SymmFusion: Injecting Multi-View-Consistent Generative Priors into Symmetry-Guided Point Cloud Completion

Supplementary Material

A. Supplementary Material

For all multi-view qualitative figures in this supplementary, we export the predicted point clouds as `.ply` files and visualize them in *CloudCompare*. We then capture screenshots from fixed camera viewpoints to ensure a consistent and interpretable comparison across methods. This avoids plot-style rendering artifacts and provides clearer geometric inspection.

A.1. Additional Multi-View Qualitative Results

Our main claim is that multi-view priors improve completion particularly when the partial scan lacks reliable symmetric evidence. To make this effect visually verifiable, we render the same completed point cloud from multiple viewpoints and inspect stability under rotation.

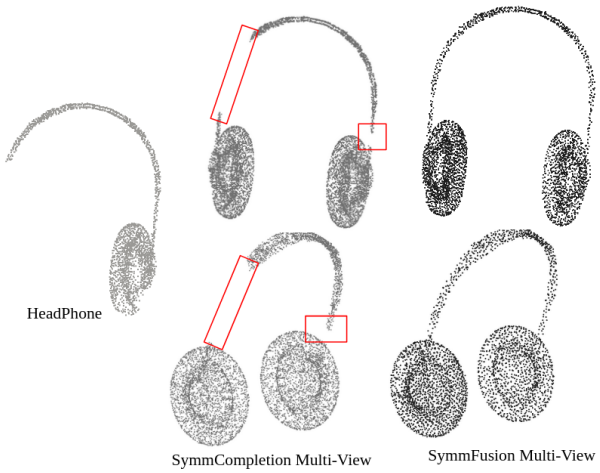

A1.Figure 1. **Multi-view qualitative comparison on ShapeNet55 (Headphone).** SymmCompletion exhibits view-dependent artifacts and incomplete regions (highlighted in red), while SymmFusion produces more coherent thin structures and fewer isolated outliers across views.

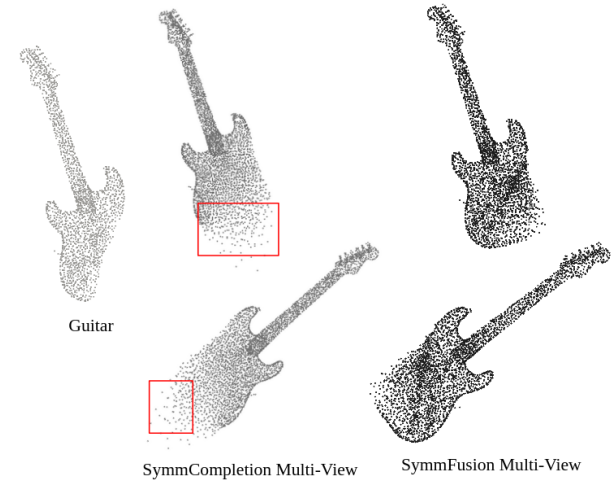

A1.Figure 2. **Multi-view qualitative comparison on ShapeNet55 (Guitar).** SymmFusion yields denser and more geometrically consistent completions across viewpoints. SymmCompletion shows missing or unstable surfaces and scattered artifacts that become more apparent under rotation (red boxes).

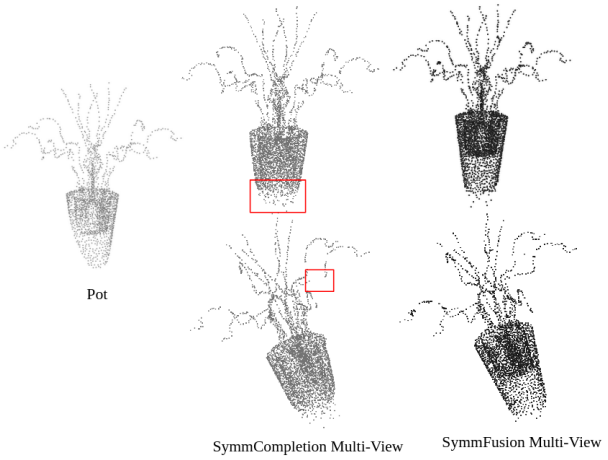

A1.Figure 3. **Multi-view qualitative comparison on ShapeNet55 (Pot).** SymmFusion reduces view-dependent artifacts and produces more coherent surfaces, while SymmCompletion contains inconsistent clusters and missing regions (red boxes).

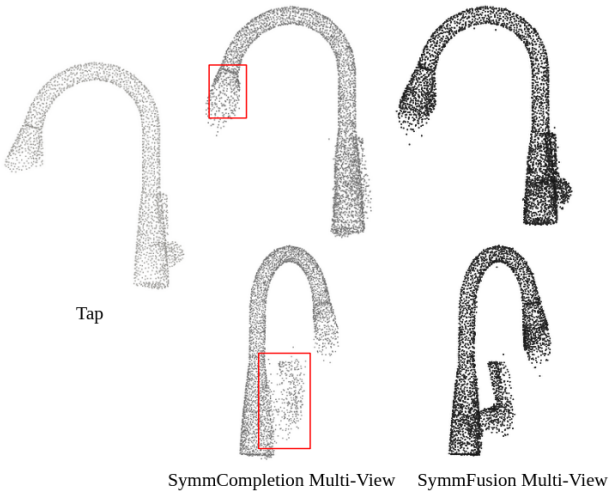

A1.Figure 4. **Multi-view qualitative comparison on ShapeNet55 (Tap).** SymmFusion reduces view-dependent artifacts and produces more coherent surfaces, while SymmCompletion contains inconsistent clusters and missing regions (red boxes).

