Quartet of Diffusions: Structure-Aware Point Cloud Generation through Part and Symmetry Guidance

Anonymous authors

Paper under double-blind review

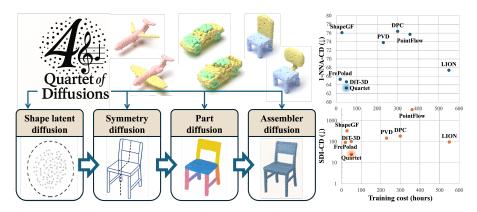


Figure 1: Our method, the Quartet of Diffusions, harmoniously orchestrates an interpretable and controllable pipeline for structure-aware, high-quality, diverse, and efficient point cloud generation guided by part and symmetry clues.

ABSTRACT

We introduce the *Quartet of Diffusions*, a structure-aware point cloud generation framework that explicitly models part composition and symmetry. Unlike prior methods that treat shape generation as a holistic process or only support part composition, our approach leverages four coordinated diffusion models to learn distributions of global shape latents, symmetries, semantic parts, and their spatial assembly. This structured pipeline ensures guaranteed symmetry, coherent part placement, and diverse, high-quality outputs. By disentangling the generative process into interpretable components, our method supports fine-grained control over shape attributes, enabling targeted manipulation of individual parts while preserving global structural consistency. A central global latent further reinforces structural coherence across assembled parts. Our experiments show that the Quartet achieves state-of-the-art performance in generating high-quality and diverse point clouds while maintaining symmetry. To our best knowledge, this is the first 3D point cloud generation framework that fully integrates and enforces both symmetry and part priors throughout the generative process.

1 Introduction

Structure-aware 3D shape generation aims to model not only the holistic surface geometry of objects but also the **underlying organizational principles** that govern their composition and formation (Chaudhuri et al., 2020). A truly structure-aware representation must comprise two essential components (Chaudhuri et al., 2020): the geometry of atomic structural elements (*e.g.*, low-level parts) and the structural patterns that dictate how these elements are assembled into coherent shapes or scenes. These patterns may manifest as hierarchical decompositions, part-whole relationships, spatial composition graphs, or local and global symmetries.

In this work, we focus on 3D point cloud generation with **explicit awareness of two fundamental structural priors** commonly observed in natural and man-made objects: **part awareness** and **symmetry awareness**. *Part awareness* captures the notion that objects are composed of semantically

meaningful, interrelated components with geometrically distinct properties. *Symmetry awareness*, in contrast, models the intrinsic geometric regularities—such as reflective or rotational symmetries—that often govern the spatial arrangement of parts. Importantly, part awareness serves as a *prerequisite* for identifying and enforcing local symmetries, which typically operate at the level of inter-part relationships. Without symmetry, shapes risk appearing disjointed or unnatural, undermining their functional performance in practical applications (*e.g.* some prior work in Fig. 4).

Compared to unstructured 3D data, structured representations offer substantial advantages. They not only enhance the *visual plausibility* and *structural coherence* of generated shapes, but also facilitate downstream tasks such as part segmentation, alignment, structure inference, and fine-grained shape editing (Schor et al., 2019; Li et al., 2020; Hertz et al., 2022; Koo et al., 2023). However, existing generative approaches largely remain structure-oblivious – treating point cloud synthesis as an unconstrained distribution learning problem. These methods often prioritize global visual fidelity and geometry-level details in the distribution, while neglecting the explicit modeling of inherent structure (Vahdat et al., 2022; Yang et al., 2019; Mo et al., 2023; Zhou et al., 2021a). As a result, they tend to produce shapes that suffer from poor part organization, broken or inconsistent symmetries, and limited generalization to unseen or complex object categories (Li et al., 2020). While a few recent efforts support part-level generation, they rarely incorporate symmetry as a *learned*, *relational prior*. To date, there remains a notable gap: the absence of unified, principled frameworks that jointly encode *part-level composition* and *symmetry structures* to guide the generation of coherent, interpretable, and generalizable 3D shapes.

To address this, we propose the *Quartet of Diffusions* (the *Quartet*), a structure-aware point cloud generation framework that explicitly leverages part and symmetry guidance. Specifically, our pipeline employs four diffusion models (Ho et al., 2020) to learn the distributions of shape latents, symmetries, semantic parts, and assemblers that assemble the full point cloud from parts. By explicitly modeling these structural distributions, the ensemble of four diffusions harmoniously orchestrates an effective pipeline for structure-aware point cloud generation through part and symmetry guidance: Our method enables the generation of high-quality, diverse point clouds with guaranteed symmetry. The disentangled structural representation makes the generation process interpretable and controllable, facilitating targeted modifications to individual parts while preserving global consistency. Structural coherence is further reinforced by a global shape latent, which anchors part assembly to the overall geometry. This approach addresses key limitations in structure modeling and represents the first 3D shape generation method to guarantee symmetry in the generated shapes.

2 RELATED WORK

Symmetry in 3D shapes Symmetry is a fundamental geometric property observed across natural and human-made objects. Extensive research exists on symmetry detection (Mitra et al., 2013a; Atallah, 1985; Illingworth & Kittler, 1988; Mitra et al., 2006; Je et al., 2024; Zhou et al., 2021b; Gao et al., 2020; Fukunaga & Hostetler, 1975; Comaniciu & Meer, 2002) (for more details, see Appendices A and B.1). Symmetry clues are widely used in various 3D vision tasks including acquisition and representation (Yang et al., 2024; Buades et al., 2008; Li et al., 2010; Pauly et al., 2005; Thrun & Wegbreit, 2005; Xu et al., 2009; Zheng et al., 2010), classification (Kazhdan et al., 2004; Martinet et al., 2006; Podolak et al., 2006), perception (Reisfeld et al., 1995), manipulation (Mitra et al., 2007; Podolak et al., 2007; Panozzo et al., 2012; Gal et al., 2009; Mehra et al., 2009), reconstruction (Phillips et al., 2016; Xu et al., 2024; Tulsiani et al., 2020), inverse rendering (Wu et al., 2020b), and refinement (Mitra et al., 2007). However, symmetry-aware 3D shape generation remains underexplored. Aside from our method, few approaches guarantee symmetry. PAGENet (Li et al., 2020) promotes symmetry via an MSE loss between a shape and its reflection, but it is limited to reflectional symmetry and does not guarantee it.

Part-based 3D shape generation Part-based 3D shape generation leverages the modular structure of objects (Mitra et al., 2014; 2013b) and emerges as a vital paradigm for modeling complex geometries (Zerroug & Nevatia, 1999; Kim et al., 2013), enhancing diversity (Schor et al., 2019; Chen et al., 2024) and facilitating downstream tasks including recognition (Hoffman & Richards, 1984), retrieval (Chang et al., 2015; Mitra et al., 2014), and manipulation (Huang et al., 2014). Early methods focus on geometric template fitting or hierarchical part assembly (Funkhouser et al., 2004; Bokeloh et al., 2010; Kalogerakis et al., 2012; Berthelot et al., 2017; Cohen-Or & Zhang, 2016; Fish et al., 2014). More recent work harnesses implicit shape representations (Koo et al., 2023; Hertz et al., 2022; Talabot et al., 2025; Hui et al., 2022; Genova et al., 2019; 2020). For example,

PartSDF (Talabot et al., 2025) models parts using implicit neural fields, enabling both continuous interpolation and discrete composition.

Deep neural networks are widely adopted for part-based shape modeling (Schor et al., 2019; Li et al., 2020; Dubrovina et al., 2019; Li et al., 2024; Wu et al., 2020a; Huang et al., 2015; Li et al., 2017; Mo et al., 2019; Zou et al., 2017; Nash & Williams, 2017; Wu et al., 2019; Wang et al., 2018; 2019; Gao et al., 2019). For instance, CompoNet (Schor et al., 2019) enhances diversity by varying both parts and their compositions, while PASTA (Li et al., 2024) generates shapes conditioned on part arrangement for fine-grained control. Similarly, our Quartet performs structure-aware 3D generation by modeling both parts and their compositions, with the added benefit of symmetry enforcement.

Diffusion models Diffusion models (Ho et al., 2020) are generative models based on Markovian diffusion processes (see Appendix B.3 for more details). They have demonstrated remarkable performance in various domains (Croitoru et al., 2023; Cao et al., 2024; Yang et al., 2023), particularly in image (Dhariwal & Nichol, 2021; Ho et al., 2020; Rombach et al., 2022; Ramesh et al., 2022; Brooks et al., 2023; Saharia et al., 2022), speech (Chen et al., 2020; Jeong et al., 2021; Liu et al., 2022), video (Ho et al., 2022; Xing et al., 2024; Luo et al., 2023; Yang et al., 2023), 3D scene (Wei et al., 2023; Zhai et al., 2023; Tang et al., 2024), and 3D object generation (Zhou et al., 2021a; Luo & Hu, 2021; Vahdat et al., 2022; Nakayama et al., 2023; Wu et al., 2023; Mo et al., 2023; Liu et al., 2019; Zhou et al., 2024; Koo et al., 2023). Accordingly, our Quartet recruits four diffusion models to learn the distributions of shape latents, symmetries, parts, and assemblers.

3 **METHOD**

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126 127 128

129

130 131

132

133

134 135

136

137

138

139

140

141

142

143

144

145

146 147

148

149

150

151

152

153

154 155

156

We aim to model the distribution of point clouds $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{3 \times N}$ as that over a collection of semantic parts $\{\mathbf{p}_j\}_{j=1}^M$ through part and symmetry guidance. Specifically, we view each point cloud $\mathbf{x} = \bigcup_{j=1}^{M} T_j \mathbf{p}_j$ as a composition of its M semantic parts and their corresponding assemblers. Each assembler T_i is a transformation comprising translation, rotation, and scaling in 3D Euclidean space, mapping the part \mathbf{p}_i to the correct position, orientation, and scale in the original point cloud.

The Quartet models four distributions using parameterized diffusion models: 1. point cloud shape latents $p_{\theta}(\mathbf{z})$ (Sec. 3.1), 2. part-wise symmetries $p_{\zeta}(\mathcal{S}_i \mid \mathbf{z})$ (Sec. 3.2), 3. point cloud parts $p_{\xi}(\mathbf{p}_i \mid \mathcal{S}_j, \mathbf{z})$ (Sec. 3.3), and 4. assemblers $p_{\psi}(T_i \mid \mathbf{w}_j, \mathbf{p}_i, \mathcal{S}_j, \mathbf{z})$, additionally conditioned on part latent w_i (Sec. 3.4). These components collectively define the Quartet of Diffusions architecture. Pipeline overviews are shown in Figs. 1 and 2.

Modeling point cloud distributions in this way offers several key advantages: 1. It introduces variability at both the shape and assembly levels, enabling combinatorially greater diversity; 2. By explicitly generating and enforcing symmetry early in the pipeline, the model produces outputs with more realistic and consistent symmetric properties; 3. The separate modeling of symmetries, parts, and assembly allows for better interpretability and fine-grained control – individual parts can be manipulated independently without compromising global structure; and 4. A central shape latent z ensures structural coherence across all components. These benefits are demonstrated through extensive experiments in Sec. 4.

