

PASTA: Part-Aware Sketch-to-3D Shape Generation with Text-Aligned Prior

Seunggwon Lee¹ Hwanhee Jung¹ Byoungsoo Koh² Qixing Huang³ Sangho Yoon⁴ Sangpil Kim^{1*}

¹Korea University ²KOCCA ³The University of Texas at Austin ⁴KAIST

{gwan7801, hwanyo14, spk7}@korea.ac.kr kbs0753@gmail.com

huangqx@cs.utexas.edu sangho@kaist.ac.kr

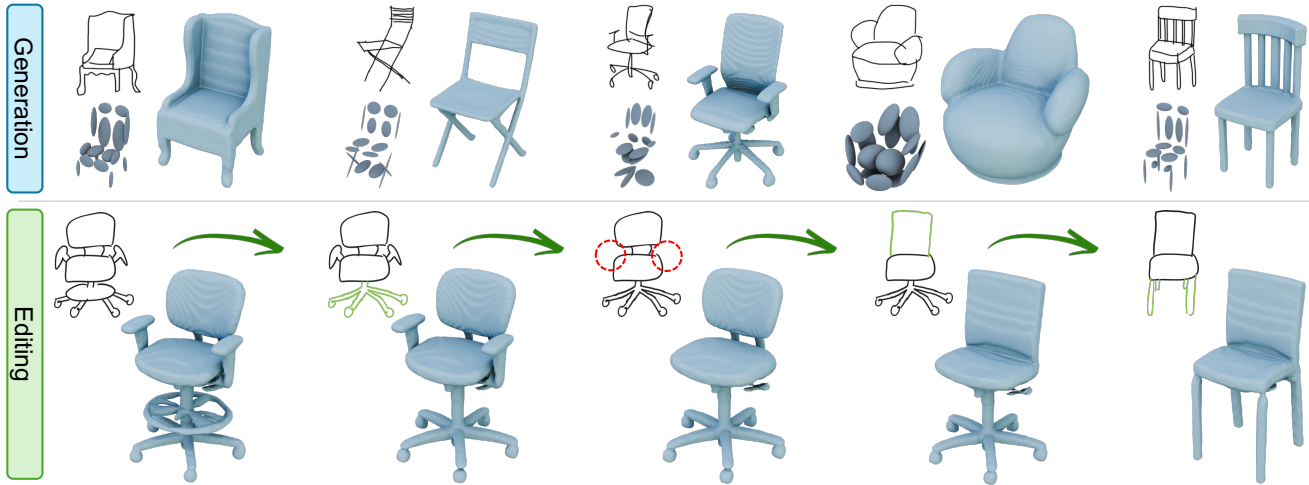


Figure 1. **PASTA** enables 3D shape generation and editing from hand-drawn sketches. It processes a sketch input using Gaussian mixture models (GMMs) for part-level decomposition, generating a corresponding 3D mesh. The framework also supports localized modifications, such as adding, deleting, or transforming parts, for precise user interactive editing.

Abstract

A fundamental challenge in conditional 3D shape generation is to minimize the information loss and maximize the intention of user input. Existing approaches have predominantly focused on two types of isolated conditional signals, i.e., user sketches and text descriptions, each of which does not offer flexible control of the generated shape. In this paper, we introduce **PASTA**, the flexible approach that seamlessly integrates a user sketch and a text description for 3D shape generation. The key idea is to use text embeddings from a vision-language model to enrich the semantic representation of sketches. Specifically, these text-derived priors specify the part components of the object, compensating for missing visual cues from ambiguous sketches. In addition, we introduce **ISG-Net** which employs two types of graph convolutional networks: *IndivGCN*, which processes fine-grained details, and *PartGCN*, which aggregates these details into parts and refines the structure of objects. Extensive experiments demonstrate that **PASTA** outperforms existing methods in part-level editing and achieves state-of-the-art results in sketch-to-3D shape generation.

*Corresponding author

1. Introduction

Conditional 3D shape generation, which generates a 3D shape from user constraints, has many applications in design, gaming, and virtual reality. So far, existing approaches have predominantly focused on two types of conditional signals, namely user sketches [31–33, 41, 70, 71] and text descriptions [13, 38, 44, 54, 65]. Most efforts focus on building machine learning models to achieve improved performance under these two conditional generation tasks. In this paper, we look at the critical question of what are the most effective and informative conditional signals for 3D shape generation? Such conditional signals should be concise, easy to specify, and offer accurate control of the shape to be generated. Unfortunately, both sketch-based and text-based conditions exhibit fundamental limitations. Text descriptions are certainly concise and offer shape semantics, but do not provide precise control of shape geometries. User sketches, on the other hand, offer better control on shape geometries, but they suffer from limited semantics and ambiguities in a single 2D sketch, c.f., [5, 6, 18]. In this paper, we study how to develop a new conditional 3D shape generation approach by seamlessly integrating user sketches and text descriptions.

To this end, we propose a novel framework called PASTA that effectively integrates complementary information to generate 3D models from sketches. Specifically, to compensate for semantic cues that the visual backbone might overlook, we leverage text embeddings extracted from the input sketches. These text embeddings, which are grounded in the rich knowledge of a pre-trained vision-language model (VLM) [36], provide detailed descriptions of object components, such as the number of legs on a chair or the presence of an armrest, that enhance the information provided by the visual backbone. By integrating both text and visual information, PASTA can more precisely understand object composition and semantics, ultimately generating accurate 3D models.

Similar to previous studies [5, 6], our model represents 3D shapes via decomposed Gaussian mixture models (GMMs) through SPAGHETTI shape decoder [21], making interactions between GMMs critically important. Accordingly, along with text priors, we introduce the Integrated Structure-Graph Network (ISG-Net), a specialized refinement module that employs two types of graph convolutional networks (GCNs) to process the sketch. The first network, *IndivGCN*, focuses on processing the details of each individual GMM, capturing subtle features within the sketch. The second network, *PartGCN*, aggregates these fine details into parts and refines the structure of the object, improving the consistency and spatial relationships between the different parts. Together, these components enable a more comprehensive structural understanding of the sketch, facilitating part-level editing with superior accuracy and flexibility compared to existing methods.

To evaluate the robustness of PASTA, We also conducted an additional experiment to extend the domain to real-world images. Our comprehensive experiments demonstrate that PASTA achieves state-of-the-art performance in qualitative and quantitative evaluations, generating more complete and structurally sound 3D models than prior approaches. The summary of our contributions is as follows:

- We propose PASTA, a novel approach for sketch-to-3D shape generation that integrates visual and text conditions from VLM to compensate for missing visual cues, enhancing the comprehension of object components in the provided user sketch.
- We introduce ISG-Net, a specialized refinement module incorporating two types of GCNs: *IndivGCN* for capturing fine-grained details and *PartGCN* for ensuring structural integrity at the part-level.
- Achieving state-of-the-art performance, PASTA outperforms existing methods in qualitative and quantitative evaluations, generating structurally accurate 3D models.
- Additionally, the scalability of PASTA is demonstrated, showing its ability to extend to real-world images and effectively handle realistic visual data beyond sketches.

2. Related Work

3D Shape Representation. Unlike 2D grid images with RGB channels, 3D data can be represented in a variety of ways. Voxel-based 3D shape generation methods [37, 56, 63] have been widely explored; however, they incur significant computational costs and suffer from limited expressive power at low resolutions. Another approach employs point cloud representations [1, 34, 42, 74], which offer a sparser and more lightweight alternative, albeit being unstructured and lacking explicit connectivity information. To address these limitations, numerous studies have focused on mesh-based 3D shape generation [17, 61, 67]. However, mesh-based approaches often necessitate complex operations to manage topological variations and high-genus structures. More recently, neural implicit representations such as signed distance functions and occupancy fields have garnered considerable attention due to their flexibility and ability to model shapes at arbitrary resolutions with compelling performance [12, 14, 29, 43, 69]. In this work, we adopt an implicit modeling strategy [21] based on a Gaussian mixture model (GMM), which enables flexible shape representation and part-level editing.

Sketch-based 3D Shape Generation. Sketch-based 3D shape generation seeks to reconstruct or generate 3D shapes from 2D sketches, bridging the gap between abstract drawings and detailed models. Prior works have focused on single-view [32, 33, 70, 71] or multi-view sketches [31, 41], yet these methods often fail to recover intricate details due to the oversimplification inherent in the sketches, inconsistencies across different views, complicating the reconstruction process. Recent view-aware approaches [69, 72] have made notable improvements by incorporating viewpoint-specific information. In addition, user interactive methods [5, 6, 18] have similarly advanced the field by enabling user-guided editing. Despite these advancements, fully reconstructing the intended shape from sketches remains challenging. Therefore, in this paper, we propose a method that leverages part-level representations [21] alongside text priors and graph structures to reinforce semantic representation and inter-part relationships, thereby generating more accurate and consistent 3D models.

