

Doodle Your 3D: From *Abstract* Freehand Sketches to *Precise* 3D Shapes

Hmrishav Bandyopadhyay¹ Subhadeep Koley^{1,2} Ayan Das^{1,2} Aneeshan Sain¹

Pinaki Nath Chowdhury¹ Tao Xiang^{1,2} Ayan Kumar Bhunia¹ Yi-Zhe Song^{1,2}

¹SketchX, CVSSP, University of Surrey, United Kingdom.

²iFlyTek-Surrey Joint Research Centre on Artificial Intelligence.

{h.bandyopadhyay, s.koley, a.das, a.sain, p.chowdhury, a.bhunias, y.song}@surrey.ac.uk




Figure 1. Unlike prior methods (LAS-D [88]), our sketch-based shape generation algorithm generalises to abstract doodles without training on human sketches. Thanks to our part-disentangled sketch and shape representations, we exhibit (a) fine-grained correspondence between sketches and generated shapes, allowing us to (b) perform highly localised shape edits through edits in sketches.

Abstract

In this paper, we democratise 3D content creation, enabling precise generation of 3D shapes from abstract sketches while overcoming limitations tied to drawing skills. We introduce a novel part-level modelling and alignment framework that facilitates abstraction modelling and cross-modal correspondence. Leveraging the same part-level decoder, our approach seamlessly extends to sketch modelling by establishing correspondence between CLIPasso edgmaps and projected 3D part regions, eliminating the need for a dataset pairing human sketches and 3D shapes. Additionally, our method introduces a seamless in-position editing process as a byproduct of cross-modal part-aligned modelling. Operating in a low-dimensional implicit space, our approach significantly reduces computational demands and processing time.

1. Introduction

We envisage a world where 3D content creation is democratised, granting everyone the liberty to freely create and modify 3D shapes. This shared vision has spurred collective efforts, initially centred on using text as a condition for 3D shape creation [43, 46]. However, the pivotal challenge arises when delving into fine-grained [6, 55, 62] creation – text can be ambiguous, lacking the nuanced cues essential for precise idea conveyance. This is where sketches step in as a complementary input modality [14, 63, 64], inherently and accurately capturing users’ intent – the promise being “what you sketch is what you get”.

All that stands in the way of fulfilling that promise is you and your sketch, or the lack of it – “I can’t sketch”, I hear you say! Prior works [7, 20, 24, 88] in this space do generate shapes given rough sketches, but the issue lies in the

resulting output – a poorly drawn sketch yields a deformed 3D shape where your lack of drawing skills is *accurately* reflected – “what you sketch is *literally* what you get” (see Fig. 1). In this paper, we aim to democratise sketch-to-3D creation by definitively addressing this very point – for the first time, enabling your *abstract* (“half-decent”) sketches to generate *precise* 3D shapes, all without any specific paired human sketch and 3D shape data!

Navigating these challenges is no small feat. First, sketches transcend mere edge maps [77]; they embody subjective abstract forms that have been proven to be difficult to model [39, 68]. Second, the hurdle lies in injecting a fine-grained understanding to capture the inherent details within sketches. Third, the challenge extends to establishing fine-grained cross-modal correspondences [13] between sketches and 3D shapes. And last but not least, we are keen on avoiding the need to collect a dataset of human sketch and 3D pairs [59, 89] to achieve all said challenges.

Our solution, we think(!), is quite elegant. It revolves around a well-known hypothesis: part-level modelling aids [16, 74] in generalising to complex (fine-grained) shapes. In this paper, we not only extend this hypothesis but also demonstrate its applicability to modelling abstraction. The intuition is straightforward – parts enable more flexibility (as they can move around) in terms of global construction, thus accommodating deformations [59] found in human sketches. It is roughly akin to using a single matchstick where you are limited to just that, but by breaking it into pieces, you can create something more meaningful.

Our part-level modelling initiates from a pre-trained auto-decoder [29], where we invert the decoder to obtain part-disentangled latent representations. However, this step alone is not sufficient. For cross-modal matching and downstream editing, we also require these parts to be aligned, roughly corresponding to the same regions in an object category. To achieve this, we align these parts, ensuring they share the same indices for similar components in the decomposition. Subsequently, we train a generative model, specifically a diffusion pipeline [31, 61], on this aligned latent space. Through sampling from the trained network, we can unconditionally generate aligned and part-disentangled latents. These latents can then be decoded [29] into tangible meshes. Note that in contrast to typical voxel-space diffusion models for 3D shape generation [43, 88], our approach involves performing diffusion in the latent space [61]. This approach, operating in a low-dimensional implicit space, enables semantic edits and interpolations with significantly reduced computational demands ($0.02\times$ params) and processing time ($3\times$ faster) when compared with prior art [88].

Curious about the role of our solution in sketch modelling? Here is where the elegance of our approach truly stands out! Surprisingly, the same decoder [29] can be conveniently leveraged to achieve part-level modelling

for sketches, all without the need to collect any sketch data. This correspondence is established by transferring the shape decompositions onto their corresponding rendered 2D edgemaps. To enhance abstraction modelling, we took an additional step by passing the aforementioned edgemaps through a CLIPasso [68] network to produce of human-like sketches before establishing the said correspondence. Finally, on the editing front, having established a part-level cross-modal correspondence, we can effortlessly determine which parts of the 3D shape the in-position sketch is editing. Subsequently, we can condition the generation of each part with individual part representations.

In summary, our contributions are: (i) Empowering abstract sketches to generate precise 3D shapes and execute in-position editing, surpassing limitations associated with drawing skills prevalent in prior methods. (ii) Introducing an intuitive part-level modelling and alignment framework that facilitates abstraction modelling and cross-modal correspondence. (iii) Conveniently leveraging the same part-level decoder for sketch modelling, achieved by establishing correspondence between CLIPasso edgemaps and projected 3D part regions. (iv) Introducing a seamless in-position editing process as a byproduct of our cross-modal part-aligned modelling.

2. Related Works

3D Representations: 3D shapes have been popularly represented in their *explicit* geometric form via (i) point clouds, (ii) voxel grids and (iii) polygonal meshes. Generally captured by 3D sensors [33], point clouds represent the 3D surface of real-world objects [5] using coordinate sets [60] on a 3D space. While representative in nature [75], point clouds are typically sparse and are converted to denser representations [4, 26] like meshes [26, 83] or used as voxel [84] grids. Voxelised representations of shapes are dense 3D grids with filled/unfilled information per grid-cell, forming the most straightforward extension of a pixel in 3D. Voxels rich in 3D spatial data, are thus commonly processed with 3D CNNs [76] for both representative [71] and generative [88] tasks. Lighter than voxels, meshes represent approximate 3D *surfaces* with *polygon-faces*, requiring only vertices and connecting faces to depict shapes. This allows easy deformation [44, 45] and manipulation by offsetting vertex positions [46] for 3D generative tasks [3, 26, 44, 46]. Explicit 3D geometrical representations are, however, compute-intensive [66, 72] and hence limited to low grid-resolutions.

