

LaS-Comp: Zero-shot 3D Completion with Latent–Spatial Consistency

Weilong Yan¹ Haipeng Li² Hao Xu³ Nianjin Ye⁴
Yihao Ai¹ Shuaicheng Liu² Jingyu Hu^{3†}

¹National University of Singapore ²University of Electronic Science and Technology of China

³The Chinese University of Hong Kong ⁴Changhong Intelligent Robot

Abstract

This paper introduces *LaS-Comp*, a zero-shot and category-agnostic approach that leverages the rich geometric priors of 3D foundation models to enable 3D shape completion across diverse types of partial observations. Our contributions are threefold: First, *LaS-Comp* harnesses these powerful generative priors for completion through a complementary two-stage design: (i) an explicit replacement stage that preserves the partial observation geometry to ensure faithful completion; and (ii) an implicit refinement stage ensures seamless boundaries between the observed and synthesized regions. Second, our framework is training-free and compatible with different 3D foundation models. Third, we introduce *Omni-Comp*, a comprehensive benchmark combining real-world and synthetic data with diverse and challenging partial patterns, enabling a more thorough and realistic evaluation. Both quantitative and qualitative experiments demonstrate that our approach outperforms previous state-of-the-art approaches. Our code and data will be available at *LaS-Comp*.

1. Introduction

Shape completion is a fundamental problem in 3D vision and graphics, aiming to reconstruct complete 3D shapes from partial observations, with broad applications in robotics [33, 34], autonomous driving [76, 77, 85], and AR/VR [3, 40, 79]. An effective shape completion approach should meet the following requirements: (i) robustly handle diverse partial patterns, from single-view scans to missing semantic parts; (ii) well generalize across broad object categories; (iii) avoid dependence on paired partial–complete datasets; and (iv) flexibly support both text guidance and automatic completion without user input for applications.

Traditional methods [50, 51, 54, 55, 58] use hand-crafted priors to infer the missing geometry from partial observations. However, such priors lack flexibility and fail to gen-




Figure 1. Our new framework supports category-agnostic shape completion across diverse partial patterns, including (a) random crops, (b) single-view scans, and (c) missing semantic parts. It further supports both unconditional and text-guided completion, offering flexible control for real-world applications, see (d).

eralize to diverse real-world shapes. Subsequently, [14, 82] propose to learn a mapping from partial to complete shapes by training a neural network. Yet, these methods rely on paired partial–complete datasets, limiting their generalizability to unseen categories. To overcome such constraints, recent works [28, 35, 38] leverage generative priors from large pre-trained models to enable category-agnostic shape completion. However, these methods assume that a partial input can be rendered into at least one complete image, where the object appears geometrically complete in the rendered view. When this assumption fails, *e.g.*, the incomplete regions are visible from any viewpoint (see Fig. 1(c)), the incomplete renderings lead to degraded results.

Recent 3D foundation models [19, 32, 41, 57, 66, 67, 84,

† Corresponding author

88], trained on large-scale datasets [15, 16], show strong cross-category and in-the-wild shape generation capabilities, enabling a wide range of downstream applications such as 3D editing [4, 24, 26, 27, 30, 31, 47] and analysis [7, 17, 71–73, 87]. In contrast, existing shape completion methods struggle to generalize across the diverse input partial patterns and shape categories. This gap motivates us to exploit the powerful geometric priors embedded in these foundation models for more robust and generalizable shape completion. Yet, unlike earlier 3D generative models [25, 30, 46, 48, 89] that operate directly in the spatial domain, *e.g.*, points or voxels, these foundation models adopt a latent-generative-based pipeline: A variational auto-encoder first maps shape to a compact latent space, and a diffusion model is trained over this space. While this design effectively compresses 3D data and supports scalable training of large foundation models, it simultaneously introduces a unique challenge for shape completion. A straightforward approach for shape completion within the latent framework is to encode the partial input into the latent space and directly use the latent codes of these observed regions as conditions to guide the generative model to predict the missing geometry. However, we observe that even when a complete shape and its partial input share identical geometry in the overlapping regions, their latent encodings in those regions differ significantly. Therefore, directly completing shapes in latent space becomes unreliable due to this domain gap.

To address this challenge, we propose a zero-shot **Latent–Spatial Consistency Completion (LaS-Comp)** framework, which bridges the gap between the latent and spatial domains of pre-trained 3D foundation models for faithful and category-agnostic shape completion. LaS-Comp operates via a complementary two-stage design: First, we introduce an **Explicit Replacement Stage (ERS)**, ensuring fidelity to the input partial shape by directly injecting its geometric information into the latent representation. Second, we adopt an **Implicit Alignment Stage (IAS)** to improve smoothness between observed and synthesized regions by refining the latent features with a geometry-alignment loss. By combining explicit geometric replacement with implicit latent refinement, LaS-Comp unleashes the power of 3D foundation models to deliver high-fidelity shape completions.

Our framework is compatible with different latent-generative-based 3D foundation models [66, 67]. By exploiting the rich geometric priors encoded by these models, our framework is capable of robustly handling diverse partial patterns, ranging from single-view scans to irregularly missing regions; see Fig. 1(a-c). Furthermore, our approach leverages the built-in classifier-free guidance (CFG) mechanism of these foundation models to support both text-guided and unconditional shape completion, offering flexible user control; see Fig. 1(d). Finally, we found the evaluation gap

in existing benchmarks [8, 28], which are mostly limited to single-view partial scans and do not reflect real-world diversity. To bridge this gap, we introduce Omni-Comp: a new benchmark composed of real-world scans and synthetic data with diverse, challenging partial patterns, enabling a more comprehensive in-the-wild evaluation.

In summary, our contributions are listed as follows:

- We introduce LaS-Comp, a novel zero-shot, category-agnostic 3D shape completion framework. LaS-Comp bridges the gap between latent and spatial domains via a two-stage design: the Explicit Replacement Stage (ERS) preserves fidelity to the partial input, while the Implicit Alignment Stage (IAS) ensures boundary coherence.
- Our framework is training-free, compatible with different 3D foundation models, and highly efficient. It completes each shape in 20 seconds, making it over $3\times$ faster than existing zero-shot methods.
- We introduce Omni-Comp, a new benchmark composed of both real-world scans and synthetic shapes with diverse partial patterns to enable more comprehensive in-the-wild evaluation. Extensive experiments on this benchmark and others confirm our method’s state-of-the-art performance.

2. Related Work

Supervised Shape Completion. Supervised shape completion methods learn to reconstruct complete 3D shapes from partial inputs by training on paired datasets of partial and complete shapes. Pioneering works [14, 82] established this paradigm by introducing paired datasets and end-to-end networks to directly learn the mapping from partial to complete shapes. Subsequent studies [29, 61, 64, 70, 86] employed coarse-to-fine refinement to better reconstruct missing regions and recover fine geometric details. Transformer-based approaches [39, 59, 62, 68, 74, 75, 78, 80, 90] leverage attention mechanisms to effectively aggregate multi-scale geometric features, leading to higher-fidelity shape completions. Meanwhile, diffusion-based methods [9, 11, 42, 49, 89] formulate shape completion as a conditional generation task, where the partial input serves as a condition for the diffusion model [21] to produce a completed shape. Although effective on in-domain data, these methods fail to generalize to unseen categories and partial patterns. Moreover, collecting large-scale paired datasets is costly, motivating the exploration of unsupervised alternatives.

Unsupervised Shape Completion. To alleviate reliance on paired data, recent approaches explore unsupervised shape completion. [6] first leverages unpaired data and adversarial losses. Follow-up works [1, 63, 65, 69, 83] further advance research through various unsupervised strategies, *e.g.*, cycle consistency and latent space alignment. Building upon these efforts, recent self-supervised approaches [10, 12, 22, 36, 44] relax the need for complete shapes




Figure 2. Overview of the LaS-Comp framework. Starting from Gaussian noise, the process iteratively refines a latent feature x_t under the guidance of the partial input S_p . At each iteration t , this refinement is performed in two stages: the *Explicit Replacement Stage (ERS)* and the *Implicit Alignment Stage (IAS)*. The ERS explicitly injects the known geometry of S_p into x_t to produce an updated latent x_t^* . The IAS then refines x_t^* using a gradient-based optimization, yielding a spatially aligned latent x_{t-dt} for the next step. After the final iteration, the completed shape S_c is obtained by decoding the refined latent.

by learning geometric priors directly from partial observations. However, these methods are trained on ShapeNet [5], which contains limited categories, leading to degraded performance on real-world data or unseen categories.

