

Nordic AI Meet 2026

P4: Machine Learning, Intelligent Systems & AI for Science and Engineering

2. Automated Quality Evaluation of 3D scenes

Authors: Joakim Lindén^{1,2}; Håkan Forsberg¹; Masoud Daneshtalab¹

Affiliations: 1. Department of Computer Science & Engineering, Mälardalen University, Sweden

2. Saab Aeronautics, Advanced Programs, Stockholm, Sweden

Abstract: Recent advances in Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have enabled highly realistic 3D scene reconstruction from images. However, reconstruction quality depends heavily on where images are captured, and existing pipelines typically rely on manually designed acquisition strategies or uniformly sampled viewpoints. These approaches often collect redundant observations while failing to capture regions where additional data would most improve the reconstructed scene.

In this work, we present an AI-driven closed-loop framework for autonomous scene refinement that evaluates reconstruction quality and guides further data collection. Rather than treating reconstruction and image acquisition as separate tasks, the proposed framework iteratively reconstructs a scene, predicts where reconstruction quality is insufficient, plans efficient data collection routes, acquires the missing supplementary image data, and updates the reconstruction.

Our reference-free quality prediction network estimates perceptual reconstruction quality directly from rendered images without requiring ground-truth references during deployment. The predictor is trained to approximate LPIPS perceptual similarity and produces quality estimates that strongly correlate with measured reconstruction errors on held-out sets. These predictions are used to rank candidate viewpoints according to their expected contribution to scene improvement.

To efficiently acquire the selected images, we formulate a multi-agent Vehicle Routing Problem (VRP), enabling multiple autonomous platforms to cooperatively capture supplementary images whilst minimizing travel time and operational cost.

Experiments performed in simulated urban environments demonstrate that the proposed quality predictor accurately identifies reconstruction deficiencies, enabling targeted refinement of 3DGS models. Our results also show an important limitation of purely greedy viewpoint selection: the highest predicted-error viewpoints frequently cluster around the same artifacts, resulting in redundant sampling. Introducing a spatial Non-Maximum Suppression (NMS) stage helps distribute selected viewpoints more effectively across the scene and consistently outperforms greedy quality selection, random sampling, and geometric coverage-based baselines.

The work demonstrates how learned quality estimation can actively guide autonomous sensing and data acquisition, illustrating a practical application of AI for closing the loop between perception and action. From a holistic point of view, our framework can be seen as a strategy for adaptive data collection in robotics and autonomous systems, where machine learning models identify knowledge gaps and direct future observations to improve system performance.