Alleviating this limitation, continuous *implicit* functions f_θ [29, 56] represent shapes by mapping coordinates to implicit values. Implicit values like (i) occupancy [51] and (ii) signed distance [56] represent the coordinate’s *presence* within the shape and distance from the surface respectively. Sampling for these implicit values along a template grid [56] of mesh vertices or voxel coordinates allows us

to build tangible shapes in the form of *meshes* and *voxels*. Thus, having function parameters θ enables us to obtain explicit shape forms of arbitrary resolution (0.1 or 0.01 etc.) by sampling any coordinate with the known continuous implicit function. Beyond implicit and explicit forms, part-level representations decompose 3D shapes into parts for part-specific editing [29, 82] and generation [29, 54]. Shapes can be hierarchically decomposed without explicit part-level supervision into hyperplanes [10], quadratic surfaces [81], superquadrics [57, 58], and GMM (Gaussian Mixture Model) based [28, 29] shape-partitions.

3D Generation: 3D shapes can be generated by sampling on a learned latent space of implicit functions like signed distance [56] and occupancy [51], or adversarially by converting to explicit voxels [38, 87], meshes [49], or point clouds [38]. The recent success of denoising diffusion models [31, 61, 65] in generating high fidelity images has led to their adoption for diverse generation of explicit 3D structures [46, 50, 88]. Denoising randomly sampled gaussians, however, requires multiple iterations resulting in slow and compute heavy inference [48], particularly in the context of explicit generation [88]. As a result, generation of high-resolution shapes becomes *impractical*. While generation of implicit functions [12, 53] or function parameters [18] with latent diffusion [61] is scalable, it limits generation diversity [34] and editability. In this work, we leverage the efficiency of implicit representations with part-level decomposition of GMM based models [29] to learn a latent diffusion model in a part-aware implicit space for high resolution shape generation, guiding and editing.

Sketches for 3D: Shape reconstruction from single-view [15, 19, 78] and multi-view [15, 69] images has been extensively studied with deep encoder-decoder architectures [15, 19] as well as through generative modelling with adversarial [17, 86] and diffusion pipelines [23, 46]. While sketches are sparser [24] representations of shapes compared to images, their expressive nature makes them an ideal input modality for interactive tasks like retrieval [59, 79], modelling [7, 41, 42, 73, 88], and editing [7, 88] of 3D shapes. However, hand drawn sketches are inaccurate in representing objects from pre-defined viewpoints [22] or with consistent style. This makes the task of sketch-to-3D more of a generative [88] one than reconstructive [7], as one sketch can correspond to multiple plausible shapes. Nevertheless, shapes can be constructed both from reconstruction and generation perspectives with deterministic encoder-decoders [7, 40] and denoising diffusion models [88] respectively. In this work, we encode sketches in a fine-grained part-aware representation for conditional shape generation through latent diffusion [61]. We explicitly imbibe part-specific knowledge in sketch representations by aligning part-disentangled sketch encodings with underlying shape parts, allowing us to capture fine-grained shape information from the sketch.

3. Proposed Methodology

Overview: The expressive [30] nature of sketch as an input modality makes it an excellent choice for spatial control over 3D generative tasks [88]. We aim to perform sketch-conditioned generation of 3D shapes via (i) diffusion-based generative modelling on implicit representations of a pre-trained neural implicit decoder [29] and (ii) mapping sketches to this implicit space for fine-grained control. Specifically, we represent part decompositions of 3D shapes in a part-disentangled latent space of a pre-trained decoder [29], and map sketches to this space via similar part-disentangled sketch representations. Evidently, this mapping allows us to (i) condition shapes based on *highly abstract doodles* [25] that are *unseen* during training, (ii) perform fine-grained and localised shape edits by matching edited regions in sketches, and (iii) generate a morph [2] of multiple shapes by naive interpolation of sketches.

3.1. Baseline Sketch-to-3D Generation:

Implicit Neural functions $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ efficiently characterised 3D shapes with the relationship between 3D coordinates (x, y, z) and their implicit values $O = f(x, y, z)$ as occupancy [51] or signed distance [56]. Neural networks, as universal function approximators [52], learn these functions for various shapes as a unified function $O_I = f_\theta(I, x, y, z)$, where implicit code $I \in \mathbb{R}^d$ describes a shape uniquely [56]. A naive way to model an explicit shape on an input sketch \mathcal{S} would be to learn a mapping from sketch to implicit I using a visual encoder E (like ResNet-18 [27]) and use a pre-trained function f_θ to sample implicit values as:

$$O_I = f_\theta(E(\mathcal{S}), X) \quad (1)$$

where, $X = (x, y, z)$ refers to sampling coordinates in 3D. The explicit shape can then be reconstructed [29, 56] from a uniform grid of coordinates and their implicit values. While such a formulation explores 3D reconstructions based on efficient neural implicit modelling [56], it has a few inherent limitations: (i) generated explicit shapes lack editability as *discovering* local edits in implicit codes that translate to explicit edits is non-trivial [29]. This is primarily due to a lack of feature disentanglement in the naive implicit codes $I \in \mathbb{R}^d$ that learn a global immutable feature per instance. Secondly, mapping of *sketch* \rightarrow *implicit*s is a one-to-many task [88] as 2D line drawings lack the *granularity* to perfectly depict a single shape, corresponding to multiple plausible shapes at once. Modelling such a problem as a deterministic one-to-one mapping *i.e.* $I = E(\mathcal{S})$ is thus ill-posed. Finally, a mapper trained on synthetic edgemap-like sketches \mathcal{S} (from scarcity of line drawings) does not generalise [88] to human-drawn doodles that suffer from geometric and perspective inaccuracies.

To address these challenges, we (i) perform *part-aware* neural implicit modelling (Sec. 3.2) by decomposing ex-

PLICIT shapes into m parts as $Z \in \mathbb{R}^{m \times d}$, and mapping them to global implicit codes I . Part-level disentanglement in the latent space allows for part-specific editing by mixing or swapping parts from multiple latents. (ii) We learn a diffusion [31] pipeline as a generative model [61] on this part-aware latent space to stochastically generate editable shape implicit codes, where we exert *fine-grained* control over shape generation with sketches. (iii) To perform local shape edits by editing on sketches (Fig. 1), we learn the same disentanglement of parts – this time to represent sketches, by matching sketch regions with pre-decomposed parts in Z . We find, this helps in generalising to freehand doodles, despite training on synthetic sketches.

3.2. Part-aware Neural Implicit Shape modelling

A pre-trained decoder \mathcal{D} [29] decomposes an explicit shape (mesh) M as $Z = \{\omega_i\}_{i=1}^m$ where $\omega_i \in \mathbb{R}^d$, and maps it to implicit codes I , to get output shapes. These part-latents (Z) are decoded to form (i) structural representations $Z^p = \mathcal{D}_p(Z) \in \mathbb{R}^{m \times d}$ and (ii) volumetric descriptors $Z^g = \mathcal{D}_g(Z) \in \mathbb{R}^{m \times 16}$ where $\mathcal{D}_p, \mathcal{D}_g$ are fully connected layers in \mathcal{D} . While part-structures (Z^p) are comprised by d -dimensional latents, part-volumes (Z^g) are represented as parametric 3D Gaussians, $\mathcal{N}(\mu_i, \Sigma_i)$ (16-parameters e.g. $\mu \in \mathbb{R}^3, \Sigma \in \mathbb{R}^{3 \times 3}$, etc. [28, 29]) under a global Gaussian mixture model (GMM) with mixing weights π_i as $p(X) = \sum_{i=1}^m \pi_i \mathcal{N}(\mu_i, \Sigma_i)$. A Gaussian representing a part in Z , captures the probability of a sampled point (X in Eq. (1)) to belong to that part. In unsupervised implicit decomposition, part-specific 3D Gaussians represent part-orientation and positions in the 3D space. They (i) enable uniform part decomposition and (ii) allow explicit part-disentanglement of overlapping and closely-placed parts. Finally, part-structures and Gaussian parameters are concatenated followed by self-attention [67] to form implicit codes I , thus extending Eq. (1) as:

$$O_I = f_\theta(\mathcal{D}(Z), X) \quad (2)$$

Accordingly, the mesh M can be constructed with Marching Cubes [47] over a grid of implicit values O_I .

