduce adaptive compression to address sequence length, improving autoregressive model efficiency. PolyGen [27] is an autoregressive method for generating 3D meshes. It employs transformer [41] and pointer networks [42] to predict sequences of vertices and faces in the mesh. Similarly, MeshGPT [36] learns and encodes quantized embeddings of triangular mesh geometry in order to autonomously generate highly compact triangular meshes. SolidGen [15] utilizes transformers and pointer networks to directly model B-rep and predict vertices, edges, and faces in an autoregressive manner.

While PolyGen [27] and SolidGen [15] share similarities with our approach, they serialize primitives (vertices, edges, faces) separately. This independent processing can accumulate errors; inaccuracies in vertices may impact edges and faces. Moreover, these methods construct primitive sequences based on coordinates, neglecting semantic relationships. In contrast, our method emphasizes exploiting these semantic connections to sequence line segments, creating coherent wireframes that are semantically rich and hierarchically organized.

3 Methodology

In 3D design, professional designers construct wireframes by drawing each line segment individually. Similarly, large language models use a sequential approach to build complex linguistic structures. Our methodology synthesizes these principles, adopting a line-by-line generation approach to craft 3D wireframes. As shown in Fig. 1, we first learn a geometric embedding vocabulary from a large collection of 3D wireframes. Each line segment is quantized into a latent space, facilitating its encoding and subsequent processing as detailed in Sec. 3.1. Following this, we employ a transformer-based model trained to predict the subsequent code within our predefined vocabulary, gradually generating the corresponding 3D wireframes 3.2.

3.1 Learning Quantized Line Segment Embeddings

As demonstrated by SolidGen [15] and MeshGPT [36], the development of a quantized feature or embedding vocabulary is foundational for effective 3D shape generation. Our approach begins with the derivation of a latent code for each line segment through an encoder, denoted as E , followed by the quantization of these codes into a geometric codebook via Residual Quantization [17]. We define a 3D wireframe \mathcal{W} as a set of line segments:

$$\mathcal{W} = (l_1, l_2, l_3, \dots, l_N), \quad (1)$$

where \mathcal{W} consists of N line segments. Each segment $l_i \in \mathbb{R}^{n_{in}}$ is characterized by: 1) Coordinates of its endpoints, 2) Length of the segment 3) Directional orientation 4) Angle between it and adjacent segments 5) Coordinates of its midpoint. All these features are quantized to an integer of range $[0, 128)$ and then embedded into a 196-dimensional vector.

As shown in Fig. 2, to enhance neighborhood data representation and simplify the geometric vocabulary, we first transform the feature of each line segment l_i into a n_z -dimensional vector z_i using an encoder $E(\mathcal{W})$, $z = 384$ in all our experiments. The encoder consists of a graph convolutional encoder $E_G(\mathcal{W})$ and an attention-based information exchanger $E_A(\mathcal{W})$. The encoder $E_G : (\mathbb{R}^{N \times n_{in}}) \rightarrow \mathbb{R}^{N \times n_z}$ treat line segments as nodes, with their adjacency defining the graph edges. We utilize SAGEConv layers [9] to project node features into latent space, capturing the local geometric feature of each segment.

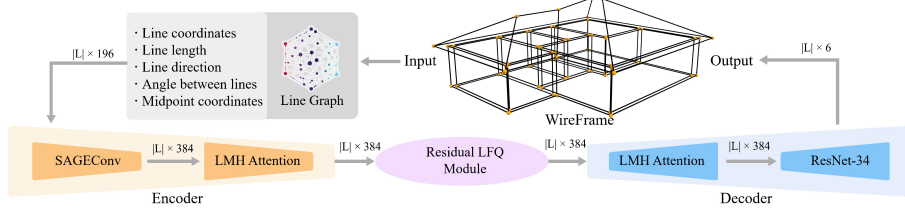


Fig. 2: Pipeline of learning the geometric vocabulary of line segments.

Given that the input wireframe \mathcal{W} may contain disconnected sub-graphs, where two nodes that are spatially close might not be topologically connected, relying solely on graph convolutions might be insufficient for information exchange.

To mitigate this, we introduce an attention-based exchanger $E_A : (\mathbb{R}^{N \times n_z}) \rightarrow \mathbb{R}^{N \times n_z}$, specifically a Local Multi Head Attention (LMH Attention) [3, 34], comprised of multiple self-attention layers. It allows for a complete information exchange across the entire structure, ensuring the extracted features are not only topologically informed but also contextually enriched.

Based on the high-dimensional feature of each segment l_i , we then learn a geometric embedding codebook through Residual Quantization (RQ). Similar to MeshGPT [36], we learn the quantized code for each vertex instead of each line segment to get better generalizability. This involves assigning the segment features to their endpoints and normalizing by the count of adjacent segments, effectively distributing the feature information across the vertices. For a given codebook C , a Residual Quantization (RQ) with D layers of depth can represent a feature z as follows:

$$(m_1, m_2, \dots, m_D) \in [M]^D = \mathcal{RQ}(z; C, D), \quad (2)$$

where, m_d is the code of z at depth d .

