G. QAL: Supplementary Material

We provide extended qualitative comparisons on the **MVP** dataset [19] using the ECG [18] backbone trained with different loss functions. Figure 8 shows representative success cases. Across diverse object categories, QAL yields completions with *fewer spurious points*, better preservation of *thin structures* such as chair legs and lamp posts, and sharper recovery of *fine geometric details* compared to CD, EMD, DCD, HCD, and InfoCD. These improvements demonstrate QAL's ability to balance recall and precision, producing reconstructions that are both complete and clean.

To complement these strengths, Figure 7 presents challenging failure cases. While QAL generally outperforms existing losses, objects with highly irregular or ambiguous geometry remain difficult: elongated parts may still be under-reconstructed, flat surfaces can appear distorted, and spurious points may emerge under severe occlusion. These examples highlight open challenges in reliably capturing fine details from severely incomplete observations and suggest promising directions for future work on stronger priors and shape reasoning.

H. Per-Category Coverage Analysis

To provide a more detailed view of reconstruction quality, we report **per-category coverage** across all models and loss functions in Table 7. While the main paper highlights average performance, these results reveal finer trends within individual object categories of the MVP dataset.

Several observations emerge. First, **QAL consistently achieves the highest coverage across most categories and models**, often improving upon CD, HCD, and InfoCD. For instance, on *airplane*, *cabinet*, and *table* objects, QAL provides clear gains over both CD and InfoCD. Second, although EMD enforces one-to-one correspondences, its coverage remains markedly lower across categories, underscoring its tendency to produce over-smoothed reconstructions. Finally, categories with thin or fine-grained structures such as *lamp* and *chair* show the greatest benefit from QAL, where balancing recall and precision is critical to capturing delicate geometric details.

These category-level results complement the aggregate analysis and further validate the robustness of QAL: it improves completeness of reconstructions without sacrificing precision, even on challenging object classes.

I. Evaluation Protocol and Trade-off Analysis

Choice of ϵ for Coverage Evaluation We report coverage at a fixed threshold of $\epsilon = 0.03$. This value was selected based on the average nearest-neighbor spacing in MVP samples at 2048 points, which is approximately 0.0156. Setting ϵ to twice this mean spacing provides a strict enough tolerance to penalize missing thin structures, while remaining flexible enough to ignore minor surface noise or small alignment discrepancies. Empirically, thresholds lower than 0.02 excessively penalized completions for surface noise, while thresholds above 0.05 blurred differences between methods.

Reporting Trade-offs Recall-oriented training with QAL naturally induces a trade-off between recall and precision. To highlight this, we present grouped bar plots in our ablations, showing Cov (↑), the complement of spurious points $\overline{\text{SP}}$ (↑), and Chamfer Distance (CD; ↓) side-by-side for each hyperparameter setting. This visualization makes trade-offs explicit without collapsing into a single score. While we report an aggregate quality metric in the main paper for completeness, we emphasize that per-metric reporting provides clearer insight into how QAL balances recall and precision.

