

GO-PRE: Requested Viewpoints and Predictive Entropy

GO-PRE directs acquisition toward better predictions at a specified set of camera viewpoints. Its goal is a camera-view manifold, such as a sector around a reconstruction, rather than a semantic object instance [1]. This distinction matters: both are goal-conditioned acquisition, but only the latter directly specifies which object’s geometry should improve.

Objective and approximation

The method estimates how a candidate observation would reduce predictive rendering entropy over the goal viewpoints. A Jensen upper bound, Monte Carlo probe poses and diagonal Laplace/Jacobian approximations make this objective tractable. The current renderer supplies the derivatives used to score candidates before their RGB images are acquired. Successive acquisitions update the reconstruction; the inspected pool-selection protocol does not establish motion-feasible robot control or optimal finite-horizon geometric return.

The requested camera sector in Figure 1 determines which uncertainty matters for selection.




Figure 1: GO-PRE combines posterior uncertainty with sensitivity of the requested rendered views. The source schematic distinguishes goal probes from candidate and historical views. Reproduced from [1, Fig. 2], with surrounding whitespace removed.

Scope and evidence

The study uses Gaussian/radiance reconstruction on Blender objects and full Mip-NeRF360 and Tanks and Temples scenes. PSNR, SSIM and LPIPS assess rendered-image quality globally or in the specified camera sector. Predictive uncertainty concerns an observable output more directly than raw parameter uncertainty, but it remains different from a measured decrease in target geometric error. Performance also depends on the chosen goal-view distribution.

Consequence for ARIA-NBV

A precise target definition is required when comparing goal-directed NBV. A desired set of rendered views, an internally chosen look-at point, and a requested object are different goals. GO-PRE is a useful comparison for goal-conditioned information, not evidence that ARIA-NBV should replace its target point-mesh objective or introduce Gaussian reconstruction.

Primary version read: arXiv:2607.29037v1, method and evaluation protocols; publication identity checked against ICML 2026, on 13 October 2026.

Bibliography

- [1] Y. Song, Z. Li, C. Li, L. He, Y. Wang, and W. Zhang, “GO-PRE: Goal-Oriented Next-Best-View Selection via Predictive Rendering Entropy for Active 3D Reconstruction,” in *Proceedings of the 43rd International Conference on Machine Learning*, in Proceedings of Machine Learning Research, vol. 306. 2026, pp. 114687–114699. [Online]. Available: <https://proceedings.mlr.press/v306/song26k.html>