This quantization process involves comparing each vertex feature, against all codes in the codebook, thus increasing the model’s computational complexity. To overcome this challenge, we employed the lookup-free quantization (LFQ) [49]. LFQ quantizes vertex features by treating them as a Cartesian product of single-dimensional variables, which eliminates the need for the codebook lookup step typically required in traditional quantization.

After the quantization, we stack the quantized vertex features according to the indices of the line segment vertices to form the features of the line segment. These features are then decoded into six distinct coordinate values via a decoder network comprising LMH Attention layers and a 1D ResNet34 [10] architecture. We guide the reconstruction process of 3D wireframes by applying cross-entropy loss to discrete 3D wireframe coordinates, coupled with a commitment loss for vertex embeddings. Additionally, we employ an entropy penalty [49] to enhance the utilization of the codebook. This strategy not only boosts the network’s confidence in its predictions but also promotes a more diverse usage of the codebook.

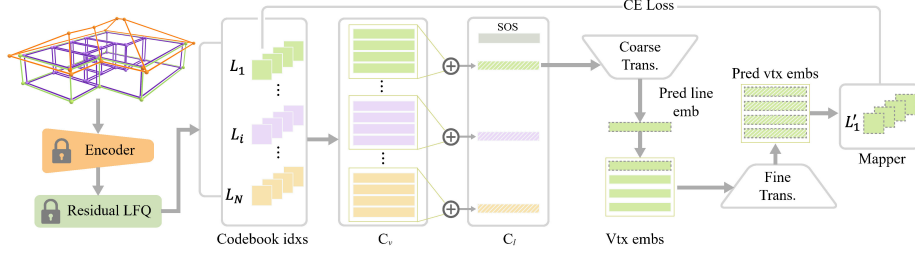


Fig. 3: Pipeline of the transformer training. Firstly, the wireframe is encoded by the encoder to extract features and undergo Residual LFQ [49] to obtain code indices of each line segment. These indices are then split and transformed into code embeddings for vertices (C_v) and are recovered to the code embeddings of line segments (C_l). The line embeddings are progressively predicted (Pred line emb) by the coarse transformer and refined to the vertex embedding (Pred vtx embs) by the fine transformer. They are finally transformed back to codebook indices through a mapper, and these indices are then optimized using a loss function to generate high-quality wireframes.

entries. More details can be found in supplementary. Once we have trained the encoder and the codebook, we can use the quantized feature of each line segment as a token sequence to further train the autoregressive model.

3.2 3D Wireframe Generation with Transformers

Leveraging the quantized features of each line segment, we employ a transformer model trained to predict the next segment in the sequence autoregressively. Autoregressive models, by design, generate sequences where each subsequent token depends on the previous ones. In the context of 3D wireframe generation with transformers, establishing a predetermined sequence for line segment generation is crucial. This sequence not only ensures the structural integrity of the wireframes but also aids the model in learning the intricate geometric shapes.

Unlike previous methods [15, 36] that rely solely on coordinate-based token sequences, our approach sequences tokens based on the semantic meaning of the line segments. Adapting PolyGen’s [27] strategy, we organize wireframe junctions and segments in a z-y-x hierarchical order. Then we treat the line segments of the wireframe as graph nodes and the junctions as graph edges, thereby transforming the 3D wireframe into a graph structure. Note that house wireframes often comprise multiple disconnected subgraphs. Thus, a Breadth-first search algorithm is applied to group segments in the same subgraph, thereby ensuring sequential generation of segments from the same object.

As shown in Fig. 3, our model utilizes a decoder-only transformer architecture to predict segment sequence indices derived from a previously established codebook (see Section 3.1). In this architecture, indices are transformed into learnable embeddings, complemented by three distinct types of encoding: 1) a learnable discrete positional encoding, which marks the position of each code in the code sequence. 2) a vertex positional encoding, indicating the position

of vertices within a line segment. 3) a quantizer level encoding to differentiate between various depths of residual quantization levels. This results in encoded vertex code sequences C_v , with a length of $|C_v| = 2 \cdot D \cdot N$.

To effectively learn the discrete code stack with depth D extracted by Residual LFQ, we trained the transformer in two stages: coarse transformer and fine transformer [17]. During the coarse transformer stage, C_v is reshaped into a $2 \cdot D \times N$. Subsequently, by merging these $2 \cdot D$ dimensions, we acquire the line code sequences C_l , whose length is $|C_l| = N$. Then, we utilize the coarse transformer to predict the line embeddings autoregressively. During the fine transformer stage, we predict each line segment’s vertex embeddings along the depth dimension based on each predicted line embeddings.

During the inference time, the transformer can generate a token sequence autoregressively. The sequence’s corresponding embeddings, fetched from the codebook, are decoded to form the 3D wireframe. Duplicate vertices are merged as a final step to construct the wireframe.