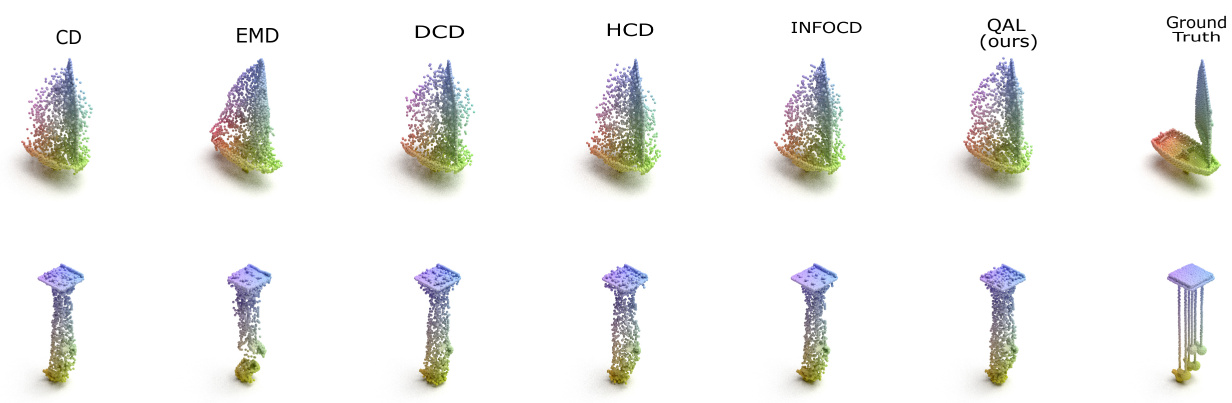

Figure 7. Failure cases of point cloud completion on challenging **MVP** samples [19]. Each row shows point cloud completions generated by ECG [18] trained with CD, EMD, DCD, HCD, InfoCD, and our proposed QAL, followed by the ground truth. While QAL generally outperforms existing losses, certain instances with complex geometry remain difficult: thin or elongated parts may be under-reconstructed, and spurious points can still appear around highly ambiguous regions. These examples highlight open challenges in reliably capturing fine details under severe incompleteness. Input partial clouds are omitted for space; all methods use identical inputs.

Figure 8. Qualitative comparisons of point cloud completion across 8 representative object categories from the **MVP** dataset [19]. Each row shows completions generated by ECG [18] trained with CD, EMD, DCD, HCD, InfoCD, and our proposed QAL, followed by the ground truth. QAL consistently recovers thin structures (e.g., chair legs, lamp posts) and suppresses spurious artifacts, producing reconstructions that are both more complete and geometrically faithful compared to existing losses. Input partial clouds are omitted for space; all methods use identical inputs.

Table 7. Per-category coverage (%) across models and loss functions on the MVP dataset. QAL consistently achieves the highest coverage, with best results highlighted in bold.(Validation ϵ=0.03)

Model	Loss	Airplane	Cabinet	Car	Chair	Lamp	Sofa	Table	Boat	Average
PCN	**CD**	77.0	55.3	62.3	50.9	56.1	53.3	56.4	67.1	59.8
	EMD	34.5	13.2	20.8	17.2	11.0	27.0	30.7	14.2	21.1
	DCD	78.5	**61.2**	67.0	54.0	56.6	58.7	59.6	70.4	63.2
	HCD	77.0	56.0	63.1	50.3	54.3	53.3	57.3	67.3	59.8
	InfoCD	77.6	58.2	65.9	55.1	59.3	57.1	58.9	70.1	62.8
	QAL	**79.7**	61.1	**68.4**	**58.1**	**62.2**	**60.7**	**62.6**	**74.5**	**65.9**
TOPNET	**CD**	75.0	49.7	57.0	45.6	49.8	47.1	52.0	65.7	55.2
	HCD	75.5	48.7	56.0	45.9	53.0	47.8	52.4	67.0	55.8
	InfoCD	74.9	49.7	57.9	46.0	55.2	47.4	52.8	66.2	56.3
	QAL	**77.7**	**51.1**	**60.0**	**49.5**	**57.9**	**50.7**	**55.3**	**70.4**	**59.1**
ECG	**CD**	85.7	56.0	60.2	61.4	71.5	56.7	68.0	73.2	66.6
	EMD	83.1	**59.5**	61.5	61.1	71.2	**59.8**	66.9	72.8	67.0
	DCD	86.0	58.6	62.8	62.7	72.1	58.6	69.8	74.5	68.1
	HCD	85.4	57.0	60.3	61.7	71.5	56.9	68.7	72.8	66.8
	InfoCD	85.3	57.0	61.8	62.1	71.3	57.7	66.7	74.0	67.0
	QAL	**87.2**	59.3	**63.1**	**63.4**	**73.3**	59.2	**70.0**	**75.8**	**68.9**
Seedformer	**CD**	86.1	67.4	66.4	69.4	77.8	66.6	72.9	77.7	73.0
	InfoCD	76.7	38.4	40.1	50.7	64.0	41.5	56.9	58.8	53.4
	QAL	**89.5**	**70.0**	**73.5**	**73.8**	**81.1**	**72.4**	**77.4**	**85.1**	**77.8**

Table 8. Completion Networks trained with CD (L2-CD $\times 1e^3$), HCD, InfoCD and **QAL** for $\epsilon = 0.03$. Cov, $\overline{SP}$ and Quality metrics are scaled as percentages.