Vision-Language Pre-trained Models. In computer vision, pre-trained vision-language models (VLMs) have exerted a profound impact by integrating textual modalities. Pre-trained on large-scale image-text pairs [11, 39, 73], these models have been successfully applied to a variety of downstream tasks, including visual question answering, text-to-image generation, and image captioning. Notably, models such as CLIP [50] and ALIGN [25] utilize contrastive learning to acquire cross-modal representations, thereby enabling the understanding of visual concepts in zero-shot scenarios. More recently, research has

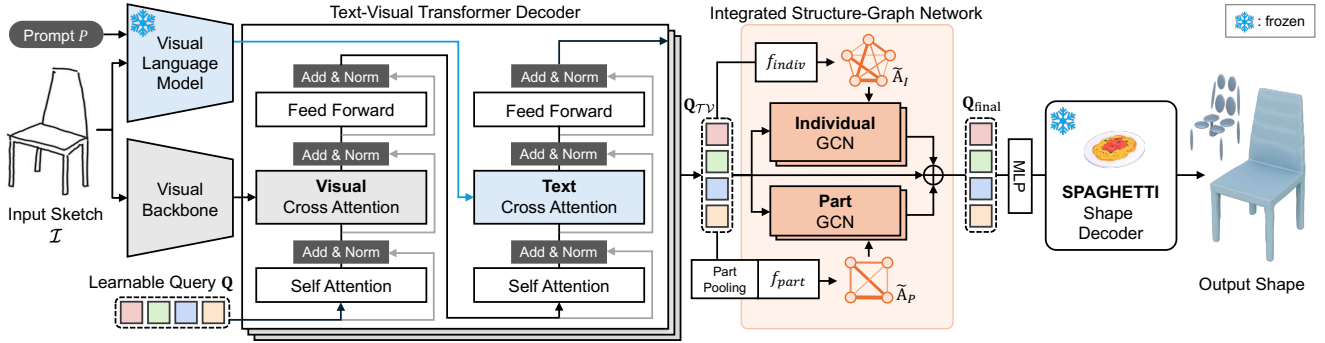


Figure 2. **Overview of PASTA.** Our framework enhances sketch-based 3D shape generation by integrating visual embeddings and text-aligned priors. A visual backbone and vision-language model (VLM) extract meaningful features from an input sketch, which are then processed by a Text-Visual Transformer Decoder with learnable queries. To refine structural details, we introduce ISG-Net, which consists of IndivGCN for fine-grained feature processing and PartGCN for aggregating part-level information. The output features are fed into the SPAGHETTI shape decoder [21], producing a more complete and structurally accurate 3D model.

leveraged the rich knowledge embedded within large language models (LLMs) [2, 7, 15, 16] to enhance visual context comprehension and image processing [3, 35, 36]. This paradigm has further extended to 3D vision tasks, such as 3D scene understanding [23, 24, 49, 51–53]. In our work, we exploit VLMs to accurately recognize object components from sketch images and to strengthen the modeling of their relationships, thereby addressing the inherent limitations of abstract sketches.

3. Method

Our proposed method, PASTA, aims to generate 3D shapes from hand-drawn sketches. To enable more flexible shape modeling, we encode sketches in the latent space of the part-level implicit shape decoder [21]. Consequently, it is crucial to map the information from the sketches accurately to aligned latent vectors, which are then represented as Gaussian mixture models (GMMs) for generating the 3D shapes. To achieve this, we utilize text-aligned priors and graph-based processing. In Fig. 2, we provide an overview of the framework. Given a sketch, a visual backbone [46] processes the input sketch image, while a vision-language model (VLM) [36] captures semantic priors from a sketch. A Text-Vision Transformer Decoder then processes these features using learnable queries to fuse and refine visual and text information. To ensure structural and part-level consistency, we introduce the Integrated Structure-Graph Network (ISG-Net), which comprises an Individual Graph Convolutional Network (IndivGCN) for local feature refinement and a Part Graph Convolutional Network (PartGCN) for maintaining part-level coherence. The refined features are finally fed into the shape decoder [21] to generate a complete 3D shape that reflects the sketch and its semantic attributes. Section 3.1 details the Text-Vision Transformer Decoder, while Section 3.2 discusses the ISG-Net module.

3.1. Text-Visual Transformer Decoder

Due to the ambiguous and oversimplified nature of the sketches, inaccurate shape generation and part omission often occur. To address these challenges, we introduce a text-aligned prior that integrates visual and text condition into the 3D shape generation process. This prior is extracted from vision-language model (VLM) [36] that analyzes the overall composition of the sketch and, through a text-attention mechanism, integrate visual cues to enrich the highly abstract characteristics of the sketch. The integration of these components ensures that the generated shapes are coherently aligned with semantic attributes derived from text embeddings, even when the sketches do not fully convey the intended details. Consequently, our approach improves the ability of the network to perceive the composition of the sketch, resulting in a more robust understanding.

Vision-Language Embedding. Starting with an input sketch, we extract feature representations from distinct modalities using a dual-branch approach. Firstly, the visual backbone [46] produces a visual embedding $\mathcal{V} = f_{\text{vis}}(\mathcal{I})$, where \mathcal{I} denotes the input sketch and $f_{\text{vis}}(\cdot)$ denotes the forward pass of the visual backbone encoder. However, relying solely on this visual information is insufficient to capture the semantic nuances from a simple sketch, leading to sub-optimal reconstruction performance.

Therefore, we leverage semantically meaningful text embeddings obtained from the pre-trained VLM [36], thereby exploiting the knowledge encapsulated within this model. Text embeddings are computed as $\mathcal{T} = f_{\text{vlm}}(\mathcal{I}, P)$, where P denotes the input prompt and $f_{\text{vlm}}(\cdot)$ represents the VLM text and image encoder. As illustrated in Fig. 3, the obtained \mathcal{T} represents the embedding immediately preceding the conversion to the output text description and contains information regarding the overall composition of the sketch (e.g., the shape of the backrest and seat, the number of legs and the absence of armrests).

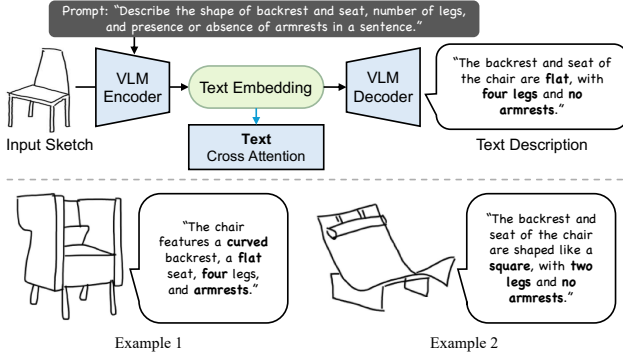


Figure 3. An example of the prompts we used for chair sketches and their various corresponding text descriptions from the VLM [36]. The figure illustrates how the generated text descriptions emphasize key semantic details of the sketch.

Text-Conditioned Query Learning. For accurate 3D shape generation, we employ a Text-Visual Transformer Decoder to integrate visual embeddings \mathcal{V} and text embeddings \mathcal{T} . As the shape decoder [21] expects N latent vectors as input to form N GMMs, we define learnable queries $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N] \in \mathbb{R}^{N \times d}$, where d denotes the dimensionality of each query. Initially, queries are processed by a self-attention mechanism to capture their internal dependencies. For the attention operation, \mathbf{Q} is linearly projected via learnable weight matrices W_Q , W_K , and W_V , which map \mathbf{Q} to query, key, and value vectors, respectively, which is expressed as:

$$\mathbf{Q} = W_Q^T \cdot \mathbf{Q}, \quad \mathbf{K} = W_K^T \cdot \mathbf{Q}, \quad \mathbf{V} = W_V^T \cdot \mathbf{Q}, \quad (1)$$

$$\mathbf{Q} = \text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q} \cdot \mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}. \quad (2)$$

The resulting queries are added to the residuals and normalized before being subjected to cross-attention with the visual embeddings \mathcal{V} . Through visual cross-attention, each of the N queries acquires visual embeddings extracted from sketch, a process that can be expressed as:

$$\mathbf{Q}_{\mathcal{V}} = \text{Attn}(W_Q^T \cdot \mathbf{Q}, W_K^T \cdot \mathcal{V}, W_V^T \cdot \mathcal{V}). \quad (3)$$

During training, the learnable queries are optimized to retrieve the most representative visual cues for each vector \mathbf{q} . Next, the queries $\mathbf{Q}_{\mathcal{V}}$ enter another round of self-attention and feed-forward transformation, further refining their relationships. Subsequently, queries are integrated with text embeddings \mathcal{T} obtained from the VLM through a text cross-attention mechanism:

$$\mathbf{Q}_{\mathcal{T}\mathcal{V}} = \text{Attn}(W_Q^T \cdot \mathbf{Q}_{\mathcal{V}}, W_K^T \cdot \mathcal{T}, W_V^T \cdot \mathcal{T}). \quad (4)$$

This operation ensures that queries $\mathbf{Q}_{\mathcal{T}\mathcal{V}}$ are conditioned not only on visual information but also on the semantic information embedded in the text representation, effectively

integrating both modalities. By incorporating text embeddings, the model enriches the visual queries with complementary linguistic context, allowing for a more comprehensive understanding of object. Finally, a feed-forward layer processes the updated queries $\mathbf{Q}_{\mathcal{T}\mathcal{V}}$ before passing them to the next iteration. This entire process is repeated 12 times, each iteration allowing the queries to progressively capture fine-grained semantic details that the visual backbone might have otherwise overlooked by integrating information from text and visual embeddings.