Shape Completion with Generative Priors. To overcome the limitations of paired-data supervision, recent works [28, 35, 38] leverage priors from large-scale pre-trained generative models for category-agnostic 3D shape completion. SDS-Complete [35] distills 2D priors from Stable Diffusion [53] to optimize the signed distance field, while ComPC [28] extends this idea by initializing partial inputs as 3D Gaussians and utilizing Zero-1-to-3 [45] to render multi-view images and iteratively refine them for shape completion. GenPC [38] renders partial input into an image, reconstructs a coarse shape via image-to-3D models [23, 56], and refines it using Zero-1-to-3 [45]. However, these methods rely on the assumption that partial inputs can be rendered into at least one reasonably complete image. When the partial input fails to yield a complete image from any viewpoint, the incomplete renderings often cause sub-optimal completion results. To address the limitations of existing approaches, we propose the first framework that fully leverages latent-generative-based 3D foundation models for high-fidelity, zero-shot, and category-agnostic shape completion. Without relying on any rendering assumptions, our method robustly handles highly diverse partial input patterns, from single-view scans to irregular occlusions. Moreover, it is fully training-free, enabling seamless deployment across a wide range of real-world scenarios.

3. Method

3.1. Overview

Given a partial 3D shape $S_p \in \mathbb{R}^{k \times 3}$, we desire to generate a completed shape $S_c \in \mathbb{R}^{k \times 3}$ that is geometrically faithful to S_p . To achieve this, we introduce the LaS-Comp

framework, which effectively leverages the powerful geometric priors of pre-trained 3D foundation models [66, 67] for zero-shot shape completion. Through the classifier-free guidance sampling mechanism [20] of 3D foundation models, our method supports both text-guided and unconditional shape completion.

As illustrated in Fig. 2, starting from a Gaussian noise, our framework iteratively refines it through multiple denoising steps under the guidance of the partial input S_p , progressively recovering the underlying complete geometry. At each iteration $t \in [0, 1]$, LaS-Comp takes the current latent feature $x_t \in \mathbb{R}^{n^3 \times c}$ as input, together with the partial shape S_p , and performs two complementary operations:

- *Explicit Replacement Stage (ERS)* explicitly injects the partial input S_p into x_t , producing an updated latent x_t^* that enforces strict fidelity to the partial input.
- *Implicit Alignment Stage (IAS)* refines x_t^* through a one-step optimization guided by a geometric-alignment loss, which encourages the smoothness between synthesized and observed regions, producing a coherent latent x_{t-dt} for the next step (dt is the length of a denoising step).

After the final iteration, the refined latent feature x_0 is decoded through the decoder \mathcal{D} to produce the completed shape: $S_c = \mathcal{D}(x_0)$. The details of ERS and IAS are demonstrated in Sec. 3.2 and Sec. 3.3, respectively.

3.2. Explicit Replacement Stage

Taking the latent feature x_t from the previous generative step and the partial shape S_p as input, the ERS performs primary completion by injecting the geometry of S_p into the completion process, producing an updated latent feature x_t^* that explicitly encodes the observed regions.

As depicted in Fig. 3, motivated by the latent decomposition in FlowDPS [37], we decompose the generation at each timestep into two branches to enhance both fidelity and diversity: (i) a clean branch for enforcing the input fidelity; and (ii) a noisy branch for improving generation diversity.

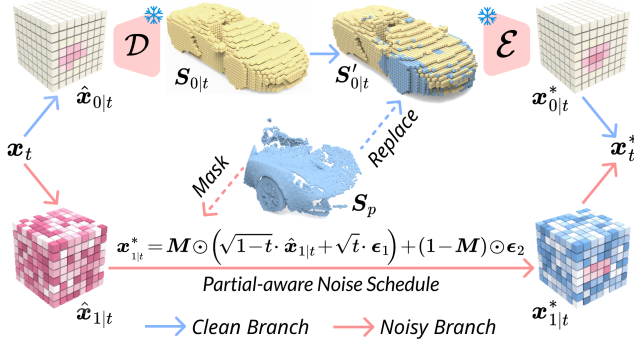


Figure 3. Overview of the Explicit Replacement Stage (ERS). At each timestep t , ERS decomposes the latent generation into two parallel branches. The clean branch (top) enforces spatial consistency, yielding $\mathbf{x}_{0|t}^*$. Concurrently, the noisy branch (bottom) enhances fidelity, producing $\mathbf{x}_{1|t}^*$. These two branch outputs are then interpolated to compute the final aligned latent \mathbf{x}_t^* .

Clean Branch. Given the current latent feature \mathbf{x}_t , the clean branch uses the generator \mathcal{G} to estimate the noise-free latent feature $\hat{\mathbf{x}}_{0|t}$ based on the linear flow path:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_t - t \cdot \mathcal{G}(\mathbf{x}_t, t), \quad (1)$$

where $\mathcal{G}(\mathbf{x}_t, t)$ is the predicted velocity from \mathbf{x}_t towards \mathbf{x}_0 . This latent feature $\hat{\mathbf{x}}_{0|t}$ is then decoded into the spatial domain to yield a complete shape prediction:

$$\mathbf{S}_{0|t} = \mathcal{D}(\hat{\mathbf{x}}_{0|t}). \quad (2)$$

Although $\mathbf{S}_{0|t} \in \mathbb{R}^{N^3}$ provides a complete shape produced by \mathcal{G} , it does not yet faithfully preserve the observed regions of the partial input \mathbf{S}_p . To enforce this fidelity, we first voxelize the input partial shape \mathbf{S}_p into an occupancy grid of size N^3 aligned with the spatial domain of $\mathbf{S}_{0|t}$. We then construct a binary spatial mask $\mathbf{M} \in \{0, 1\}^{N^3}$, where $\mathbf{M} = 1$ indicates occupied voxels observed in \mathbf{S}_p . Using this mask, we perform a spatial replacement to inject the known geometry into the generative prediction:

$$\mathbf{S}'_{0|t} = \mathbf{S}_p \odot \mathbf{M} + \mathbf{S}_{0|t} \odot (1 - \mathbf{M}), \quad (3)$$

where \odot is element-wise multiplication. This operation explicitly replaces the geometry with \mathbf{S}_p for fidelity. The missing regions are then completed by the generative prediction $\mathbf{S}_{0|t}$. Next, we feed $\mathbf{S}'_{0|t}$ into the encoder \mathcal{E} to produce a latent feature $\mathbf{x}_{0|t}^*$ that incorporates the geometries from the partial input as follows:

$$\mathbf{x}_{0|t}^* = \mathcal{E}(\mathbf{S}'_{0|t}). \quad (4)$$

Noisy Branch. Concurrently, the noisy branch estimates the noisy latent $\hat{\mathbf{x}}_{1|t}$ using the same generator prediction $\mathcal{G}(\mathbf{x}_t, t)$ as in Eq. (1):

$$\hat{\mathbf{x}}_{1|t} = \mathbf{x}_t + (1 - t) \cdot \mathcal{G}(\mathbf{x}_t, t). \quad (5)$$

Rather than employing $\hat{\mathbf{x}}_{1|t}$ for latent update, we introduce a novel *Partial-aware Noise Schedule (PNS)* to modulate the stochasticity during denoising. The motivation is that $\hat{\mathbf{x}}_{1|t}$ treats all spatial regions equally, whereas shape completion is asymmetric: observed regions should remain stable to preserve the input geometry, while missing regions should permit higher stochasticity to explore plausible completions. The final noisy latent $\mathbf{x}_{1|t}^*$ is then composed as:

$$\mathbf{x}_{1|t}^* = \mathbf{M} \odot \left(\sqrt{1-t} \cdot \hat{\mathbf{x}}_{1|t} + \sqrt{t} \cdot \epsilon_1 \right) + (1 - \mathbf{M}) \odot \epsilon_2, \quad (6)$$

where $\epsilon_1, \epsilon_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and \mathbf{M} is downsampled to size of n^3 to match the dimensions of the latent features. This design applies the following mechanisms:

- For the observed regions ($\mathbf{M} = 1$): Since these regions correspond to the reliable partial input, their latent features should remain largely stable during denoising, with only minimal stochastic disturbance. To achieve this, we blend the model-predicted noisy latent $\hat{\mathbf{x}}_{1|t}$ with Gaussian noise ϵ_1 using the coefficients $\sqrt{1-t}$ and \sqrt{t} . This schedule imposes a time-dependent perturbation magnitude: early iterations inject greater randomness to allow adjustments that maintain overall coherence, while later iterations gradually reduce stochasticity and preserve the geometry of the observed regions.
- For the missing regions ($\mathbf{M} = 0$): No reliable observations are available to constrain the denoising process. To encourage broad exploration of the possible completions, we replace the model-predicted noisy latent with pure Gaussian noise ϵ_2 . This promotes diversity in the generated geometry for the missing areas.

