Semantic and Geometry-Aware Depth Discretization for Multiplane Image Construction: A Comparative Study of Depth Sampling Strategies in MPI-Based Novel View Synthesis

Author: L. Paul
Email: L.B.Paul@student.tudelft.nl
Supervisor: P. Kellnhofer

Introduction:

Multiplane Images (MPIs) are a popular representation for novel view synthesis.

Rendering quality strongly depends on how continuous scene depth is discretized into a finite number of depth planes.

This work investigates whether geometry and semantic information can improve MPI depth-plane allocation.

Method:

Habitat-Matterport 3D Semantic dataset (HM3DSem). Scanned and reconstructed building interiors.

Construct MPIs using four depth allocation strategies:

- Uniform sampling
- Importance-based sampling
- Semantic-statistics sampling (min, max, mean, median)
- Edge-aware semantic sampling

Render novel viewpoints using model-view- and projection-matrix provided by blender.

Results:

Semantic methods generally outperform geometric baselines.

Sem-Median achieved the best overall quality. Ranking 1st in SSIM and LPIPS score as well as 2nd in PSNR score

Uniform sampling achieved the highest PSNR, but lower perceptual quality.

Importance sampling did not consistently improve results.

Strategy	SSIM↑	PSNR↑	LPIPS↓
uniform	0.916321	**31.181276**	0.046051
importance	0.916397	29.615104	0.063322
sem-min	0.916374	30.264391	0.046438
sem-max	0.910540	29.305154	0.049441
sem-mean	0.929213	31.069450	0.042421
sem-median	**0.931271**	31.174671	**0.041641**
sem-edge-min	0.917462	30.339033	0.045957
sem-edge-max	0.910660	29.315854	0.049381
sem-edge-mean	0.927480	30.931549	0.043076
sem-edge-median	0.928905	31.059272	0.042913

Table 1: Average values per metric for each investigated strategy

Limitations:

HM3DSem contains reconstruction artifacts, missing geometry, and texture misalignments.

Semantic labels are not available for all scenes.

Camera poses were manually selected.

Evaluation uses reconstructed scenes rather than perfect ground-truth geometry.

Importance sampling introduces non-determinism unless a fixed random seed is used.

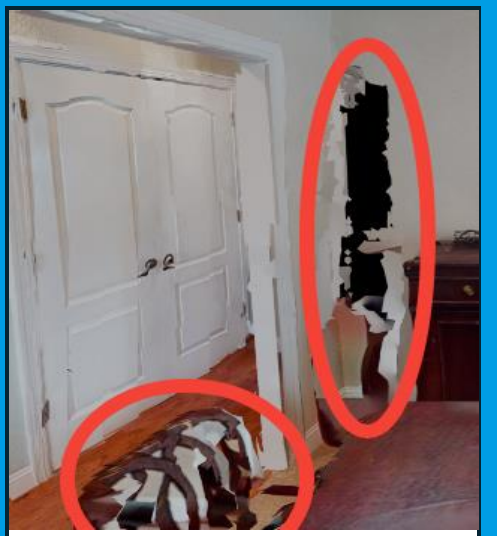

Figure 1: Example of mismatched texture and missing geometry

Conclusions:

Depth-plane allocation significantly affects MPI rendering quality.

Semantic information provides useful structural cues for depth discretization.

Semantic-statistics approaches outperform uniform and importance-based strategies on perceptual metrics.

The Sem-Median strategy achieved the strongest overall performance

Future Work:

Learn depth-plane placement automatically using optimization or deep learning.

Incorporate scene-level context and object relationships.

Test on higher-quality synthetic datasets with accurate geometry.