3.2. ISG-Net

To enhance the structural coherence and geometric accuracy of the generated 3D shapes, we introduce ISG-Net, a GCN-based module that refines the queries $\mathbf{Q}_{\mathcal{T}\mathcal{V}}$ through relational reasoning. ISG-Net consists of two key components: IndivGCN, which refines fine-grained local features, and PartGCN, which aggregates part-level structural information to enforce global consistency. By jointly leveraging these two networks, ISG-Net encodes both local geometric details and high-level structural relationships, ensuring a faithful generation of the target 3D shape.

3.2.1. IndivGCN: Fine-Grained Feature Processing

The IndivGCN module considers the relationships among individual queries to integrate features from neighboring nodes. For graph-based operations, obtaining an adjacency matrix is essential; therefore, we introduce an individual adjacency predictor f_{indiv} composed of a multilayer perceptron (MLP) to define the interquery relationships. As illustrated in Fig. 4 (a), the predicted $\mathbf{Q}_I = f_{\text{indiv}}(\mathbf{Q}_{\mathcal{T}\mathcal{V}})$ is used to compute the individual adjacency matrix $\tilde{\mathbf{A}}_I = \mathbf{Q}_I^T \mathbf{Q}_I \in \mathbb{R}^{N \times N}$, which is supervised by a pseudo-ground truth \mathbf{A}_I . The loss $\mathcal{L}_{\text{indiv}}$ is computed using the mean squared error (MSE) as follows:

$$\mathcal{L}_{\text{indiv}} = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \left\| \tilde{\mathbf{A}}_I(i, j) - \mathbf{A}_I(i, j) \right\|_2^2. \quad (5)$$

Regarding the pseudo-ground truth, \mathbf{A}_I is derived based on the positions $\mu \in \mathbb{R}^{N \times 3}$ of the N GMMs. The 3D shape is converted into GMMs by inversion of SPAGHETTI [21], and the distances between all Gaussian means μ are computed. Note that these distances are normalized to the interval between 0 and 1, then inverted, as smaller distances correspond to stronger connections. Through this, f_{indiv} is optimized to predict the distances between the actual GMM centers, and the derived adjacency matrix is then employed to perform graph convolution as follows:

$$\mathbf{Q}_{\text{indiv}} = \sigma(\tilde{\mathbf{A}}_I \mathbf{Q}_{\mathcal{T}\mathcal{V}} \mathbf{W}_I), \quad (6)$$

where σ is a non-linear activation function and \mathbf{W}_I is the learnable weight matrix for graph convolution. This formulation allows the network to dynamically learn relationships, refining feature representations of individual nodes.

3.2.2. PartGCN: Structural aggregate at the part-level

To further refine the details through IndivGCN and enforce structural consistency, we perform graph convolution by grouping GMMs into subsets at the part-level. First, hierarchical clustering [68] is used to produce indices that assign the N queries to the K parts ($K < N$), with further details provided in the supplementary material. The indices obtained are used to perform an average pooling operation on the queries, resulting in part-level queries $\mathbf{Q}_P \in \mathbb{R}^{K \times d}$, which in turn are used by f_{part} to predict $\tilde{\mathbf{A}}_P \in \mathbb{R}^{K \times K}$. The loss $\mathcal{L}_{\text{part}}$ is also calculated using MSE as follows:

$$\mathcal{L}_{\text{part}} = \frac{1}{K^2} \sum_{i=1}^K \sum_{j=1}^K \left\| \tilde{\mathbf{A}}_P(i, j) - \mathbf{A}_P(i, j) \right\|_2^2, \quad (7)$$

where \mathbf{A}_P denotes the adjacency matrix of the pseudo-ground truth part. As shown in Fig. 4 (b), it is constructed based on the part coordinates obtained by applying clustering on the distances of the GMMs, similar to \mathbf{A}_T . With the queries $\mathbf{Q}_{\mathcal{T}\mathcal{V}}$ and the predicted $\tilde{\mathbf{A}}_P$, we perform the following operations:

$$\mathbf{Q}_{\text{part}} = \sigma(\tilde{\mathbf{A}}_P \mathbf{Q}_{\mathcal{T}\mathcal{V}} \mathbf{W}_P). \quad (8)$$

After obtaining the part-level representations, the features are unpooled to match the size of the individual queries, ensuring compatibility with IndivGCN. The final representation is obtained by fusing the output from IndivGCN and PartGCN with the residual connection:

$$\mathbf{Q}_{\text{final}} = \text{norm}(\alpha \mathbf{Q}_{\text{indiv}} + (1 - \alpha) \mathbf{Q}_{\text{part}} + \mathbf{Q}_{\mathcal{T}\mathcal{V}}), \quad (9)$$

where α is a weighting factor that balances the impact of two GCNs. The result is subsequently normalized to produce $\mathbf{Q}_{\text{final}}$, which is processed by an MLP to infer latent vectors $\tilde{\mathbf{z}}$. These vectors are then passed through the shape decoder, which predicts occupancy fields and generates a 3D mesh via the marching cubes algorithm [40].

3.2.3. Loss Function

To leverage the part-level implicit shape decoder [21], that is crucial that the vectors predicted by our network are accurately aligned with the latent space of the pre-trained shape decoder. To ensure that the predicted vectors $\tilde{\mathbf{z}}$ remain close to the ground truth latent vectors \mathbf{z} in the feature space, we optimize the network using the following loss function:

$$\mathcal{L}_{\text{align}} = \frac{1}{N} \sum_{i=1}^N \|\tilde{\mathbf{z}}_i - \mathbf{z}_i\|_1, \quad (10)$$

where \mathbf{z} represents the latent vectors extracted from the 3D shape using the inversion of SPAGHETTI [21]. The total loss function \mathcal{L} is defined as the sum of three components:

$$\mathcal{L} = \lambda_{\text{align}} \mathcal{L}_{\text{align}} + \lambda_{\text{indiv}} \mathcal{L}_{\text{indiv}} + \lambda_{\text{part}} \mathcal{L}_{\text{part}}. \quad (11)$$

In this equation, λ denotes the weighting factors for each loss term. By jointly optimizing all three losses, PASTA achieves structural coherence in 3D shape generation.

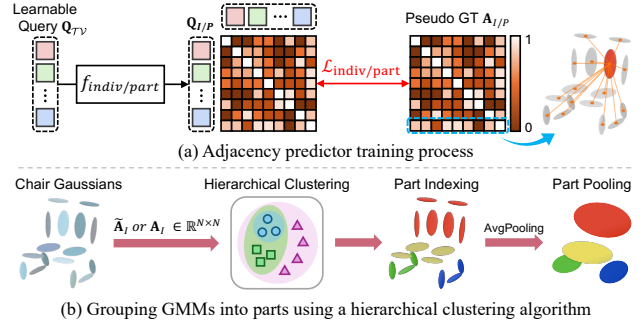


Figure 4. Detailed illustration of ISG-Net. (a) Shows the training process of the adjacency predictor, which learns the relationships between queries. (b) Demonstrates the process of grouping GMMs into parts using a hierarchical clustering algorithm to better capture structural dependencies.

4. Experiments

4.1. Experimental Setup

Datasets. We employ a subset of the ShapeNet dataset [10] along with the corresponding sketches. Specifically, the chair dataset comprises 6,755 shapes, the airplane dataset contains 1,575, and the lamp dataset includes 733. Each shape is normalized to a unit cube scale and rendered from 6 viewpoints. Abstract sketches are obtained using CLIPasso [60] from these rendered images. For the chair category, we additionally leverage non-photo-realistic renderings [9], which map photographs into line drawings.

For evaluation, we conduct two experiments for the chair category. First, we utilize the AmateurSketch-3D [48] dataset containing 1,000 chair shapes along with corresponding hand-drawn sketches from three different viewpoints. Second, we evaluate 500 chair shapes on the expert-drawn sketch dataset, ProSketch-3D[70], which provides more detailed representations. For the airplane and lamp categories, sketches corresponding to 200 airplane shapes and 100 lamp shapes are obtained with CLIPasso. Following the protocols adopted in prior work [5, 6, 47], the latent vectors used during training are derived from 3D shapes through the inversion functionality provided by SPAGHETTI [21].

Evaluation Metrics. We evaluate the quality of the 3D shapes generated using three metrics, with lower values indicating better results. First, Chamfer Distance (CD) measures the average distance between the corresponding points in two-point clouds. Similarly, Earth Mover Distance (EMD) quantifies the minimal cost of transforming one point cloud into another, capturing structural differences. Lastly, Fréchet Inception Distance (FID) [22] compares the feature distributions of the rendered images from generated 3D shapes with those of actual 3D shape data using a pre-trained Inception network. Detailed information and equations for the three evaluation metrics are provided in the supplementary material.