With the updated noise-free latent $\mathbf{x}_{0|t}^*$ and the noisy latent $\mathbf{x}_{1|t}^*$, we reconstruct the updated latent state \mathbf{x}_t^* for the current timestep t via the forward flow interpolation:

$$\mathbf{x}_t^* = (1 - t) \cdot \mathbf{x}_{0|t}^* + t \cdot \mathbf{x}_{1|t}^*, \quad (7)$$

the resulting latent \mathbf{x}_t^* encodes the partial input geometry via the ERS, where the observed regions are directly imposed in the spatial domain. Although this explicit spatial replacement preserves fidelity to the input, it may introduce inconsistencies or discontinuities near the boundaries between observed and synthesized regions. To mitigate these artifacts and improve local coherence, we feed \mathbf{x}_t^* into the Implicit Alignment Stage (IAS) for a refinement step.

3.3. Implicit Alignment Stage

To improve coherence between observed and synthesized regions, we introduce the Implicit Alignment Stage (IAS), which refines the ERS output \mathbf{x}_t^* before the next iteration. The refinement begins by estimating a noise-free latent from \mathbf{x}_t^* . Refer to the Eq. (1), we reuse the notation $\hat{\mathbf{x}}_{0|t}$ to denote the noise-free latent predicted from \mathbf{x}_t^* :

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_t^* - t \cdot \mathcal{G}(\mathbf{x}_t^*, t), \quad (8)$$

then this predicted noise-free latent $\hat{\mathbf{x}}_{0|t}$ is then decoded into its spatial domain to obtain $\mathbf{S}_{0|t}$:

$$\mathbf{S}_{0|t} = \mathcal{D}(\hat{\mathbf{x}}_{0|t}). \quad (9)$$

To reduce the boundary artifacts introduced by the ERS, we define geometry-alignment loss $\mathcal{L}_{\text{align}}$ to refine the latent feature. This loss provides a localized correction within the reliable masked regions, helping smooth the discontinuities. Specifically, we use the mask \mathbf{M} to select the masked regions and compute a BCE loss between the predicted occupancy and the corresponding voxels of the partial input:

$$\mathcal{L}_{\text{align}} = \text{BCE}(\mathbf{S}_{0|t} \odot \mathbf{M}, \mathbf{S}_p \odot \mathbf{M}). \quad (10)$$

Notice that this loss is not used to update the model parameters. Instead, we compute the gradient of the loss with respect to the latent feature, and optimize it via a **single-step** gradient update to refine it:

$$\mathbf{x}_{0|t}^{\text{aligned}} = \hat{\mathbf{x}}_{0|t} - \eta \cdot \nabla_{\hat{\mathbf{x}}_{0|t}} \mathcal{L}_{\text{align}}, \quad (11)$$

where η is the learning rate. By refining the feature $\hat{\mathbf{x}}_{0|t}$ to better match the reliable masked regions, we obtain an updated latent feature $\mathbf{x}_{0|t}^{\text{aligned}}$ that is more coherent. This updated feature is then used to compute the latent state \mathbf{x}_{t-dt} for the next iteration as follows:

$$\mathbf{x}_{t-dt} = \mathbf{x}_{0|t}^{\text{aligned}} + (t-dt) \cdot \mathcal{G}(\mathbf{x}_t^*, t). \quad (12)$$

The latent feature \mathbf{x}_{t-dt} then serves as the input for the next iteration of the completion process.

4. Experiments

4.1. Datasets and Implementation details

Evaluation Datasets. We conduct a comprehensive evaluation across multiple benchmarks. (i) **Redwood** [8]: We follow [28, 35, 38], utilizing 10 reconstructed meshes from real-world RGB-D scans, with single scans as the partial input. (ii) **Synthetic dataset** [28]: Following [28, 43], this contains 12 objects from different categories, with virtually-rendered single scans as the partial input. (iii) **KITTI** [18]: We follow [6, 83], using cars from real-world LiDAR scans, with extremely sparse points as the partial input. (iv) **ScanNet** [13]: We follow [6, 12, 83], cropping 48 chairs and 49 tables from real-world single scans (no ground truths).

Proposed Omni-Comp Benchmark. These aforementioned benchmarks, while valuable, highlight clear limitations for comprehensive evaluation, including limited scales (e.g., about 10 meshes in Redwood/Synthetic), restricted category diversity (e.g., ≤ 2 categories in KITTI/ScanNet), and confinement to a single partial pattern (e.g., depth/LiDAR scans). To address these limitations, we introduce **Omni-Comp**, a new benchmark designed for

a more comprehensive and robust evaluation of 3D shape completion. Our benchmark features a challenging set of 30 objects, each from a distinct category, curated from diverse sources: 10 real-world scans from Redwood [8] (chosen for complex geometry), 10 real-world everyday objects from YCB [2] (motivated by downstream applications, e.g., robotic grasping), and 10 synthetic shapes from [60] (chosen for rich semantic structure). Critically, inspired by [83], our benchmark generates three distinct partiality patterns for each object: (i) *Single Scan*: using the projection of the captured depth map for real-world data, and simulating a standard depth camera capture for synthetic data; (ii) *Random Crop*: Representing arbitrary occlusions by randomly cropping a portion; and (iii) *Semantic Part*: Keeping a semantic component and removing other parts. By creating two samples for each pattern per object, the benchmark comprises 180 partial samples with ground truths. More details are in the supplementary material.

Evaluation Metrics. We follow previous works [28, 35, 38, 83] to assess completion quality from three perspectives: (i) Chamfer Distance (CD) and Earth Mover’s Distance (EMD), reported both $\times 10^2$, which measures completion accuracy by the distance between the prediction and the ground truth; (ii) Unidirectional Chamfer Distance (UCD) and Unidirectional Hausdorff Distance (UHD), reported $\times 10^4$ and $\times 10^2$, which measure the fidelity of the prediction compared with the partial input; and (iii) Minimum Matching Distance (MMD) and Total Mutual Difference (TMD), reported both $\times 10^2$, which assess completion diversity across multiple outputs. More details are in the supplementary material.

Baselines. We compare our method against a comprehensive set of recent approaches, categorized into three main groups, including supervised methods [62, 81, 90], unsupervised methods [12, 83], and generative prior-based (zero-shot) methods [28, 35, 38]. To ensure a fair comparison, we adopt the same experiment settings following previous zero-shot 3D completion methods [28, 35, 38, 83].

Implementation Details. We implement our framework based on two pre-trained 3D generative foundation models, TRELIS [67] and Direct3D-S2 [66]. Following [28], we sample point clouds from the output meshes with farthest point sampling for quantitative evaluation. The learning rate for the one-step IAS refinement step is set to 1×10^{-5} . More details are in the supplementary material.

4.2. Comparison with State-of-the-art Methods

We present comprehensive analyses across multiple benchmarks for assessment in this section. For fair visualization, all point clouds are converted to meshes using [52].

Evaluation on Completion Correctness. We first evaluate the completion correctness on single-scan partial shapes from the real-world Redwood dataset [8] and synthetic

Table 1. Quantitative comparisons on Redwood [8]. We highlight the **best** and **second-best** results.