The Quartet is trained by sequentially optimizing the parameters of four distributions to fit the given dataset \mathcal{X} . Point cloud generation is performed in two phases: part generation and part assembly. In part generation, a shape latent z is first sampled from the latent distribution $p_{\theta}(\mathbf{z})$. Conditioned on \mathbf{z} , symmetry groups \mathcal{S}_j are sampled from $p_{\zeta}(\mathcal{S}_j \mid \mathbf{z})$, and parts \mathbf{p}_j are drawn from $p_{\xi}(\mathbf{p}_i \mid \mathcal{S}_j, \mathbf{z})$ for $j = 1, 2, \dots, M$; In part assembly, parts \mathbf{p}_i are encoded into latent representations \mathbf{w}_i using an encoder $q_{\phi}(\mathbf{w}_i \mid \mathbf{p}_i, \mathbf{z})$, and the corresponding assemblers T_i are sampled from $p_{\psi}(T_j \mid \mathbf{w}_j, \mathbf{p}_j, \mathcal{S}_j, \mathbf{z})$. These assemblers are then applied to the parts to reconstruct the full point cloud $\tilde{\mathbf{x}} := \bigcup_{i=0}^{M} T_i \mathbf{p}_i$. Mathematically, the generation process can thus be formulated as

$$p\left(\left\{\left(\mathbf{p}_{j}, T_{j}, \mathcal{S}_{j}, \mathbf{w}_{j}\right)\right\}_{j=0}^{M}, \mathbf{z}\right) = p_{\theta}\left(\mathbf{z}\right) p\left(\left\{\left(\mathbf{p}_{j}, T_{j}, \mathcal{S}_{j}, \mathbf{w}_{j}\right)\right\}_{j=0}^{M} \middle| \mathbf{z}\right) = p_{\theta}\left(\mathbf{z}\right) \prod_{i=0}^{M} p(\mathbf{p}_{j}, T_{j}, \mathbf{w}_{j} \mid \mathbf{z}\right)$$
(1)

$$p\left(\left\{\left(\mathbf{p}_{j}, T_{j}, \mathcal{S}_{j}, \mathbf{w}_{j}\right)\right\}_{j=0}^{M}, \mathbf{z}\right) = p_{\boldsymbol{\theta}}\left(\mathbf{z}\right) p\left(\left\{\left(\mathbf{p}_{j}, T_{j}, \mathcal{S}_{j}, \mathbf{w}_{j}\right)\right\}_{j=0}^{M} \middle| \mathbf{z}\right) = p_{\boldsymbol{\theta}}\left(\mathbf{z}\right) \prod_{j=0}^{M} p(\mathbf{p}_{j}, T_{j}, \mathbf{w}_{j} \mid \mathbf{z}\right)$$
(1)
$$= p_{\boldsymbol{\theta}}\left(\mathbf{z}\right) \prod_{j=0}^{M} p_{\boldsymbol{\zeta}}\left(\mathcal{S}_{j} \mid \mathbf{z}\right) p_{\boldsymbol{\xi}}\left(\mathbf{p}_{j} \mid \mathcal{S}_{j}, \mathbf{z}\right) q_{\boldsymbol{\phi}}\left(\mathbf{w}_{j} \mid \mathbf{p}_{j}, \mathbf{z}\right) p_{\boldsymbol{\psi}}\left(T_{j} \mid \mathbf{w}_{j}, \mathbf{p}_{j}, \mathcal{S}_{j}, \mathbf{z}\right),$$
(2)

166

168

169

170

171

172

173

174

175

176

177

178

187

189

190

191

192

193

194

195

196

197

198

199

201

202

203

205 206

207

208

209

210

211

212

213

214

215

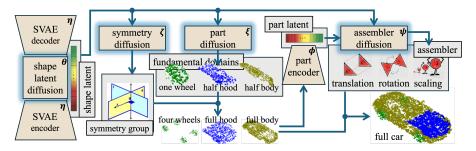


Figure 2: Overview of the Quartet's architecture. Four diffusions are employed to learn the distributions of shape latents, symmetries, parts, and assemblers. By explicitly modeling different distributions, the Quartet provides an interpretable and controllable framework to generate high-quality, diverse 3D shapes with guaranteed symmetry. Beige blocks denote learnable modules; gray blocks indicate outputs directly generated from the models.

where the second equality in (1) follows from the conditional independence of parts given the shape latent z. Therefore, the overall point cloud distribution can be modeled as

$$p(\mathbf{x}) = p\left(\left\{\left(\mathbf{p}_{j}, T_{j}\right)\right\}_{j=0}^{M}\right) = \int \cdots \int p\left(\left\{\left(\mathbf{p}_{j}, T_{j}, \mathcal{S}_{j}, \mathbf{w}_{j}\right)\right\}_{j=0}^{M}, \mathbf{z}\right) \left(\prod_{j=1}^{M} d\mathcal{S}_{j}\right) \left(\prod_{j=1}^{M} d\mathbf{w}_{j}\right) d\mathbf{z}$$
(3)
$$\stackrel{(2)}{=} \int p_{\theta}(\mathbf{z}) \prod_{j=0}^{M} \left(\int \int p_{\xi}\left(\mathbf{p}_{j} \mid \mathcal{S}_{j}, \mathbf{z}\right) q_{\phi}\left(\mathbf{w}_{j} \mid \mathbf{p}_{j}, \mathbf{z}\right) p_{\psi}\left(T_{j} \mid \mathbf{w}_{j}, \mathbf{p}_{j}, \mathcal{S}_{j}, \mathbf{z}\right) d\mathcal{S}_{j} d\mathbf{w}_{j}\right) d\mathbf{z}.$$
(4)

$$\stackrel{(2)}{=} \int p_{\boldsymbol{\theta}} \left(\mathbf{z} \right) \prod_{j=0}^{M} \left(\int \int p_{\boldsymbol{\xi}} \left(\mathbf{p}_{j} \mid \mathcal{S}_{j}, \mathbf{z} \right) q_{\boldsymbol{\phi}} \left(\mathbf{w}_{j} \mid \mathbf{p}_{j}, \mathbf{z} \right) p_{\boldsymbol{\psi}} \left(T_{j} \mid \mathbf{w}_{j}, \mathbf{p}_{j}, \mathcal{S}_{j}, \mathbf{z} \right) d\mathcal{S}_{j} d\mathbf{w}_{j} \right) d\mathbf{z}. \tag{4}$$

3.1 Shape Latent Diffusion

Although we model each distribution separately, they must remain aware of the underlying point cloud to coordinate coherently. Directly conditioning on the full point cloud is computationally expensive due to its high dimensionality. Instead, we condition them on shape latents obtained from a variational autoencoder (VAE) (Kingma & Welling, 2013; Zhou et al., 2024; Vahdat et al., 2022) (see Appendix B.2 for background).

To effectively encode point clouds while preserving semantic information, we propose a novel sparse variational autoencoder (SVAE) with a latent diffusion modeling its latent distribution. SVAE builds on a VAE architecture implemented with point-voxel convolutional neural networks (PVCNNs) (Liu et al., 2019; Zhou et al., 2021a), drawing inspiration from sparse autoencoders (Ng et al., 2011). Prior work suggests that enforcing sparsity in activation layers improves interpretability (Cunningham et al., 2023; Makelov et al., 2024; Marks et al., 2024) and semantic disentanglement (Bricken et al., 2023; O'Neill et al., 2024), which is critical for our downstream tasks. To this end, we introduce a sparsity constraint on the final activation layer $a_{\eta}(\mathbf{x})$ of the encoder $q_{\eta}(\mathbf{z} \mid \mathbf{x})$ while maximizing the VAE evidence lower bound (ELBO) \mathcal{L}_{ELBO} (Eq. (19) in Appendix B.2):

$$\max_{\boldsymbol{\eta}} \mathbb{E}_{\mathbf{x}} \left[\mathcal{L}_{\text{ELBO}}(\boldsymbol{\eta}; \mathbf{x}) \right] \text{ subject to } \mathbb{E}_{\mathbf{x}} \left[\| a_{\boldsymbol{\eta}}(\mathbf{x}) \|_{1} \right] < \delta, \tag{5}$$

where $\delta > 0$ controls the sparsity strength. Leveraging Karush–Kuhn–Tucker (KKT) methods (Kuhn & Tucker, 1951; Karush, 1939), we reformulate the constrained optimization problem into a Lagrangian $\mathcal{F}(\eta, \lambda; \mathbf{x})$ whose optimal point is a global maximum over the domain of η and obtain its lower bound:

$$\mathcal{F}(\boldsymbol{\eta}, \lambda; \mathbf{x}) := \mathcal{L}_{\text{ELBO}}(\boldsymbol{\eta}; \mathbf{x}) - \lambda \left(\|a_{\boldsymbol{\eta}}(\mathbf{x})\|_{1} - \delta \right)$$
(6)

$$\geq \mathcal{L}_{\text{ELBO}}(\boldsymbol{\eta}; \mathbf{x}) - \lambda \|a_{\boldsymbol{\eta}}(\mathbf{x})\|_{1} =: \mathcal{L}_{\text{SVAE}}(\boldsymbol{\eta}; \mathbf{x}), \tag{7}$$

where λ is the KKT multiplier. The SVAE is trained by maximizing this lower bound \mathcal{L}_{SVAE} . The hyperparameter λ controls the trade-off between reconstruction quality and sparsity, with $\lambda=0$ recovering the standard ELBO.

While a simple Gaussian prior is commonly used for VAE latent distributions, evidence suggests that such a restricted prior cannot accurately capture complex latent distributions (i.e. the prior hole problem (Vahdat et al., 2021; Zhou et al., 2023)) and can degrade VAE performance (Chen et al., 2016). To address this, we follow prior work (Zhou et al., 2024) and employ a diffusion model $p_{\theta}(\mathbf{z})$, implemented with a U-Net backbone, to learn the latent distribution. Once trained, the shape latents can be directly sampled from it. This design enables effective and interpretable feature extraction via SVAE while enhancing latent expressiveness and maintaining runtime efficiency.

3.2 Symmetry Enforcement

The Quartet guarantees symmetry by explicitly learning its distribution. Once learned, the model can generate and enforce appropriate symmetries during point cloud generation. For a 3D shape \mathbf{p} , a finite group of rigid transformations $\mathcal{S} = \langle \mathcal{T} \rangle \subseteq \mathrm{E}(3)$, generated by $\mathcal{T} = \{S_1, S_2, \ldots, S_n\}$, is said to be a *symmetry group* over a subset $\mathbf{d} \subseteq \mathbf{p}$ if the full shape \mathbf{p} can be reconstructed by $\mathcal{S}\mathbf{d}$, the application of \mathcal{S} on \mathbf{d} , which is defined to be the union of all images under the transformations in \mathcal{S} :

$$\bigcup_{S \in S} S\mathbf{d} = \mathbf{p},\tag{8}$$

where the group generation is defined as $\langle \mathcal{T} \rangle := \{ \prod_{i=1}^n S_i \, | \, S_i \in \mathcal{T}, n \in \mathbb{N} \}$. The minimal such subset \mathbf{d} is called the *fundamental domain* for \mathcal{S} . Figure 3 illustrates fundamental domains for various parts in point cloud airplanes, cars, and chairs.