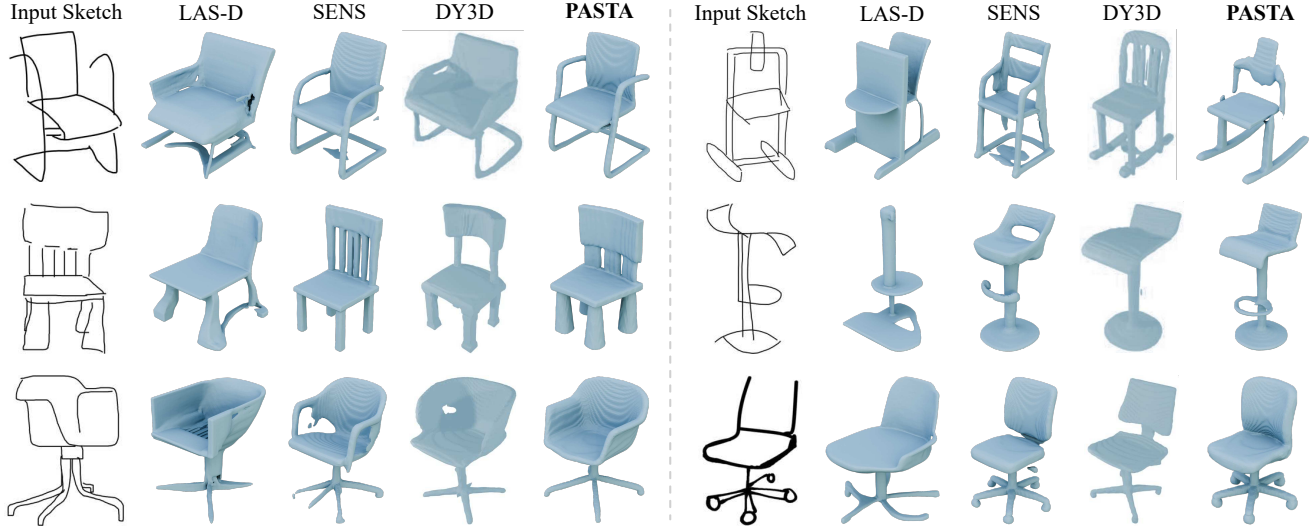


Figure 5. Qualitative comparison of chair reconstructions. The figure demonstrates the effectiveness of our methods in generating high-quality 3D chair models, accurately capturing intricate details such as armrests, backrests, and leg structures, while also preserving overall global structures, compared to LAS-D [69], SENS [6], and DY3D [5].

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
Sketch2Mesh [18]	0.257	0.211	392.2	0.228	0.171	297.8
LAS-D [69]	0.159	0.128	197.5	0.195	0.147	193.5
SENS [6]	0.121	0.096	171.3	0.116	0.076	160.5
DY3D [5]	0.109	0.091	-	0.093	0.087	-
PASTA (Ours)	0.090	0.071	143.9	0.055	0.049	112.2

Table 1. Quantitative comparison of different methods on the AmateurSketch-3D [48] and ProSketch-3D [70] datasets. The performance is evaluated using CD, EMD, and FID, where lower values indicate better performance.

4.2. Quantitative Analysis

In this section, we evaluate the performance of PASTA, on the AmateurSketch-3D [48] and ProSketch-3D [70]. When evaluating the AmateurSketch-3D and ProSketch-3D dataset, we train our model using two sketch styles [9, 60]. We compare our method with several state-of-the-art approaches [5, 6, 18, 69] that take user sketches as input. As shown in Tab. 1, PASTA consistently outperforms all competing methods in both datasets. Specifically, our method achieves the lowest CD and EMD values, demonstrating superior geometric accuracy in the 3D shape reconstruction task. Moreover, PASTA also achieves the best FID scores, further confirming its ability to generate highly realistic and accurate 3D shapes. Additionally, as shown in Tab. 2, ours performs particularly well in categories such as airplane and lamp, showcasing its versatility across various object types. Notably, PASTA excels in preserving missing visual cues in oversimplified sketches, demonstrating improvements over existing methods, and highlighting its effectiveness in capturing detailed 3D structures.

Methods	Airplane			Lamp		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
SENS [6]	0.240	0.162	97.2	0.253	0.194	152.8
PASTA (Ours)	0.188	0.113	90.5	0.195	0.147	125.6

Table 2. Quantitative comparison of SENS [6] on the airplane and lamp datasets obtained with CLIPasso [60]. PASTA demonstrates superior result across all three metrics in both categories.

4.3. Qualitative Analysis

To further evaluate the performance of our model, we present qualitative results across three categories of objects, including chairs, airplanes, and lamps, as well as its ability to edit shapes. Visual comparisons demonstrate that our model excels in capturing fine-grained details and complex structures more accurately than competing methods.

Chairs. We compare PASTA with state-of-the-art methods, including LAS-D [69], SENS [6], and DY3D [5], to demonstrate superior performance in reconstructing diverse chair designs with accurately capture intricate details and global structures. As shown in Fig. 5, PASTA effectively preserves subtle design elements such as armrests, backrests, and leg structures, which are often oversimplified in sketches. Our 3D shape generations exhibit more refined geometries, confirming the robustness of our approach.

Airplanes. For the airplane category, the qualitative results in Fig. 6 show that our method produces more accurate and well-structured reconstructions than the previous model [6]. PASTA effectively captures essential structural components, including wings, tails, and engines, also maintaining fine-grained details, while SENS often struggles with preserving these details.




Figure 6. 3D generation results of SENS [6] and PASTA for airplane and lamp sketches. The shapes generated by our method accurately capture structural details for both object types.

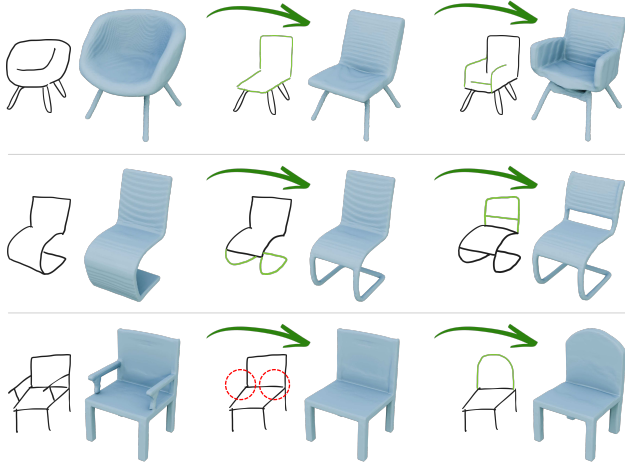


Figure 7. Qualitative results of shape editing. Our model enables intuitive modifications, allowing for controlled adjustments while preserving structural consistency.

Lamps. We also evaluate our model in lamp sketches, where it successfully reconstructs a variety of lamp styles, from simple desk lamps to chandelier-like designs. As depicted in Fig. 6, PASTA generates plausible 3D shapes that closely align with the input sketches while preserving intricate design elements such as the head shape, lamp base, and delicate frame structures.

Shape Editing. Our method provides interactive editing capability that enables users to precisely localize and modify the intended parts, thanks to the GMM-based decomposed part-level representation. As illustrated in Fig. 7, PASTA obviously reflects changes in appearance, as well as modifications, additions, and deletions. For example, when the shapes of the backrest, seat, and legs are altered across all rows, the intended modifications are accurately reflected while maintaining the shape of other parts. Moreover, in the first row, the addition of armrests is clearly depicted, in the second row, the legs and backrest are modified into the desired shapes, and in the third row, their removal is effectively captured. These experiments demonstrate that our model effectively learns the composition of objects and the relationships of the parts.

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
Baseline	0.117	0.095	169.2	0.109	0.076	157.9
w/ IndivGCN	0.103	0.087	158.8	0.089	0.065	148.4
w/ PartGCN	0.109	0.090	163.1	0.093	0.071	152.0
w/ ISG-Net	0.098	0.083	155.7	0.068	0.054	129.2

Table 3. Ablation study to examine the impact of different GCN configurations, including “w/ IndivGCN”, “w/ PartGCN”, and “w/ ISG-Net”, which refers to the model with both IndivGCN and PartGCN applied.

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
Baseline	0.117	0.095	169.2	0.109	0.076	157.9
w/ Text	0.101	0.086	157.6	0.081	0.063	135.5
w/ ISG-Net	0.098	0.083	155.7	0.068	0.054	129.2
PASTA(Ours)	0.090	0.071	143.9	0.055	0.049	112.2

Table 4. Evaluating the impact of text embedding and ISG-Net on model performance. The “w/ Text” and “w/ ISG-Net” rows show the effects of adding each component individually, while the final PASTA model combines both for improved performance.

4.4. Ablation Study

In this study, we conducted two ablation studies to assess the impact of key contributions in our framework by comparing them against a “Baseline” without any of these components. In particular, we examine the effectiveness of IndivGCN and PartGCN in capturing structural information from input sketches. Additionally, we evaluate the contributions of text embedding and ISG-Net to the overall performance, analyzing how they enhance the quality of the generated 3D shapes.