CD ↓ / EMD ↓	Table	Exe-Chair	Out-Chair	Old-Chair	Vase	Off-Can	Vespa	Tricycle	Trash	Couch	Average
SVDFormer [90]	5.48/6.68	3.20/5.85	0.79/1.45	3.79/5.95	5.70/6.98	5.13/6.71	3.10/4.91	2.73/4.96	3.67/5.05	1.86/2.92	3.54/5.15
AdaPoinTr [81]	5.02/6.25	2.58/4.80	0.82/1.37	3.62/5.64	5.14/6.50	4.47/6.35	1.96/3.54	1.83/3.66	1.21/3.08	1.01/2.09	2.77/4.33
Shape-Inv [83]	1.58/2.84	3.59/5.75	1.36/2.12	4.55/7.39	3.91/7.40	3.10/4.77	4.36/7.24	5.05/7.40	2.48/3.83	1.95/2.76	3.19/5.15
P2C [12]	1.57/2.64	3.87/6.42	1.28/2.10	3.72/5.82	4.54/7.04	3.36/4.82	6.75/10.7	7.78/11.2	3.28/4.21	2.55/3.16	3.87/5.81
SDS-Comp [35]	1.35/2.30	1.96/2.65	2.51/3.92	2.77/3.77	3.00/5.25	3.79/4.28	3.36/5.73	3.18/3.49	2.69/3.21	2.95/4.56	2.74/3.93
PCDremer [62]	0.82/2.36	1.43/3.56	1.19/2.07	2.61/5.34	2.41/3.64	2.62/3.89	2.20/4.30	2.53/3.82	1.55/2.47	1.47/2.61	1.88/3.41
GenPC [38]	1.28/2.07	1.43/2.29	1.16/1.68	1.36/2.20	2.86/4.85	2.72/4.36	1.36/2.47	1.38/2.97	2.31/3.17	1.58/2.78	1.74/2.88
ComPC [28]	1.73/3.29	1.29/1.85	1.35/1.94	1.14/1.63	2.89/4.13	3.55/3.92	1.21/1.86	2.65/2.48	2.15/2.34	1.59/2.58	1.95/2.59
Ours (Direct3D-S2)	0.75/1.29	1.54/2.16	0.96/ 1.37	1.11/1.69	2.61/3.89	2.90/2.92	1.35/2.08	1.96/2.41	1.21/1.85	1.95/2.26	1.64/2.19
Ours (TRELLIS)	0.82/1.45	1.25/1.65	0.96/1.38	1.07/1.37	2.09/3.61	2.57/2.53	1.17/1.78	2.07/ 1.70	1.32/ 1.76	0.88/1.14	1.42/1.84

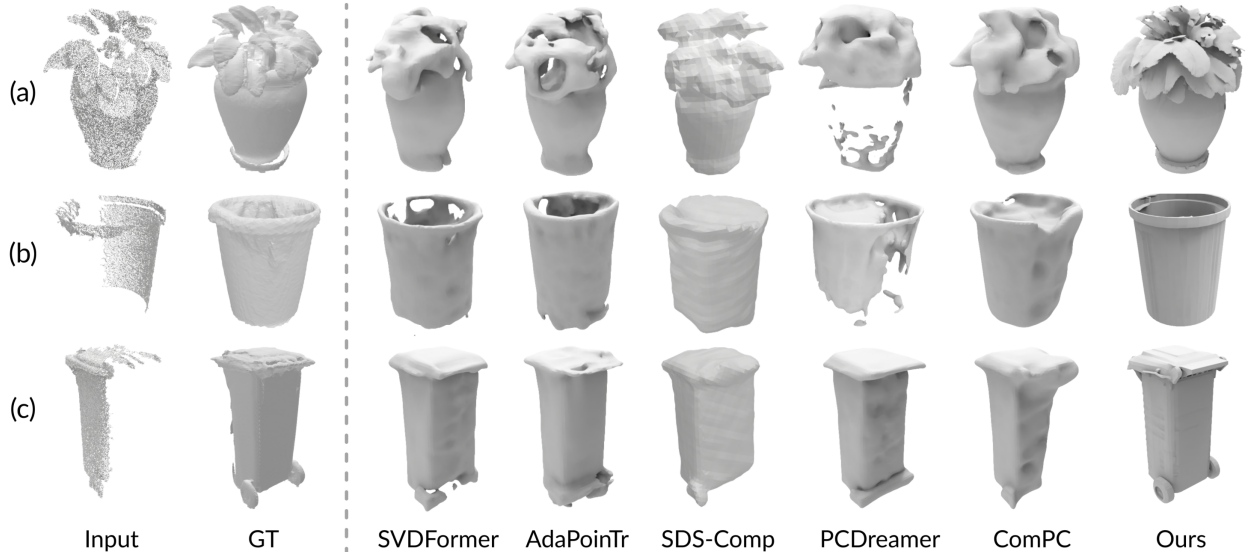


Figure 4. Qualitative comparison on Redwood dataset [8]. We compare with various supervised and unsupervised methods [28, 35, 62, 81, 90], and visualize the output as meshes utilizing the commonly-used mesh reconstruction method [52].

data [28, 43], as shown in Tab. 1 and Tab. 2. Our method achieves state-of-the-art performance in most categories and demonstrates significant improvements compared to recent zero-shot methods. Specifically, we outperform ComPC [28] by 27.2% in CD and 29.0% in EMD, and surpass GenPC [38] by 18.4% in CD and 36.1% in EMD, respectively. We attribute this significant gap to a fundamental difference in methodology. Prior methods [28, 38] either rely on generating multiple 2D views for supervision or lift a 2D rendered image to 3D with post-processing. In contrast, our approach directly conducts completion by leveraging both the strong generative priors and the rich 3D information within the partial input, enabling a more robust and accurate inference of the missing geometry.

We show qualitative comparisons in Fig. 4. For specific examples, such as the leaf structures of the plant (see Row (a)) and the rim and wheels of the trash bin (see Rows (b-c)), our method achieves high fidelity in the observed regions while generating missing parts with precise surface geometry and coherent topology. In contrast, existing methods either distort the visible geometry or generate over-smoothed,

implausible structures. This demonstrates that enforcing latent-spatial consistency with strong 3D generative priors yields superior geometric realism and structural integrity.

Evaluation on Completion Fidelity. To test generalization and robustness to more challenging inputs, we evaluate on datasets with complex partiality, including sparse LiDAR scans from KITTI [18] and noisy depth scans from ScanNet [13]. As shown in Tab. 5, our method outperforms the latest zero-shot method ComPC [28] by a large margin in almost all categories, confirming its robustness and superior fidelity on challenging real-world data. Some qualitative examples can be found in Fig. 5.

Evaluation on Omni-Comp. Our proposed Omni-Comp dataset specifically tests generalization across diverse partial patterns in both real-world and synthetic scenarios. As shown in Tab. 4, we report the average metrics on all shapes according to different partial patterns. Previous unsupervised approaches [28, 35] are heavily optimized for single-scan patterns, suffer a consistent and significant performance drop when generalizing to other partiality types. In contrast, our method leverages a pattern-agnostic 3D prior,

Table 2. Quantitative comparisons on the synthetic data [28, 43]. We highlight the **best** and second-best results.

