

Neural Visibility Field: Uncertainty and Active Mapping

A reconstruction model can be uncertain because a region has not been observed or because its appearance remains ambiguous. Neural Visibility Field (NVF) adds learned visibility and colour covariance to a radiance field so that these sources contribute to a candidate view’s predicted image uncertainty [1].

Selection before acquisition

The current field predicts an entropy score for each candidate image. Interpreting this as information gain assumes that conditional measurement entropy is constant; the method also corrects independent-pixel aggregation for spatial correlation using depth. The selected view is acquired and the field is trained again. Candidate poses can be sampled or refined through gradients. Repeating this cycle for a capture budget does not make the acquisition score a learned multi-step return.

Figure 1 places uncertainty scoring between reconstruction fitting and the next acquisition.




Figure 1: NVF alternates fitting the current field, estimating candidate-image uncertainty and acquiring the selected observation. Reproduced from [1, Fig. 2], with surrounding whitespace removed.

Reconstruction scope and evaluation

Experiments include isolated assets and a two-room mapping environment. The scene experiment values the whole reconstruction, rather than one externally requested object. Rendering quality and geometric accuracy, completeness and coverage assess the outcome. They are evaluation measures, not the entropy quantity optimized during selection. Their reported improvement supports the tested acquisition method, not equality between entropy reduction and geometric reconstruction improvement.

Relation to ARIA-NBV

NVF motivates distinguishing unobserved space from uncertain learned predictions. It queries the current model before acquiring the candidate image, respecting the temporal distinction important for candidate scoring. Adopting a radiance field or its uncertainty score is optional: ARIA-NBV’s target error and finite-horizon supervision remain separate scientific choices. A useful comparison would test whether an uncertainty feature helps rank candidates against measured reconstruction improvement, without relabelling that feature as the target value.

Primary version read: arXiv:2406.06948v1, named method and evaluation sections, on 13 October 2026.

Bibliography

- [1] S. Xue, J. Dill, P. Mathur, F. Dellaert, P. Tsiotras, and D. Xu, “Neural Visibility Field for Uncertainty-Driven Active Mapping,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [Online]. Available: <https://arxiv.org/abs/2406.06948>