Impact of IndivGCN and PartGCN. To assess the effectiveness of two GCNs, we compare the following settings in Tab. 3: (i) Baseline, (ii) IndivGCN applied only, (iii) PartGCN applied only, and (iv) ISG-Net (IndivGCN and PartGCN applied together). This result shows that using IndivGCN enhances performance by capturing detailed information from individual queries, while the application of PartGCN improves global structure by leveraging part awareness. Combining both modules results in even better performance, highlighting their complementary roles in feature extraction.

Effect of Text-Aligned Prior and ISG-Net. We examine the influence of the text-aligned prior and ISG-Net by comparing Tab. 4: (i) Baseline, (ii) text embedding applied only, (iii) ISG-Net applied only, and (iv) PASTA (text embedding and ISG-Net applied together). When applying only the text-aligned prior, a significant performance improvement is observed over the baseline across all three metrics, highlighting the importance of the text condition. Additionally, using only ISG-Net also shows improved perfor-



Figure 8. Examples of sketches from CLIPasso [60], and non-photo-realistic rendering [9], along with their converted versions using ControlNet [66].



Figure 9. Comparison of outputs when using real-world photos as input. We compare our model with SENS [6], showing the improved accuracy and detail in the output. Results demonstrate the ability of PASTA to extend the domain over the sketch.

mance, which validates the effectiveness of learning structural relationships through the two GCN modules. Ultimately, combining both methods achieves the highest performance, demonstrating that their integration leads to optimal results.

4.5. Qualitative Results on Real-Image

To assess the general robustness of the proposed method, we extend qualitative comparisons to real-world image domains. Specifically, ControlNet [66] is utilized to convert grayscale sketches from training datasets [9, 60] into RGB images while generating colored outputs that preserve the original sketch’s shapes and enhance realism, as illustrated in Fig. 8. To ensure a fair comparison, PASTA is trained on these RGB images alongside SENS [6]. During inference, as shown in Fig. 9, real-world photographs with varying shapes in different places are utilized to further evaluate the model’s generalization ability. The results demonstrate that PASTA achieves superior accuracy and detail in 3D shape generation from real images compared to previous method. These findings underscore the effectiveness of the proposed approach in structural analysis and interpretation, validating its applicability beyond grayscale sketch data to RGB real-world images.

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
Baseline	0.117	0.095	169.2	0.109	0.076	157.9
w/ Part Type	0.103	0.088	159.5	0.092	0.069	142.8
w/ Verbose	0.105	0.088	161.8	0.085	0.066	134.1
w/ Single Sent.	0.101	0.086	157.6	0.081	0.063	135.5

Table 5. Evaluating the impact of various text description formats for text condition, including Part Type, Verbose (more extended sentences), and Single Sentences.

4.6. Effect of Description Types in VLM

In this section, we further evaluate the impact of text embeddings by comparing results across various description types. As shown in Tab. 5, using embeddings that describe only the type of parts (e.g., backrest, armrest, legs) results in improved performance, which demonstrates the effectiveness of text embeddings. This approach helps the model recognize basic structures and parts. However, listing only the parts is not enough to capture detailed features or complex structures. Conversely, providing excessively verbose explanations yields some improvement but often leads to overstatements or inaccuracies such as hallucinations, which introduce noise that degrades performance. Therefore, we use embeddings that specify the exact number and appearance of parts in a sentence, as shown in Fig. 3. This approach effectively compensates for missing visual cues and achieves optimal performance on both benchmarks [48, 70]. More detailed examples of part type and verbose descriptions are provided in the supplementary material.

5. Conclusion

We present PASTA as a significant advancement in sketch-based 3D shape generation, addressing key limitations of prior methods [5, 6, 69]. Leveraging text conditions from VLM [36] enhances the understanding of object composition, while ISG-Net improves fine-grained detail processing and refines structural coherence through IndivGCN and PartGCN. Additionally, PASTA also enables more precise and flexible part-level editing. Moreover, PASTA proves its versatility by enabling 3D shape generation not only from sketches but also from real-world images, emphasizing its scalability. Comprehensive experiments, including both qualitative and quantitative analyses, PASTA is established as a significant advancement in 3D shape generation, demonstrating its superiority in structural coherence, detail preservation, and part-level flexibility.

Limitations & Future Work. Our approach uses a shape decoder, SPAGHETTI [21], enabling fast training and flexible generation. However, it is limited to a few classes and a fixed number of Gaussians. Future work may explore removing these constraints and incorporating a dedicated implicit decoder to support a wider range of categories.

References

- [1] 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. 2
- [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023. 3
- [3] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023. 3
- [4] Inhwon Bae, Jin-Hwi Park, and Hae-Gon Jeon. Learning pedestrian group representations for multi-modal trajectory prediction. In *European Conference on Computer Vision*, pages 270–289. Springer, 2022. 4
- [5] Hmrishav Bandyopadhyay, Subhadeep Koley, Ayan Das, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Doodle your 3d: From abstract freehand sketches to precise 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9795–9805, 2024. 1, 2, 5, 6, 8
- [6] Alexandre Binniger, Amir Hertz, Olga Sorkine-Hornung, Daniel Cohen-Or, and Raja Giryes. Sens: Part-aware sketch-based implicit neural shape modeling. In *Computer Graphics Forum*, page e15015. Wiley Online Library, 2024. 1, 2, 5, 6, 7, 8
- [7] 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. 3
- [8] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013. 4
- [9] Caroline Chan, Frédo Durand, and Phillip Isola. Learning to generate line drawings that convey geometry and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7915–7925, 2022. 5, 6, 8, 1
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [11] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. 2
- [12] 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. 2
- [13] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G. Schwing, and Liangyan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 4456–4465. IEEE, 2023. 1
- [14] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465, 2023. 2
- [15] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6, 2023. 3
- [16] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024. 3
- [17] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems*, 35:31841–31854, 2022. 2
- [18] Benoit Guillard, Edoardo Remelli, Pierre Yvernavy, and Pascal Fua. Sketch2mesh: Reconstructing and editing 3d shapes from sketches. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13023–13032, 2021. 1, 2, 6
- [19] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017. 4
- [20] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015. 4
- [21] Amir Hertz, Or Perel, Raja Giryes, Olga Sorkine-Hornung, and Daniel Cohen-Or. Spaghetti: Editing implicit shapes through part aware generation. *ACM Transactions on Graphics (TOG)*, 41(4):1–20, 2022. 2, 3, 4, 5, 8
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5, 2
- [23] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023. 3
- [24] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22157–22167, 2023. 3

- [25] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 2
- [26] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. 4
- [27] Sungjune Kim, Hadam Baek, Seungwan Lee, Hyung-gun Chi, Hyerin Lim, Jinkyu Kim, and Sangpil Kim. Enhanced motion forecasting with visual relation reasoning. In *European Conference on Computer Vision*, pages 311–328. Springer, 2024. 4
- [28] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 4
- [29] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial generation of continuous implicit shape representations. *arXiv preprint arXiv:2002.00349*, 2020. 2
- [30] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018. 4
- [31] Changjian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. Robust flow-guided neural prediction for sketch-based freeform surface modeling. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. 1, 2
- [32] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. Sketch2cad: Sequential cad modeling by sketching in context. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 2
- [33] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. Free2cad: Parsing freehand drawings into cad commands. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 1, 2
- [34] Chun-Liang Li, Manzil Zaheer, Yang Zhang, Barnabas Poczos, and Ruslan Salakhutdinov. Point cloud gan. *arXiv preprint arXiv:1810.05795*, 2018. 2
- [35] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 3
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 2, 3, 4, 8
- [37] Jerry Liu, Fisher Yu, and Thomas Funkhouser. Interactive 3d modeling with a generative adversarial network. In *2017 International Conference on 3D Vision (3DV)*, pages 126–134. IEEE, 2017. 2
- [38] Zhengzhe Liu, Yi Wang, Xiaojuan Qi, and Chi-Wing Fu. Towards implicit text-guided 3d shape generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022*, pages 17875–17885. IEEE, 2022. 1
- [39] Siqu Long, Feiqi Cao, Soyeon Caren Han, and Haiqin Yang. Vision-and-language pretrained models: A survey. *arXiv preprint arXiv:2204.07356*, 2022. 2
- [40] 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. 5
- [41] Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhansu Maji, and Rui Wang. 3d shape reconstruction from sketches via multi-view convolutional networks. In *2017 International Conference on 3D Vision (3DV)*, pages 67–77. IEEE, 2017. 1, 2
- [42] 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. 2
- [43] 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. 2
- [44] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 1
- [45] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pages 2014–2023. PMLR, 2016. 4
- [46] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3
- [47] 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. 5
- [48] Anran Qi, Yulia Gryaditskaya, Jifei Song, Yongxin Yang, Yonggang Qi, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Toward fine-grained sketch-based 3d shape retrieval. *IEEE transactions on image processing*, 30:8595–8606, 2021. 5, 6, 8, 1
- [49] Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. In *European Conference on Computer Vision*, pages 214–238. Springer, 2024. 3
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3
- [51] Wonseok Roh, Seung Hyun Lee, Won Jeong Ryoo, Jakyung Lee, Gyeongrok Oh, Sooyeon Hwang, Hyung-Gun Chi,