3.3 Implementation Details

During the learning process of the line segment vocabulary, our residual quantization use a depth of $D = 2$, representing each line segment by $2 \cdot D = 4$ embeddings. Our codebook size is 8192. Since we use Residual LFQ, each embedding vector has a dimension of $\log_2^{8192} = 13$. We employ two distinct, non-shared codebooks for richer information learning. Subsequently, the decoder predicts the coordinates of line segments, distributed across 128 categories, achieving spatial discretization with a resolution of 128^3 . This encoder-decoder network model was trained on 8 RTX 3090 GPUs for approximately 2 days. For the transformer, we employed a 12+2 layers coarse-to-fine decoder-only transformer model. The model was trained for about 5 days on 8 RTX 3090 GPUs.

Both the autoencoder and transformer were written in PyTorch and trained using the ADAM optimizer, with a learning rate of 1×10^{-4} . For the transformer, the learning rate linearly increases to a maximum of $lr_{\max} = 1 \times 10^{-4}$ during the initial $t_0 = 10$ epochs. It then gradually decreases following a cosine decay schedule $lr_{\max} * 0.5 * (1 + \cos(\frac{t-t_0}{T-t_0}\pi))$ until it reaches a minimum of 1×10^{-6} .

4 Experiments

We compare our method with the state-of-the-art generative model [27, 36] to evaluate the performance of our method. To the best of our knowledge, we are the first to propose a method for generating 3D house wireframes with semantic enrichment. While PolyGen [27] and MeshGPT [36] aim for generating 3D models, they do not explicitly construct the wireframe structure. We modified PolyGen and MeshGPT to shift their focus from predicting triangle or polygon vertices to line segments, allowing for a fair comparison with our method. We also conduct a series of ablation studies to verify the effectiveness of our technical contributions and provide an in-depth analysis of our approach.

4.1 Dataset and Metrics

Since there is no existing dataset for 3D wireframe generation, we created a new dataset for this task. We first extracted corners from the house layouts in the RPLAN dataset [44] to obtain the basic wireframes of houses. Then we used the straight skeleton [1, 2] to construct wireframes of the house roofs. This process resulted in a 3D house wireframe dataset with approximately 78,000 wireframes. Various data augmentation techniques, including rotation, random shifts, and scaling, were employed to facilitate training. To ensure the wireframe fit within the transformer’s context window, we selected only wireframes with fewer than 400 line segments for our training set. Detailed information on the augmentation processes and data splits is provided in the supplementary.

Evaluating 3D wireframe generation is challenging due to the absence of ground truth. We adopted established metrics from previous works [22, 51, 53], including Minimum Matching Distance (MMD), Coverage (COV), and 1-Nearest Neighbor (1-NN), based on Chamfer Distance (CD) and Earth Mover’s Distance (EMD). To assess the structural integrity of the generated wireframe, we also examined vertex-line segment connectivity with Two-Line-Connected Vertex Proportion (2L-CVP) and Three-Line-Connected Vertex Proportion (3L-CVP). These metrics measure the proportion of vertices connected to at least two or three line segments, respectively, indicating wireframe structural validity. Additionally, we use KL divergence (KLD) to compare the distributions of the number of connected components in the generated samples with the test data.

4.2 Quantitative Evaluations

As shown in Tab. 1, our method surpasses the comparative approaches across all evaluation metrics. Regarding COV, our method achieves high scores in both the CD-based and EMD-based measures, indicating a strong spatial correspondence between our wireframes and the test set. In terms of MMD, the differences between our wireframes and the test set are minimized. The results of 1-NN further validate the superiority of our method, as it maintains a high level of consistency with the test set wireframes. As for Structural Validity, our wireframes exhibit high coherence while preserving a spatial distribution highly consistent with the test set. With respect to KLD, our results more closely align with the real data distribution compared to other methods. These results confirm the effectiveness and superiority of our method in generating 3D house wireframes.

To further evaluate the quality of the generated 3D house wireframes, we conducted a user study comparing wireframes produced by different methods. Participants used an intuitive interface to view pairwise 3D wireframe models and selected the best-quality wireframe based on topological rationality. 60 participants each made 24 sets of choices.

As shown in Tab. 2, participants showed a preference for the wireframes generated by our method compared to other methods. Specifically, our method received a quality rating of 0.84 compared to PolyGen and 0.75 compared to MeshGPT. These ratings were based on the frequency with which participants

Table 1: Quantitative comparison for unconditional wireframe generation on the 3D house wireframe dataset. CD is multiplied by 10^3 and EMD is multiplied by 10^2 .

Model	COV (% , \uparrow)		MMD (\downarrow)		1-NN (%)		Struct. Valid. (% , \uparrow)		KLD
	CD	EMD	CD	EMD	CD	EMD	2L-CVP	3L-CVP	
PolyGen [27]	38.67	47.95	8.67	6.43	74.43	67.65	81.47	75.80	12.75
MeshGPT [36]	54.78	54.29	9.13	6.27	64.61	61.70	80.91	70.77	8.98
Ours	56.15	58.64	8.11	5.75	55.21	51.35	99.53	99.26	0.73

Table 2: Results of the user study on evaluating wireframes generated by PolyGen, MeshGPT, our method, and ground truth.

