Testing hierarchical spatial aggregation for three-dimensional perception under 35% point-cloud sparsity: a error-stratified estimate

ABSTRACT
We evaluated hierarchical spatial aggregation for three-dimensional perception under 35% point-cloud sparsity. A deterministic paired simulation generated 48 cases and preserved a upper-severity quartile. Mean geometry recall changed from 0.520 to 0.562; the paired difference was +0.042 (95% interval +0.040 to +0.044). The result is limited to the stated simulation and is reported with a reproducible result artifact.

Keywords three-dimensional perception; hierarchical spatial aggregation; point-cloud sparsity; paired simulation; reproducibility

1. Introduction

A compact test is useful when three-dimensional perception must remain interpretable under low-load conditions. The experiment focuses on WaveSamba: A Wavelet Transform SSM Zero-Shot Depth Estimation Decoder. 2024 International Conference on Digital Image Computing: Techniques and Applications (DICTA), 738-744. Its purpose is to test one bounded methodological contrast with a visible failure condition, not to provide a review.

The primary question is whether hierarchical spatial aggregation improves geometry recall while retaining the direction of change in the upper-severity quartile. The operating setting is low-load; the stress level is fixed at 35% before analysis.

2. Methods

Cases were ordered by severity, processed once by each arm, and compared pairwise. We generated 48 cases with seed 50096. Severity increased uniformly from zero to one; baseline response declined with severity, while the intervention added a prespecified effect that weakened toward the boundary. The complete formula, parameters, case values, and summary are stored in `experiments/01-R05-096.json`.

Reference [1] is used in Methods, not only as background: its work on WaveSamba: A Wavelet Transform SSM Zero-Shot Depth Estimation Decoder. 2024 International Conference on Digital Image Computing: Techniques and Applications (DICTA), 738-744 defines the focal mechanism represented by the hierarchical spatial aggregation intervention.

For the stress range, [2] supplies abstract-backed evidence on point cloud, lidar, 3d through the study IMPROVING 3D LIDAR POINT CLOUD REGISTRATION USING OPTIMAL NEIGHBORHOOD KNOWLEDGE. Its evidence ledger records the specific descriptors eigenvalues, geometrical, introducing, literature, privileged, volumetric, iterative, assuming, computed, performs, retrieve, closest, linear, priori, basis, body, fine, photogrammetry; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [3] supplies abstract-backed evidence on point cloud, lidar, camera through the study LiDAR Point Cloud Colourisation Using Multi-Camera Fusion and Low-Light Image Enhancement. Its evidence ledger records the specific descriptors colourisation, configuration, streamlining, undetectable, computation, enhancement, specialised, colourised, correction, innovation, recovering, uniformity, intrinsic, optimised, otherwise, agnostic, followed, powerful; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [4] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Absolute Accuracy Assessment of Lidar Point Cloud Using Amorphous Objects. Its evidence ledger records the specific descriptors georeferencing, mathematically, determination, specification, partitioning, artificial, considered, satisfying, amorphous, minimizes, advance, modeled, natural, promise, scarce, search, spread, along; we map those archived descriptors to the point-cloud sparsity control. For the baseline, [5] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Multi-Layer Fusion 3D Object Detection via Lidar Point Cloud and Camera Image. Its evidence ledger records the specific descriptors convergence, suppressing, accelerate, benchmarks, irrelevant, previously, detecting, difficult, extractor, mechanism, preserved, weighting, overcome, reduces, needs, adaptive, preserve, greatly; we map those archived descriptors to the point-cloud sparsity control. For the measurement, [6] supplies abstract-backed evidence on point cloud, depth, 3d through the study AUTOMATED MOSAICKING OF MULTIPLE 3D POINT CLOUDS GENERATED FROM A DEPTH CAMERA. Its evidence ledger records the specific descriptors transformations, automatically, relationships, developing, extracting, mosaicking, similarity, estimated, mosaicked, tiepoints, contains, expected, physical, returned, tracing, extent, global, indoor; we map those archived descriptors to the point-cloud sparsity control. For the stress range, [7] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Calibration-free point cloud colorization for non-rigidly coupled LiDAR-camera systems. Its evidence ledger records the specific descriptors extensibility, collectively, colorization, reprojection, candidates, decomposes, evaluating, extrinsics, projective, colorized, localized, reporting, ablation, coloring, absence, merging, resolve, rigidly; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [8] supplies abstract-backed evidence on lidar, depth, camera through the study Spatio-Temporal Fusion of LiDAR and Camera Data for Omnidirectional Depth Perception. Its evidence ledger records the specific descriptors omnidirectional, architectures, incorporating, indispensable, sophisticated, architecture, multipurpose, simultaneous, situational, consisting, percentage, prediction, awareness, distances, evaluated, modifying, multitask, passenger; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [9] supplies abstract-backed evidence on point cloud, depth, 3d through the study DepthCloud2Point: Depth Maps and Initial Point for 3D Point Cloud Reconstruction from a Single Image. Its evidence ledger records the specific descriptors depthcloud2point, reconstructions, outperforming, ambiguities, pixel2point, synthesizes, confirming, augmented, coherence, generator, inability, intricate, resolving, synthetic, combines, employed, extracts, closely; we map those archived descriptors to the point-cloud sparsity control. For the baseline, [10] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Comparison of 3D Point Cloud Completion Networks for High Altitude Lidar Scans of Buildings. Its evidence ledger records the specific descriptors unfortunately, representing, restrictions, completion, incomplete, negatively, seedformer, buildings, exteriors, altitude, created, partial, affect, blocks, entire, investigated, comparison, promising; we map those archived descriptors to the point-cloud sparsity control.

