Testing hierarchical spatial aggregation for three-dimensional perception under 30% point-cloud sparsity: a error-stratified estimate

ABSTRACT
We evaluated hierarchical spatial aggregation for three-dimensional perception under 30% point-cloud sparsity. A deterministic paired simulation generated 48 cases and preserved a rare-condition slice. Mean geometry recall changed from 0.480 to 0.516; the paired difference was +0.036 (95% interval +0.034 to +0.038). 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 sparse-observation 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 rare-condition slice. The operating setting is sparse-observation; the stress level is fixed at 30% before analysis.

2. Methods

Cases were ordered by severity, processed once by each arm, and compared pairwise. We generated 48 cases with seed 50040. 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-R05-040.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 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, [3] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Deep Learned Quantization-Based Codec for 3D Airborne LiDAR Point Cloud Images. Its evidence ledger records the specific descriptors approximately, decomposition, decompressed, decompresses, quantization, compressing, compression, established, instruction, transformed, compresses, stochastic, execution, regarding, constant, gradient, descent, dlqcpcd; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [4] supplies abstract-backed evidence on point cloud, depth, 3d through the study 3D Head Pose Estimation via Normal Maps: A Generalized Solution for Depth Image, Point Cloud, and Mesh. Its evidence ledger records the specific descriptors generalized, interaction, topological, procrustes, embedding, headspace, phenotype, primarily, abundant, directly, landmark, offering, analyze, herein, kinect, limits, manner, normal; 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 A Point-Wise LiDAR and Image Multimodal Fusion Network (PMNet) for Aerial Point Cloud 3D Semantic Segmentation. Its evidence ledger records the specific descriptors observational, proliferation, governmental, policymakers, nonetheless, permutation, supervision, augmenting, geospatial, invariance, management, respecting, companies, considers, currently, imbalance, insurance, producing; we map those archived descriptors to the point-cloud sparsity control. For the measurement, [6] supplies abstract-backed evidence on point cloud, lidar, 3d through the study General External Uncertainty Models of Three-Plane Intersection Point for 3D Absolute Accuracy Assessment of Lidar Point Cloud. Its evidence ledger records the specific descriptors intersections, calculating, substantial, analytical, checkpoint, positional, densities, elevation, guideline, qualities, routinely, accounts, identify, involves, practice, smallest, surveyed, unwanted; 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 Towards a Meaningful 3D Map Using a 3D Lidar and a Camera. Its evidence ledger records the specific descriptors subsequently, incremental, meaningful, registered, integrate, processed, sequences, surveying, consists, labeling, odometry, operated, utilized, removes, towards, burden, latest, traces; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [8] supplies abstract-backed evidence on point cloud, 3d, camera through the study 3D Point Cloud Verification and Scatter Removal Using a Dual Camera/Projector Setup. Its evidence ledger records the specific descriptors discrimination, triangulation, investigates, verification, trustworthy, turbidities, boundaries, geometries, performing, properties, scattering, underwater, unreliable, usefulness, motivated, ourselves, projector, uncertain; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [9] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Aerial Lidar Point Cloud Voxelization with its 3D Ground Filtering Application. Its evidence ledger records the specific descriptors simultaneously, discontinuous, coefficients, triangulated, reconstruct, publicized, separating, adjacency, axelsson, implicit, indicate, unground, benefit, samples, binary, lowest, notion, eight; we map those archived descriptors to the point-cloud sparsity control. For the baseline, [10] supplies abstract-backed evidence on point cloud, depth, 3d through the study Acquisition and Registration of Human Point Cloud Based on Depth Camera. Its evidence ledger records the specific descriptors optimize, acquire, perspective, implement, matching, least, registration, square, human, solve, pose, effectiveness, experimental, reconstruction, acquisition, step, proposes, feature; 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 rare-condition slice. 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, sparse-observation 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 0.480; the intervention mean was 0.516. The paired change was +0.036, with a 95% interval from +0.034 to +0.038. In the prespecified adverse slice, the change was +0.031.

The machine-readable artifact `experiments/01-R05-040.json` contains all 48 paired records for three-dimensional perception. Consequently, the +0.036 result can be recomputed directly under the sparse-observation setting rather than inferred from prose.

5. Discussion

The estimate is a mechanism check, not evidence of field superiority. For three-dimensional perception under sparse-observation, the sign of the geometry recall change remained positive in both the full set and the rare-condition slice. 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 rare-condition slice, 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 sparse-observation generator. A real-data study should retain point-cloud sparsity and the rare-condition slice 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. 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
3. Tamilmathi, A. C., & Chithra, P. L. (2021). Deep Learned Quantization-Based Codec for 3D Airborne LiDAR Point Cloud Images. Frontiers in Robotics and AI, 8, 606770. https://doi.org/10.3389/frobt.2021.606770
4. Wu, J., & Chen, H. (2024). 3D Head Pose Estimation via Normal Maps: A Generalized Solution for Depth Image, Point Cloud, and Mesh. Advanced Intelligent Systems, 6(11), 2400159. https://doi.org/10.1002/aisy.202400159
5. Poliyapram, V., Wang, W., & Nakamura, R. (2019). A Point-Wise LiDAR and Image Multimodal Fusion Network (PMNet) for Aerial Point Cloud 3D Semantic Segmentation. Remote Sensing, 11(24), 2961. https://doi.org/10.3390/rs11242961
6. Kim, M., Park, S., Danielson, J., Irwin, J., Stensaas, G., Stoker, J., & Nimetz, J. (2019). General External Uncertainty Models of Three-Plane Intersection Point for 3D Absolute Accuracy Assessment of Lidar Point Cloud. Remote Sensing, 11(23), 2737. https://doi.org/10.3390/rs11232737
7. Jeong, J., Yoon, T. S., & Park, J. B. (2018). Towards a Meaningful 3D Map Using a 3D Lidar and a Camera. Sensors, 18(8), 2571. https://doi.org/10.3390/s18082571
8. Thielemann, J. T., & Thorstensen, J. (2025). 3D Point Cloud Verification and Scatter Removal Using a Dual Camera/Projector Setup. Sensors, 25(23), 7250. https://doi.org/10.3390/s25237250
9. Wang, L., Xu, Y., & Li, Y. (2017). Aerial Lidar Point Cloud Voxelization with its 3D Ground Filtering Application. Photogrammetric Engineering & Remote Sensing, 83(2), 95-107. https://doi.org/10.14358/pers.83.2.95
10. Wang, H., & Wang, S. (2018). Acquisition and Registration of Human Point Cloud Based on Depth Camera. MATEC Web of Conferences, 173, 03001. https://doi.org/10.1051/matecconf/201817303001