3.3. Part-level Alignment in the INR Latent Space

Constructing a part-aligned latent representation allows better control for conditional generation and editing [32]. Although Z is decomposed into m parts, it lacks part-level alignment across different shapes since the pre-trained decoder \mathcal{D} does not enforce positional encoding [29] while using self-attention blocks. Hence, we align Z such that its part-indices $i \in [1, m]$ correspond (Fig. 2) to similar parts (e.g. chair’s leg) across all shapes. For this, we first pre-compute Z for all shapes by (i) using the occupancy values O_I and 3D coordinates X to invert our pre-trained decoder \mathcal{D} and implicit function f_θ as $\mathcal{D}_\phi(\cdot) = \text{Inv}(f_\theta(\mathcal{D}(\cdot), X))$,

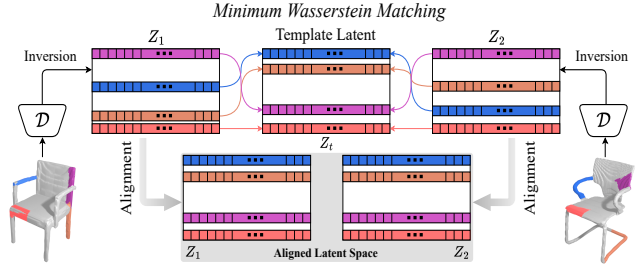


Figure 2. Decomposing shapes into latents, we shuffle m part indices of each latent $Z \in \mathbb{R}^{m \times d}$ for minimal Wasserstein distance [35] with corresponding parts in template latent $Z_t \in \mathbb{R}^{m \times d}$.

(ii) optimising for $Z = \mathcal{D}_\phi(O_I)$ to match its corresponding implicit code $I = \mathcal{D}(Z)$. Next, we randomly select the part-latents (Z) from one shape as “*template part-latents*” $Z_t \in \mathbb{R}^{m \times d}$ and align all other Z by minimising the distance between their part-volumes $Z^g = \mathcal{D}_g(Z)$ [46]. Specifically, given the parameters $\{\mu_i, \Sigma_i\}_{i=1}^m \in Z^g$, we compute the Wasserstein distance (W_{ij}) [35] between two 3D Gaussians – $\mathcal{N}(\mu_i, \Sigma_i)$ for i^{th} part of Z and $\mathcal{N}(\mu_j, \Sigma_j)$ for j^{th} part of Z_t as:

$$W_{i,j}^2 = \|\mu_i - \mu_j\|_2^2 + \text{Tr}(\Sigma_i + \Sigma_j - 2(\Sigma_i^{\frac{1}{2}} \Sigma_j \Sigma_i^{\frac{1}{2}})^{\frac{1}{2}}) \quad (3)$$

For alignment, we replace the n^{th} part of Z with its \hat{n}^{th} part that has the minimum Wasserstein distance from n^{th} part (same part-indices) of template part-latent Z_t , as:

$$\hat{n} = \arg \min_{1 \leq i \leq m} W_{i,n} \quad (4)$$

3.4. Unsupervised Part Discovery for Sketches

To enable fine-grained control of shape generation using sketches $\mathcal{S} \in \mathbb{R}^{H \times W}$, we aim to learn an m -level part-aligned sketch representation $f_s(\mathcal{S}) \in \mathbb{R}^{m \times d_s}$. First, we handle the scarcity of paired sketch and 3D shapes [59, 89], by rendering 2D projections $\mathcal{I} = \mathcal{R}_\nu(M)$ of shape (M) [70] from multiple views $\nu \in \mathcal{V}$ with rendering function $\mathcal{R}_\nu(\cdot)$. The 2D projections (i.e., images) \mathcal{I} are used to generate synthetic line-drawings $\mathcal{S} = \text{sketch}(\mathcal{I})$ as augmented edgmaps [8] following [68]. Next, we learn the m -level part-aligned sketch representation $f_s(\mathcal{S})$ by discovering parts corresponding to part-latents $Z \in \mathbb{R}^{m \times d}$. Specifically, we use the part-volume parameters $\{\mu_i, \Sigma_i, \pi_i\} \in Z^g$ to represent the probability of a 3D coordinate X belonging to the i^{th} part using a GMM with 3D Gaussian $\mathcal{N}(\mu_i, \Sigma_i)$ and mixing coefficient π_i as:

$$p_i(X) = \pi_i \cdot \mathcal{N}(X | \mu_i, \Sigma_i) \quad (5)$$

We use $p_i(X)$ to identify parts from shape coordinates X and construct corresponding m 3D segment maps $\{\mathcal{M}_i^{3D}\}_{i=1}^m$. Next, we overlay these segmentation maps over synthetic sketches (Fig. 3) by rendering them from same viewpoint (as sketches) ν as $\mathcal{M}_{1:m}^{2D} = \{\mathcal{R}_\nu(\mathcal{M}_i^{3D})\}_{i=1}^m$, to

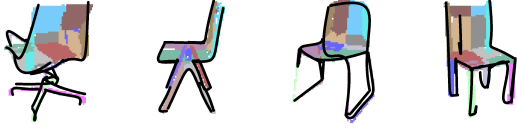


Figure 3. Part-level segmentation maps are created by segmenting 3D shapes into parts with part Gaussians and rendering individual 3D shape parts on synthetic sketches. This segments sketch regions based on shape parts of their corresponding 3D shapes.

construct segmentation maps for sketch regions that correspond to m shape parts. The annotations $\mathcal{M}_{1:m}^{2D}$ allow us to obtain part-disentangled sketch representations $f_s(\mathcal{S})$, aligned with part-latents $Z \in \mathbb{R}^{m \times d}$ of corresponding shapes to enhance control over generation and editing.

3.5. Latent Diffusion for 3D Generation

We use a pre-trained decoder \mathcal{D} to represent shapes in m parts as part-latents $Z \in \mathbb{R}^{m \times d}$. Z is aligned across all shapes in a category, such that similar parts have the same part indices $i \in [1, m]$. Rendering the (a) shape and (b) its individual parts on a 2D plane separately, we (a) construct sketches from shape boundaries and (b) overlay sketch regions with rendered parts (as part-level annotations). Next, we train a diffusion model on part-latents Z for generative modelling, which we condition with part-disentangled sketch representations from part-level annotations.

The diffusion pipeline consists of (i) a predefined forward process $q(z_{0:T})$, where noise is added to z_0 progressively from $0 \rightarrow T$ till $q(z_T) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and a (ii) reverse process $p_\theta(z_{T:0})$ where a network estimates denoising conditionals $p_\theta(z_{t-1}|z_t)$ for each step $t(< T)$. Specifically, given a noisy sample $z_t = \sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}\epsilon$ at time step t with constants $\{\alpha_t\}_{t=0}^1$ and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the network estimates the noise ϵ as $\epsilon_\theta(z_t, t)$. The *simplified* loss is:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{z_t \sim q(z_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), t} [\|\epsilon - \epsilon_\theta(z_t, t)\|_2^2] \quad (6)$$

During inference, z_T is sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and is iteratively denoised to z_0 . We learn the underlying distribution of part-latents $Z \in \mathbb{R}^{m \times d}$ as z_0 with latent diffusion, where our denoiser ϵ_θ consists of fully-connected layers f_d and multi-head attention module \mathcal{C} (Fig. 4(a)).