CD ↓ / EMD ↓	Horse	Max-Planck	Armadillo	Cow	Homer	Teapot	Bunny	Nefertiti	Bimba	Ogre	Lucy	Dragon	Average
SVDFormer [90]	4.33/5.17	8.84/7.95	4.97/6.13	3.50/4.39	2.30/3.20	5.00/5.91	9.61/9.19	5.46/5.94	7.54/7.18	4.90/5.55	2.08/2.93	2.88/4.35	5.12/5.66
AdaPoinTr [81]	4.88/5.45	8.60/8.52	5.14/5.97	3.48/4.51	2.28/3.30	3.92/4.53	9.33/8.86	5.54/6.16	8.16/7.62	4.53/5.41	1.85/2.79	3.07/4.54	5.07/5.64
Shape-Inv [83]	6.55/10.4	4.94/5.72	4.79/7.22	4.74/7.83	2.36/3.75	3.29/4.53	5.79/6.58	4.39/5.70	5.29/6.85	5.76/9.51	2.99/4.07	4.48/7.91	4.62/6.68
PCDreamer [62]	2.52/4.34	5.20/3.85	2.78/4.53	1.79/3.06	1.54/2.97	1.95/3.57	5.48/5.66	2.72/3.43	3.83/5.39	2.57/3.97	1.50/2.61	2.41/4.97	2.86/4.11
ComPC [28]	1.29/1.77	1.22/1.51	2.18/3.23	1.78/1.90	1.32/1.65	1.09/1.34	1.46/1.76	1.84/2.13	1.65/1.90	1.56/2.11	1.97/2.70	1.99/3.10	1.61/2.09
Ours (Direct3D-S2)	0.96/1.42	1.11/1.63	1.15/1.67	1.21/1.63	0.87/1.26	0.88/1.24	1.55/1.97	0.96/1.40	1.11/1.58	1.03/1.51	0.96/1.43	1.51/2.12	1.11/1.57
Ours (TRELIS)	0.97/ 1.31	0.89/1.20	1.06/1.48	1.23/ 1.54	0.62/0.87	0.61/0.82	1.60/ 1.73	0.81/1.12	1.50/ 1.55	0.86/1.16	1.30/1.87	1.73/2.36	1.11/1.41

Table 3. Completion fidelity on ScanNet [13] and KITTI [18].

UCD ↓ / UHD ↓	ScanNet-Chair	ScanNet-Table	KITTI-Car
SVDFormer [90]	1.4 / 2.4	1.4 / 2.4	1.8 / 5.0
AdaPoinTr [81]	1.4 / 2.4	1.3 / 2.4	1.6 / 4.9
Shape-Inv [83]	4.0 / 9.3	6.6 / 11.0	5.3 / 12.5
P2C [12]	3.4 / 5.0	2.4 / 4.9	3.9 / 5.8
PCDreamer [62]	1.5 / 3.5	1.3 / 3.7	2.7 / 6.6
ComPC [28]	2.0 / 5.3	3.0 / 7.0	1.1 / 5.7
Ours	0.8 / 2.0	0.9 / 2.0	1.4 / 4.5

Table 4. Quantitative comparisons on our proposed Omni-Comp.

CD ↓ / EMD ↓	Single Scan	Random Crop	Semantic Part
SVDFormer [90]	5.32 / 6.84	5.35 / 6.73	5.66 / 7.19
AdaPoinTr [81]	5.27 / 6.87	5.33 / 6.68	5.59 / 7.23
Shape-Inv [83]	4.98 / 7.31	5.39 / 6.79	5.37 / 6.91
SDS-Comp [35]	5.42 / 6.27	5.12 / 6.35	5.60 / 6.75
PCDreamer [62]	3.32 / 4.74	3.78 / 4.55	3.88 / 4.98
ComPC [28]	4.24 / 4.61	5.48 / 5.95	6.37 / 6.18
Ours	2.21 / 3.15	2.60 / 3.31	3.30 / 3.68




Figure 5. Visual examples on ScanNet [13] and KITTI [18] real-world datasets, which only contain real scans of table, chair, and car, with very sparse points.

enabling it to maintain robust and superior performance across all categories and patterns. On average, we achieve improvements of 49.6% in CD and 39.4% in EMD over ComPC [28], demonstrating the effectiveness of our idea.

We provide some qualitative comparisons in Fig. 6. Specifically, baselines’ reliance on rendered 2D views that contain the complete object contour [28, 62] fails on complex inputs where contours are ill-posed (see Rows (c-d)). Furthermore, [28, 62] fail to complete the random crop and semantic part due to out-of-distribution categories and partial patterns (see Rows (a-b) and (e-f)). In stark contrast, our method proves robust across all these challenging cases,

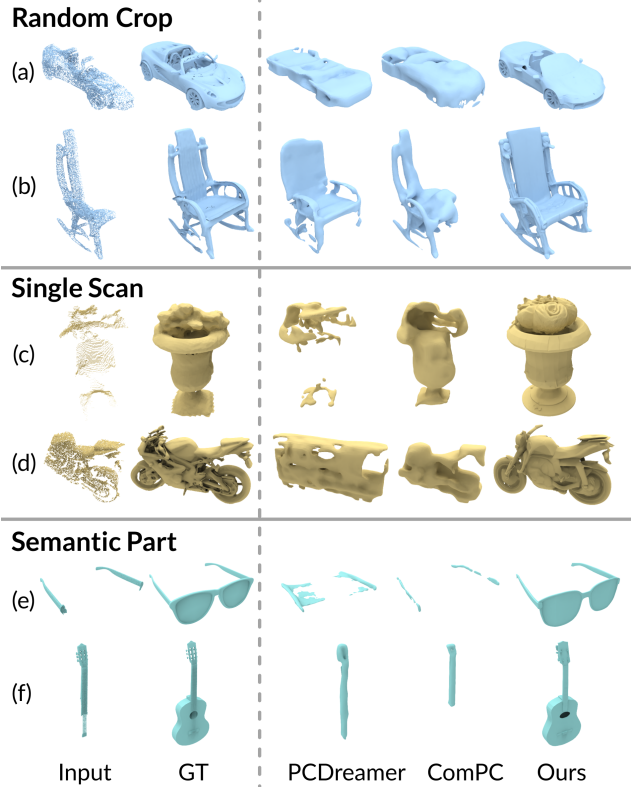


Figure 6. Qualitative comparisons on cases with different partial patterns from our Omni-Comp benchmark. Our approach produces more reasonable results than the latest methods [28, 62].

leveraging its pattern-agnostic 3D prior to reconstruct structurally plausible and geometrically detailed completions, demonstrating its strong generalizability.

Evaluation on Completion Diversity. Finally, we assess completion diversity against the latest generative-based baselines [28, 62] on Redwood [8] and synthetic [28, 43] datasets, reporting average metrics over 5 random runs per object for each method. As shown in Tab. 5, our approach achieves consistently better MMD and TMD. Notably, while PCDreamer [62] utilizes generative models, its reliance on varying multi-view depth maps produces limited geometric variation, resulting in low TMD scores. Our strong performance on both metrics verifies that our method can generate diverse shape completions, as shown in Fig. 7.

Table 5. Completion diversity evaluation on Redwood [8] and synthetic data [28, 43].

MMD ↓ / TMD ↑	Redwood	Synthetic
PCDreamer [62]	1.88 / 0.13	2.86 / 0.11
ComPC [28]	2.01 / 0.72	1.65 / 0.62
Ours (Direct3D-S2)	1.81 / 1.35	1.23 / 0.89
Ours (TRELIS)	1.62 / 0.99	1.20 / 0.94

Table 6. Ablation studies. The baseline uses only latent replacement for completion.

CD ↓ / EMD ↓	Redwood	Synthetic
(a) Baseline (latent replacement)	2.15 / 2.93	2.33 / 2.78
(b) Full w/o ERS	3.42 / 4.94	3.53 / 4.85
(c) Full w/o PNS	1.94 / 2.56	2.27 / 2.67
(d) Full w/o IAS	1.88 / 2.14	1.17 / 1.56
(e) IAS w/ 10 optimization steps	1.42 / 1.87	1.20 / 1.48
(f) Full pipeline (Ours)	1.42 / 1.84	1.11 / 1.41

4.3. Ablation Studies

We conduct ablation studies to validate the contribution of each component in our framework, and the quantitative and qualitative results are presented in Tab. 6 and Fig. 8.

Naive Baseline. We first establish a naive baseline that only utilizes the latents from the partial shape for replacement. Due to the latent discrepancy of the corresponding regions of partial and ground truth samples, it produces poor completion results both quantitatively and qualitatively, as shown in Tab. 6 (comparing Rows (a) and (f)) and Fig. 8.

Analysis of Key Components. We analyze the contribution of our three key components, Explicit Replacement Stage (ERS), Partial-aware Noise Schedule (PNS), and Implicit Alignment Stage (IAS), by ablating them individually.