Method (Network+Loss)	CD (↓)	F1 (↑)	Cov. (↑)	$\overline{SP}$ (↑)	Quality (↑)
PCN + CD	17.0	0.30	60.8	**64.6**	62.7
PCN + EMD	44.0	0.11	21.9	41.3	31.6
PCN + HCD	17.2	**0.31**	60.4	64.5	62.4
PCN + INFO	**16.9**	0.30	63.6	62.7	63.2
PCN + QAL	16.9	**0.31**	**64.8**	63.1	**64.0**
TOPNET + CD	**17.7**	**0.28**	56.2	**60.9**	58.6
TOPNET + HCD	17.8	**0.28**	56.8	60.1	58.4
TOPNET + INFO	**17.7**	**0.28**	57.4	60.3	58.8
TOPNET + QAL	17.8	0.27	**60.5**	57.9	**59.2**
ECG + CD	6.2	**0.82**	90.9	80.6	85.8
ECG + EMD	11.8	0.49	82.8	77.3	80.0
ECG + DCD	6.1	0.80	91.0	79.9	85.4
ECG + HCD	6.0	**0.82**	**91.6**	**81.0**	**86.3**
ECG + INFO	6.6	0.81	89.6	79.7	84.6
ECG + QAL	**5.8**	0.81	**91.6**	80.9	86.2
SEEDFORMER + CD	9.4	0.62	87.8	87.8	**87.8**
SEEDFORMER + INFO	**5.9**	0.65	75.1	**95.8**	85.4
SEEDFORMER + QAL	8.4	**0.66**	**88.3**	87.1	87.7

J. Effect of Loss on Reconstruction.

Using the same trained models from Section 4.2 (Table 1), we further evaluate reconstruction performance when ground-truth point clouds are provided as input on the MVP dataset (1200 samples). This controlled setting removes the challenge of inferring missing regions and highlights how the choice of loss function alone shapes the fidelity of the reconstructed geometry. Table 8 shows that while CD and F1 remain standard, they often provide incomplete or even misleading signals. For example, in PCN, QAL improves Coverage by +4.5 pts over CD (65.3 vs. 60.8) while CD and F1 remain nearly unchanged, masking the recovery of fine details. In TOPNet, CD and F1 suggest negligible differences across losses, yet Coverage and Quality reveal that QAL yields denser and more uniform reconstructions (+4.3 pts Coverage over CD). In ECG, HCD and QAL achieve near-identical Cov (91.6) and $\overline{\mathrm{SP}}$ ($\sim$81), yet CD differs (6.0 vs. 5.8), illustrating that CD alone can produce misleading signals even when recall and precision are effectively the same. Finally, in SeedFormer, F1 stays nearly flat across CD, InfoCD, and QAL (0.62–0.66), while Coverage rises from 87.8 (CD) to 88.3 (QAL), and Quality remains competitive despite InfoCD reporting the lowest CD.

This confirms that CD and F1 alone cannot capture the nuanced trade-offs between recall and precision. To complement these quantitative results, we also include 3–4 representative visualizations per model (Figs. 9 to 12). These examples qualitatively confirm the same trends: QAL consistently recovers sharper structures and reduces holes relative to CD/EMD, while avoiding the spurious clusters sometimes induced by InfoCD or HCD. Such visualizations reinforce that the geometric advantages highlighted by Cov and $\overline{\mathrm{SP}}$ are clearly visible in the reconstructions.

These cases highlight that CD favors small average distances, while F1 is threshold-sensitive and may improve even if reconstructions cluster unnaturally, leading to contradictory interpretations. By contrast, the additional dimensions introduced here—Cov and $\overline{\mathrm{SP}}$—capture complementary geometric aspects: whether fine details are covered and whether spurious points are suppressed. Together, they resolve the ambiguity of CD/F1 and provide a more faithful picture of reconstruction quality, confirming that QAL not only benefits partial-to-complete completion but also enhances detail preservation in autoencoding scenarios.