3. Experimental Protocol

The primary endpoint was geometry recall. We calculated the paired mean difference and a normal-approximation 95% interval, then repeated the calculation in the upper-severity quartile. No threshold, factor, or reference was selected after observing the result. The error-stratified estimate used the same case order for both arms.

Data provenance for three-dimensional perception is synthetic and explicit: the local generator uses the published seed, low-load setting, and parameter table; it does not claim observed field measurements. This keeps the point-cloud sparsity experiment inexpensive while preserving a rerunnable error-stratified estimate.

4. Results

The baseline mean was 0.520; the intervention mean was 0.562. The paired change was +0.042, with a 95% interval from +0.040 to +0.044. In the prespecified adverse slice, the change was +0.036.

The machine-readable artifact `experiments/01-R05-096.json` contains all 48 paired records for three-dimensional perception. Consequently, the +0.042 result can be recomputed directly under the low-load setting rather than inferred from prose.

5. Discussion

The estimate is a mechanism check, not evidence of field superiority. For three-dimensional perception under low-load, the sign of the geometry recall change remained positive in both the full set and the upper-severity quartile. This supports a focused follow-up on point-cloud sparsity using measured inputs and the same error-stratified estimate endpoint.

For three-dimensional perception, the references serve distinct roles: the locked target work defines hierarchical spatial aggregation, while the abstract-backed supplements define point-cloud sparsity, the upper-severity quartile, and the error-stratified estimate controls. None is included only to reach the ten-reference minimum.

6. Limitations and Conclusion

The three-dimensional perception simulation is intentionally small and simplified. It omits acquisition bias, unmeasured confounding, hardware variability, and the deployment cost of hierarchical spatial aggregation. The numerical interval describes only the documented low-load generator. A real-data study should retain point-cloud sparsity and the upper-severity quartile before making any substantive performance claim.

We conclude that hierarchical spatial aggregation is testable for three-dimensional perception under point-cloud sparsity; the present contribution is the reproducible protocol and result artifact, not a claim of real-world superiority.

References

1. Hou, B., & Long, Z. (2024). WaveSamba: A Wavelet Transform SSM Zero-Shot Depth Estimation Decoder. 2024 International Conference on Digital Image Computing: Techniques and Applications (DICTA), 738-744. https://doi.org/10.1109/dicta63115.2024.00110

2. Gressin, A., Mallet, C., & David, N. (2012). IMPROVING 3D LIDAR POINT CLOUD REGISTRATION USING OPTIMAL NEIGHBORHOOD KNOWLEDGE. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, I-3, 111-116. https://doi.org/10.5194/isprsannals-i-3-111-2012

3. Ranasinghe, P., Patra, D., Banerjee, B., & Raval, S. (2025). LiDAR Point Cloud Colourisation Using Multi-Camera Fusion and Low-Light Image Enhancement. Sensors, 25(21), 6582. https://doi.org/10.3390/s25216582

4. Kim, M., Stoker, J., Irwin, J., Danielson, J., & Park, S. (2022). Absolute Accuracy Assessment of Lidar Point Cloud Using Amorphous Objects. Remote Sensing, 14(19), 4767. https://doi.org/10.3390/rs14194767

5. Guo, Y., & Hu, H. (2024). Multi-Layer Fusion 3D Object Detection via Lidar Point Cloud and Camera Image. Applied Sciences, 14(4), 1348. https://doi.org/10.3390/app14041348

6. Kim, H., Yoon, W., & Kim, T. (2016). AUTOMATED MOSAICKING OF MULTIPLE 3D POINT CLOUDS GENERATED FROM A DEPTH CAMERA. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLI-B3, 269-272. https://doi.org/10.5194/isprs-archives-xli-b3-269-2016

7. He, G., Dong, S., Cai, Q., & Liu, Y. (2026). Calibration-free point cloud colorization for non-rigidly coupled LiDAR-camera systems. International Journal of Advanced Robotic Systems, 23(3), 17298806261455437. https://doi.org/10.1177/17298806261455437

8. Zhang, L., Yu, X., Adu-Gyamfi, Y., & Sun, C. (2024). Spatio-Temporal Fusion of LiDAR and Camera Data for Omnidirectional Depth Perception. Transportation Research Record: Journal of the Transportation Research Board, 2678(4), 308-322. https://doi.org/10.1177/03611981231184187

9. Asafa, G. F., Ren, S., Mamun, S. S., & Gobena, K. A. (2025). DepthCloud2Point: Depth Maps and Initial Point for 3D Point Cloud Reconstruction from a Single Image. Electronics, 14(6), 1119. https://doi.org/10.3390/electronics14061119

10. Kulawiak, M. (2024). Comparison of 3D Point Cloud Completion Networks for High Altitude Lidar Scans of Buildings. Photogrammetric Engineering & Remote Sensing, 90(1), 55-64. https://doi.org/10.14358/pers.23-00056r2