- **w/o ERS:** Comparing Rows (b) and (f) in Tab. 6, removing the ERS causes the most significant performance drop. Without this explicit geometric conditioning, the model fails to preserve the input structure and instead hallucinates a novel shape based on the generative prior (e.g., a completely different horse pose in Fig. 8(c)). This confirms that ERS is critical for ensuring input fidelity.
- **w/o PNS:** Comparing Rows (c) and (f) in Tab. 6, disabling the PNS leads to noticeable stripe-like artifacts on the surface. We attribute this to the mismatch in the denoising process between the clean (known) partial input and the noisy (unknown) regions being generated. PNS is thus essential for seamlessly blending these two regions.
- **w/o IAS:** Comparing Rows (d) and (f) in Tab. 6, removing the IAS, which harmonizes the latent context via optimization, results in small holes and boundary inconsistencies, as highlighted by red boxes in Fig. 8. This demonstrates that ERS alone is insufficient; IAS is necessary to




Figure 7. Visual examples of the completion diversity from our method on the real-world data.

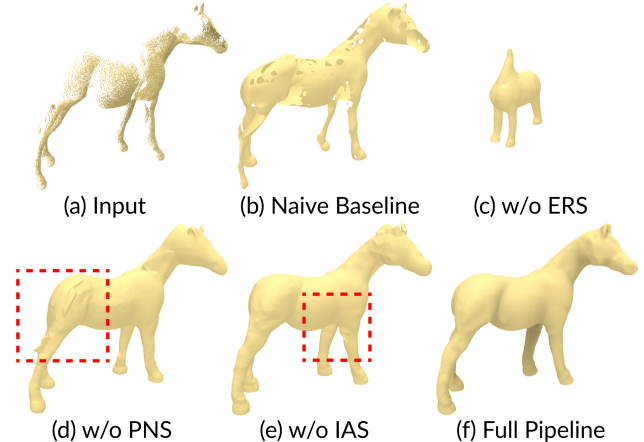


Figure 8. Visual comparison of the ablation studies. The red boxes highlight the artifacts and holes.

ensure the final latent representation is coherent.

- **Optimization steps in IAS:** We also do ablation about the optimization steps in IAS. As can be seen in Tab. 6 (e), using 10 steps for optimization at each timestep does not show obvious performance gain compared with (f).

The full pipeline successfully combines these components to achieve geometrically accurate and fidelity-preserving completions, with the best performance in Tab. 6. The horse, as shown in Fig. 8 (f), exhibits excellent geometric quality without noticeable artifacts.

5. Conclusion

We presented LaS-Comp, a zero-shot and category-agnostic framework for 3D shape completion that leverages the rich geometric priors of 3D foundation models. The method integrates two complementary components, an Explicit Replacement Stage and an Implicit Alignment Stage, to jointly ensure high-fidelity reconstruction and global geometric coherence. We also introduce Omni-Comp, a benchmark combining real-world scans and synthetic shapes with diverse partial patterns for comprehensive evaluation. Experiments on Omni-Comp and standard benchmarks show that our method generalizes well across diverse categories and partial patterns, surpassing existing approaches.

Acknowledgement This work was supported in part by the National Natural Science Foundation of China (NSFC) under grant 62372091 and in part by the Hainan Province Science and Technology Plan Project under Grant ZDYF2024(LALH)001.

References

- [1] Yingjie Cai, Kwan-Yee Lin, Chao Zhang, Qiang Wang, Xiaogang Wang, and Hongsheng Li. Learning a Structured Latent Space for Unsupervised Point Cloud Completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5543–5553, 2022. 2
- [2] Berk Çalli, Aaron Walsman, Arjun Singh, Siddhartha S. Srinivasa, Pieter Abbeel, and Aaron M. Dollar. Benchmarking in Manipulation Research: The YCB Object and Model Set and Benchmarking Protocols. In *International Conference on Advanced Robotics (ICAR)*, 2015. 5
- [3] Xiao Cao, Beibei Lin, Bo Wang, Zhiyong Huang, and Robby T Tan. Ssnerf: Sparse View Semi-supervised Neural Radiance Fields with Augmentation. *arXiv preprint arXiv:2408.09144*, 2024. 1
- [4] Xiao Cao, Beibei Lin, Bo Wang, Zhiyong Huang, and Robby T Tan. 3DOT: Texture Transfer for 3DGS Objects from a Single Reference Image. *arXiv preprint arXiv:2503.18853*, 2025. 2
- [5] Angel X. Chang, Thomas Funkhouser, Leonidas J. 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. 3
- [6] Xuelin Chen, Baoquan Chen, and Niloy J Mitra. Unpaired Point Cloud Completion on Real Scans using Adversarial Training. In *International Conference on Learning Representations (ICLR)*, 2020. 2, 5
- [7] Ying-Hao Chen, Jian Li, Shi-Peng Xie, and Qin Wang. Single-photon 3D Imaging with a Multi-stage Network. *Optics Express*, 30(16):29173–29188, 2022. 2
- [8] Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun. A Large Dataset of Object Scans. *arXiv:1602.02481*, 2016. 2, 5, 6, 7, 8
- [9] Jisheng Chu, Wenrui Li, Xingtao Wang, Kanglin Ning, Yidan Lu, and Xiaopeng Fan. Digging into Intrinsic Contextual Information for High-fidelity 3D Point Cloud Completion. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 2573–2581, 2025. 2
- [10] Lei Chu, Hao Pan, and Wenping Wang. Unsupervised Shape Completion via Deep Prior in the Neural Tangent Kernel Perspective. *ACM Transactions on Graphics*, 40(3):1–17, 2021. 2
- [11] Ruihang Chu, Enze Xie, Shentong Mo, Zhenguo Li, Matthias Nießner, Chi-Wing Fu, and Jiaya Jia. DiffComplete: Diffusion-based Generative 3D Shape Completion. In *Conference on Neural Information Processing Systems (NeurIPS)*, pages 75951–75966, 2023. 2
- [12] Ruikai Cui, Shi Qiu, Saeed Anwar, Jiawei Liu, Chaoyue Xing, Jing Zhang, and Nick Barnes. P2C: Self-supervised Point Cloud Completion from Single Partial Clouds. In *IEEE International Conference on Computer Vision (ICCV)*, pages 14351–14360, 2023. 2, 5, 6, 7
- [13] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2443, 2017. 5, 6, 7
- [14] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape Completion using 3D-Encoder-Predictor CNNs and Shape Synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5868–5877, 2017. 1, 2
- [15] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-XL: A Universe of 10M+ 3D Objects. *arXiv preprint arXiv:2307.05663*, 2023. 2
- [16] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A Universe of Annotated 3D Objects. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13142–13153, 2023. 2
- [17] Keyu Du, Jingyu Hu, Haipeng Li, Hao Xu, Haibin Huang, Chi-Wing Fu, and Shuaicheng Liu. Hierarchical Neural Semantic Representation for 3D Semantic Correspondence. In *Proceedings of SIGGRAPH Asia*, pages 1–11, 2025. 2
- [18] Andreas Geiger. Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page 3354–3361, 2012. 5, 6, 7
- [19] Xianglong He, Zi-Xin Zou, Chia-Hao Chen, Yuan-Chen Guo, Ding Liang, Chun Yuan, Wanli Ouyang, Yan-Pei Cao, and Yangguang Li. SparseFlex: High-Resolution and Arbitrary-Topology 3D Shape Modeling. In *IEEE International Conference on Computer Vision (ICCV)*, 2025. 1
- [20] Jonathan Ho and Tim Salimans. Classifier-free Diffusion Guidance. *arxiv preprint arxiv:2207.12598*, 2022. 3
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Conference on Neural Information Processing Systems (NeurIPS)*, pages 6840–6851, 2020. 2
- [22] Sangmin Hong, Mohsen Yavartanoo, Reyhaneh Neshatavar, and Kyoung Mu Lee. ACL-SPC: Adaptive Closed-Loop system for Self-supervised Point Cloud Completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9435–9444, 2023. 2
- [23] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large Reconstruction Model for Single Image to 3D. In *International Conference on Learning Representations (ICLR)*, 2024. 3
- [24] Jingyu Hu*, Ka-Hei Hui*, Zhengzhe Liu, Hao Zhang, and Chi-Wing Fu. CLIPXPlore: Coupled CLIP and Shape Spaces for 3D Shape Exploration. In *Proceedings of SIGGRAPH Asia*, pages 1–12, 2023. 2
- [25] Jingyu Hu, Ka-Hei Hui, Zhengzhe Liu, Ruihui Li, and Chi-Wing Fu. Neural Wavelet-domain Diffusion for 3D Shape Generation, Inversion, and Manipulation. *ACM Transactions on Graphics (TOG)*, 42(6), 2024. 2
- [26] Jingyu Hu, Ka-Hei Hui, Zhengzhe Liu, Hao Zhang, and Chi-Wing Fu. CNS-Edit: 3D Shape Editing via Coupled Neural