	PolyGen	MeshGPT	Ours	Ground Truth
PolyGen [27]	—	-0.67	-0.84	-0.93
MeshGPT [36]	0.67	—	-0.75	-0.83
Ours	0.84	0.75	—	-0.13
Ground Truth	0.93	0.83	0.13	—

Table 3: Quantitative comparison for unconditional generation on the ABC dataset.

Model	COV (% , \uparrow)	MMD (\downarrow)	1-NN (%)
PolyGen [27]	39.94	25.53	75.87
MeshGPT [36]	42.38	24.62	67.65
Ours	44.10	22.12	62.96

ranked our method as higher quality than the other methods, accounting for 92% and 87% respectively. Note that compared to ground truth wireframes, our method still achieved a relatively objective score of -0.13, indicating a good approximation to real wireframe models. We also conducted a quantitative experiment on a subset of the ABC dataset [16], comprising 24,634 samples of planar shapes (see Tab. 3). Compared with the house dataset, the ABC dataset exhibited a greater diversity of shapes. As a result, the performance of all methods suffered some decline. However, our method still outperformed the others and produced the best results. More details can be found in the supplementary.

4.3 Qualitative Evaluations

In Fig. 4, we provide a qualitative comparison between 3D house wireframes generated by our approach and those from existing methods. PolyGen often results in floating vertices and line segments due to its sequential process of generating vertices before connections, which suffers from an error accumulation and can negatively impact the coherence of the wireframe. MeshGPT also shows similar issues, attributed to its neglect of semantic relationships between wireframe components like exterior walls, rooms, and roofs. Conversely, our method

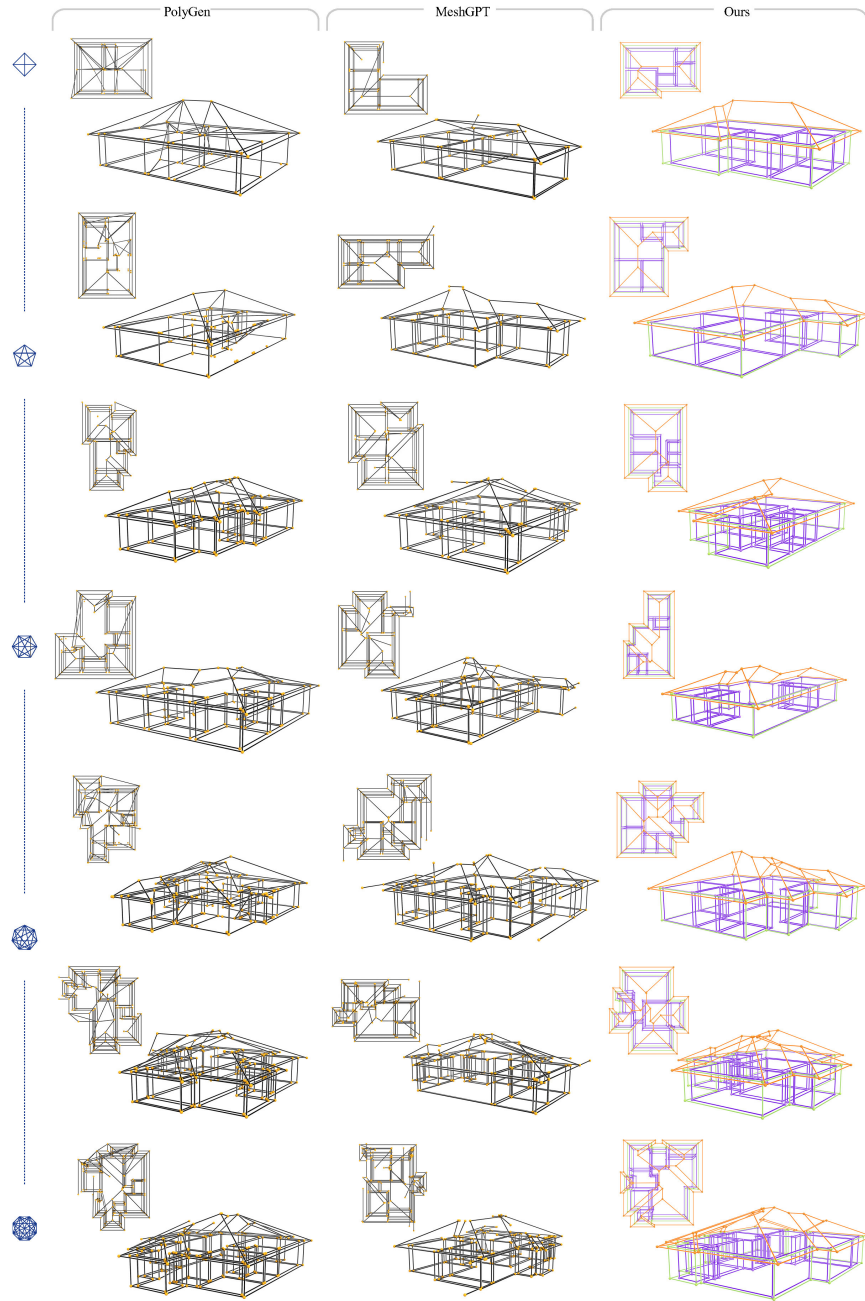


Fig. 4: Qualitative comparison of 3D house wireframes. Compared to baselines, our method produces valid wireframes with high geometric fidelity and greater simplicity.