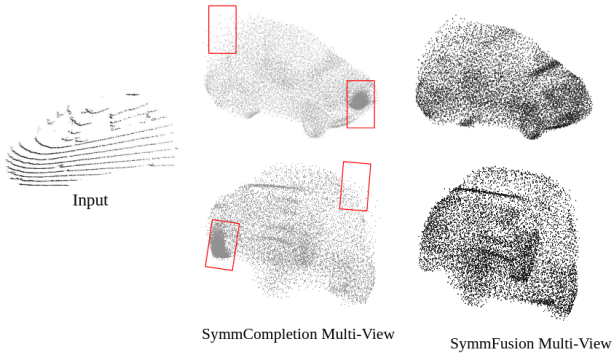

A1.Figure 5. **Multi-view qualitative comparison on KITTI Cars.** We visualize the same completed car from multiple viewpoints. SymmFusion reduces view-dependent artifacts and produces more coherent surfaces, while SymmCompletion contains inconsistent clusters and missing regions (red boxes). These qualitative examples complement proxy metrics reported in the main paper.

A.2. Multi-View Prior Comparison: Shap-E vs TripoSR

Our ablation indicates that multi-view priors help most when they encode cross-view geometric consistency. To illustrate why TripoSR priors outperform alternatives in our setting, we visualize the multi-view reconstructions produced by Shap-E and TripoSR from the same input images.

A2.Figure 1. **Multi-view prior visual comparison between Shap-E and TripoSR.** For each input image (center), we show representative novel views from Shap-E (left) and TripoSR (right). TripoSR tends to maintain more stable geometry across viewpoints, whereas Shap-E can exhibit view-to-view inconsistency and less reliable structure. This qualitative behavior aligns with our quantitative prior-source ablation.

A.3. Failure Cases and Difficulty Split Analysis on ShapeNet55

ShapeNet55 provides difficulty splits that reflect the severity of partiality. To complement the quantitative results in the main paper, we visualize representative examples from the easy, medium, and hard settings to highlight where completion succeeds and where it still fails. All views are rendered by loading the saved `.ply` point clouds in *CloudCompare* and capturing fixed camera screenshots.

Observations: In the easy setting, the partial input retains most of the global structure, and the completion primarily fills small holes or restores minor missing surfaces. In the medium setting, larger missing regions appear, and the model must infer extended parts of the object. In the hard setting, failure cases become more frequent: the partial observation may contain only a small visible subset, which can lead to incomplete recovery of thin structures, local overfilling, or missing appendages.

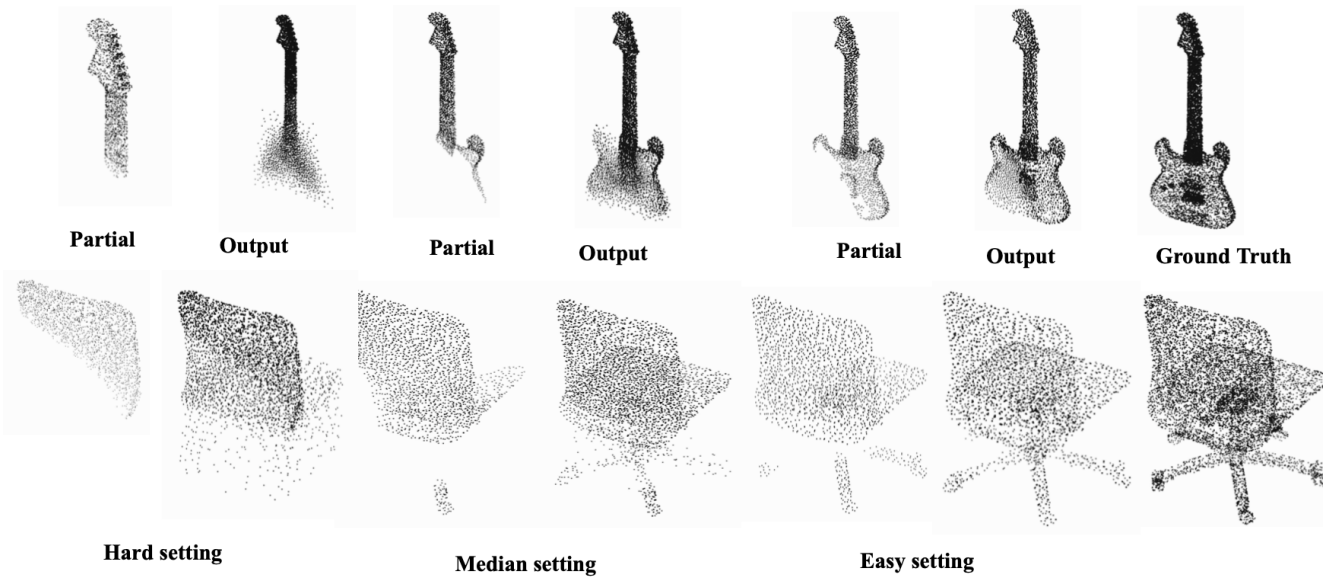

A3.Figure 1. **Qualitative comparison across ShapeNet55 difficulty settings.** We show representative examples from hard, medium, and easy settings. Each group visualizes the partial input, the predicted completion, and the ground truth. As partiality becomes more severe (hard setting), the input provides fewer geometric cues and completion becomes more underconstrained.