- and Sangpil Kim. Functional hand type prior for 3d hand pose estimation and action recognition from egocentric view monocular videos. In *BMVC*, page 193, 2023. 3
- [52] Wonseok Roh, Hwanhee Jung, Jong Wook Kim, Seung-gwan Lee, Innfarn Yoo, Andreas Lugmayr, Seunggeun Chi, Karthik Ramani, and Sangpil Kim. Catsplat: Context-aware transformer with spatial guidance for generalizable 3d gaussian splatting from a single-view image. *arXiv preprint arXiv:2412.12906*, 2024.
- [53] Wonseok Roh, Hwanhee Jung, Giljoo Nam, Jinseop Yeom, Hyunje Park, Sang Ho Yoon, and Sangpil Kim. Edge-aware 3d instance segmentation network with intelligent semantic prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20644–20653, 2024. 3
- [54] Aditya Sanghi, Rao Fu, Vivian Liu, Karl DD Willis, Hooman Shayani, Amir H Khasahmadi, Srinath Sridhar, and Daniel Ritchie. Clip-sculptor: Zero-shot generation of high-fidelity and diverse shapes from natural language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18339–18348, 2023. 1, 3
- [55] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1): 61–80, 2008. 4
- [56] Edward J Smith and David Meger. Improved adversarial systems for 3d object generation and reconstruction. In *Conference on Robot Learning*, pages 87–96. PMLR, 2017. 2
- [57] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, pages 369–386. SPIE, 2019. 1
- [58] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 2
- [59] Damien Teney, Lingqiao Liu, and Anton van Den Hengel. Graph-structured representations for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2017. 4
- [60] Yael Vinker, Ehsan Pajouheshgar, Jessica Y Bo, Roman Christian Bachmann, Amit Haim Bermano, Daniel Cohen-Or, Amir Zamir, and Ariel Shamir. Clipasso: Semantically-aware object sketching. *ACM Transactions on Graphics (TOG)*, 41(4):1–11, 2022. 5, 6, 8, 1
- [61] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 52–67, 2018. 2
- [62] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Learning combinatorial embedding networks for deep graph matching. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3056–3065, 2019. 4
- [63] 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. 2
- [64] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5419, 2017. 4
- [65] Biao Zhang, Jiapeng Tang, Matthias Nießner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Trans. Graph.*, 42(4):92:1–92:16, 2023. 1
- [66] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 8, 1, 2
- [67] Song-Hai Zhang, Yuan-Chen Guo, and Qing-Wen Gu. Sketch2model: View-aware 3d modeling from single free-hand sketches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6012–6021, 2021. 2
- [68] Ying Zhao, George Karypis, and Usama Fayyad. Hierarchical clustering algorithms for document datasets. *Data mining and knowledge discovery*, 10:141–168, 2005. 5, 4
- [69] Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. Locally attentional sdf diffusion for controllable 3d shape generation. *ACM Transactions on Graphics (ToG)*, 42(4):1–13, 2023. 2, 6, 8
- [70] Yue Zhong, Yonggang Qi, Yulia Gryaditskaya, Honggang Zhang, and Yi-Zhe Song. Towards practical sketch-based 3d shape generation: The role of professional sketches. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(9):3518–3528, 2020. 1, 2, 5, 6, 8
- [71] Yue Zhong, Yulia Gryaditskaya, Honggang Zhang, and Yi-Zhe Song. A study of deep single sketch-based modeling: View/style invariance, sparsity and latent space disentanglement. *Computers & Graphics*, 106:237–247, 2022. 1, 2
- [72] Jie Zhou, Zhongjin Luo, Qian Yu, Xiaoguang Han, and Hongbo Fu. Ga-sketching: Shape modeling from multi-view sketching with geometry-aligned deep implicit functions. In *Computer Graphics Forum*, page e14948. Wiley Online Library, 2023. 2
- [73] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 2
- [74] 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. 2

PASTA: Part-Aware Sketch-to-3D Shape Generation with Text-Aligned Prior

Supplementary Material

Overview

The supplementary material provides a detailed description of our proposed method, PASTA, for sketch-to-3D shape generation. It begins with the experimental setup in Section 1, which outlines the data sets [9, 48, 60, 70] used for training and evaluation, the details of the implementation, and the evaluation metrics used to assess the quality of the 3D shapes generated. In Section 2, we elaborate details on the Text-Visual Transformer Decoder, which plays a crucial role in integrating text and visual information to enhance the semantic understanding of input hand-drawn sketches. In Section 3, we analyze Integrated Structure-Graph Network (ISG-Net), a graph-based refinement module designed to improve the structural coherence of the generated 3D models. Finally, in Section 4, we present more qualitative results, demonstrating the effectiveness of PASTA in generating high-quality 3D shapes that faithfully preserve the intended design elements from sketches.

1. Experimental Setup

1.1. Dataset

To evaluate the effectiveness of our proposed method, we utilize multiple datasets [9, 48, 60, 70] representing diverse sketching styles and their corresponding 3D shapes. Specifically, the datasets used in this work include CLIPasso [60] and non-photo-realistic renderings [9], both of which provide diverse stylistic representations of objects. CLIPasso generates abstract, highly simplified depictions, while non-photo-realistic renderings produce more realistic sketches compared to CLIPasso. For evaluation, we used the AmateurSketch-3D [48] and ProSketch-3D [70] datasets. AmateurSketch-3D consists of freehand sketches drawn by non-experts, often exhibiting variability in proportions and details, whereas ProSketch-3D comprises highly detailed, expert-drawn sketches that adhere closely to object structures. Each dataset contributes a distinct perspective on object depiction, spanning a spectrum from abstract simplifications to highly refined artistic renderings. Fig. 1 presents sample sketches from these datasets [9, 48, 60, 70], highlighting differences in abstraction level, detail, and style of chair sketches.

In addition to these datasets [9, 60], we convert realistic images using ControlNet [66] to further analyze our approach on real-world photo based 3D shape generation. Specifically, we apply ControlNet to sketches from CLIPasso and non-photo-realistic renderings, producing enhanced versions that preserve the original structural essence

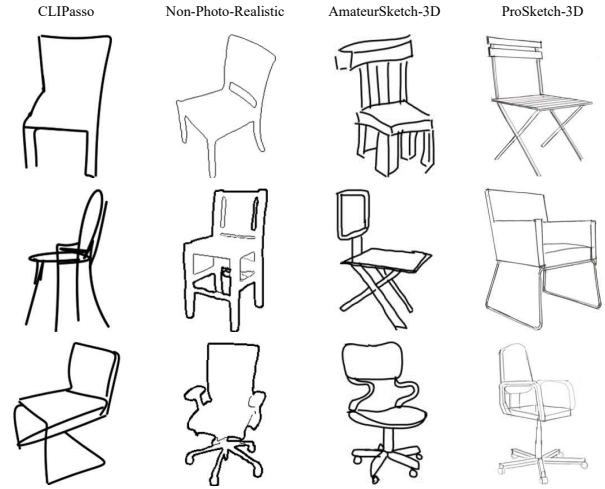


Figure 1. Sample images from the datasets used in our experiments, including CLIPasso [60], non-photo-realistic renderings [9], AmateurSketch-3D [48], and ProSketch-3D [70]. These datasets capture a range of sketching styles, from highly abstract representations to detailed, expert-drawn sketches.

of the sketches while incorporating greater realism. Using ControlNet to guide image generation with fine-grained control, we ensure that the generated results retain the fundamental characteristics of the input sketches while resembling realistic RGB images. This allows us to bridge the gap between abstract grayscale representations and photo-realistic RGB images, offering a richer set of inputs for 3D generation. The generated results are displayed in Fig. 2, demonstrating how ControlNet enhances the visual fidelity of sketches while maintaining their structural shape through examples in the realistic RGB images.

1.2. Implementation Details

Training Configuration. For the training and evaluation of our model, we adopt a carefully designed experimental setup to ensure robust performance and fair comparisons with existing methods. Our model is trained on a single RTX 3090 GPU for approximately 38 hours for the chair category. The training process follows a batch size of 16 and employs an initial learning rate of 10^{-4} , which is dynamically adjusted using the OneCycle learning rate scheduler [57] to facilitate stable convergence. This scheduler gradually increases the learning rate in the early stages of training before decaying it towards the end, preventing premature convergence and improving generalization. The model is optimized using the Adam optimizer, and we train for 650 epochs, ensuring sufficient iterations for convergence while mitigating the risk of overfitting.



Figure 2. Illustration of sketches processed using ControlNet [66]. The transformed sketches retain the structural essence of the originals while incorporating enhanced realism.