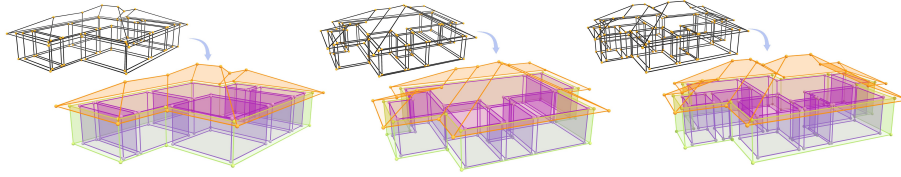


Fig. 5: The resulting wireframe can be easily converted into a mesh model.

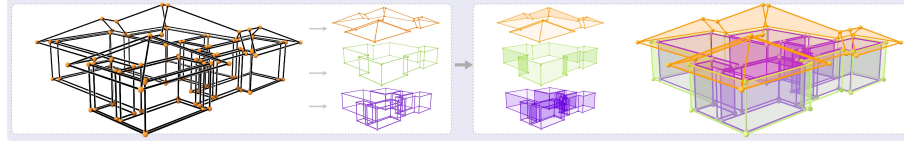


Fig. 6: The reconstructed wireframe model can be easily split into several components. We also show their corresponding mesh on the right.

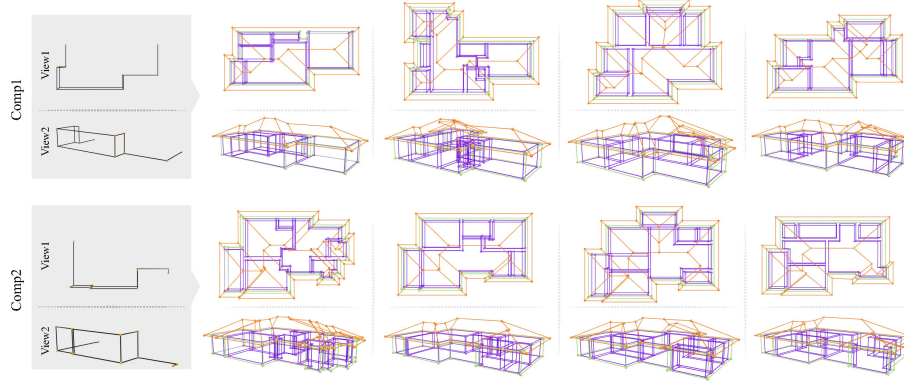


Fig. 7: Our method has the capability to generate multiple potential completions for a given partial wireframe.

excels in generating structurally sound and semantically rich wireframes, highlighting clear distinctions between exterior walls, rooms, and roofs. This not only improves the interpretability of the wireframes but also significantly reduces floating vertices and extraneous line segments, enhancing both accuracy and structural integrity. The resulting wireframe can be easily split into multiple components, such as walls, roofs, and different rooms, based on the connectivity of the line segments, as shown in Fig. 6. Moreover, our wireframes are readily convertible into mesh models, as illustrated in Fig. 5. Also, it can produce multiple possible completions for a given partial wireframe since our model generates wireframe sequences in a probabilistic manner, as shown in Fig. 7. We included text-conditioned results in Fig. 8. Features from the input text were incorporated via cross-attention, enriching the wireframe generation process. Fig. 9 demonstrates our method’s capacity to generate complex objects.

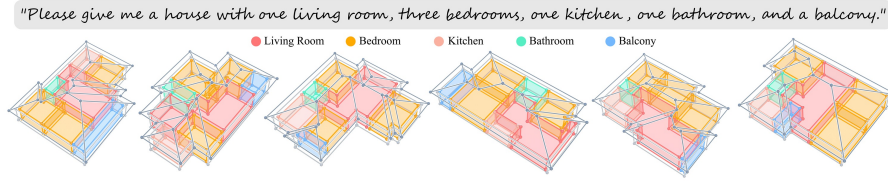


Fig. 8: Samples of 3D house wireframes generated based on text conditions.

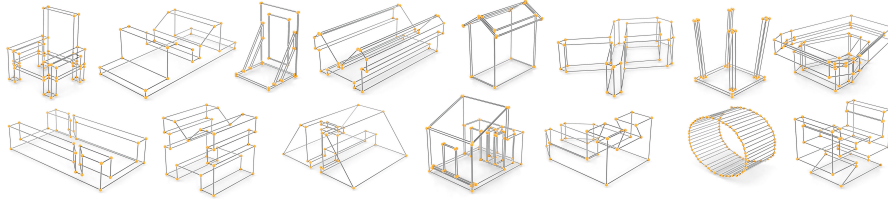


Fig. 9: Samples generated by our method, trained on the ABC dataset [16].

Table 4: Ablations of our design choices on the 3D house wireframe dataset. The results indicate a significant decrease in performance when any of them are removed.