As translational symmetry implies an infinite shape, for tractability we restrict symmetry groups to those generated by at most two transformations – reflections or rotations with angles $\alpha \geq \frac{\pi}{18}$ such that $\frac{2\pi}{\alpha} \in \mathbb{Z}$. By a special case of the Cartan–Dieudonné theorem (Gallier & Gallier, 2011), any 3D rotation can be expressed as a composition of two reflections across planes intersecting along the axis of rotation, where the rotation angle is twice the angle between the two planes. Consequently, we just need to search for the symmetry groups generated by at most three reflections. Each reflection is represented using the Hesse normal form (Bôcher, 1915; Duda & Hart, 1972), enabling efficient and parallelizable symmetry search (Je et al., 2024). When multiple symmetries exist, we select the one corresponding to the fundamental domain with the smallest cardinality. In practice, if a symmetry group is generated by fewer than three reflections, we pad the remaining slots with a special symbol to maintain consistent data dimensionality.

Since each part is generated separately, symmetry is enforced at the part level. During part generation, we sample only the fundamental domain d and recover the full part by applying the learned symmetry group as defined in Eq. (8). This approach reduces the number of points to generate and, more importantly, guarantees symmetry in the resulting shapes.

To learn the distribution of symmetry groups present in the dataset, we first construct a metric space $(\mathcal{M}, d_{\mathcal{M}})$ of reflectional symmetries following Je *et al.* (Je et al., 2024). We then obtain the ground truth for symmetry groups \mathcal{S} using mean-shift clustering (Mitra et al., 2006; Fukunaga & Hostetler, 1975; Comaniciu & Meer, 2002), a nonparametric method based on gradient ascent over a density function in \mathcal{M} (see Appendix B.1 for more details). Finally, we leverage a diffusion model (Ho et al., 2020) to learn $p(\mathcal{S})$.

More specifically, we apply the diffusion process (Eq. (21) in Appendix B.3) to \mathcal{S} , and note that each transition kernel $q(\mathcal{S}_t \mid \mathcal{S}_{t-1}) \sim \mathcal{N}\left(\mu_t \mathcal{S}_{t-1}, \sigma_t I\right)$ is Gaussian. Therefore, intermediate samples \mathcal{S}_t can be expressed in closed form as

$$S_t = \left(\prod_{i=1}^t \mu_t\right) S_0 + \sqrt{\sum_{i=1}^t \sigma_i^2 \prod_{j=i+1}^t \mu_j^2} \epsilon, \tag{9}$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, I)$. Setting $\mu_t \coloneqq 1$ and $\gamma_t \coloneqq \sqrt{\sum_{i=1}^t \sigma_i^2}$ simplifies this to $\mathcal{S}_t = \mathcal{S}_0 + \gamma_t \epsilon$. Thus, the distribution of \mathcal{S}_t becomes a convolution:

$$p(\mathcal{S}_{t}) = (p_{0} * \varphi_{\mathbf{0},\gamma_{t}})(\mathcal{S}_{t}) = \int p_{0}(\mathcal{S}) \varphi_{\mathbf{0},\gamma_{t}}(\mathcal{S}_{t} - \mathcal{S}) d\mathcal{S} = \int p_{0}(\mathcal{S}) \varphi_{\mathcal{S},\gamma_{t}}(\mathcal{S}_{t}) d\mathcal{S}, \quad (10)$$

where p_0 is the distribution of S_0 , φ_{S,γ_t} is the probability density function for $\mathcal{N}\left(S,\gamma_t I\right)$, and * denotes convolution. With this, we train a diffusion model (Ho et al., 2020) $s_{\zeta}\left(S,t\right)$ parametrized by ζ to approximate the time-dependent score function (Song et al., 2020b) (*i.e.* the gradient of the log-density of noisy data) via an empirical approximation (Je et al., 2024):

$$s_{\zeta}(\mathcal{S}, t) \approx \nabla_{\mathcal{S}} \log p(\mathcal{S}_t) \approx \frac{1}{\gamma_t^2} \left(\frac{\sum_{R \in \mathcal{M}} \varphi_{R, \gamma_t}(\mathcal{S}) R}{\sum_{R \in \mathcal{M}} \varphi_{R, \gamma_t}(\mathcal{S})} - \mathcal{S} \right).$$
 (11)

Since symmetry can vary across shapes, in practice, we condition the diffusion model on the shape latent \mathbf{z} , modeling the symmetry group distribution as $p_{\mathcal{C}}(\mathcal{S} \mid \mathbf{z})$.

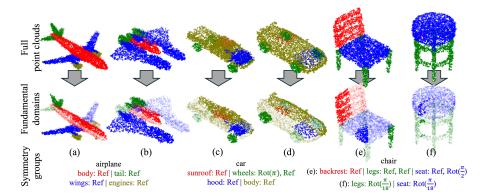


Figure 3: Point cloud airplanes, cars, and chairs with identified symmetry groups and corresponding fundamental domains for each color-coded part. Ref denotes reflection; $Rot(\alpha)$ denotes rotation by angle α . Our symmetry formulation allows greater flexibility by supporting symmetries composed of multiple transformations, such as two reflections (car wheels) or a reflection followed by a rotation (chair seat (e)). Circular symmetry is approximated via small-angle rotations (chair seat (f)).

After training, symmetries can be sampled from $s_{\zeta}(\mathcal{S},t)$ through annealed stochastic gradient Langevin dynamics (Song & Ermon, 2019): We initialize $\mathcal{S}_{\tau}^{(0)} \sim \mathcal{N}(\mathbf{0},I)$ and sequentially sample from noise-perturbed distributions $s_{\zeta}(\mathcal{S},t)$ for $t=\tau,\tau-1,\ldots,1$ using L Langevin steps:

$$S_t^{(0)} \leftarrow \begin{cases} \boldsymbol{\epsilon}_{\tau}^{(0)} & \text{if } t = \tau \\ S_{t+1}^{(L)} & \text{otherwise} \end{cases} ; \tag{12}$$

$$\mathcal{S}_{t}^{(i+1)} \leftarrow \mathcal{S}_{t}^{(i)} + \beta_{t} s_{\zeta} \left(\mathcal{S}_{t}^{(i)}, t \right) + \sqrt{2\beta_{t}} \boldsymbol{\epsilon}_{t}^{(i)}, \quad i = 0, 1, \dots, L,$$

$$(13)$$

where $\epsilon_t^{(i)} \sim \mathcal{N}(0, I)$, and β_t is the step size. The final sample $\mathcal{S}_1^{(L)}$ is the generated symmetry.

Notably, substituting Eq. (11) into Eq. (13) and setting $\beta_t := \gamma_t^2$ yields the following update rule:

$$\mathcal{S}_{t}^{(i+1)} \leftarrow \mathcal{S}_{t}^{(i)} + \frac{\beta_{t}}{\gamma_{t}^{2}} \left(\frac{\sum_{T \in \mathcal{M}} \varphi_{T,\gamma_{t}} \left(\mathcal{S}_{t}^{(i)} \right) T}{\sum_{T \in \mathcal{M}} \varphi_{T,\gamma_{t}} \left(\mathcal{S}_{t}^{(i)} \right)} - \mathcal{S}_{t}^{(i)} \right) + \sqrt{2\beta_{t}} \boldsymbol{\epsilon} \stackrel{\beta_{t} := \gamma_{t}^{2}}{=} \frac{\sum_{T \in \mathcal{M}} \varphi_{T,\gamma_{t}} \left(\mathcal{S}_{t}^{(i)} \right) T}{\sum_{T \in \mathcal{M}} \varphi_{T,\gamma_{t}} \left(\mathcal{S}_{t}^{(i)} \right)} + \sqrt{2\gamma_{t}} \boldsymbol{\epsilon}.$$

$$(14)$$

This update rule resembles mean-shift clustering (Eq. (17) in Appendix B.1), except it assumes an infinite neighborhood $B(\mathcal{S}_t) = \mathcal{M}$ and uses a Gaussian kernel with bandwidth γ_t : $K_t(\mathcal{S}) := \frac{1}{\sqrt{2\pi}\gamma_t}e^{-\frac{\|\mathcal{S}\|^2}{2}}$. Another difference is the injected noise $\sqrt{2}\gamma_t\epsilon$ in Eq. (14), which has been shown to improve robustness, sample quality, and mitigate mode collapse (e.g. Je et al., 2024).

Figure 3 shows examples of identified symmetry group generators. Most parts exhibit symmetry of a single reflection. In addition, we note that classical symmetries, such as reflectional and rotational, are special cases of our formulation, where the symmetry group S is generated by a single transformation. Our approach generalizes this by allowing symmetries composed of multiple transformations, such as sequential reflections (*e.g.*, car wheels) or a reflection followed by a rotation (*e.g.*, the chair seat in (e)). While our model does not explicitly capture circular symmetry, it can be approximated using discrete rotations with a minimum angle of $\frac{\pi}{18}$ (*e.g.*, chair seat in (f)).

3.3 PART GENERATION

Given the identified symmetry groups S_j and corresponding fundamental domains \mathbf{d}_j for each part \mathbf{p}_j , we model the part distribution via the distribution of fundamental domains. Since each full part \mathbf{p}_j can be reconstructed by applying S_j to \mathbf{d}_j (Eq. (8)), we have

$$p_{\xi}\left(\mathbf{p}_{i} \mid \mathcal{S}_{j}, \mathbf{z}\right) = p_{\xi}\left(\mathbf{d}_{j} \mid \mathcal{S}_{j}, \mathbf{z}\right). \tag{15}$$

To learn p_{ξ} ($\mathbf{d}_j \mid \mathcal{S}_j, \mathbf{z}$), we employ a transformer-based diffusion model (Mo et al., 2023; Peebles & Xie, 2023; Vaswani et al., 2017), conditioned on the shape latent \mathbf{z} and symmetry group \mathcal{S}_j . The model operates directly on voxelized point clouds, using 3D positional and patch embeddings to capture spatial context. Unlike U-Net architectures (Ho et al., 2020), which face limitations

Table 1: Quantitative comparison of point cloud generation. PA denotes part awareness; SA denotes symmetry awareness. Our Quartet is the only model that supports both, achieving significant improvements over most baselines and setting a new state of the art.