Architectural Specifications. In addition to these experimental settings, we provide detailed specifications regarding the architectural components of PASTA. We set the number of queries and Gaussian mixture models to $N = 16$. Within the PartGCN module, the number of graph clustering groups is set to $K = 4$, ensuring effective part-wise feature aggregation. The overall model architecture is designed with two layers in the graph-convolution network (GCN), allowing for a balance between expressiveness and computational efficiency. As described in Equation (9) of the main paper, the weighting factor α in IndivGCN and PartGCN is set to 0.8 based on experimental results. Furthermore, the loss weights for different components, including λ_{align} , λ_{indiv} , and λ_{part} , are set to 1.0, 0.1, and 0.1, respectively, following the formulation in Equation (11) of the main paper, ensuring optimal performance in various sketch styles and object categories.

1.3. Evaluation Metrics

For evaluation, we employ three metrics to assess the quality of the 3D shapes generated. Chamfer distance (CD) quantifies the accuracy of the point-wise reconstruction by measuring the discrepancy between the predicted point clouds set P and the ground-truth point clouds set G :

$$\text{CD}(P, G) = \frac{1}{|P|} \sum_{p \in P} \min_{g \in G} \|p - g\|_2 + \frac{1}{|G|} \sum_{g \in G} \min_{p \in P} \|p - g\|_2. \quad (12)$$

The Earth Mover’s Distance (EMD) quantifies structural differences by determining the minimum cost required to transform one point cloud into another, based on an optimal correspondence $\pi \in \Pi(P, G)$. Here, $\Pi(P, G) \in \mathbb{R}^{n \times m}$ consists of elements in the range between 0 and 1, such that

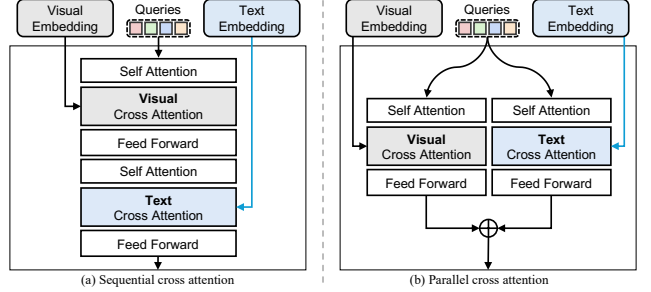


Figure 3. Comparison of two architectures for the Text-Visual Transformer Decoder: (a) sequential cross-attention and (b) parallel cross-attention mechanisms.

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
(a) Sequential	0.090	0.071	143.9	0.055	0.049	112.2
(b) Parallel	0.095	0.074	150.2	0.071	0.061	120.7

Table 1. Quantitative comparison of the two Text-Visual Transformer Decoder architectures. (a) Sequential method consistently outperforms (b) parallel method across multiple evaluation metrics, highlighting the benefits of first enriching visual embeddings before merging them with textual information.

the sum of each row and each column equals one.

$$\text{EMD}(P, G) = \min_{\pi \in \Pi(P, G)} \sum_{i=1}^n \sum_{j=1}^m \pi_{i,j} \|p_i - g_j\|. \quad (13)$$

For the calculation of CD and EMD, 2,048 points are sampled from both the ground truth mesh and the generated mesh. Finally, the Fréchet Inception Distance (FID) [22] is used to assess the realism of the generated shapes by comparing the feature distributions of the rendered 3D models with those of real-world reference data:

$$\text{FID} = \frac{1}{20} \sum_{i=1}^{20} \left(\|\mu_i - \mu'_i\|_2^2 + \text{Tr} \left(\Sigma_i + \Sigma'_i - 2\sqrt{\Sigma_i \Sigma'_i} \right) \right). \quad (14)$$

To compute the FID score, we first sample 20 different views and render both the ground truth shape S the generated shape S' . The features are then extracted from these images using the Inception-V3 network [58], which maps each image to a probability distribution across 1,000 classes. From this distribution, we compute the mean μ_i and the covariance matrix Σ_i for each image i . These statistics are then used to calculate the final FID score. To improve the accuracy of the calculation, we employ the shading-image-based FID metric.

By systematically integrating these experimental configurations, our approach is designed to generate high-fidelity 3D shape generations while maintaining structural consistency across different types of input sketches. The following sections further analyze our results, highlighting details in qualitative and quantitative improvements.



Input Sketch	Part Type	Single Sentence	Verbose
	"Backrest, seat, legs"	"The chair has a curved backrest, a flat seat, four legs, and no armrests ."	"The chair in the image is a modern, minimalist design with a simple and elegant appearance. It features a single, curved backrest and a single, curved seat. The chair has a single leg, which adds to its sleek..."
	"Backrest, seat, legs, armrests"	"The chair has a curved backrest and seat, four legs, and armrests ."	"The chair in the image is a simple, modern design with a backrest and a seat. The backrest is designed to provide support and comfort for the user's back, while the seat is designed to provide a comfortable and ergonomic ..."

Figure 4. Examples of different text description styles used in the VLM [36]. The three types include Part Type, Single Sentence, and Verbose descriptions.

Methods	AmateurSketch-3D			ProSketch-3D		
	CD ↓	EMD ↓	FID ↓	CD ↓	EMD ↓	FID ↓
CLIP [50]	0.101	0.080	152.1	0.066	0.053	117.7
BLIP2 T5 [35]	0.095	0.073	147.4	0.061	0.051	115.6
LLaVa-13B [36]	0.091	0.071	144.3	0.054	0.050	113.9
LLaVa-7B [36]	0.090	0.071	143.9	0.055	0.049	112.2

Table 2. Performance comparison of different VLMs [35, 36, 50] in our framework. LLaVa-7B [36] achieves the best balance between CD, EMD, and FID that making it the most suitable choice for enhancing text-guided 3D shape generation.

2. Details of Text-Visual Transformer Decoder

The Text-Visual Transformer Decoder is a key contribution in our framework, responsible for integrating text priors with visual features to enhance 3D shape generation from sketches. In this section, we present the architectural design of the decoder, analyze the impact of different vision-language models (VLMs) [35, 36, 50], and evaluate the role of input text descriptions in guiding the generation process.

2.1. Architectural Design of Text-Visual Transformer Decoder

Fig. 3 presents two different architectures to integrate text and visual condition within our Text-Visual Transformer Decoder. In method (a), the visual embedding is first processed through a cross-attention mechanism before being fused with text embeddings via a second cross-attention operation. In contrast, method (b) applies cross-attention separately to both visual and text embeddings before merging them later in the process. Tab. 1 compares the performance of these two approaches, showing that method (a) consistently outperforms method (b). The key advantage of method (a) is that the initial cross-attention incorporates visual condition into the query, allowing the subsequent text cross-attention to extract more relevant information. Since the text embeddings extracted from the VLM [36] are aligned with the visual latent space, integrating the sketch embedding into the query first aids in the consolidation of corresponding part-specific information. In contrast, as in

method (b), applying cross-attention between a query devoid of visual cues and the text embeddings may lead to unintended information exchange. This results in improved alignment between sketches and their corresponding 3D shapes, leading us to adopt method (a) in our framework.

2.2. Impact of Vision-Language Models

To assess the influence of different VLMs [35, 36, 50] on our framework, we conduct a comparative analysis using various VLMs, as summarized in Tab. 2. When employing CLIP [54], only marginal performance improvements are observed, likely due to its categorical training and alignment, which are specifically optimized for classification tasks. In contrast, models that excel in image captioning and visual question answering (VQA) yield more substantial enhancements. Notably, BLIP2 [35], lacking explicit fine-tuning on instruction data, frequently produces suboptimal outcomes compared to LLaVA-7B/13B. For instance, it sometimes generates responses that are irrelevant to the user’s instructions (e.g., “a chair is a piece of furniture with a seat and backrest.”) or omits critical components, such as the number of legs and the presence of armrests (e.g., “a chair with a backrest and seat that are shaped in the form of a horseshoe.”). In contrast, the LLaVA-7B model consistently provides responses that comprehensively incorporate the necessary elements. While LLaVA-13B delivers descriptions similar to those of the 7B model, its performance is marginally inferior. Considering model size, computational cost, and overall performance balance, we therefore adopt the 7B model in this work.

Effect of Description Types on Text Embeddings. The quality and structure of text descriptions play a valuable role in guiding the 3D generation process. Fig. 4 provides examples of different description styles, corresponding to part type based descriptions, single sentence descriptions, and verbose descriptions. For instance, in the **Verbose**, the description of such discrepancies arises due to overemphasis on stylistic attributes like “modern”, “minimalist”, and “elegant” which can mislead the 3D generation model. On the other hand, **Part Type**, such as “Backrest, seat, legs, armrests,” provides a fundamental structural understanding. However, they lack details about the shape and number of components, which are crucial for precise shape generation. In contrast, **Single Sentence** strikes a balance between clarity and informativeness. They specify the number and form of object components while avoiding unnecessary complexity. Our analysis, supported by Table 5 in the main paper, reveals that verbose descriptions often introduce hallucinations and excessive modifiers leading to inconsistencies in shape generation. As a result, our framework adopts single sentence descriptions to maximize the accuracy and reliability of text-guided 3D shape generation, by maintaining specificity while avoiding hallucinations.