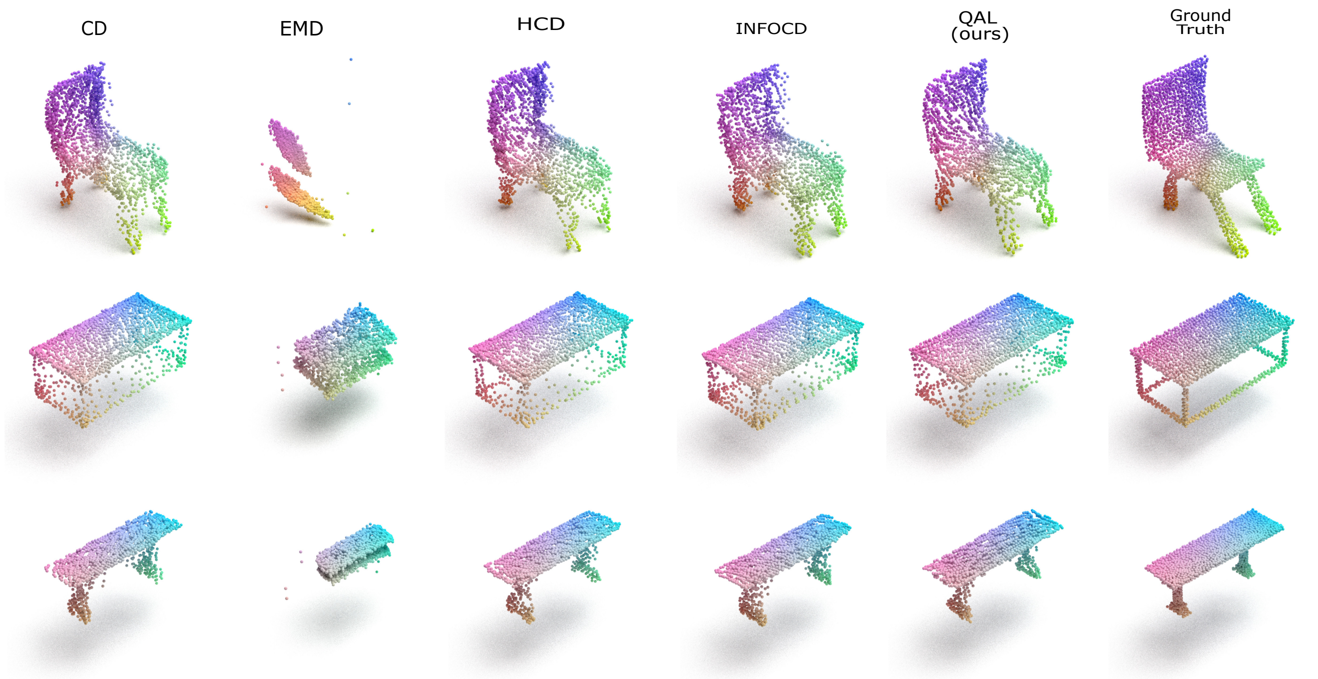

Figure 9. Representative reconstructions with PCN. CD and HCD reconstructions leave gaps in thin structures, while EMD collapses coverage with poor fidelity. InfoCD improves coverage modestly but introduces spurious clusters. QAL (ours) yields denser and more uniform point sets (+4.5 pts Cov. over CD), visibly recovering sharper chair backs and object contours.