| Model | | | | Air | plane | | | C | ar | | Chair | | | | |
|--------------------------------|---|----|-----------|------|---------|------|-----------|------|---------------|------|-----------|------|---------------|------|--|
| | | SA | 1-NNA (↓) | | SDI (↓) | | 1-NNA (↓) | | SDI (\bigcup) | | 1-NNA (↓) | | SDI (\bigcup) | | |
| | | | CD | EMD | CD | EMD | CD | EMD | CD | EMD | CD | EMD | CD | EMD | |
| Training set | | | 64.4 | 64.1 | 0.954 | 4.90 | 51.3 | 54.8 | 7.49 | 1.18 | 51.7 | 50.0 | 5.56 | 1.68 | |
| PointFlow (Yang et al., 2019) | Х | Х | 75.7 | 70.7 | 3410 | 782 | 62.8 | 60.6 | 679 | 347 | 58.1 | 56.3 | 7290 | 530 | |
| ShapeGF (Cai et al., 2020) | X | X | 81.2 | 80.9 | 332 | 98.6 | 58.0 | 61.3 | 645 | 40.9 | 61.8 | 57.2 | 1100 | 101 | |
| DPF-Net (Klokov et al., 2020) | Х | X | 75.2 | 65.6 | 4256 | 245 | 62.0 | 58.5 | 827 | 452 | 62.4 | 54.5 | 5234 | 245 | |
| SetVAE (Kim et al., 2021) | X | X | 76.5 | 67.7 | 2830 | 824 | 58.8 | 60.6 | 1240 | 327 | 59.9 | 59.9 | 5320 | 673 | |
| DPC (Luo & Hu, 2021) | Х | X | 76.4 | 86.9 | 187 | 44.2 | 60.1 | 74.8 | 217 | 30.3 | 68.9 | 80.0 | 335 | 50.6 | |
| PVD (Zhou et al., 2021a) | X | X | 73.8 | 64.8 | 150 | 42.0 | 56.3 | 53.3 | 213 | 31.2 | 54.6 | 53.8 | 275 | 58.4 | |
| LION (Vahdat et al., 2022) | Х | X | 67.4 | 61.2 | 97.2 | 40.6 | 53.7 | 52.3 | 168 | 30.8 | 53.4 | 51.1 | 201 | 55.2 | |
| SPAGHETTI (Hertz et al., 2022) | 1 | X | 78.2 | 77.0 | 1530 | 529 | 72.3 | 71.0 | 581 | 284 | 70.7 | 69.0 | 5930 | 582 | |
| DiT-3D (Mo et al., 2023) | Х | X | 64.7 | 60.3 | 105 | 42.4 | 52.7 | 50.2 | 206 | 327 | 52.5 | 53.1 | 235 | 49.0 | |
| SALAD (Koo et al., 2023) | 1 | X | 73.9 | 71.1 | 198 | 45.1 | 59.2 | 57.2 | 236 | 29.4 | 57.8 | 58.4 | 308 | 52.6 | |
| FrePolad (Zhou et al., 2024) | Х | X | 65.3 | 62.1 | 94.1 | 38.1 | 52.4 | 53.2 | 173 | 29.6 | 51.9 | 50.3 | 252 | 50.9 | |
| Quartet (ours) | 1 | 1 | 63.3 | 59.7 | 25.7 | 1.87 | 50.1 | 51.8 | 25.7 | 2.28 | 51.6 | 53.7 | 28.9 | 2.86 | |

in scalability and spatial coherence, our transformer backbone incorporates 3D window attention, reducing computational cost while preserving local structure. This design yields higher efficiency, scalability, and fidelity in point cloud generation (Mo et al., 2023).

3.4 PART ASSEMBLY

To assemble the full point cloud, we learn assemblers T_j that correctly transform each part \mathbf{p}_j . We first train a part encoder $q_{\phi}\left(\mathbf{w}_j\mid\mathbf{p}_j,\mathbf{z}\right)$ to obtain compact representations \mathbf{w}_j , then train a diffusion model $p_{\psi}\left(T_j\mid\mathbf{w}_j,\mathbf{p}_j,\mathcal{S}_j,\mathbf{z}\right)$ to model the assembler distribution $p_{\psi}\left(T_j\mid\mathbf{w}_j,\mathbf{p}_j,\mathcal{S}_j,\mathbf{z}\right)$. To simplify learning, we assume that \mathbf{w}_j captures all necessary information about \mathbf{p}_j and \mathcal{S}_j , making T_j conditionally independent of them given \mathbf{w}_j . Consequently, the diffusion model is conditioned only on \mathbf{w}_j and the global shape latent \mathbf{z} : $p_{\psi}\left(T_j\mid\mathbf{w}_j,\mathbf{z}\right)\approx p_{\psi}\left(T_j\mid\mathbf{w}_j,\mathbf{p}_j,\mathcal{S}_j,\mathbf{z}\right)$. This assumption decouples geometric modeling from transformation learning, reducing complexity and allowing the model to focus on learning flexible, generalizable spatial configurations.

Our assembler diffusion model (Tang et al., 2024) employs a U-Net backbone (Ronneberger et al., 2015) augmented with skip connections and cross-attention layers (Vaswani et al., 2017). The cross-attention modules integrate contextual information from both part latents and the global shape latent, enabling the model to generate transformations that are coherent at both local and global levels. This architecture effectively handles diverse part configurations and generalizes well to novel shapes. The hierarchical structure of U-Net further supports multi-scale spatial reasoning, which is essential for accurate 3D part placement.

For the part encoder, we adopt the SVAE architecture (see Sec. 3.1) with *equivariance fine-tuning* (*EFT*): During training, random rigid transformations, including translations, rotations, and reflections, are applied to each part, and the encoder is encouraged to produce the latent transformed accordingly. This promotes geometry-aware but pose-invariant part embeddings, enhancing their suitability for the part assembly task. As a result, the diffusion model receives more stable and semantically meaningful latents, leading to more accurate and coherent assembler predictions.

4 EXPERIMENT

4.1 DATASET

We evaluate our method on the ShapeNetPart dataset (Yi et al., 2016), a subset of ShapeNet (Chang et al., 2015) with semantic part annotations. We focus on three representative categories: airplanes (body, tail, wings, engines), cars (sunroof, wheels, hood, body), and chairs (backrest and armrests, legs, seat). Segmentation examples are shown in Fig. 3. We use the official train/validation/test split provided with the dataset.

Following common practice, we use point clouds with 2048 points. Each part is resized to match the mean number of points per part within its category. During preprocessing, parts with more points

Table 2: Per-part SDI-CD (\downarrow) for point cloud generation. With explicit symmetry enforcement, the Quartet achieves significantly lower SDI scores, indicating stronger symmetry in generated parts.

| Model | | Aiı | plane | | | Са | ır | Chair | | | |
|--------------------------------|-------|-------|-------|---------|------|--------|------|-------|------|------|------|
| 1.10401 | body | tail | wings | engines | roof | wheels | hood | body | back | legs | seat |
| Training set | 0.877 | 0.700 | 0.883 | 0.717 | 5.25 | 3.06 | 3.15 | 5.27 | 4.11 | 2.65 | 3.04 |
| SPAGHETTI (Hertz et al., 2022) | 318 | 268 | 417 | 125 | 258 | 339 | 127 | 417 | 2857 | 2948 | 1358 |
| SALAD (Koo et al., 2023) | 39.1 | 26.2 | 30.5 | 18.5 | 98.2 | 68.5 | 152 | 256 | 163 | 93.5 | 104 |
| Quartet (ours) | 7.86 | 4.08 | 7.76 | 5.84 | 9.72 | 4.10 | 5.76 | 9.90 | 10.3 | 6.19 | 10.4 |

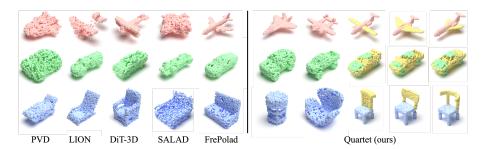


Figure 4: Point cloud generation. Samples from the Quartet are visually appealing, diverse, and exhibit strong structural consistency. The last three columns illustrate targeted manipulation.

are randomly subsampled, while those with fewer are upsampled via random duplication. This normalization ensures consistent input dimensions across training batches, enhancing the stability and efficiency of both the encoder and generative models.

4.2 EVALUATION METRICS

Following prior work (e.g. Zhou et al., 2024), we assess the quality and diversity of generated point clouds using 1-nearest neighbor (1-NNA) (Lopez-Paz & Oquab, 2016), a retrieval-based metric computed with Chamfer Distance (CD) and Earth Mover's Distance (EMD). 1-NNA measures how often a shape's nearest neighbor belongs to the same distribution (Yang et al., 2019). A balanced score near 50% indicates close alignment between generated and real shape distributions.

To assess whether a generated 3D shape is symmetric, we introduce the *symmetry discrepancy index* (SDI), which quantifies how well a 3D shape $\bf p$ conforms to a given symmetry group $\cal S$: For a normalized shape $\bf p$, SDI is defined as the distance d—either CD or EMD—between $\bf p$ and the shape reconstructed from its fundamental domain $\bf d$ under $\cal S$: $L_{\rm SDI}(\bf p):=d(\bf p,\cal Sd)$. Lower SDI values indicate stronger symmetry. For methods without explicit symmetry modeling, we compute SDI using the simplest yet most common symmetry: reflection across the vertical bisector plane. For readability, we report SDI-CD scaled by 10 and SDI-EMD scaled by 10^3 .

4.3 POINT CLOUD GENERATION

We benchmark the Quartet against several competitive 3D generative models. Implementation and training details are provided in Appendix C, and additional experimental results can be found in Appendix D. Quantitative results based on 1-NNA and SDI are reported in Tab. 1. Notably, the Quartet is the only method that explicitly models both part structure and symmetry. Per-part SDI scores for part-aware models are presented in Tab. 2. The Quartet consistently outperforms state-of-the-art methods, achieving high fidelity and diversity while closely matching the real shape distribution. Remarkably, due to explicit symmetry enforcement, the Quartet achieves significantly lower SDI scores, indicating strong alignment with ideal symmetric structures. Qualitative results in Fig. 4 show that the Quartet generates visually coherent and diverse point clouds obeying symmetry across all three object categories, demonstrating its effectiveness.

Targeted manipulation The Quartet models the distribution of each part separately, enabling targeted part-level manipulation while guaranteeing disentanglement. In Fig. 1 and in the last three columns of Fig. 4, we present point cloud samples in which the yellow-highlighted parts vary while the remaining structure is held fixed.