Figure 5. This figure presents prompts used for airplane and lamp sketches, along with the text descriptions produced by the VLM [36]. The generated descriptions capture the key structural features of each sketch, providing valuable semantic information.

Text-Based Descriptions and Generated Outputs. We extend our analysis beyond the chair category to include airplane and lamp, as shown in Fig. 5. The results demonstrate that VLM [36] effectively captures fine-grained structural details across different object types, showcasing its ability to perform effectively across diverse shape categories. By utilizing text-aligned priors, PASTA successfully distinguishes between distinct design elements, such as the wing and fuselage in airplanes or the lamp head and base in lamps. These findings highlight the robustness of text-visual integration strategy, confirming its effectiveness in enhancing the semantic understanding of sketches across multiple categories.

3. Analysis of ISG-Net

The Integrated Structure-Graph Network (ISG-Net) is designed to refine the structural consistency of 3D shape generation by incorporating graph-based reasoning. The network consists of two key modules: IndivGCN, which focuses on fine-grained feature extraction, and PartGCN, which aggregates and refines part-level information. In this section, we first explore related work on Graph Neural Networks (GNNs) and then analyze the effects of the clustering method used in PartGCN, the impact of the number of clusters K , and the balance between IndivGCN and PartGCN using the parameter α .

3.1. Related Work on Graph Neural Networks

Graph neural networks (GNNs) [55] model complex structured data as nodes and edges, thus effectively learning relationships within the data. Early work focused on message-passing mechanisms to aggregate information among nodes, and later convolutional neural networks (CNNs) were generalized to graph-structured data, leading to the development of graph convolutional networks (GCNs) that facilitate information exchange between neighboring nodes. GCNs can be classified into spatial-based methods [8, 45], which apply trainable filters directly to

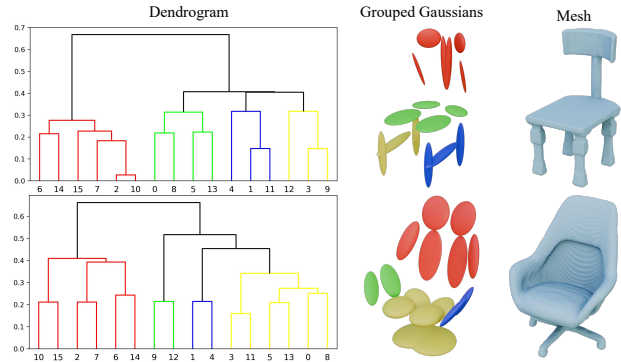


Figure 6. Dendrogram [68] and visualization of Gaussians with the mesh illustrating the clustering process in PartGCN, where $N = 16$ nodes are grouped into $K = 4$ clusters to enhance part-wise feature aggregation.

connected nodes, and spectral-based methods [20, 28], which define locality through spectral analysis. Such graph-based neural networks have been applied not only to social networks [19] but also to various computer vision tasks [4, 26, 27, 30, 59, 62, 64]. In this work, we apply GCNs to Gaussian mixture models (GMMs) for 3D shape representation. We define the graph structure by producing an adjacency matrix based on the distances between Gaussians in 3D space, thereby enabling information exchange between adjacent nodes for more structurally accurate shape generation and reconstruction. Furthermore, by grouping similar nodes and performing additional graph operations, we achieve robust structural learning.

3.2. Clustering Method for PartGCN

To perform PartGCN operations, N queries and Gaussians must be assigned to K subsets. To achieve this, we employ a hierarchical clustering technique [68] to partition the parts. This clustering is applied to both the pseudo-ground truth and the predicted queries. For pseudo-ground truth, clustering is performed based on $\mathbf{A}_I \in \mathbb{R}^{N \times N}$ to form K groups, and the coordinates of the Gaussians within each group are aggregated using average pooling to obtain a representative coordinate. The distances between these coordinates are then computed to produce the final pseudo-ground truth adjacency matrix $\mathbf{A}_P \in \mathbb{R}^{K \times K}$. Simultaneously, predicted queries are pooled for each group to obtain $\mathbf{Q}_P \in \mathbb{R}^{K \times d}$, which is processed through f_{part} and subsequently subjected to an dot product operation to produce the predicted adjacency matrix $\hat{\mathbf{A}}_P \in \mathbb{R}^{K \times K}$. As illustrated in Fig. 6, the clustering process is depicted through a dendrogram [68], where N nodes are grouped into K clusters. In particular, it flexibly groups dynamically varying Gaussians (e.g., those located near the legs versus those near the armrests). This method enables PartGCN to accurately identify and select relevant structural parts, ensuring that the network focuses efficiently on the details of the part level while maintaining overall structural integrity.

Methods	AmateurSketch-3D		
	CD ↓	EMD ↓	FID ↓
$K = 2$	0.098	0.081	160.2
$K = 4$	0.090	0.071	143.9
$K = 6$	0.093	0.076	149.0
$K = 8$	0.094	0.078	156.4

Table 3. Performance comparison for different cluster numbers K . Too few clusters weaken structural representation, while too many reduce synergy with IndivGCN.

3.3. Effect of Cluster Number K

The experimental results on varying the number of clusters K reveal important insights into the impact of cluster selection on the performance of our model. When K is too small, there is insufficient granularity in the clustering process, causing PartGCN to lose critical structural relationships. As a result, the overall structural integrity of the output suffers and the model does not adequately represent the coherence of the object. On the other hand, when K is excessively large, the model faces a different challenge. Although more clusters may provide a higher resolution of detail, it also introduces redundancy and excessive complexity. In this case, PartGCN becomes overloaded with unnecessary information, leading to a loss of synergy with IndivGCN. This imbalance occurs because excessive detail hampers the effective integration of local features, causing IndivGCN and PartGCN to work against each other rather than complement each other. The model’s performance, therefore, suffers from a lack of coherence between the fine-grained details and the global structure. As shown in Tab. 3, $K=4$ ensures that both components work harmoniously, leading to the generation of high-quality 3D shapes without sacrificing performance or structural integrity. This process allows PartGCN to contribute effectively to the overall structure while maintaining the balance with IndivGCN.

3.4. Influence of α in IndivGCN and PartGCN

In our experiments, we use the parameter α as a weighting coefficient for the outputs of IndivGCN and PartGCN, as described in Equation (9) of the main paper. Specifically, α scales the output of IndivGCN, which captures fine-grained local features, while the complement $(1 - \alpha)$ scales the output of PartGCN, which encodes global structural elements. The experimental results, shown in Tab. 4, demonstrate that the best performance is achieved when α is set to 0.8, striking the optimal synergy between these two components. In this configuration, IndivGCN is weighted more heavily, which allows it to preserve fine-grained local details that are crucial for the accurate representation of intricate features. At the same time, PartGCN is still capable of effectively contributing to the integrity of the global structure by encoding broader spatial relationships and consistency at the

Methods	AmateurSketch-3D		
	CD ↓	EMD ↓	FID ↓
$\alpha = 0.0$	0.095	0.084	157.0
$\alpha = 0.2$	0.095	0.081	155.4
$\alpha = 0.4$	0.092	0.077	151.8
$\alpha = 0.6$	0.094	0.077	152.1
$\alpha = 0.8$	0.090	0.071	143.9
$\alpha = 1.0$	0.092	0.071	145.3

Table 4. Effect of α on IndivGCN and PartGCN balance. The best performance is achieved at $\alpha = 0.8$, optimizing detail and structure integration.

part level, as it is scaled by $(1 - \alpha)$. This balance is key to ISG-Net success, while IndivGCN focuses on refining local geometry and intricate details, PartGCN provides essential structural support, preventing shape distortions, and preserving overall coherence. Consequently, this leads to the best overall performance of ISG-Net, enabling it to generate 3D shapes that are rich in fine-grained detail and structurally sound, preserving oversimplified sketches.

4. Qualitative Results

Additional qualitative results across different object categories, including chair, airplane, and lamp, are provided to further support our experimental findings. For chair, further results can be found in Fig. 7 and Fig. 8, where our model successfully preserves key structural elements such as detailed leg structure, backrest curvature, and armrest presence, even in abstract sketches. In the case of airplane, additional results are shown in Fig. 9, demonstrating that PASTA accurately generates aircraft components such as wings, fuselage, engines, and tail sections, with a high level of structural coherence. For lamp, additional 3D shape generations are also presented in Fig. 9, showcasing how our model captures intricate details such as delicate frame structures, lampshades, and support bases, producing models that closely align with input sketches. These qualitative results confirm that our method outperforms existing approaches across multiple object categories, offering superior generalization and the ability to handle complex and varied shapes with high fidelity and structural consistency.



Figure 7. Qualitative results for chair reconstructions, demonstrating the preservation of key structural elements such as leg orientation, backrest curvature, and armrest presence.



Figure 8. Qualitative results for chair reconstructions, demonstrating the preservation of key structural elements such as leg orientation, backrest curvature, and armrest presence.




Figure 9. Additional qualitative results for airplanes and lamps, showcasing accurate reconstruction of aircraft components and intricate lamp structures.