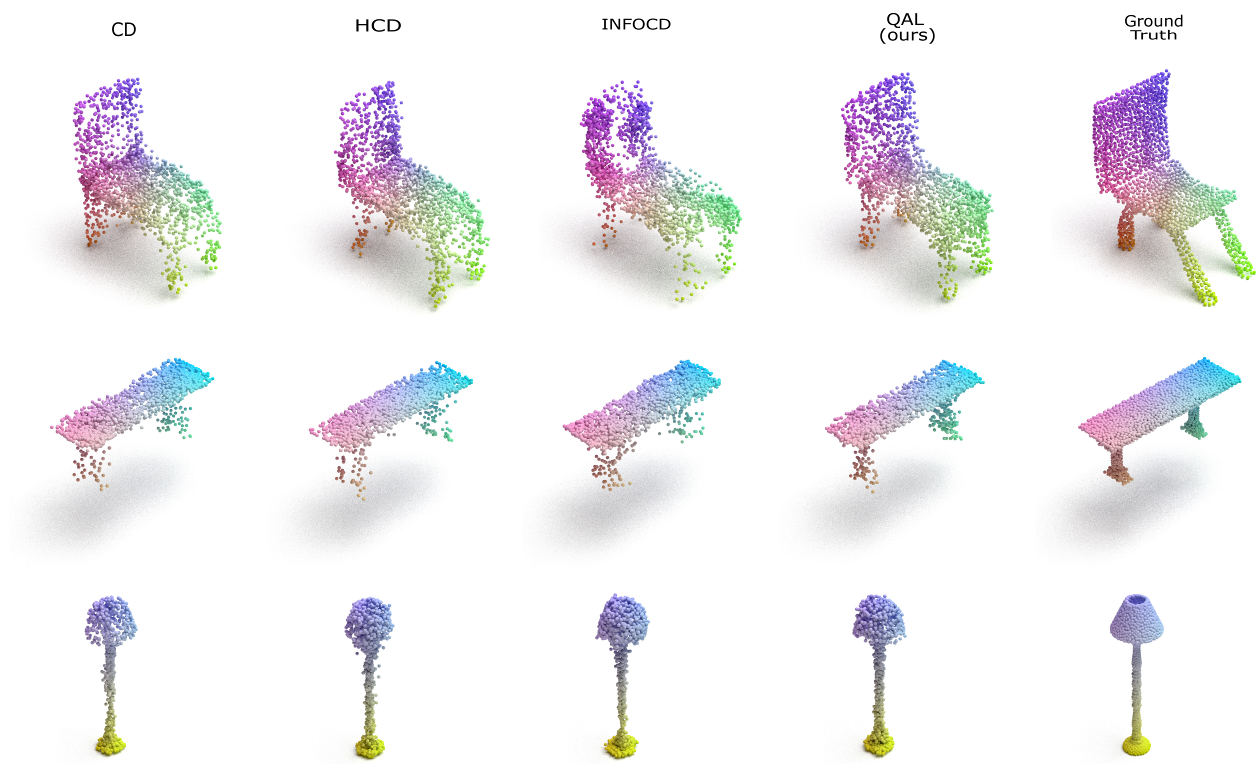

Figure 10. Representative reconstructions with TOPNet. CD, HCD, and InfoCD produce completions that look visually similar, consistent with flat CD/F1 values. However, QAL improves coverage (+4.3 pts Cov. over CD), filling in missing areas and producing more uniform reconstructions. These examples illustrate how CD/F1 miss improvements that are evident with coverage-aware evaluation.