Method	COV \uparrow	MMD \downarrow	1-NN	2L-CVP \uparrow	3L-CVP \uparrow
w/o Encoder LMH Attention	49.07	10.98	68.72	64.97	59.20
w/o Residual LFQ	50.02	9.24	69.87	69.17	63.47
w/o Coarse-to-Fine	52.56	9.08	66.20	73.97	68.77
w/o Semantic Order	51.33	9.07	67.27	72.42	67.69
Ours	56.15	8.11	55.21	99.53	99.26

4.4 Wireframe Novelty Analysis

Inspired by previous methods [7, 11, 36], we analyzed the novelty of the generated wireframes compared to the training dataset samples. Using our method, we generated 4096 wireframes and found the most similar samples in the training set (with the lowest Chamfer distance, CD). Fig. 10 shows a histogram of the generated wireframes. When the CD is small, our method’s wireframes cover the training samples well. As CD values increase, our wireframes show more significant differences from the most similar training structures, demonstrating our method’s ability to create novel wireframes. Even in the histogram’s middle region, where most generated wireframes are located, our wireframes exhibit structural differences from training samples, indicating significant diversity. More details can be found in the supplementary material.

4.5 Ablation Studies

To verify our technical contributions, we conducted ablation studies on four major components of our model: 1) The LMH attention layers in our autoencoder;

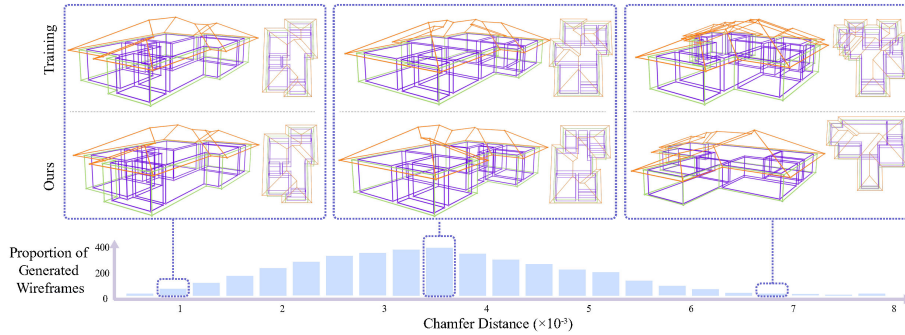


Fig. 10: Wireframe novelty analysis on the 3D house wireframe dataset. We plot the distribution of 4096 samples generated by our method and assessed their similarity to the distribution of the training dataset based on the Chamfer Distance (CD). Our method is capable of generating wireframes that are both similar (low CD) and novel (high CD) compared to the training distribution.

2) The geometric codebook implemented by a Residual LFQ; 3) A coarse-to-fine transformer when decoding the sequence; 4) The semantic sequence reordering when training the transformer. As shown in Tab. 4, removing LMH Attention layers resulted in intersecting line segments, highlighting their role in encoding spatial relationships. Replacing Residual LFQ with VQ reduced model quality due to VQ’s limited effectiveness with large vocabularies. Discontinuing the coarse-to-fine strategy increased sequence length, raising training difficulty and affecting quality. Removing semantic sorting of line segments in the training data and sorting only in the z-y-x sequence also degraded quality. These findings demonstrate that each component significantly impacts overall performance. LMH Attention layers extract geometric features, while Residual LFQ, a coarse-to-fine strategy, and semantic ordering enhance generation quality.

5 Conclusion

In summary, we introduce a novel autoregressive model that significantly enhances the generation of semantically enriched 3D house wireframes. Leveraging a unified, wire-based representation and semantic sequence reordering, our approach outperforms existing methods in accuracy, novelty, and semantic integrity. Our experiments on a diverse 3D house dataset underline our model’s superior capability, positioning it as a valuable asset for applications in 3D modeling, computer-aided design, and virtual reality.

In the future, we aim to explore the deeper integration of semantic information related to house and architectural wireframes through various conditional generation methods, including controllable text, images, and point clouds.

Acknowledgements

We thank all the anonymous reviewers for their insightful comments. This work was supported in parts by NSFC (U21B2023, U2001206, 62161146005), Guangdong Basic and Applied Basic Research Foundation (2023B1515120026), DEGP Innovation Team (2022KCXTD025), Shenzhen Science and Technology Program (KQTD20210811090044003, RCJC20200714114435012, JCYJ20210324120213036), Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), and Scientific Development Funds from Shenzhen University.

