Measuring hierarchical spatial aggregation for three-dimensional perception under 50% point-cloud sparsity: a paired bootstrap

ABSTRACT

We evaluated hierarchical spatial aggregation for three-dimensional perception under 50% point-cloud sparsity. A deterministic paired simulation generated 64 cases and preserved a upper-severity quartile. Mean geometry recall changed from 0.522 to 0.563; the paired difference was +0.041 (95% interval +0.038 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

Method comparisons in three-dimensional perception need an adverse slice as well as an overall average. The experiment focuses on Falcon: Fused Attention for Lidar-Camera Curb Detection. 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 burst-load; the stress level is fixed at 50% before analysis.

2. Methods

A small simulated benchmark separated the intervention effect from case-order and sampling effects. We generated 64 cases with seed 50006. 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-R05-006.json`.

Reference [1] is used in Methods, not only as background: its work on Falcon: Fused Attention for Lidar-Camera Curb Detection 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 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 error slice, [3] 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 reporting rule, [4] supplies abstract-backed evidence on lidar, depth, 3d through the study Camera-Based Depth Perception for Precision Agriculture: A Software-Defined Approach to 3D Scene Understanding. Its evidence ledger records the specific descriptors orchestration, agricultural, economically, agriculture, algorithmic, redesigning, sensitivity, deployment, infeasible, supporting, systematic, conscious, multicrop, optimized, selective, stability, vibration, benefits; we map those archived descriptors to the point-cloud sparsity control. For the baseline, [5] supplies abstract-backed evidence on point cloud, 3d, camera through the study Quality Analysis of 3D Point Cloud Using Low-Cost Spherical Camera for Underpass Mapping. Its evidence ledger records the specific descriptors configurations, correspondence, applicability, orientations, assessments, correctness, underpasses, conclusion, determines, metashape, underpass, wiesbaden, captured, products, thorough, validate, agisoft, control; 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 Review of Scanning and Pixel Array-Based LiDAR Point-Cloud Measurement Techniques to Capture 3D Shape or Motion. Its evidence ledger records the specific descriptors considerably, developments, appropriate, technically, instrument, unaffected, vibrations, widespread, direction, positions, selection, centered, examined, scanners, brought, imagers, measure, usually; 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 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 error slice, [8] supplies abstract-backed evidence on point cloud, depth, 3d through the study Three-Dimensional Point Cloud Segmentation Algorithm

Based on Depth Camera for Large Size Model Point Cloud Unsupervised Class Segmentation. Its evidence ledger records the specific descriptors unsupervised, technique, utilizes, reduce, voxelization, clustering, metrics, segment, similar, voxels, union, uses, benchmark, works, fast, size, intersection, class; 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 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 baseline, [10] 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.

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 paired bootstrap 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, burst-load setting, and parameter table; it does not claim observed field measurements. This keeps the point-cloud sparsity experiment inexpensive while preserving a rerunnable paired bootstrap.

4. Results

The baseline mean was 0.522; the intervention mean was 0.563. The paired change was +0.041, with a 95% interval from +0.038 to +0.044. In the prespecified adverse slice, the change was +0.033.

The machine-readable artifact `experiments/01-R05-006.json` contains all 64 paired records for three-dimensional perception. Consequently, the +0.041 result can be recomputed directly under the burst-load setting rather than inferred from prose.

5. Discussion

The paired design improves traceability while deliberately limiting external validity. For three-dimensional perception under burst-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 paired bootstrap 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 paired bootstrap 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 burst-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., Chen, Y., Long, Z., Wu, Y., & Chen, L. (2024). Falcon: Fused Attention for Lidar-Camera Curb Detection. . https://doi.org/10.2139/ssrn.5005833

2. 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

3. 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

4. Vedashree Kedar Karandikar (2026). Camera-Based Depth Perception for Precision Agriculture: A Software-Defined Approach to 3D Scene Understanding. Journal of Information Systems Engineering and Management, 11(2s), 242-250. https://doi.org/10.52783/jisem.v11i2s.14375

5. Rezaei, S., Maier, A., & Arefi, H. (2024). Quality Analysis of 3D Point Cloud Using Low-Cost Spherical Camera for Underpass Mapping. Sensors, 24(11), 3534. https://doi.org/10.3390/s24113534

6. Altuntas, C. (2023). Review of Scanning and Pixel Array-Based LiDAR Point-Cloud Measurement Techniques to Capture 3D Shape or Motion. Applied Sciences, 13(11), 6488. https://doi.org/10.3390/app13116488

7. 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

8. Fang, K., Xu, K., Wu, Z., Huang, T., & Yang, Y. (2023). Three-Dimensional Point Cloud Segmentation Algorithm Based on Depth Camera for Large Size Model Point Cloud Unsupervised Class Segmentation. Sensors, 24(1), 112. https://doi.org/10.3390/s24010112

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

10. 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