Part-level Sketch Representations: For fine-grained control over generated shapes, we use part-disentangled representations $f_s(\mathcal{S}) \in \mathbb{R}^{m \times d_s}$ from sketches \mathcal{S} matching with similar disentanglement in part-latents $Z \in \mathbb{R}^{m \times d}$. For this, we utilise part annotations of sketch regions corresponding to m shape parts and train for segmentation to predict parts as sketch segment-maps. This segmentation is facilitated by encoding sketches as $\eta = f_s(\mathcal{S}) \in \mathbb{R}^{m \times d_s}$ with a ResNet-18 [27] encoder f_s , and decoding to segment-map predictions of individual parts with auxiliary decoder f'_s consisting of upsampling blocks (Fig. 4(b)). To optimize for disen-

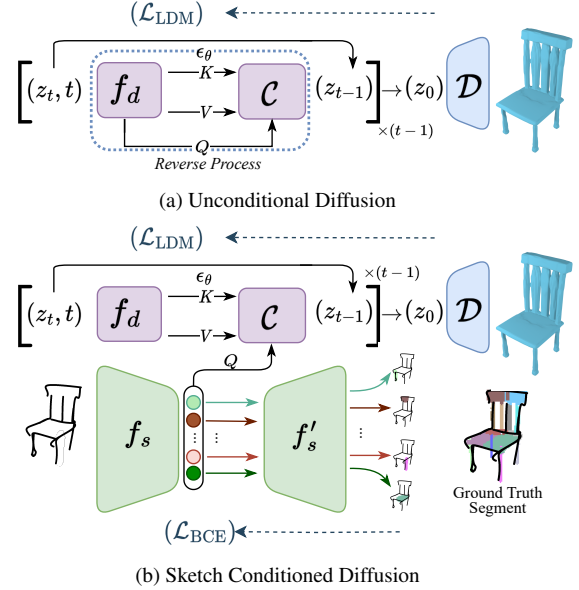


Figure 4. Model overview: (a) The diffusion pipeline denoises latent vector $z_t \in \mathbb{R}^{m \times d}$ to z_{t-1} with fully connected layers f_d and a multi-head attention module \mathcal{C} at time step t . After $t = 0$, the fully denoised vector z_0 corresponds to a generated part-latent Z . (b) We encode sketches as part-disentangled representations with encoder f_s by segmenting them into segment maps of individual parts with shared decoder f'_s . These sketch representations are fed to the attention module \mathcal{C} as a Query with intermediate diffusion outputs (from f_d) as Key-Value pairs.

tanglement of part features in sketch representations $f_s(\mathcal{S})$, representation of each part of size $\in \mathbb{R}^{d_s}$ is decoded individually with common decoder f'_s , making disentanglement indispensable for accurate part-segmentation.

Importantly through part-aware representations, we align encoding $f_s(\mathcal{S})$ to represent the same explicit parts from sketch \mathcal{S} irrespective of its viewpoint ν . As such, we can naively aggregate information on different parts from multi-view $\{\nu_1, \dots, \nu_n\}$ sketches $\{\mathcal{S}_{\nu_j}\}_{j=1}^n$, allowing us to reconstruct the shape more accurately (Fig. 8) with an aggregated representation $f_s(\mathcal{V}) = \frac{1}{n} \cdot \sum_{j=1}^n f_s(\mathcal{S}_{\nu_j})$

Sketch Conditioning: Sketches are encoded as $\eta = f_s(\mathcal{S}) \in \mathbb{R}^{m \times d}$ where sketch-features are disentangled into m parts-features as $\{\eta_i\}_{i=1}^m$. Part-features have one-to-one correspondence with pre-defined explicit shape parts, and hence part-indexing holds positional significance to uphold this correspondence. We utilise this property for sketch conditioned generation by concatenating sinusoidal positional embeddings of part-indices to both the condition vector $f_s(\mathcal{S})$ from sketch \mathcal{S} and the noised input z_t . We then use our multi-head attention block \mathcal{C} as a cross-attention module [67] by using the sketch conditioning $f_s(\mathcal{S})$ and intermediate diffusion outputs $f_s(\mathcal{S})$ as Query and Key-Value pairs respectively to compute attention as:

$$\text{Attention}(Q, K, V) = \text{softmax}(QK^T / \sqrt{d_k}) \cdot V \quad (7)$$



Figure 5. From left to right: 3D shape; an edgemap of a 2D render of the shape; corresponding sketches from the ProSketch3D [89] and AmateurSketch3D [59] datasets; an abstract CLIPasso [68] sketch of the shape.

with $Q = W^Q \cdot f_s(\mathcal{S})$, $K = W^K \cdot f_d(z_t, t)$, $V = W^V \cdot f_d(z_t, t)$ where W^Q , W^K , W^V represent the projection matrices for Queries, Keys, and Values respectively. The output from \mathcal{C} (used for cross-attention) is projected back to the input dimension with fully connected layers and skip connections. The loss function is extended from Eq. (6) as:

$$\mathcal{L}_C = \mathbb{E}_{z \sim q(z_0), \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} [\|\epsilon - \epsilon_\theta(z_t, t, f_s(\mathcal{S}))\|_2^2] \quad (8)$$

3.6. Inference Pipeline

Inference with the trained diffusion pipeline is performed as a reverse process, by (i) randomly sampling $z_T \in \mathbb{R}^{m \times d}$ from the normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and (ii) estimating the denoising conditionals at each step $t : T \rightarrow 0$ as $\epsilon_\theta(z_t, t, f_s(\mathcal{S})) \in \mathbb{R}^{m \times d}$ with conditioning signal $f_s(\mathcal{S}) \in \mathbb{R}^{m \times d_s}$ from sketch \mathcal{S} . The intermediate denoised latent z_t is now denoised to z_0 with estimated noise iteratively as $z_t \rightarrow z_{t-1} \rightarrow z_{t-2} \dots \rightarrow z_0$ giving us part-latents $Z = z_0 \in \mathbb{R}^{m \times d}$ that can be decoded to obtain implicit codes $I = \mathcal{D}(Z)$ with pre-trained \mathcal{D} . Implicit values (occupancies) are sampled with a coordinate grid as Eq. (2) to construct output mesh M with marching-cubes [47].

4. Experiments

Our latent diffusion pipeline is trained on edgemaps [8] and synthetic sketches [68] from 2D projections of 3D shapes. For generalisation to human sketches, we evaluate primarily on hand-drawn doodles for generative control and diversity. **Datasets:** We use a subset of the ShapeNet [9] dataset, consisting of 6755 ‘chair’ shapes [7]. All shapes are normalised to unit cube and rendered from 6 distinct views $\nu \in \mathcal{V}$ with azimuthal angles distributed across $[-\pi, \pi)$ and at a constant elevation of $\frac{\pi}{10}$ from a distance of 2.07 units. These renders are used to create non photo-realistic renderings [8] and abstract sketches using CLIPasso [68] similar to [7] (Fig. 5). Thus, we obtain a total of $\sim 80\text{K}$ sketch samples synthetically generated from our shape dataset. We invert pre-trained \mathcal{D}_s for these 6755 shapes, obtaining their latent representations Z , which we align and use as ground truth latent codes following the evaluation split in [56]. Without training on *any* human sketches, we evaluate our model on sketches from the *AmateurSketch-3D* dataset [59] consisting of 3000 hand-drawn sketches of 1000 chairs drawn from viewpoints at angles of 0° , 30° , and 75° . We further evaluate on the *ProSketch* dataset [89] containing 1500 sketches

of 500 chairs drawn at 0° , 45° , and 90° . Our evaluation set thus consists of (i) unseen hand-drawn sketches drawn from (ii) unseen azimuthal angles, allowing us to demonstrate the robustness of our algorithm to sketch style and viewpoint perturbation, respectively. Finally, we demonstrate further robustness to abstraction by generating shapes from the chair category in the *Quick-Draw!* dataset [25] of highly abstract diverse sketches from 15 million people.