References

1. Aichholzer, O., Aurenhammer, F.: Straight skeletons for general polygonal figures in the plane. In: *Comput. Combin.: Annu. Int. Conf.* pp. 117–126 (1996)
2. Aichholzer, O., Aurenhammer, F., Alberts, D., Gärtner, B.: *A novel type of skeleton for polygons.* Springer (1996)
3. Beltagy, I., Peters, M.E., Cohan, A.: Longformer: The long-document transformer (2020)
4. Cao, L., Xu, Y., Guo, J., Liu, X.: Wireframenet: A novel method for wireframe generation from point cloud. *Comput. Graph.* **115**, 226–235 (2023)
5. Chen, Q., Wu, Q., Tang, R., Wang, Y., Wang, S., Tan, M.: Intelligent home 3d: Automatic 3d-house design from linguistic descriptions only. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 12622–12631 (2020)
6. Cheng, A.C., Li, X., Liu, S., Sun, M., Yang, M.H.: Autoregressive 3d shape generation via canonical mapping. In: *Eur. Conf. Comput. Vis.* pp. 89–104. Springer (2022)
7. Erkoç, Z., Ma, F., Shan, Q., Nießner, M., Dai, A.: Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In: *Int. Conf. Comput. Vis.* pp. 14254–14264 (2023)
8. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 12873–12883 (2021)
9. Hamilton, W.L., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: *Adv. Neural Inform. Process. Syst.* pp. 1024–1034 (2017)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 770–778 (2016)
11. Hu, J., Hui, K.H., Liu, Z., Li, R., Fu, C.W.: Neural wavelet-domain diffusion for 3d shape generation, inversion, and manipulation. *ACM Trans. Graph.* **43**(2), 16:1–16:18 (2024)
12. Hu, R., Huang, Z., Tang, Y., van Kaick, O., Zhang, H., Huang, H.: Graph2plan: learning floorplan generation from layout graphs. *ACM Trans. Graph.* **39**(4), 118:1–118:14 (2020)
13. Huang, K., Wang, Y., Zhou, Z., Ding, T., Gao, S., Ma, Y.: Learning to parse wireframes in images of man-made environments. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 626–635 (2018)
14. Ibing, M., Kobsik, G., Kobbelt, L.: Octree transformer: Autoregressive 3d shape generation on hierarchically structured sequences. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2698–2707 (2023)

15. Jayaraman, P.K., Lambourne, J.G., Desai, N., Willis, K.D.D., Sanghi, A., Morris, N.J.W.: Solidgen: An autoregressive model for direct b-rep synthesis. *Trans. Mach. Learn. Res.* **2023** (2023)
16. Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M., Zorin, D., Panozzo, D.: Abc: A big cad model dataset for geometric deep learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (June 2019)
17. Lee, D., Kim, C., Kim, S., Cho, M., Han, W.S.: Autoregressive image generation using residual quantization. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 11523–11532 (2022)
18. Leng, S., Zhou, Y., Dupty, M.H., Lee, W.S., Joyce, S., Lu, W.: Tell2design: A dataset for language-guided floor plan generation. In: Rogers, A., Boyd-Graber, J.L., Okazaki, N. (eds.) *Assoc. Comput. Linguist.* pp. 14680–14697 (2023)
19. Li, Z., Liu, F., Yang, W., Peng, S., Zhou, J.: A survey of convolutional neural networks: Analysis, applications, and prospects. *IEEE Trans. Neural Netw. Learn. Syst.* **33**(12), 6999–7019 (2022)
20. Lin, Y., Pinte, S.L., van Gemert, J.C.: Deep hough-transform line priors. In: *Eur. Conf. Comput. Vis.* pp. 323–340 (2020)
21. Liu, Y., D’Aronco, S., Schindler, K., Wegner, J.D.: PC2WF: 3d wireframe reconstruction from raw point clouds. In: *Int. Conf. Learn. Represent.* (2021)
22. Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2837–2845 (2021)
23. Luo, Y., Ren, J., Zhe, X., Kang, D., Xu, Y., Wonka, P., Bao, L.: Learning to construct 3d building wireframes from 3d line clouds. In: *Brit. Mach. Vis. Conf.* p. 91 (2022)
24. Ma, W., Tan, B., Xue, N., Wu, T., Zheng, X., Xia, G.: How-3d: Holistic 3d wireframe perception from a single image. In: *Int. Conf. 3D Vision.* pp. 596–605 (2022)
25. Matveev, A., Artemov, A., Zorin, D., Burnaev, E.: 3d parametric wireframe extraction based on distance fields. In: *Int. Conf. Artif. Intell. Pattern Recog.* pp. 316–322 (2021)
26. Mittal, P., Cheng, Y.C., Singh, M., Tulsiani, S.: Autosdf: Shape priors for 3d completion, reconstruction and generation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 306–315 (2022)
27. Nash, C., Ganin, Y., Eslami, S.A., Battaglia, P.: Polygen: An autoregressive generative model of 3d meshes. In: *Int. Conf. Mach. Learn.* pp. 7220–7229 (2020)
28. Nauata, N., Chang, K.H., Cheng, C.Y., Mori, G., Furukawa, Y.: House-gan: Relational generative adversarial networks for graph-constrained house layout generation. In: *Eur. Conf. Comput. Vis.* pp. 162–177 (2020)
29. Nauata, N., Hosseini, S., Chang, K.H., Chu, H., Cheng, C.Y., Furukawa, Y.: House-gan++: Generative adversarial layout refinement network towards intelligent computational agent for professional architects. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13632–13641 (2021)
30. van den Oord, A., Kalchbrenner, N., Espeholt, L., Kavukcuoglu, K., Vinyals, O., Graves, A.: Conditional image generation with pixelcnn decoders. In: *Adv. Neural Inform. Process. Syst.* pp. 4790–4798 (2016)
31. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2536–2544 (2016)
32. Pautrat, R., Barath, D., Larsson, V., Oswald, M.R., Pollefeys, M.: Deeplsd: Line segment detection and refinement with deep image gradients. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 17327–17336 (2023)

33. Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with VQ-VAE-2. In: *Adv. Neural Inform. Process. Syst.* pp. 14837–14847 (2019)
34. Roy, A., Saffar, M., Vaswani, A., Grangier, D.: Efficient content-based sparse attention with routing transformers. *Trans. Assoc. Comput. Linguistics* **9**, 53–68 (2021)
35. Shabani, M.A., Hosseini, S., Furukawa, Y.: Housediffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5466–5475 (2023)
36. Siddiqui, Y., Alliegro, A., Artemov, A., Tommasi, T., Sirigatti, D., Rosov, V., Dai, A., Nießner, M.: Meshgpt: Generating triangle meshes with decoder-only transformers (2023)
37. Sun, J., Wu, W., Liu, L., Min, W., Zhang, G., Zheng, L.: Wallplan: synthesizing floorplans by learning to generate wall graphs. *ACM Trans. Graph.* **41**(4), 92:1–92:14 (2022)
38. Sun, Y., Wang, Y., Liu, Z., Siegel, J., Sarma, S.: Pointgrow: Autoregressively learned point cloud generation with self-attention. In: *Winter Conf. Appl. Comput. Vis.* pp. 61–70 (2020)
39. Tan, X., Zhang, D., Tian, L., Wu, Y., Chen, Y.: Coarse-to-fine pipeline for 3d wireframe reconstruction from point cloud. *Comput. Graph.* **106**, 288–298 (2022)
40. Upadhyay, A., Dubey, A., Arora, V., Kuriakose, M.S., Agarawal, S.: Flnet: Graph constrained floor layout generation. In: *Int. Conf. Multimedia Expo Workshops.* pp. 1–6 (2022)
41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Adv. Neural Inform. Process. Syst.* pp. 5998–6008 (2017)
42. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In: *Adv. Neural Inform. Process. Syst.* pp. 2692–2700 (2015)
43. Wang, S., Zeng, W., Chen, X., Ye, Y., Qiao, Y., Fu, C.: Actfloor-gan: Activity-guided adversarial networks for human-centric floorplan design. *IEEE Trans. Vis. Comput. Graph.* **29**(3), 1610–1624 (2023)
44. Wu, W., Fu, X., Tang, R., Wang, Y., Qi, Y., Liu, L.: Data-driven interior plan generation for residential buildings. *ACM Trans. Graph.* **38**(6), 234:1–234:12 (2019)
45. Xue, N., Bai, S., Wang, F., Xia, G.S., Wu, T., Zhang, L.: Learning attraction field representation for robust line segment detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1595–1603 (2019)
46. Xue, N., Tan, B., Xiao, Y., Dong, L., Xia, G.S., Wu, T.: Volumetric wireframe parsing from neural attraction fields (2023)
47. Xue, N., Wu, T., Bai, S., Wang, F., Xia, G., Zhang, L., Torr, P.H.S.: Holistically-attracted wireframe parsing: From supervised to self-supervised learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(12), 14727–14744 (2023)
48. Yan, X., Lin, L., Mitra, N.J., Lischinski, D., Cohen-Or, D., Huang, H.: Shapeformer: Transformer-based shape completion via sparse representation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 6239–6249 (2022)
49. Yu, L., Lezama, J., Gundavarapu, N.B., Versari, L., Sohn, K., Minnen, D., Cheng, Y., Gupta, A., Gu, X., Hauptmann, A.G., Gong, B., Yang, M.H., Essa, I., Ross, D.A., Jiang, L.: Language model beats diffusion – tokenizer is key to visual generation (2023)
50. Yue, Y., Kontogianni, T., Schindler, K., Engelmann, F.: Connecting the Dots: Floorplan Reconstruction Using Two-Level Queries. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2023)

51. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: LION: latent point diffusion models for 3d shape generation. In: *Adv. Neural Inform. Process. Syst.* (2022)
52. Zhang, S., Tong, H., Xu, J., Maciejewski, R.: Graph convolutional networks: a comprehensive review. *Comput. Soc. Netw.* **6**(1), 1–23 (2019)
53. Zhou, L., Du, Y., Wu, J.: 3d shape generation and completion through point-voxel diffusion. In: *Int. Conf. Comput. Vis.* pp. 5826–5835 (2021)
54. Zhou, Y., Qi, H., Ma, Y.: End-to-end wireframe parsing. In: *Int. Conf. Comput. Vis.* pp. 962–971 (2019)
55. Zhou, Y., Qi, H., Zhai, Y., Sun, Q., Chen, Z., Wei, L., Ma, Y.: Learning to reconstruct 3d manhattan wireframes from a single image. In: *Int. Conf. Comput. Vis.* pp. 7697–7706 (2019)