Table 3: Ablation study. L denotes latent diffusion, S symmetry diffusion, and P part and assembler diffusions. All four members in the Quartet are essential; removing any degrades performance.

| | | | | Airplane | | | | | C | ar | | Chair | | | | |
|-------------|---|---|---|------------|---------------------|------|------|------|------|------|-------------------|-------|---------------------|------|-------------------|--|
| Variation | L | S | P | 1-NI CD | 1-NNA (↓) CD EMD | | | | (4) | | SDI (↓) CD EMD | | 1-NNA (↓) CD EMD | | SDI (↓) CD EMD | |
| Solo | Х | Х | Х | 69.2 | 64.2 | 154 | 44.2 | 55.5 | 53.6 | 241 | 351 | 53.2 | 53.8 | 295 | 52.8 | |
| Duet Var. 1 | Х | Х | 1 | 76.3 | 82.1 | 927 | 614 | 70.2 | 72.9 | 522 | 460 | 68.3 | 67.2 | 3261 | 615 | |
| Duet Var. 2 | 1 | Х | Х | 66.8 | 63.3 | 103 | 42.4 | 52.3 | 52.7 | 192 | 28.1 | 53.2 | 51.6 | 184 | 49.2 | |
| Trio Var. 1 | Х | 1 | 1 | 85.1 | 85.8 | 3516 | 1450 | 83.9 | 83.6 | 836 | 326 | 92.5 | 85.2 | 6286 | 562 | |
| Trio Var. 2 | 1 | Х | 1 | 63.1 | 63.6 | 95 | 39.2 | 49.8 | 52.4 | 205 | 32.3 | 52.3 | 53.9 | 174 | 52.5 | |
| NoSVAE | 1 | 1 | 1 | 64.3 | 61.8 | 25.2 | 2.51 | 51.6 | 52.0 | 25.9 | 2.55 | 52.6 | 53.9 | 29.3 | 3.15 | |
| NoEFT | 1 | 1 | 1 | 63.9 | 62.4 | 32.7 | 9.27 | 51.8 | 52.0 | 27.2 | 5.98 | 52.5 | 54.0 | 30.2 | 4.15 | |
| Quartet | 1 | 1 | ✓ | 63.3 | 59.7 | 25.7 | 1.87 | 50.1 | 51.8 | 25.7 | 2.28 | 51.6 | 53.7 | 28.9 | 2.86 | |

4.4 ABLATION STUDY

Our Quartet consists of four diffusion models responsible for shape latents, symmetries, parts, and assemblers. As an ablation study, in Tab. 3, we evaluate several simplified variants with some diffusions or key techniques removed:

- NoSVAE: replaces SVAEs (Sec. 3.1) with standard VAEs for full point cloud and part encoding.
- NoEFT: removes the equivariance fine-tuning (Sec. 3.4) applied to the part SVAE.
- Trio Var. 1: replaces shape latent diffusion with a simple Gaussian prior;
- Trio Var. 2: removes the symmetry diffusion with part diffusion generating the full part directly;
- Duet Var. 1: removes both shape latent and symmetry diffusions;
- Duet Var. 2: removes symmetry diffusion and merges part and assembler diffusions into a single full-shape diffusion;
- Solo: uses a single diffusion generating full point clouds.

Our results show that all four diffusion models are indispensable members of the Quartet, each contributing to the generation of high-quality, diverse point clouds with strong symmetry. Two key observations emerge: First, comparing Solo and Duet Var. 1, we see that generating parts without a central shape latent significantly degrades performance. This echoes previous findings that a simple Gaussian prior is insufficient for capturing the structural complexity of 3D shapes (Vahdat et al., 2021; Tomczak & Welling, 2018; Rosca et al., 2018; Zhou et al., 2024; Vahdat et al., 2022). The performance drops further in Trio Var. 1, where symmetry diffusion is added, likely compounding the mismatch between prior and part representations. Second, comparing Trio Var. 2 and the full Quartet, we observe that while symmetry enforcement substantially improves SDI, it does not consistently improve 1-NNA – even with latent diffusion present. This may be due to the added constraint or increased learning complexity introduced by enforcing symmetry.

4.5 RUNTIME ANALYSIS

Training the Quartet on each object category takes approximately 50 hours on a single GPU. Figure 1 presents the generation performance vs. training time for models trained on the airplane category. Some baselines require longer training due to joint optimization (Yang et al., 2019), operation on full point clouds (Zhou et al., 2021a; Luo & Hu, 2021), or the use of complex, high-dimensional latent spaces (Vahdat et al., 2022). In contrast, despite using four diffusions, the Quartet trains faster than most baselines. This is because the shape latent, symmetry, and assembler diffusions operate on low-dimensional representations (124-, 12M-, and 9M-dimensional, respectively, where M is the number of parts). The majority of training time is spent on part diffusion, which is accelerated using an efficient transformer-based diffusion (Mo et al., 2023) applied to fundamental domains that are typically only half or a quarter the size of the full parts.

5 CONCLUSION AND DISCUSSION

We presented the Quartet of Diffusions, a structure-aware framework for 3D point cloud generation that explicitly models part composition and symmetry. At its current stage, our work focuses on unconditional generation, but it highlights the importance of integrating symmetry- and part-based reasoning into structure-aware models. The model's ability to disentangle and manipulate individual parts makes it well-suited for interactive shape editing and user-guided design applications.

REFERENCES

- Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, pp. 40–49. PMLR, 2018.
 - Atallah. On symmetry detection. *IEEE Transactions on Computers*, 100(7):663–666, 1985.
 - David Berthelot, Thomas Schumm, and Luke Metz. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017.
- Maxime Bôcher. Plane analytic geometry: with introductory chapters on the differential calculus.
 H. Holt, 1915.
 - Martin Bokeloh, Michael Wand, and Hans-Peter Seidel. A connection between partial symmetry and inverse procedural modeling. In *ACM SIGGRAPH 2010 papers*, pp. 1–10. ACM New York, NY, USA, 2010.
 - Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. https://transformer-circuits.pub/2023/monosemantic-features/index.html.
 - Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18392–18402, 2023.
 - Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. Nonlocal image and movie denoising. *International journal of computer vision*, 76:123–139, 2008.
 - Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snavely, and Bharath Hariharan. Learning gradient fields for shape generation. In *Computer Vision–ECCV* 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16, pp. 364–381. Springer, 2020.
 - Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. A survey on generative diffusion models. *IEEE Transactions on Knowledge and Data Engineering*, 2024.
 - 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.
 - Siddhartha Chaudhuri, Daniel Ritchie, Jiajun Wu, Kai Xu, and Hao Zhang. Learning generative models of 3d structures. *Computer Graphics Forum*, 39(2):643–666, 2020.
 - Minghao Chen, Roman Shapovalov, Iro Laina, Tom Monnier, Jianyuan Wang, David Novotny, and Andrea Vedaldi. Partgen: Part-level 3d generation and reconstruction with multi-view diffusion models. *arXiv preprint arXiv:2412.18608*, 2024.
 - Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020.
 - Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
 - Daniel Cohen-Or and Hao Zhang. From inspired modeling to creative modeling. *The Visual Computer*, 32:7–14, 2016.

- Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002.
- Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9): 10850–10869, 2023.
 - Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pp. 146–158, 1975.
 - Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*, 2023.
 - Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
 - Anastasia Dubrovina, Fei Xia, Panos Achlioptas, Mira Shalah, Raphaël Groscot, and Leonidas J Guibas. Composite shape modeling via latent space factorization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8140–8149, 2019.
 - Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972.
 - Vassiliy A Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory of Probability & Its Applications*, 14(1):153–158, 1969.
 - Noa Fish, Melinos Averkiou, Oliver Van Kaick, Olga Sorkine-Hornung, Daniel Cohen-Or, and Niloy J Mitra. Meta-representation of shape families. *ACM Transactions on Graphics (TOG)*, 33(4):1–11, 2014.
 - Keinosuke Fukunaga and Larry Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on information theory*, 21(1):32–40, 1975.
 - Thomas Funkhouser, Michael Kazhdan, Philip Shilane, Patrick Min, William Kiefer, Ayellet Tal, Szymon Rusinkiewicz, and David Dobkin. Modeling by example. *ACM transactions on graphics (TOG)*, 23(3):652–663, 2004.
 - Ran Gal, Olga Sorkine, Niloy J Mitra, and Daniel Cohen-Or. iwires: An analyze-and-edit approach to shape manipulation. In *ACM SIGGRAPH 2009 papers*, pp. 1–10. ACM New York, NY, USA, 2009.
 - Jean Gallier and Jean Gallier. The cartan–dieudonné theorem. *Geometric Methods and Applications:* For Computer Science and Engineering, pp. 231–280, 2011.
 - Lin Gao, Jie Yang, Tong Wu, Yu-Jie Yuan, Hongbo Fu, Yu-Kun Lai, and Hao Zhang. Sdm-net: Deep generative network for structured deformable mesh. *ACM Transactions on Graphics (TOG)*, 38 (6):1–15, 2019.
 - Lin Gao, Ling-Xiao Zhang, Hsien-Yu Meng, Yi-Hui Ren, Yu-Kun Lai, and Leif Kobbelt. Prs-net: Planar reflective symmetry detection net for 3d models. *IEEE transactions on visualization and computer graphics*, 27(6):3007–3018, 2020.
 - Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 7154–7164, 2019.
 - Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4857–4866, 2020.
 - Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

- 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.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- Donald D Hoffman and Whitman A Richards. Parts of recognition. Cognition, 18(1-3):65–96, 1984.
- Haibin Huang, Evangelos Kalogerakis, and Benjamin Marlin. Analysis and synthesis of 3d shape families via deep-learned generative models of surfaces. *Computer Graphics Forum*, 34(5):25–38, 2015.
- Qixing Huang, Fan Wang, and Leonidas Guibas. Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG)*, 33(4):1–11, 2014.
- Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural template: Topology-aware reconstruction and disentangled generation of 3d meshes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18572–18582, 2022.
- Aapo Hyvärinen, Jarmo Hurri, Patrik O Hoyer, Aapo Hyvärinen, Jarmo Hurri, and Patrik O Hoyer. Estimation of non-normalized statistical models. *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*, pp. 419–426, 2009.
- John Illingworth and Josef Kittler. A survey of the hough transform. *Computer vision, graphics, and image processing*, 44(1):87–116, 1988.
- Jihyeon Je, Jiayi Liu, Guandao Yang, Boyang Deng, Shengqu Cai, Gordon Wetzstein, Or Litany, and Leonidas Guibas. Robust symmetry detection via riemannian langevin dynamics. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, 2024.
- Myeonghun Jeong, Hyeongju Kim, Sung Jun Cheon, Byoung Jin Choi, and Nam Soo Kim. Diff-tts: A denoising diffusion model for text-to-speech. *arXiv preprint arXiv:2104.01409*, 2021.
- Penglei Ji and Xinguo Liu. A fast and efficient 3d reflection symmetry detector based on neural networks. *Multimedia Tools and Applications*, 78(24):35471–35492, 2019.
- Evangelos Kalogerakis, Siddhartha Chaudhuri, Daphne Koller, and Vladlen Koltun. A probabilistic model for component-based shape synthesis. *Acm Transactions on Graphics (TOG)*, 31(4):1–11, 2012.
- William Karush. Minima of functions of several variables with inequalities as side constraints. M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago, 1939.
- Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Symmetry descriptors and 3d shape matching. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pp. 115–123, 2004.
- Hyeongju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow: Probabilistic framework for normalizing flow on manifolds. *Advances in Neural Information Processing Systems*, 33:16388–16397, 2020.
- Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical composition for generative modeling of set-structured data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15059–15068, 2021.
- Vladimir G Kim, Wilmot Li, Niloy J Mitra, Siddhartha Chaudhuri, Stephen DiVerdi, and Thomas Funkhouser. Learning part-based templates from large collections of 3d shapes. *ACM Transactions on Graphics (TOG)*, 32(4):1–12, 2013.

- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint* arXiv:1312.6114, 2013.
 - Roman Klokov, Edmond Boyer, and Jakob Verbeek. Discrete point flow networks for efficient point cloud generation. In *European Conference on Computer Vision*, pp. 694–710. Springer, 2020.
 - Juil Koo, Seungwoo Yoo, Minh Hieu Nguyen, and Minhyuk Sung. Salad: Part-level latent diffusion for 3d shape generation and manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14441–14451, 2023.
 - H Kuhn and A Tucker. Nonlinear programming. In *Proceedings of 2nd Berkeley symposium (pp. 481–492)*, 1951.
 - Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and Fujie Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
 - Guo Li, Ligang Liu, Hanlin Zheng, and Niloy J Mitra. Analysis, reconstruction and manipulation using arterial snakes. *ACM Trans. Graph.*, 29(6):152, 2010.
 - Jun Li, Kai Xu, Siddhartha Chaudhuri, Ersin Yumer, Hao Zhang, and Leonidas Guibas. Grass: Generative recursive autoencoders for shape structures. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017.
 - Jun Li, Chengjie Niu, and Kai Xu. Learning part generation and assembly for structure-aware shape synthesis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 11362–11369, 2020.
 - Shidi Li, Miaomiao Liu, and Christian Walder. Editvae: Unsupervised parts-aware controllable 3d point cloud shape generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 1386–1394, 2022a.
 - Shidi Li, Christian Walder, and Miaomiao Liu. Spa-vae: Similar-parts-assignment for unsupervised 3d point cloud generation. *arXiv preprint arXiv:2203.07825*, 2022b.
 - Songlin Li, Despoina Paschalidou, and Leonidas Guibas. Pasta: Controllable part-aware shape generation with autoregressive transformers. *arXiv preprint arXiv:2407.13677*, 2024.
 - Songxiang Liu, Dan Su, and Dong Yu. Diffgan-tts: High-fidelity and efficient text-to-speech with denoising diffusion gans. *arXiv* preprint arXiv:2201.11972, 2022.
 - Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. *Advances in Neural Information Processing Systems*, 32, 2019.
 - David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. *arXiv preprint arXiv:1610.06545*, 2016.
 - 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*, pp. 2837–2845, 2021.
 - Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. *arXiv* preprint arXiv:2303.08320, 2023.
 - Aleksandar Makelov, George Lange, and Neel Nanda. Towards principled evaluations of sparse autoencoders for interpretability and control. *arXiv preprint arXiv:2405.08366*, 2024.
 - Luke Marks, Alasdair Paren, David Krueger, and Fazl Barez. Enhancing neural network interpretability with feature-aligned sparse autoencoders. *arXiv preprint arXiv:2411.01220*, 2024.
 - Aurélien Martinet, Cyril Soler, Nicolas Holzschuch, and François X Sillion. Accurate detection of symmetries in 3d shapes. *ACM Transactions on Graphics (TOG)*, 25(2):439–464, 2006.

- Ravish Mehra, Qingnan Zhou, Jeremy Long, Alla Sheffer, Amy Gooch, and Niloy J Mitra. Abstraction of man-made shapes. In *ACM SIGGRAPH Asia 2009 papers*, pp. 1–10. ACM New York, NY, USA, 2009.
 - Niloy J Mitra, Leonidas J Guibas, and Mark Pauly. Partial and approximate symmetry detection for 3d geometry. *ACM Transactions on Graphics (TOG)*, 25(3):560–568, 2006.
 - Niloy J Mitra, Leonidas J Guibas, and Mark Pauly. Symmetrization. *ACM Transactions on Graphics* (*TOG*), 26(3):63–es, 2007.
 - Niloy J Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. Symmetry in 3d geometry: Extraction and applications. *Computer graphics forum*, 32(6):1–23, 2013a.
 - Niloy J Mitra, Michael Wand, Hao Zhang, Daniel Cohen-Or, and Martin Bokeloh. Structure-aware shape processing. In *Computer Graphics Forum (Eurographics State-of-the-art Report)*, pp. 175–197. Wiley and Eurographics, 2013b.
 - Niloy J Mitra, Michael Wand, Hao Zhang, Daniel Cohen-Or, Vladimir Kim, and Qi-Xing Huang. Structure-aware shape processing. In *ACM SIGGRAPH 2014 Courses*, pp. 1–21. Wiley and Eurographics, 2014.
 - Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas J Guibas. Structurenet: Hierarchical graph networks for 3d shape generation. *arXiv preprint arXiv:1908.00575*, 2019.
 - Shentong Mo, Enze Xie, Ruihang Chu, Lanqing Hong, Matthias Niessner, and Zhenguo Li. Dit-3d: Exploring plain diffusion transformers for 3d shape generation. *Advances in neural information processing systems*, 36:67960–67971, 2023.
 - George Kiyohiro Nakayama, Mikaela Angelina Uy, Jiahui Huang, Shi-Min Hu, Ke Li, and Leonidas Guibas. Difffacto: Controllable part-based 3d point cloud generation with cross diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14257–14267, 2023.
 - Charlie Nash and Christopher KI Williams. The shape variational autoencoder: A deep generative model of part-segmented 3d objects. *Computer Graphics Forum*, 36(5):1–12, 2017.
 - Andrew Ng et al. Sparse autoencoder. CS294A Lecture notes, 72(2011):1–19, 2011.
 - Charles O'Neill, Christine Ye, Kartheik Iyer, and John F Wu. Disentangling dense embeddings with sparse autoencoders. *arXiv preprint arXiv:2408.00657*, 2024.
 - Daniele Panozzo, Yaron Lipman, Enrico Puppo, and Denis Zorin. Fields on symmetric surfaces. *ACM Transactions on Graphics (TOG)*, 31(4):1–12, 2012.
 - Mark Pauly, Niloy J Mitra, Joachim Giesen, Markus H Gross, and Leonidas J Guibas. Example-based 3d scan completion. In *Symposium on geometry processing*, pp. 23–32, 2005.
 - Mark Pauly, Niloy J Mitra, Johannes Wallner, Helmut Pottmann, and Leonidas J Guibas. Discovering structural regularity in 3d geometry. In *ACM SIGGRAPH 2008 papers*, pp. 1–11. ACM New York, NY, USA, 2008.
 - William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
 - Cody J Phillips, Matthieu Lecce, and Kostas Daniilidis. Seeing glassware: from edge detection to pose estimation and shape recovery. In *Robotics: Science and Systems*, volume 3, pp. 3. Michigan, USA, 2016.
 - Joshua Podolak, Philip Shilane, Aleksey Golovinskiy, Szymon Rusinkiewicz, and Thomas Funkhouser. A planar-reflective symmetry transform for 3d shapes. In *ACM SIGGRAPH 2006 Papers*, pp. 549–559. ACM New York, NY, USA, 2006.

- Joshua Podolak, Aleksey Golovinskiy, and Szymon Rusinkiewicz. Symmetry-enhanced remeshing of surfaces. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pp. 235–242, 2007.
 - Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
 - Daniel Reisfeld, Haim Wolfson, and Yehezkel Yeshurun. Context-free attentional operators: The generalized symmetry transform. *International Journal of Computer Vision*, 14:119–130, 1995.
 - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.
 - Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pp. 234–241. Springer, 2015.
 - Mihaela Rosca, Balaji Lakshminarayanan, and Shakir Mohamed. Distribution matching in variational inference. *arXiv preprint arXiv:1802.06847*, 2018.
 - Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4713–4726, 2022.
 - Nadav Schor, Oren Katzir, Hao Zhang, and Daniel Cohen-Or. Componet: Learning to generate the unseen by part synthesis and composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8759–8768, 2019.
 - Zeyun Shi, Pierre Alliez, Mathieu Desbrun, Hujun Bao, and Jin Huang. Symmetry and orbit detection via lie-algebra voting. *Computer graphics forum*, 35(5):217–227, 2016.
 - Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learn-ing*, pp. 2256–2265. PMLR, 2015.
 - Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv* preprint arXiv:2010.02502, 2020a.
 - Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
 - Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv* preprint *arXiv*:2011.13456, 2020b.
 - Kevin Swersky, Marc'Aurelio Ranzato, David Buchman, Nando D Freitas, and Benjamin M Marlin. On autoencoders and score matching for energy based models. In *Proceedings of the 28th international conference on machine learning*, pp. 1201–1208, 2011.
 - Nicolas Talabot, Olivier Clerc, Arda Cinar Demirtas, Doruk Oner, and Pascal Fua. Partsdf: Partbased implicit neural representation for composite 3d shape parametrization and optimization. *arXiv preprint arXiv:2502.12985*, 2025.
 - Jiapeng Tang, Yinyu Nie, Lev Markhasin, Angela Dai, Justus Thies, and Matthias Nießner. Diffuscene: Denoising diffusion models for generative indoor scene synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 20507–20518, 2024.
 - Sebastian Thrun and Ben Wegbreit. Shape from symmetry. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pp. 1824–1831. IEEE, 2005.
 - Jakub Tomczak and Max Welling. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pp. 1214–1223. PMLR, 2018.

- Shubham Tulsiani, Or Litany, Charles R Qi, He Wang, and Leonidas J Guibas. Object-centric multiview aggregation. *arXiv preprint arXiv:2007.10300*, 2020.
 - Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems*, 34:11287–11302, 2021.
 - Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, Karsten Kreis, et al. Lion: Latent point diffusion models for 3d shape generation. *Advances in Neural Information Processing Systems*, 35:10021–10039, 2022.
 - Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
 - Hao Wang, Nadav Schor, Ruizhen Hu, Haibin Huang, Daniel Cohen-Or, and Hui Huang. Global-to-local generative model for 3d shapes. *ACM Transactions on Graphics (TOG)*, 37(6):1–10, 2018.
 - Lei Wang, Yuchun Huang, Pengjie Tao, Yaolin Hou, and Yuxuan Liu. Learning geometry-image representation for 3d point cloud generation. *arXiv preprint arXiv:2011.14289*, 2020.
 - Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5): 1–12, 2019.
 - Qiuhong Anna Wei, Sijie Ding, Jeong Joon Park, Rahul Sajnani, Adrien Poulenard, Srinath Sridhar, and Leonidas Guibas. Lego-net: Learning regular rearrangements of objects in rooms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19037–19047, 2023.
 - Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning*, pp. 681–688. Citeseer, 2011.
 - Jan D Wolter, Tony C Woo, and Richard A Volz. Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer*, 1:37–48, 1985.
 - Lemeng Wu, Dilin Wang, Chengyue Gong, Xingchao Liu, Yunyang Xiong, Rakesh Ranjan, Raghuraman Krishnamoorthi, Vikas Chandra, and Qiang Liu. Fast point cloud generation with straight flows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9445–9454, 2023.
 - Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen. Pq-net: A generative part seq2seq network for 3d shapes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 829–838, 2020a.
 - Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1–10, 2020b.
 - Zhijie Wu, Xiang Wang, Di Lin, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Sagnet: Structure-aware generative network for 3d-shape modeling. *ACM Transactions on Graphics* (*TOG*), 38(4):1–14, 2019.
 - Zhen Xing, Qijun Feng, Haoran Chen, Qi Dai, Han Hu, Hang Xu, Zuxuan Wu, and Yu-Gang Jiang. A survey on video diffusion models. *ACM Computing Surveys*, 57(2):1–42, 2024.
 - Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024.
 - Kai Xu, Hao Zhang, Andrea Tagliasacchi, Ligang Liu, Guo Li, Min Meng, and Yueshan Xiong. Partial intrinsic reflectional symmetry of 3d shapes. In *ACM SIGGRAPH Asia 2009 papers*, pp. 1–10. ACM New York, NY, USA, 2009.

- Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4541–4550, 2019.
- Jing Yang, Kyle Fogarty, Fangcheng Zhong, and Cengiz Oztireli. Sym3d: Learning symmetric triplanes for better 3d-awareness of gans. *arXiv preprint arXiv:2406.06432*, 2024.
- Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *Entropy*, 25(10):1469, 2023.
- Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016.
- Mourad Zerroug and Ramakant Nevatia. Part-based 3d descriptions of complex objects from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(9):835–848, 1999.
- Guangyao Zhai, Evin Pınar Örnek, Shun-Cheng Wu, Yan Di, Federico Tombari, Nassir Navab, and Benjamin Busam. Commonscenes: Generating commonsense 3d indoor scenes with scene graph diffusion. *Advances in Neural Information Processing Systems*, 36:30026–30038, 2023.
- Qian Zheng, Andrei Sharf, Guowei Wan, Yangyan Li, Niloy J Mitra, Daniel Cohen-Or, and Baoquan Chen. Non-local scan consolidation for 3d urban scenes. *ACM Trans. Graph.*, 29(4):94–1, 2010.
- Chenliang Zhou, Fangcheng Zhong, and Cengiz Öztireli. Clip-pae: Projection-augmentation embedding to extract relevant features for a disentangled, interpretable and controllable text-guided face manipulation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–9, 2023.
- Chenliang Zhou, Fangcheng Zhong, Param Hanji, Zhilin Guo, Kyle Fogarty, Alejandro Sztrajman, Hongyun Gao, and Cengiz Oztireli. Frepolad: Frequency-rectified point latent diffusion for point cloud generation. In *European Conference on Computer Vision*, pp. 434–453. Springer, 2024.
- 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*, pp. 5826–5835, 2021a.
- Yichao Zhou, Shichen Liu, and Yi Ma. Nerd: Neural 3d reflection symmetry detector. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 15940–15949, 2021b.
- Chuhang Zou, Ersin Yumer, Jimei Yang, Duygu Ceylan, and Derek Hoiem. 3d-prnn: Generating shape primitives with recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 900–909, 2017.

Quartet of Diffusions: Structure-Aware Point Cloud Generation through Part and Symmetry Guidance

Appendix

A RELATED WORK ON SYMMETRY DETECTION

Symmetry is a fundamental geometric property commonly observed in both natural and humanmade objects. As a result, detecting and leveraging symmetry has long been a central problem in the fields of computer vision and computer graphics (Mitra et al., 2013a). Early work in this area focused on detecting exact symmetries in 2D or 3D planar point sets (Atallah, 1985; Wolter et al., 1985). However, the reliance on exact symmetry limits the practicality of these methods, as realworld objects often exhibit only approximate symmetry due to noise, occlusion, or design variations.

Traditional voting-based approaches, such as those based on the Hough transform (Illingworth & Kittler, 1988), attempt to accumulate votes for potential symmetries from point correspondences. While effective in idealized settings, these methods are known to be sensitive to noise and often produce unreliable results when applied to imperfect or incomplete data.

To address these limitations, Mitra *et al.* (Mitra et al., 2006) proposed a more robust technique that replaces voting with mean-shift clustering (Fukunaga & Hostetler, 1975; Comaniciu & Meer, 2002) in a transformation space (for more details see Appendix B.1). This method searches for modes in the space of rigid transformations, making it more resilient to noise and better suited for partial symmetries. Their approach laid the groundwork for several subsequent improvements (Pauly et al., 2008; Shi et al., 2016), which further enhanced the robustness and generality of symmetry detection.

In recent years, learning-based methods have also been introduced (Ji & Liu, 2019; Zhou et al., 2021b; Gao et al., 2020), leveraging neural networks to detect symmetry directly from 3D data. These methods benefit from data-driven representations and have shown improved generalization across object categories and varying conditions.

Building on these trends, Je *et al.* (Je et al., 2024) proposed a hybrid approach that combines elements of both traditional and learning-based techniques. Their method redefines the symmetry space and applies Langevin dynamics—a sampling technique from generative modeling—to iteratively refine symmetry estimates. This formulation provides both robustness and efficiency and serves as a strong foundation for the symmetry detection component in our framework.

B BACKGROUND

B.1 SYMMETRY DETECTION BY MEAN-SHIFT CLUSTERING

A common approach to detecting symmetry in 3D objects involves using the Hough transform (Illingworth & Kittler, 1988) to vote on the parameters of potential symmetry planes (Mitra et al., 2013a; 2006). In this framework, each pair of points in the shape casts a vote for a candidate symmetry in a predefined transformation space \mathcal{M} . For example, in the case of reflectional symmetry, a point pair \mathbf{a} and \mathbf{b} casts a vote for the plane that passes through their midpoint $\frac{\mathbf{a}+\mathbf{b}}{2}$ and has a normal vector given by $\frac{\mathbf{a}-\mathbf{b}}{\|\mathbf{a}-\mathbf{b}\|}$ – the direction that would reflect one point onto the other.

In an ideal setting, where the object and its symmetries are exact, votes from all symmetric point pairs would concentrate on discrete points in \mathcal{M} . In such cases, detecting symmetry reduces to identifying the peak with the highest vote count. However, most real-world objects and 3D datasets only exhibit approximate symmetry. As a result, the votes form a smooth, continuous distribution in \mathcal{M} rather than sharp peaks. Detecting symmetry in this context thus requires identifying clusters in \mathcal{M} that correspond to the most prominent approximate symmetries.

Mitra et al. (Mitra et al., 2006) proposes to use the mean-shift clustering (Fukunaga & Hostetler, 1975; Comaniciu & Meer, 2002), a nonparametric clustering method based on gradient ascent on a

density function p(S) in \mathcal{M} defined as

$$p(S) := \sum_{R \in \mathcal{M}} K\left(\frac{S - R}{h}\right),\tag{16}$$

where K is a kernel function (e.g., Gaussian or Epanechnikov kernel (Epanechnikov, 1969)) with bandwidth h. The significant modes of p, and hence the significant symmetries, can be determined using gradient ascent: The algorithm first initializes from any of the candidate transformation $S_0 \in \mathcal{M}$ and performs the following iteration until convergence:

$$S_{t+1} \leftarrow \frac{\sum_{R \in B(S_t)} K\left(\frac{S_t - R}{h}\right) R}{\sum_{R \in B(S_t)} K\left(\frac{S_t - R}{h}\right)},\tag{17}$$

where $B(S_t)$ is a neighborhood of S_t .

While it has shown promising results in detecting symmetries (e.g., Chang et al., 2015), it may fail under noisy shapes or noisy transformation space (Je et al., 2024). Inspired by Je et al. (Je et al., 2024), we leverage a diffusion process (Ho et al., 2020) to establish the connection between the iteration in Eq. (17) and the stochastic gradient Langevin dynamics (Welling & Teh, 2011), where stochastic noise is injected to improve the sample quality and robustness.

B.2 Variational Autoencoder

We use *variational autoencoders* (*VAEs*) (Kingma & Welling, 2013) as our latent distribution model as it provides access to a low-dimensional latent space and has been successfully applied to generate point clouds (Li et al., 2022a; Wang et al., 2020; Li et al., 2022b; Zhou et al., 2024). VAEs are probabilistic generative models that can model a probability distribution of a given dataset \mathcal{X} .

Starting with a known prior distribution $p(\mathbf{z})$ of shape latents $\mathbf{z} \in \mathbb{R}^{\mathbf{z}}$, the parametric decoder of a VAE models the conditional distribution $p_{\eta}(\mathbf{x} \mid \mathbf{z})$ parametrized by η . However, training the decoder to maximize the likelihood of data is not possible as

$$p(\mathbf{x}) = \int p_{\eta}(\mathbf{x} \mid \mathbf{z}) p(\mathbf{z}) d\mathbf{z}$$
(18)

is intractable. Instead, a parametric encoder $p_{\eta}(\mathbf{x} \mid \mathbf{z})$ is used to approximates the posterior distribution. Both networks are jointly trained to maximize a lower bound on the likelihood called the *evidence lower bound (ELBO)*:

$$\mathcal{L}_{\text{ELBO}}(\boldsymbol{\eta}; \mathbf{x}) := \mathbb{E}_{q_{\boldsymbol{\eta}}(\mathbf{z}|\mathbf{x})} \left[\log p_{\boldsymbol{\eta}}(\mathbf{x} \mid \mathbf{z}) \right] - \mathcal{D}_{\text{KL}} \left(q_{\boldsymbol{\eta}}(\mathbf{z} \mid \mathbf{x}), p(\mathbf{z}) \right), \tag{19}$$

where \mathcal{D}_{KL} is the Kullback-Leibler divergence between the two distributions (Csiszár, 1975).

B.3 Denoising Diffusion Probabilistic Model

Our generative framework leverages four denoising diffusion probabilistic models (DDPMs or diffusion models) (Sohl-Dickstein et al., 2015; Ho et al., 2020) to model distinct data distributions. Given a data sample $\mathbf{z} \sim p(\mathbf{z})$, diffusion models progressively corrupt $\mathbf{z} = \mathbf{z}_0$ into a noisy sample \mathbf{z}_{τ} through a Markovian forward diffusion process. At each time step $t = 1, 2, ..., \tau$, Gaussian noise is added according to a predefined variance schedule $\{\sigma_t\}_t$:

$$q\left(\mathbf{z}_{1:\tau} \mid \mathbf{z}_{0}\right) := \prod_{t=1}^{\prime} q(\mathbf{z}_{t} | \mathbf{z}_{t-1})$$

$$(20)$$

$$q(\mathbf{z}_t \mid \mathbf{z}_{t-1}) \sim \mathcal{N}\left(\mu_t \mathbf{z}_{t-1}, \sigma_t I\right),$$
 (21)

where $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ denotes multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}$. In practice, we set $\mu_t \coloneqq \sqrt{1 - \sigma_t}$. If τ is sufficiently large (e.g., 1000 steps), $p(\mathbf{z}_{\tau})$ will approach the standard Gaussian distribution $\mathcal{N}(0, I)$.

Diffusion models learn a reverse process $p_{\theta}(\mathbf{z}_{t-1} \mid \mathbf{z}_t)$, parameterized by θ , which defines a Markovian denoising chain that inverts the forward diffusion. This process gradually transforms a sample

 of standard Gaussian noise \mathbf{z}_{τ} back into a data sample \mathbf{z}_0 :

$$p_{\boldsymbol{\theta}}(\mathbf{z}_{0:\tau}) := p(\mathbf{z}_{\tau}) \prod_{t=1}^{\tau} p_{\boldsymbol{\theta}}(\mathbf{z}_{t-1} \mid \mathbf{z}_{t});$$
(22)

$$p_{\theta}(\mathbf{z}_{t-1} \mid \mathbf{z}_t) \sim \mathcal{N}\left(\boldsymbol{\mu}_{\theta}(\mathbf{z}_t, t), \varsigma_t^2 I\right),$$
 (23)

where $\mu_{\theta}(\mathbf{z}_t, t)$ represents the predicted mean for the Gaussian distribution at time step t and $\{\varsigma_t\}_t$ is another variance schedule.