Implementation details: Noisy sample $z_t \in \mathbb{R}^{m \times 512}$ is concatenated with sinusoidal positional embeddings ($\in \mathbb{R}^{m \times 224}$) of timestep t , and projected to $\mathbb{R}^{16 \times 128}$. The intermediate outputs from these layers are concatenated with part-index positional embeddings $\in \mathbb{R}^{16 \times 16}$ and timestep positional embeddings $\mathbb{R}^{16 \times 16}$ again to form Key-Value pairs for cross-attention. Sketch encodings $f_s(\mathcal{S}) \in \mathbb{R}^{16 \times 64}$ are likewise projected and concatenated with part-index embeddings to $\mathbb{R}^{16 \times 256}$ forming Query for multi-head attention block \mathcal{C} that generate part-latents $z_0 \in \mathbb{R}^{m \times 512}$.

Training: We train the diffusion pipeline with a DPM solver [48] for 1000 timesteps for 10M iterations. We use a learning rate of $1e-4$ with the AdamW optimizer and a batch size of 128 on a NVIDIA RTX 3060 Ti. The sketch conditioning module is trained on the same GPU with a learning rate of $2e-4$ with an Adam optimizer and a batch size of 64. Training the unconditional diffusion model, the sketch conditioning, and the conditional diffusion pipeline takes 39, 41, and 44 hours on the GPU respectively.

Competitors: (i) Unconditional shape generation algorithms: **SDF-StyleGAN** [87] appropriates StyleGAN2 [36] from 2D to 3D with 3D convolutions and learns a shape in a voxelised SDF space through locally and globally critical discriminators. **Diffusion-SDF** [12] learns an MLP-based diffusion model to predict the implicit function for generating a shape. Rather than the ‘coarse’ implicit space, denoising is performed in a ‘smooth’ variational [37] latent space from which an implicit latent code is decoded. **Mesh-Diffusion** [46] parameterises meshes with tetrahedral grids, representing them with 3D convolutions and generating shapes from deformations of uniform tetrahedrals. Recent volumetric diffusion models including **LAS-D** [88], **SD-Fusion** [11] and **Wavelet-Diffusion** [34] perform diffusion in the voxelized occupancy space, SDF space, and SDF-decomposed wavelet space respectively to generate shapes. (ii) Conditional generation of shapes has been explored through language prompts [46], images [21], point-clouds [50] and sketches [7, 88]. We compare our sketch-based shape generation with *LAS-D* [88] as well as image-based shape generation models like *SDFusion* [11]. We also compare with deterministic sketch to shape reconstruction algorithms. **Sketch2Model** [85] builds shapes from sketches adversarially with sketch-viewpoint aware generative modelling. **Sketch2Mesh** [24] renders shapes differentially and optimizes the mesh by comparing (a) shape silhouette



Figure 6. Qualitative comparisons of our method with LAS-D [88] and SENS [7] on sketches of different levels of abstraction from (i) highly detailed sketches by artists (first 2 from ProSketch-3D [89]) and (ii) sketches by amateurs with perspective distortions (next 6 from AmateurSketch-3D [59]) to (iii) Highly abstract sketches drawn in <20s (last from *Quick-Draw!* [25]). While neither LAS-D nor our algorithm has seen hand-drawn doodles during training, SENS[7] was trained on ProSketch-3D sketches.

Methods	AmateurSketch-3D		ProSketch-3D		Inference	
	CD ↓	EMD ↓	CD ↓	EMD ↓	Time ↓	Params ↓
Sketch2Model [85]	0.913	0.631	1.050	0.301	1.47s	85M
SketchSampler [20]	0.615	0.537	0.582	0.240	1.64s	46M
Sketch2Mesh [24]	0.257	0.211	0.228	0.171	90s	9M
SENS [7]	0.121	0.096	-	-	3.33s	177M
SDFusion [11]	0.632	0.483	0.375	0.259	25s	1099M
LAS-D [88]	0.159	0.128	0.195	0.147	6s	767M
Ours	0.109	0.091	0.093	0.087	2s	19M
Ours-Multi-View	0.097	0.089	0.085	0.082	2.5s	19M

Table 1. Comparison of conditional generation and model efficiency on the AmateurSketch-3D [59] and ProSketch-3D [89] datasets (unit for CD here is 10^{-1}). Performance of SENS [7] not reported for ProSketch-3D as it is in their training set.

with sketches and (b) segmentation masks from shape renderers with segmented sketches. SENS [7] learns a Vision-Transformer mapping of sketch patches to latent space of \mathcal{D} , forging a *one-to-one* relationship between sketches and corresponding shape implicit codes.

Comparative Analysis: We include a qualitative performance analysis of our algorithm with SOTA LAS-D [88] and SENS [7] on different sketch datasets in Fig. 6. For quantitative evaluation of conditional shape generation, we follow recent works [88] to sample 2048 points on both generated and ground truth meshes and use point-cloud based metrics like (i) **Chamfer Distance** (CD) computed as the squared distances between nearest points in predicted and ground truth point clouds and (ii) **Earth Mover’s Distance** (EMD) computed as the average point-to-point distance under a global match of the prediction with the ground truth. From Tab. 1, our algorithm outperforms previous SOTA LAS-D [88] and SENS [7] in sketch-conditioned generation

by 0.0076/0.049 and 0.0012/0.05 on average CD/EMD respectively without any training on human sketches. For unconditional shape generation (Tab. 2), we evaluate diversity with shading-image based **Fréchet Inception Distance** (FID) [87] where images are rendered both from the predicted and ground truth meshes in 20 uniform views, from which view-specific FIDs are averaged. We find that our method outperforms Mesh-Diffusion [46], SDF-StyleGAN [87], and Wavelet Diffusion [34] in FID by 23.57, 14.16, and 6.32 respectively (lower is better) while performing at par with frameworks like LAS-D [88] and Diffusion-SDF [12]. We also use point-cloud based metrics [1, 80] like (i) **Coverage** (Cov.) measuring the fraction of generated shape point clouds that match (based on CD/EMD) with the ground truth shape point clouds, (ii) **Minimum Matching Distance** (MMD) measuring the average minimum matching distances (Chamfer/Earth Mover Distances) between point clouds from generated and ground truth shapes and (iii) **1-Nearest Neighbour Accuracy** (1-NNA) measuring the similarity between generated and ground truth point clouds using a 1-NN classifier. While for Cov. and MMD, a higher and lower value respectively is better, for 1-NNA, the best model should be closest to 50 % denoting that the 1-NN classifier is confused whether a sample is real or generated. We evaluate with these metrics on the ShapeNet [9] dataset where our model performs (Tab. 2) at par with other diffusion-based 3D generative frameworks.