- Shape Optimization. In *Proceedings of SIGGRAPH*, pages 1–12, 2024. 2
- [27] Jingyu Hu, Bin Hu, Ka-Hei Hui, Haipeng Li, Zhengzhe Liu, Daniel Cohen-Or, and Chi-Wing Fu. PEGAsus: 3D Personalization of Geometry and Appearance. *arXiv preprint arXiv:2602.08198*, 2026. 2
- [28] Tianxin Huang, Zhiwen Yan, Yuyang Zhao, and Gim H Lee. ComPC: Completing a 3D Point Cloud with 2D Diffusion Priors. In *International Conference on Learning Representations (ICLR)*, pages 51765–51784, 2025. 1, 2, 3, 5, 6, 7, 8
- [29] Zitian Huang, Yikuan Yu, Jiawen Xu, Feng Ni, and Xinyi Le. PF-Net: Point Fractal Network for 3D Point Cloud Completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7662–7670, 2020. 2
- [30] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural Wavelet-domain Diffusion for 3D Shape Generation. In *Proceedings of SIGGRAPH Asia*, pages 1–9, 2022. 2
- [31] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural Template: Topology-aware Reconstruction and Disentangled Generation of 3D Meshes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18572–18582, 2022. 2
- [32] Team Hunyuan3D, Bowen Zhang, Chunchao Guo, Haolin Liu, Hongyu Yan, Huiwen Shi, Jingwei Huang, Junlin Yu, Kunhong Li, Penghao Wang, et al. Hunyuan3D-Omni: A Unified Framework for Controllable Generation of 3D Assets. *arXiv preprint arXiv:2509.21245*, 2025. 1
- [33] Shun Iwase, Katherine Liu, Vitor Guizilini, Adrien Gaidon, Kris Kitani, Rareş Ambruş, and Sergey Zakharov. Zero-Shot Multi-Object Scene Completion. In *ECCV*, 2024. 1
- [34] Shun Iwase, Muhammad Zubair Irshad, Katherine Liu, Vitor Guizilini, Robert Lee, Takuya Ikeda, Amma Ayako, Koichi Nishiwaki, Kris Kitani, Rareş Ambruş, and Sergey Zakharov. ZeroGrasp: Zero-Shot Shape Reconstruction Enabled Robotic Grasping. In *CVPR*, 2025. 1
- [35] Yoni Kasten, Ohad Rahamim, and Gal Chechik. Point Cloud Completion with Pretrained Text-to-image Diffusion Models. In *Conference on Neural Information Processing Systems (NeurIPS)*, pages 12171–12191, 2023. 1, 3, 5, 6, 7
- [36] Jihun Kim, Hyeokjun Kwon, Yunseo Yang, and Kuk-Jin Yoon. Learning Point Cloud Completion without Complete Point Clouds: A Pose-aware Approach. In *IEEE International Conference on Computer Vision (ICCV)*, pages 14157–14167, 2023. 2
- [37] Jeongsol Kim, Bryan Sangwoo Kim, and Jong Chul Ye. FlowDPS : Flow-Driven Posterior Sampling for Inverse Problems. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12328–12337, 2025. 3
- [38] An Li, Zhe Zhu, and Mingqiang Wei. GenPC: Zero-shot Point Cloud Completion via 3D Generative Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1308–1318, 2025. 1, 3, 5, 6
- [39] Shanshan Li, Pan Gao, Xiaoyang Tan, and Mingqiang Wei. ProxyFormer: Proxy Alignment Assisted Point Cloud Completion with Missing Part Sensitive Transformer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9466–9475, 2023. 2
- [40] Shuwei Li, Lei Tan, and Robby T. Tan. Bridging Day and Night: Target-Class Hallucination Suppression in Unpaired Image Translation, 2026. 1
- [41] Yanguang Li, Zi-Xin Zou, Zexiang Liu, Dehu Wang, Yuan Liang, Zhipeng Yu, Xingchao Liu, Yuan-Chen Guo, Ding Liang, Wanli Ouyang, et al. TripoSG: High-Fidelity 3D Shape Synthesis using Large-Scale Rectified Flow Models. *arXiv preprint arXiv:2502.06608*, 2025. 1
- [42] Zijun Li, Hongyu Yan, Shijie Li, Kunming Luo, Li Lu, Xulei Yang, and Weisi Lin. DiffPCN: Latent Diffusion Model Based on Multi-view Depth Images for Point Cloud Completion. *arXiv preprint arXiv:2509.23723*, 2025. 2
- [43] Yaron Lipman, David Levin, and Daniel Cohen-Or. Green Coordinates. *ACM Transactions on Graphics*, 27(3):1–10, 2008. 5, 6, 7, 8
- [44] Mengya Liu, Ajad Chhatkuli, Janis Postels, Luc Van Gool, and Federico Tombari. Self-supervised Shape Completion via Involution and Implicit Correspondences. In *European Conference on Computer Vision (ECCV)*, pages 212–229, 2024. 2
- [45] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot One Image to 3D Object. In *IEEE International Conference on Computer Vision (ICCV)*, pages 9298–9309, 2023. 3
- [46] Zhengzhe Liu, Jingyu Hu, Ka-Hei Hui, Xiaojuan Qi, Daniel Cohen-Or, and Chi-Wing Fu. Exim: A hybrid explicit-implicit representation for text-guided 3d shape generation. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 42(6): 1–12, 2023. 2
- [47] Zishan Liu, Zecong Tang, RuoCheng Wu, Xinzhe Zheng, Jingyu Hu, Ka-Hei Hui, Haoran Xie, Bo Dai, and Zhengzhe Liu. Imagine a city: Citygenagent for procedural 3d city generation. *arXiv preprint arXiv:2602.05362*, 2026. 2
- [48] Shitong Luo and Wei Hu. Diffusion Probabilistic Models for 3D Point Cloud Generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, 2021. 2
- [49] Zhaoyang Lyu, Zhifeng Kong, Xudong Xu, Liang Pan, and Dahua Lin. A Conditional Point Diffusion-Refinement Paradigm for 3D Point Cloud Completion. In *International Conference on Learning Representations (ICLR)*, 2022. 2
- [50] Mark Pauly, Niloy J Mitra, Joachim Giesen, Markus H Gross, and Leonidas J Guibas. Example-Based 3D Scan Completion. In *Eurographics Symposium on Geometry Processing (SGP)*, 2005. 1
- [51] Mark Pauly, Niloy J Mitra, Johannes Wallner, Helmut Pottmann, and Leonidas J Guibas. Discovering Structural Regularity in 3D Geometry. In *ACM Transactions on Graphics (SIGGRAPH)*, pages 1–11, 2008. 1
- [52] Songyou Peng, Chiyu “Max” Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape As Points: A Differentiable Poisson Solver. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 5, 6