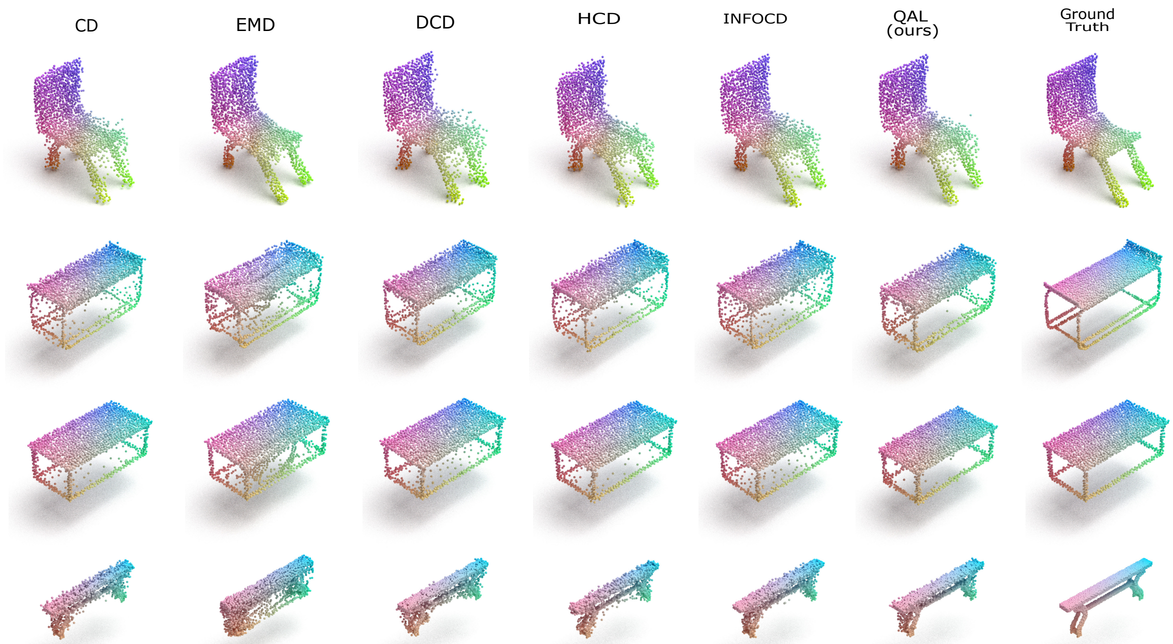

Figure 11. Representative reconstructions with ECG. CD and DCD achieve good overall alignment but leave subtle missing details. HCD and QAL both reach high coverage (91.6), yet CD differs (6.0 vs. 5.8), showing how CD alone can give misleading signals. Qualitatively, QAL better balances recall and precision, producing sharper and more consistent reconstructions without adding spurious points.

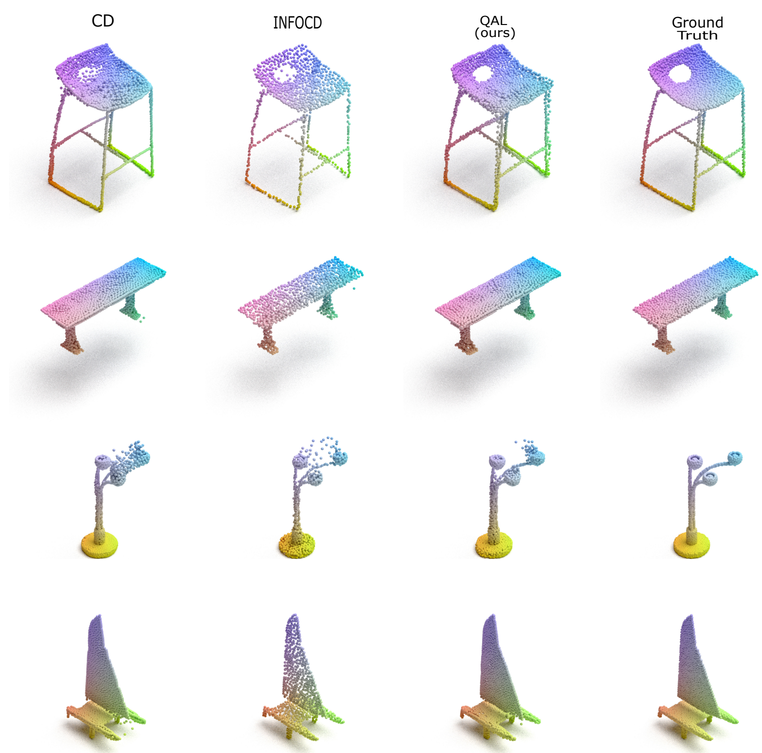

Figure 12. Representative reconstructions with SeedFormer. CD reconstructions miss fine chair details and introduce spurious points around thin lamp structures, while InfoCD collapses coverage (75.1) and clusters points densely in certain regions, leaving other areas sparse. In contrast, QAL restores coverage to 88.3 while maintaining balance between detail recovery and noise suppression, yielding more faithful reconstructions across both thin and coarse structures.