DDPMs are trained by maximizing the variational lower bound of log-likelihood of the data \mathbf{z}_0 under $q(\mathbf{z}_0)$:

$$\mathbb{E}_{q(\mathbf{z}_0)}\left[\log p_{\boldsymbol{\theta}}(\mathbf{z}_0)\right] \ge \mathbb{E}_{q(\mathbf{z}_{0:\tau})}\left[\log \frac{p_{\boldsymbol{\theta}}(\mathbf{z}_{0:\tau})}{q(\mathbf{z}_{1:\tau} \mid \mathbf{z}_0)}\right]. \tag{24}$$

Expanding Eq. (24) with Eq. (22) and noticing that $p(\mathbf{z}_{\tau})$ and $q(\mathbf{z}_{1:\tau} \mid \mathbf{z}_0)$ are constant with respect to $\boldsymbol{\theta}$, we obtain our objective function to maximize:

$$\mathbb{E}_{q(\mathbf{z}_0), q(\mathbf{z}_{1:\tau}|\mathbf{z}_0)} \left[\sum_{t=1}^{\tau} \log p_{\boldsymbol{\theta}}(\mathbf{z}_{t-1} \mid \mathbf{z}_t) \right]. \tag{25}$$

Since we can factor the joint posterior

$$q(\mathbf{z}_{1:\tau} \mid \mathbf{z}_0) = \prod_{t=1}^{\tau} q(\mathbf{z}_{t-1} \mid \mathbf{z}_t, \mathbf{z}_0)$$
 (26)

and both $q(\mathbf{z}_{t-1} \mid \mathbf{z}_t, \mathbf{z}_0)$ and $p_{\theta}(\mathbf{z}_{t-1} \mid \mathbf{z}_t)$ are Gaussian, maximizing Eq. (25) is equivalent as minimizing the following score matching objective for a parametric model $\epsilon_{\theta}(\mathbf{z}_t, t)$:

$$\mathcal{L}_{\text{diffusion}}(\boldsymbol{\theta}) := \mathbb{E}_{p(\mathbf{z}_0), t \sim \mathcal{U}(1, \tau), \boldsymbol{\epsilon} \sim \mathcal{N}(0, I)} \left[\| \boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{z}_t, t) \|^2 \right], \tag{27}$$

where $\mathcal{U}(1,\tau)$ is the uniform distribution on $\{1,2,\ldots,\tau\}$. Intuitively, minimizing this loss corresponds to learning to predict the noise ϵ required to denoise the diffused sample \mathbf{z}_t . Notably, the training objective of diffusion models closely aligns with estimating the gradient of the log data density -i.e., the score function - as used in score-based energy models (Song et al., 2020a; Swersky et al., 2011; Hyvärinen et al., 2009; LeCun et al., 2006).

During inference, the network allows sampling through an iterative procedure since the learned distribution can be factorized as

$$p_{\boldsymbol{\theta}}\left(\left\{\mathbf{z}_{t}\right\}_{t=0}^{\tau}\right) = p\left(\mathbf{z}_{\tau}\right)p_{\boldsymbol{\theta}}\left(\mathbf{z}_{t-1} \mid \mathbf{z}_{t}\right) = p\left(\mathbf{z}_{\tau}\right)\prod_{t=1}^{\tau}p_{\boldsymbol{\theta}}\left(\mathbf{z}_{t-1} \mid \mathbf{z}_{t}\right)$$
(28)

for $p(\mathbf{z}_{\tau}) := \mathcal{N}(0, I)$.

C IMPLEMENTATION AND TRAINING DETAILS

Tables 4 and 5 provide the architecture and training details for the various VAE and diffusion models used in the Quartet. The hyperparameter λ in Eq. (7), which controls the activation sparsity in the SVAE, is set to 0.005. The hyperparameter for the equivariance fine-tuning in the part VAE, discussed in Sec. 3.4, is set to 0.01. For all models, we employ a learning rate scheduler with a reduce-on-plateau policy, which decreases the learning rate by a factor of 10 if the loss does not improve for 10 consecutive epochs. All training runs converged successfully by the end of training.

D More Experimental Results

Table 6 and Figs. 5 and 6 present more results for the point cloud generation experiment in Sec. 4.3.

Table 4: Model architecture details for the Quartet. M denotes the number of parts for each object category (M=4 for airplanes and cars; M=3 for chairs).

| Model | Backbone | Data dimensionality | Latent dimensionality |
|------------------------|------------------------------------|----------------------------------|-----------------------|
| Point cloud SVAE | PVCNN (Liu et al., 2019) | 2048×3 | 128 |
| Shape latent diffusion | U-Net (Ronneberger et al., 2015) | 128 | - |
| Symmetry diffusion | ResNet (He et al., 2016) | $12 \times M$ | - |
| Part diffusion | Transformer (Vaswani et al., 2017) | 2048×3 | - |
| Part VAE | PVCNN (Liu et al., 2019) | 57×3 to 1024×3 | 128 |
| Assembler diffusion | U-Net (Ronneberger et al., 2015) | $9 \times M$ | |

Table 5: Training details for the Quartet.

| Model | Batch size | Number of epoch | Learning rate |
|------------------------|------------|-----------------|---------------|
| Point cloud SVAE | 64 | 1000 | 10^{-3} |
| Shape latent diffusion | 32 | 2000 | 10^{-4} |
| Symmetry diffusion | 32 | 2000 | 10^{-4} |
| Part diffusion | 16 | 2000 | 10^{-3} |
| Part VAE | 64 | 1000 | 10^{-5} |
| Assembler diffusion | 32 | 2000 | 10^{-4} |

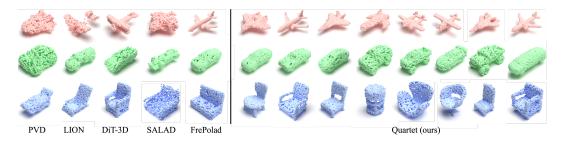


Figure 5: Full point cloud generation results. Samples from the Quartet are visually appealing, diverse, and exhibit strong structural and symmetry consistency.

Table 6: Quantitative comparison of point cloud generation. PA denotes part awareness; SA denotes symmetry awareness. Our Quartet is the only model that supports both, achieving significant improvements over most baselines and setting a new state of the art.

| | | | Airplane | | | | | C | ar | | Chair | | | |
|-------------------------------------|----|----|----------|--------|-------|-------------|------|--------|------|----------|-------|-------|------|----------|
| Model | PA | SA | 1-NN | √A (↓) | SDI | [(\lambda) | 1-NN | √A (↓) | SD | I (\psi) | 1-NN | √ (↓) | SD | I (\psi) |
| | | | | EMD | CD | EMD | CD | EMD | CD | EMD | CD | EMD | CD | EMD |
| Training set | | | 64.4 | 64.1 | 0.954 | 4.90 | 51.3 | 54.8 | 7.49 | 1.18 | 51.7 | 50.0 | 5.56 | 1.68 |
| r-GAN (Achlioptas et al., 2018) | Х | Х | 98.4 | 96.8 | 4519 | 1410 | 83.7 | 99.7 | 1053 | 362 | 94.5 | 99.0 | 7249 | 619 |
| 1-GAN/CD (Achlioptas et al., 2018) | X | X | 87.3 | 94.0 | 3629 | 1353 | 68.6 | 83.8 | 918 | 352 | 66.5 | 88.8 | 7972 | 582 |
| 1-GAN/EMD (Achlioptas et al., 2018) | Х | X | 89.5 | 76.9 | 4129 | 914 | 71.9 | 64.7 | 982 | 313 | 71.2 | 66.2 | 7184 | 521 |
| PointFlow (Yang et al., 2019) | Х | X | 75.7 | 70.7 | 3410 | 782 | 62.8 | 60.6 | 679 | 347 | 58.1 | 56.3 | 7290 | 530 |
| SoftFlow (Kim et al., 2020) | | X | 76.1 | 65.8 | 3284 | 529 | 59.2 | 60.1 | 1549 | 428 | 64.8 | 60.1 | 5628 | 420 |
| ShapeGF (Cai et al., 2020) | Х | X | 81.2 | 80.9 | 332 | 98.6 | 58.0 | 61.3 | 645 | 40.9 | 61.8 | 57.2 | 1100 | 101 |
| DPF-Net (Klokov et al., 2020) | Х | X | 75.2 | 65.6 | 4256 | 245 | 62.0 | 58.5 | 827 | 452 | 62.4 | 54.5 | 5234 | 245 |
| SetVAE (Kim et al., 2021) | X | X | 76.5 | 67.7 | 2830 | 824 | 58.8 | 60.6 | 1240 | 327 | 59.9 | 59.9 | 5320 | 673 |
| DPC (Luo & Hu, 2021) | Х | X | 76.4 | 86.9 | 187 | 44.2 | 60.1 | 74.8 | 217 | 30.3 | 68.9 | 80.0 | 335 | 50.6 |
| PVD (Zhou et al., 2021a) | X | X | 73.8 | 64.8 | 150 | 42.0 | 56.3 | 53.3 | 213 | 31.2 | 54.6 | 53.8 | 275 | 58.4 |
| LION (Vahdat et al., 2022) | Х | X | 67.4 | 61.2 | 97.2 | 40.6 | 53.7 | 52.3 | 168 | 30.8 | 53.4 | 51.1 | 201 | 55.2 |
| SPAGHETTI (Hertz et al., 2022) | 1 | X | 78.2 | 77.0 | 1530 | 529 | 72.3 | 71.0 | 581 | 284 | 70.7 | 69.0 | 5930 | 582 |
| DiT-3D (Mo et al., 2023) | Х | Х | 64.7 | 60.3 | 105 | 42.4 | 52.7 | 50.2 | 206 | 327 | 52.5 | 53.1 | 235 | 49.0 |
| SALAD (Koo et al., 2023) | 1 | X | 73.9 | 71.1 | 198 | 45.1 | 59.2 | 57.2 | 236 | 29.4 | 57.8 | 58.4 | 308 | 52.6 |
| FrePolad (Zhou et al., 2024) | Х | X | 65.3 | 62.1 | 94.1 | 38.1 | 52.4 | 53.2 | 173 | 29.6 | 51.9 | 50.3 | 252 | 50.9 |
| Quartet (ours) | ✓ | ✓ | 63.3 | 59.7 | 25.7 | 1.87 | 50.1 | 51.8 | 25.7 | 2.28 | 51.6 | 53.7 | 28.9 | 2.86 |

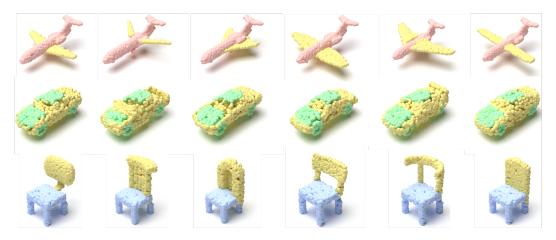


Figure 6: More targeted manipulation results. We vary the yellow-highlighted parts while holding the remaining structure fixed.