- [53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 3
- [54] Chao-Hui Shen, Hongbo Fu, Kang Chen, and Shi-Min Hu. Structure Recovery by Part Assembly. *ACM Transactions on Graphics (SIGGRAPH)*, 31(6):1–11, 2012. 1
- [55] Minhyuk Sung, Vladimir G Kim, Roland Angst, and Leonidas Guibas. Data-driven Structural Priors for Shape Completion. *ACM Transactions on Graphics (SIGGRAPH)*, 34(6):1–11, 2015. 1
- [56] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. LGM: Large Multi-View Gaussian Model for High-Resolution 3D Content Creation. In *European Conference on Computer Vision (ECCV)*, pages 1–18, 2024. 3
- [57] Tencent Hunyuan3D Team. Hunyuan3D 2.1: From Images to High-Fidelity 3D Assets with Production-Ready PBR Material. *arXiv preprint arXiv:2506.15442*, 2025. 1
- [58] Sebastian Thrun and Ben Wegbreit. Shape from Symmetry. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1824–1831, 2005. 1
- [59] Jun Wang, Ying Cui, Dongyan Guo, Junxia Li, Qingshan Liu, and Chunhua Shen. PointAttN: You only Need Attention for Point Cloud Completion. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 5472–5480, 2024. 2
- [60] Penghao Wang, Yiyang He, Xin Lv, Yukai Zhou, Lan Xu, Jingyi Yu, and Jiayuan Gu. PartNeXt: A Next-Generation Dataset for Fine-Grained and Hierarchical 3D Part Understanding. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2025. 5
- [61] Xiaogang Wang, Marcelo H Ang Jr, and Gim Hee Lee. Cascaded Refinement Network for Point Cloud Completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 790–799, 2020. 2
- [62] Guangshun Wei, Yuan Feng, Long Ma, Chen Wang, Yuanfeng Zhou, and Changjian Li. PCDreamer: Point Cloud Completion Through Multi-view Diffusion Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27243–27253, 2025. 2, 5, 6, 7, 8
- [63] Xin Wen, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. Cycle4Completion: Unpaired Point Cloud Completion using Cycle Transformation with Missing Region Coding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13080–13089, 2021. 2
- [64] Xin Wen, Peng Xiang, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. PMP-Net: Point Cloud Completion by learning Multi-step Point Moving Paths. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7443–7452, 2021. 2
- [65] Rundi Wu, Xuelin Chen, Yixin Zhuang, and Baoquan Chen. Multimodal Shape Completion via Conditional Generative Adversarial Networks. In *European Conference on Computer Vision (ECCV)*, pages 281–296, 2020. 2
- [66] Shuang Wu, Youtian Lin, Feihu Zhang, Yifei Zeng, Yikang Yang, Yajie Bao, Jiachen Qian, Siyu Zhu, Philip Torr, Xun Cao, and Yao Yao. Direct3D-S2: Gigascale 3D Generation Made Easy with Spatial Sparse Attention. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2025. 1, 2, 3, 5
- [67] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3D Latents for Scalable and Versatile 3D Generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21469–21480, 2025. 1, 2, 3, 5
- [68] Peng Xiang, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. SnowflakeNet: Point Cloud Completion by Snowflake Point Deconvolution with Skip-Transformer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5499–5509, 2021. 2
- [69] Chulin Xie, Chuxin Wang, Bo Zhang, Hao Yang, Dong Chen, and Fang Wen. Style-based Point Generator with Adversarial Rendering for Point Cloud Completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4619–4628, 2021. 2
- [70] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Jiageng Mao, Shengping Zhang, and Wenxiu Sun. GRNet: Gridding Residual Network for Dense Point Cloud Completion. In *European Conference on Computer Vision (ECCV)*, pages 365–381, 2020. 2
- [71] Chongyang Xu, Shen Cheng, Haipeng Li, Haoqiang Fan, Ziliang Feng, and Shuaicheng Liu. HeRo: Hierarchical 3D Semantic Representation for Pose-aware Object Manipulation, 2026. 2
- [72] Chongyang Xu, Haipeng Li, Shen Cheng, Jingyu Hu, Haoqiang Fan, Ziliang Feng, and Shuaicheng Liu. Action-Geometry prediction with 3D Geometric Prior for Bimanual Manipulation, 2026.
- [73] Hao Xu, Haipeng Li, Yinqiao Wang, Shuaicheng Liu, and Chi-Wing Fu. HandBooster+: Boosting 3D Hand-Mesh Reconstruction From Data Synthesis to Progressive Multi-Hypothesis Aggregation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(12):11201–11219, 2025. 2
- [74] Jingjing Wang, Hang Du, Zhihong Wu, Di Xie, Shiliang Pu, Li Lu, Xuejun Yan, Hongyu Yan. FBNet: Feedback Network for Point Cloud Completion. In *European Conference on Computer Vision*, 2022. 2
- [75] Hongyu Yan, Zijun Li, Kunming Luo, Li Lu, and Ping Tan. SymmCompletion: High-Fidelity and High-Consistency Point Cloud Completion with Symmetry Guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9094–9102, 2025. 2
- [76] Weilong Yan, Robby T. Tan, Bing Zeng, and Shuaicheng Liu. Deep Homography Mixture for Single Image Rolling Shutter Correction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9868–9877, 2023. 1
- [77] Weilong Yan, Ming Li, Haipeng Li, Shuwei Shao, and Robby T. Tan. Synthetic-to-Real Self-supervised Robust

- Depth Estimation via Learning with Motion and Structure Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21880–21890, 2025. 1
- [78] Xingguang Yan, Liqiang Lin, Niloy J. Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. ShapeFormer: Shapelet Transformer for Multivariate Time Series Classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6239–6249, 2022. 2
- [79] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T. Freeman, and Jiajun Wu. WonderWorld: Interactive 3D Scene Generation from a Single Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5916–5926, 2025. 1
- [80] Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. PoinTr: Diverse Point Cloud Completion with Geometry-Aware Transformers. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12498–12507, 2021. 2
- [81] Xumin Yu, Yongming Rao, Ziyi Wang, Jiwen Lu, and Jie Zhou. AdaPoinTr: Diverse Point Cloud Completion with Adaptive Geometry-Aware Transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12): 14114–14130, 2023. 5, 6, 7
- [82] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. PCN: Point Completion Network. In *International Conference on 3D Vision (3DV)*, pages 728–737, 2018. 1, 2
- [83] Junzhe Zhang, Xinyi Chen, Zhongang Cai, Liang Pan, Haiyu Zhao, Shuai Yi, Chai Kiat Yeo, Bo Dai, and Chen Change Loy. Unsupervised 3D Shape Completion through GAN Inversion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 5, 6, 7
- [84] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. CLAY: A Controllable Large-scale Generative Model for Creating High-quality 3D Assets. *ACM Transactions on Graphics (SIGGRAPH)*, 43(4):1–20, 2024. 1
- [85] Mengtan Zhang, Yi Feng, Qijun Chen, and Rui Fan. DCPI-Depth: Explicitly Infusing Dense Correspondence Prior to Unsupervised Monocular Depth Estimation. *IEEE Transactions on Image Processing*, 34:4258–4272, 2025. 1
- [86] Wenxiao Zhang, Qingan Yan, and Chunxia Xiao. Detail Preserved Point Cloud Completion via Separated Feature Aggregation. In *European Conference on Computer Vision (ECCV)*, pages 512–528, 2020. 2
- [87] Xindan Zhang, Weilong Yan, Yufei Shi, Xuerui Qiu, Tao He, Ying Li, Ming Li, and Hehe Fan. 4DPC²hat: Towards Dynamic Point Cloud Understanding with Failure-Aware Bootstrapping. *arXiv preprint arXiv:2602.03890*, 2026. 2
- [88] Zibo Zhao, Zeqiang Lai, Qingxiang Lin, Yunfei Zhao, Haolin Liu, Shuhui Yang, Yifei Feng, Mingxin Yang, Sheng Zhang, Xianghui Yang, et al. Hunyuan3D 2.0: Scaling Diffusion Models for High Resolution Textured 3D Assets Generation. *arXiv preprint arXiv:2501.12202*, 2025. 2
- [89] Linqi Zhou, Yilun Du, and Jiajun Wu. 3D Shape Generation and Completion through Point-Voxel Diffusion. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5826–5835, 2021. 2
- [90] Zhe Zhu, Honghua Chen, Xing He, Weiming Wang, Jing Qin, and Mingqiang Wei. SVDFormer: Complementing Point Cloud via Self-view Augmentation and Self-structure Dual-Generator. In *IEEE International Conference on Computer Vision (ICCV)*, pages 14508–14518, 2023. 2, 5, 6, 7