Abstraction robustness: We use CLIPasso [68] to simulate sketches from edgmaps of 100 random rendered shapes with 8, 12, 16, 24, and 32 strokes to represent increasing abstraction levels following their assumption [68] of abstraction as a function of number of strokes. Sketches of vary-

Methods	COV(%) \uparrow		MMD \downarrow		1-NNA (%)		FID \downarrow
	CD	EMD	CD	EMD	CD	EMD	
SDF-StyleGAN [87]	45.60	45.50	0.158	0.184	63.25	67.80	36.48
Diffusion-SDF [12]	65.35	59.22	0.106	0.133	51.18	54.3	21.07
Mesh-Diffusion [46]	46.00	46.71	0.132	0.173	53.69	57.63	39.62
LAS-D [88]	53.76	52.43	0.138	0.175	64.53	65.15	20.45
Wavelet-Diffusion [34]	52.88	47.64	0.133	0.173	61.14	66.92	28.64
Ours	63.39	61.2	0.103	0.149	54.25	57.12	22.32

Table 2. Comparison of unconditional shape generation from training on ShapeNet Chairs (unit for CD here is 10^{-1}).

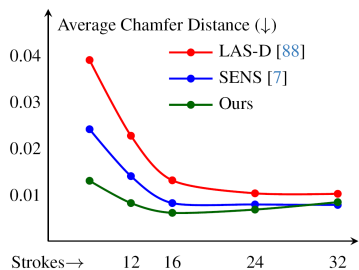


Figure 7. Drop in quality of generated shapes in LAS-D [88] and SENS [7] from increasingly abstract sketches generated with CLIPasso [68].

Methods	CD \downarrow
<i>Ours</i>	0.109
w/o alignment	0.171
part-disentangled sketch encodings	0.162
with	
1 Att ⁿ layer	0.127
2 Att ⁿ layers	0.114
4 Att ⁿ layers	0.110

Table 3. Ablative studies on the AmateurSketch-3D dataset [59] (unit for CD here is 10^{-1}).

ing abstraction levels are then used to generate 3D shapes whose Chamfer Distance (CD) from the ground truth shape is plotted in Fig. 7. Despite both our method and SENS [7] being trained on CLIPasso sketches, we find that SENS performs worse at lower abstraction levels than our algorithm, with the average CD increasing from 0.0078 to 0.0241 compared to our CD increase (0.0084 to 0.0130). Algorithms like LAS-D [88] train on edgmaps only, resulting in a lower performance (CD increases from 0.0102 to 0.0390).

Shape editing: We explore the editability of generated shapes through (i) edits in the input sketch and (ii) interpolation between shapes. Representing sketches through local part-aware encodings allows us to perform local edits [29] in the shape without disturbing the global structure, with stroke-level sketch edits. Given a sketch edit $\mathcal{S} \rightarrow \mathcal{S}'$, we identify parts affected in the edit, by the Euclidean distance between part encodings. Next, we generate corresponding part-latents Z and Z' , and then replace part information in Z corresponding to edited parts in \mathcal{S}' with respective part information from Z' . As observed in Fig. 1 (right), we are able to ensure local edits while maintaining fine-grained control over generated shapes. Fig. 9 demonstrates edits by morphing shapes into one another via interpolation of sketch representations only.

Ambiguity in Sketches: Single-view sketches of 3D shapes suffer from ambiguity in representing complex structures due to (i) the sparse nature of line drawings and (ii) the limited representative power of a single viewpoint. Deterministic reconstruction of shapes from sketches is thus inherently *biased* as they selectively form one-to-one sketch-to-

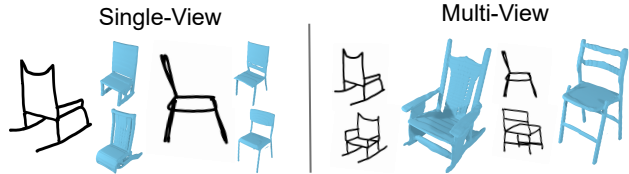


Figure 8. Ambiguity in the input sketch, owing to occlusion from specific views or poor quality sketches leads to (slight) variance in generated shapes that can be fixed with multi-view sketches.




Figure 9. Generated shapes can be smoothly morphed into one another by simple interpolation of sketch representations.

shape correspondences in a one-to-many setting. Recognising this ambiguity, we demonstrate shape variations from single-view sketches from AmateurSketch-3D in Fig. 8. To reduce ambiguity in generated shapes, we perform multi-view sketch to shape reconstruction by simply aggregating part representations from multiple views. This increases shape reconstruction accuracy (Fig. 8) in AmateurSketch-3D and ProSketch-3D sketches, dropping CD/EMD from 0.0109/0.093 to 0.0097/0.085 respectively (Tab. 1).

Ablation Experiments: We perform ablative studies to analyse the contribution of individual elements in our proposed approach. Using Chamfer Distance as a metric on the AmateurSketch-3D [59] dataset, we summarise our results in Tab. 3. We find a increase in CD of (i) 0.0062 on removing latent alignment (Sec. 3.3), (ii) of 0.0053 on replacing part-aware sketch encodings (Sec. 3.5) with naive ImageNet features, and (iii) by 0.0018/0.0005/0.0001 with 1/2/4 multi-head attention layers respectively. Removing part-aware sketch features, in particular (Tab. 3 - w/o part-disentangled sketch encodings), results in much worse generalisation with +0.0053 CD on human-drawn sketches compared to synthetic sketches (+0.0031 CD).

Time and memory constraints: Towards practical generation of 3D shapes, we evaluate the compute constraints of our algorithm against conditional shape generative networks like LAS-D, SENS, and SDFusion (Tab. 1). Our model not only outperforms conditional networks in parameter count, but is also lighter than unconditional generative models (which do not have a conditioning network) like Diffusion-SDF (123M), and Wavelet-Diffusion (99M). As a lightweight network, our model naturally has minimal inference time in comparison to SOTA generative networks.

5. Conclusion

We present a latent diffusion pipeline to generate precise 3D shapes from abstract sketches with part-disentangled sketch representations. We demonstrate fine-grained control over

generated shapes with hand-drawn doodles on a variety of abstraction levels – from highly abstract *Quick-Draw!* sketches to artistic *ProSketch-3D* diagrams. Our generated shapes can be automatically edited with fine-grained edits on the conditioning sketch and can be further improved with sketches from additional viewpoints by a simple aggregation of sketch features. Finally, we demonstrate our pipeline to be much more efficient than SOTA 3D generative models.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *ICLR*, 2018. 7
- [2] Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. Break-a-scene: Extracting multiple concepts from a single image. In *SIGGRAPH Asia*, 2023. 3
- [3] Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, and Yaser Sheikh. Modeling facial geometry using compositional vaes. In *CVPR*, 2018. 2
- [4] Matthew Berger, Joshua A Levine, Luis Gustavo Nonato, Gabriel Taubin, and Claudio T Silva. A benchmark for surface reconstruction. *ACM TOG*, 2013. 2
- [5] Matthew Berger, Andrea Tagliasacchi, Lee M Seversky, Pierre Alliez, Gael Guennebaud, Joshua A Levine, Andrei Sharf, and Claudio T Silva. A survey of surface reconstruction from point clouds. In *Computer graphics forum*, 2017. 2
- [6] Ayan Kumar Bhunia, Yongxin Yang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Sketch less for more: On-the-fly fine-grained sketch-based image retrieval. In *CVPR*, 2020. 1
- [7] Alexandre Binniger, Amir Hertz, Olga Sorkine-Hornung, Daniel Cohen-Or, and Raja Giryes. Sens: Sketch-based implicit neural shape modeling. *arXiv preprint arXiv:2306.06088*, 2023. 1, 3, 6, 7, 8
- [8] Caroline Chan, Frédo Durand, and Phillip Isola. Learning to generate line drawings that convey geometry and semantics. In *CVPR*, 2022. 4, 6
- [9] 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. 6, 7, 1, 2
- [10] Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. Bsp-net: Generating compact meshes via binary space partitioning. In *CVPR*, 2020. 3
- [11] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, 2023. 6, 7
- [12] Gene Chou, Yuval Bahat, and Felix Heide. Diffusionsdf: Conditional generative modeling of signed distance functions. In *ICCV*, 2023. 3, 6, 7, 8
- [13] Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. Democratising 2d sketch to 3d shape retrieval through pivoting. In *ICCV*, 2023. 2
- [14] Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. SceneTrilogy: On Human Scene-Sketch and its Complementarity with Photo and Text. In *CVPR*, 2023. 1
- [15] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016. 3
- [16] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *CVPR*, 2021. 2
- [17] Aysegul Dundar, Jun Gao, Andrew Tao, and Bryan Catanzaro. Progressive learning of 3d reconstruction network from 2d gan data. *arXiv preprint arXiv:2305.11102*, 2023. 3
- [18] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *ICCV*, 2023. 3
- [19] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, 2017. 3
- [20] Chenjian Gao, Qian Yu, Lu Sheng, Yi-Zhe Song, and Dong Xu. Sketchsampler: Sketch-based 3d reconstruction via view-dependent depth sampling. In *ECCV*, 2022. 1, 7
- [21] 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. In *NeurIPS*, 2022. 6
- [22] Yulia Gryaditskaya, Mark Sypesteyn, Jan Willem Hoftijzer, Sylvia C Pont, Frédo Durand, and Adrien Bousseau. Opensketch: a richly-annotated dataset of product design sketches. *ACM TOG*, 2019. 3
- [23] Jiatao Gu, Qingzhe Gao, Shuangfei Zhai, Baoquan Chen, Lingjie Liu, and Josh Susskind. Control3diff: Learning controllable 3d diffusion models from single-view images. In *3DV*, 2024. 3
- [24] Benoit Guillard, Edoardo Remelli, Pierre Yvernay, and Pascal Fua. Sketch2mesh: Reconstructing and editing 3d shapes from sketches. In *ICCV*, 2021. 1, 3, 6, 7
- [25] David Ha and Douglas Eck. A neural representation of sketch drawings. In *arXiv preprint arXiv:1704.03477*, 2018. 3, 6, 7
- [26] Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. Point2mesh: A self-prior for deformable meshes. In *SIGGRAPH*, 2020. 2
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3, 5
- [28] Amir Hertz, Rana Hanocka, Raja Giryes, and Daniel Cohen-Or. Pointgmm: A neural gmm network for point clouds. In *CVPR*, 2020. 3, 4
- [29] Amir Hertz, Or Perel, Raja Giryes, Olga Sorkine-Hornung, and Daniel Cohen-Or. Spaghetti: Editing implicit shapes through part aware generation. *ACM TOG*, 2022. 2, 3, 4, 8, 1
- [30] Aaron Hertzmann. Why do line drawings work? a realism hypothesis. *Perception*, 2020. 3

- [31] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3, 4
- [32] Lianghua Huang, Di Chen, Yu Liu, Shen Yujun, Deli Zhao, and Zhou Jingren. Composer: Creative and Controllable Image Synthesis with Composable Conditions. In *CHI*, 2023. 4
- [33] Shengyu Huang, Zan Gojcic, Jiahui Huang, Andreas Wieser, and Konrad Schindler. Dynamic 3d scene analysis by point cloud accumulation. In *ECCV*, 2022. 2
- [34] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH*, 2022. 3, 6, 7, 8
- [35] Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 1960. 4
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 6
- [37] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 6
- [38] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial generation of continuous implicit shape representations. In *Eurographics*, 2020. 3
- [39] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Picture that Sketch: Photorealistic Image Generation from Abstract Sketches. In *CVPR*, 2023. 2
- [40] 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 TOG*, 2018. 3
- [41] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. Sketch2cad: Sequential cad modeling by sketching in context. *ACM TOG*, 2020. 3
- [42] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. Free2cad: Parsing freehand drawings into cad commands. *ACM TOG*, 2022. 3
- [43] Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *CVPR*, 2023. 1, 2
- [44] Or Litany, Alex Bronstein, Michael Bronstein, and Ameesh Makadia. Deformable shape completion with graph convolutional autoencoders. In *CVPR*, 2018. 2
- [45] Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. Deepmetahandles: Learning deformation meta-handles of 3d meshes with biharmonic coordinates. In *CVPR*, 2021. 2
- [46] Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. In *ICLR*, 2023. 1, 2, 3, 4, 6, 7, 8
- [47] 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*. Association for Computing Machinery, 1998. 4, 6
- [48] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *NeurIPS*, 2022. 3, 6
- [49] Andrew Luo, Tianqin Li, Wen-Hao Zhang, and Tai Sing Lee. Surfgen: Adversarial 3d shape synthesis with explicit surface discriminators. In *ICCV*, 2021. 3
- [50] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*, 2021. 3, 6
- [51] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2018. 2, 3
- [52] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 3
- [53] Gimin Nam, Mariem Khelifi, Andrew Rodriguez, Alberto Tono, Linqi Zhou, and Paul Guerrero. 3d-ldm: Neural implicit 3d shape generation with latent diffusion models. *arXiv preprint arXiv:2212.00842*, 2022. 3
- [54] Charlie Nash and Christopher KI Williams. The shape variational autoencoder: A deep generative model of part-segmented 3d objects. In *Computer Graphics Forum*, 2017. 3
- [55] Kaiyue Pang, Ke Li, Yongxin Yang, Honggang Zhang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Generalising Fine-Grained Sketch-Based Image Retrieval. In *CVPR*, 2019. 1
- [56] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 2, 3, 6
- [57] Despoina Paschalidou, Ali Osman Ulusoy, and Andreas Geiger. Superquadrics revisited: Learning 3d shape parsing beyond cuboids. In *CVPR*, 2019. 3
- [58] Despoina Paschalidou, Luc Van Gool, and Andreas Geiger. Learning unsupervised hierarchical part decomposition of 3d objects from a single rgb image. In *CVPR*, 2020. 3
- [59] 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 TIP*, 2021. 2, 3, 4, 6, 7, 8, 1
- [60] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 2
- [61] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 4
- [62] Aneeshan Sain, Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. CLIP for All Things Zero-Shot Sketch-Based Image Retrieval, Fine-Grained or Not. In *CVPR*, 2023. 1
- [63] Patsorn Sangkloy, Wittawat Jitkrittum, Diyi Yang, and James Hays. A Sketch Is Worth a Thousand Words: Image Retrieval with Text and Sketch. In *ECCV*, 2022. 1
- [64] Jifei Song, Yi-Zhe Song, Tao Xiang, and Timothy Hospedales. Fine-Grained Image Retrieval: the Text/Sketch Input Dilemma. In *BMVC*, 2017. 1
- [65] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 3

- [66] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *ICCV*, 2017. 2
- [67] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4, 5
- [68] 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 TOG*, 2022. 2, 4, 6, 7, 8
- [69] Dan Wang, Xinrui Cui, Xun Chen, Zhengxia Zou, Tianyang Shi, Septimiu Salcudean, Z Jane Wang, and Rabab Ward. Multi-view 3d reconstruction with transformers. In *ICCV*, 2021. 3
- [70] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. 4
- [71] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM TOG*, 2017. 2
- [72] Peng-Shuai Wang, Yang Liu, and Xin Tong. Dual octree graph networks for learning adaptive volumetric shape representations. *ACM TOG*, 2022. 2
- [73] Tuanfeng Y Wang, Duygu Ceylan, Jovan Popovic, and Niloy J Mitra. Learning a shared shape space for multimodal garment design. In *SIGGRAPH Asia*, 2018. 3
- [74] Xin Wei, Xiang Gu, and Jian Sun. Learning generalizable part-based feature representation for 3d point clouds. In *NeurIPS*, 2022. 2
- [75] Huikai Wu, Junge Zhang, and Kaiqi Huang. Point cloud super resolution with adversarial residual graph networks. In *BMVC*, 2020. 2
- [76] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 2
- [77] Peng Xu, Timothy M Hospedales, Qiyue Yin, Yi-Zhe Song, Tao Xiang, and Liang Wang. Deep learning for free-hand sketch: A survey. *IEEE T-PAMI*, 2022. 2
- [78] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *NeurIPS*, 2019. 3
- [79] Rui Xu, Zongyan Han, Le Hui, Jianjun Qian, and Jin Xie. Domain disentangled generative adversarial network for zero-shot sketch-based 3d shape retrieval. In *AAAI*, 2022. 3
- [80] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *ICCV*, 2019. 7
- [81] Fenggen Yu, Zhiqin Chen, Manyi Li, Aditya Sanghi, Hooman Shayani, Ali Mahdavi-Amiri, and Hao Zhang. Capri-net: Learning compact cad shapes with adaptive primitive assembly. In *CVPR*, 2022. 3
- [82] Fenggen Yu, Qimin Chen, Maham Tanveer, Ali Mahdavi Amiri, and Hao Zhang. D²csg: Unsupervised learning of compact csg trees with dual complements and dropouts. In *NeurIPS*, 2023. 3
- [83] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022. 2
- [84] Cheng Zhang, Haocheng Wan, Xinyi Shen, and Zizhao Wu. Pvt: Point-voxel transformer for point cloud learning. *IJIS*, 2022. 2
- [85] Song-Hai Zhang, Yuan-Chen Guo, and Qing-Wen Gu. Sketch2model: View-aware 3d modeling from single free-hand sketches. In *CVPR*, 2021. 6, 7
- [86] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. In *ICLR*, 2021. 3
- [87] Xinyang Zheng, Yang Liu, Pengshuai Wang, and Xin Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. In *Computer Graphics Forum*, 2022. 3, 6, 7, 8
- [88] 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. In *SIGGRAPH*, 2023. 1, 2, 3, 6, 7, 8
- [89] 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 T-CSVT*, 2020. 2, 4, 6, 7

Doodle Your 3D: From *Abstract* Freehand Sketches to *Precise* 3D Shapes

Supplementary Material

A. Further Details on Shape Decomposition

Here, we clarify Sec. 3.2, elaborating on the decomposition of shapes (as meshes) into part-latents Z . To represent a ground truth shape M with decoder \mathcal{D} , (i) the ground truth occupancy values for coordinates $X = (x, y, z)$ in and around the shape are recorded along with the coordinates themselves. (ii) The decoder \mathcal{D} is trained to decode a randomly initialised part-latent Z to an implicit code as $I = \mathcal{D}(Z)$. (iii) Finally I is used in implicit function f_θ to *predict* occupancy values for known coordinates X , as $O_I = f_\theta(I, X)$. During pre-training, \mathcal{D} , Z , and f_θ are optimized together with binary cross entropy loss (\mathcal{L}_{BCE}) against the recorded (ground truth) occupancy values.

The decomposition of shape M occurs through its representation as part-latent $Z \in \mathbb{R}^{m \times d}$ using the decoder \mathcal{D} . Ideally, after disentanglement each latent code ω_i in part-latent $Z = \{\omega_i\}_{i=1}^m$, sufficiently and independently represents individual components of shape M , thus successfully breaking down M into m parts represented as $\{\omega_i\}_{i=1}^m$. This disentanglement of part-latents is necessary for independent representation of shape parts. However, \mathcal{L}_{BCE} is not enough for this disentanglement, as it only encourages the final output shape to match shape M , thus ignoring *part-level* correspondence.

To optimize for disentanglement, part-latent Z is projected to part structural representation $Z \rightarrow Z_p$ and part volumetric descriptor $Z \rightarrow Z_g$. Particularly important for this representation, each part’s volumetric descriptor is a parametric 3D Gaussian that captures the probability of a 3D coordinate X belonging to that part. This establishes a relationship between coordinates in the 3D space and part-latent Z , thereby representing the volume of each part in 3D. For decomposition and disentanglement respectively, this relation of part and 3D coordinates is pivotal for (i) dissipating 3D Gaussians Z_g over the entire shape volume (M) and (ii) specifically disentangling overlapped, or closely-placed parts (as information is commonly entangled here [29]), by computing distance between part-Gaussians in 3D.

B. Segmenting Hand-Drawn Sketches

We demonstrate the generalisation of sketch-based shape generation to hand-drawn sketches after being trained on synthetic sketches only (Fig. 6). To further explore this generalisation, we include qualitative results from our auxiliary segmentation task on hand-drawn sketches from the AmateurSketch-3D dataset [59] in Fig. 10. We note that despite being trained on synthetic sketches, we can generalise and segment hand-drawn sketches fairly accurately.

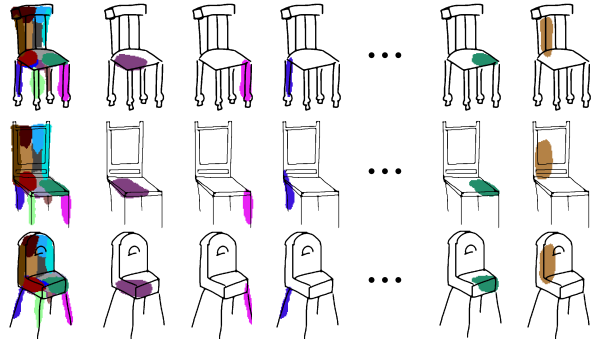


Figure 10. Segmentation results on the AmateurSketch-3D [59] dataset. (On left) We show the predicted semantic segmentation map for all 16 parts. (On right) For clarity, we show the segmentation map individually for each part.

C. Shape Interpolation




Figure 11. Interpolating shapes among four different sketches (at corners) from left to right and top to bottom.

D. Generating Shapes from Other Categories

In addition to chairs (Fig. 6), we perform sketch-based 3D generation of airplanes from the ShapeNet [9] dataset. We include qualitative results of the same in Fig. 12.




Figure 12. Single view sketch to shape generation for airplanes from the ShapeNet [9] dataset.

E. Additional Qualitative Results (Chairs)



Figure 13. Generation of chair shapes with sketches from the AmateurSketch-3D dataset [59].



Figure 14. Generation of chair shapes with sketches from the AmateurSketch-3D dataset [59] (continued from Fig. 13).