Stress-testing hierarchical spatial aggregation for three-dimensional perception under 45% point-cloud sparsity: a blocked comparison

ABSTRACT

We evaluated hierarchical spatial aggregation for three-dimensional perception under 45% point-cloud sparsity. A deterministic paired simulation generated 64 cases and preserved a median slice. Mean geometry recall changed from 0.558 to 0.602; the paired difference was +0.044 (95% interval +0.042 to +0.046). 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

This study isolates one operational question in three-dimensional perception: whether a transparent control survives low-load inputs. 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 median slice. The operating setting is low-load; the stress level is fixed at 45% before analysis.

2. Methods

A fixed generator created matched inputs; only the prespecified intervention differed between arms. We generated 64 cases with seed 50122. 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-R05-122.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 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, [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 lidar, depth, camera through the study Fuzzy Logic Based Fusion of 2D LIDAR and Depth Camera Data for Robust Perception of Obstacle Distance. Its evidence ledger records the specific descriptors misinformation, advancements, reflectivity, transparency, encouraging, implemented, perceiving, perceived, relatable, corrects, decision, integral, compact, finding, color, fused, fuzzy, given; 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 Methodology of real-time 3D point cloud mapping with UAV lidar. Its evidence ledger records the specific descriptors visualization, availability, identifying, independent, preparation, containing, facilitate, laboratory, simulation, instantly, interface, planning, platform, shortens, changes, change, stored, timely; we map those archived descriptors to the point-cloud sparsity control. For the measurement, [6] 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 stress range, [7] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Real-time Track Obstacle Detection from 3D LIDAR Point Cloud. Its evidence ledger records the specific descriptors contributions, interfered, dependent, launching, affected, equation, follows, serious, track, clustering, driverless, millimeter, obstacles, potential, obstacle, greatly, unable, great; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [8] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Automatic Tunnel Steel Arches Extraction Algorithm Based on 3D LiDAR Point Cloud. Its evidence ledger records the specific descriptors construction, assistance, automation, experiment, minimizing, researched, researches, completed, installed, shotcrete, changsha, directed, sequence, variance, growing, highway, section, tunnels; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [9] supplies abstract-backed evidence on point cloud, lidar, depth through the study Hyper-source Enhanced Retrospective Attention Fusion for Camera-LiDAR 3D Object Detection. Its evidence ledger records the specific descriptors geometrically, retrospective, interactions, dynamically, selectively, aggregates, consistent, correlates, modalities, multiscale, validating, ambiguity, balancing, learnable, hyper, bias, bird, distributions; 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 A Staged Real-Time Ground Segmentation Algorithm of 3D LiDAR Point Cloud. Its evidence ledger records the specific descriptors undersegmentation, semantickitti, concentric, flowerbeds, addresses, reflected, exhibits, flatness, judgment, validity, filters, fitting, fitted, repair, slopes, staged, grass, stage; 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 median slice. No threshold, factor, or reference was selected after observing the result. The blocked comparison 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 blocked comparison.

4. Results

The baseline mean was 0.558; the intervention mean was 0.602. The paired change was +0.044, with a 95% interval from +0.042 to +0.046. In the prespecified adverse slice, the change was +0.036.

The machine-readable artifact `experiments/01-R05-122.json` contains all 64 paired records for three-dimensional perception. Consequently, the +0.044 result can be recomputed directly under the low-load setting rather than inferred from prose.

5. Discussion

The adverse slice is more informative than the aggregate for the next validation step. For three-dimensional perception under low-load, the sign of the geometry recall change remained positive in both the full set and the median slice. This supports a focused follow-up on point-cloud sparsity using measured inputs and the same blocked comparison 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 median slice, and the blocked comparison 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 median 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., 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. 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

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. Saksena, H., & Anwar, S. (2022). Fuzzy Logic Based Fusion of 2D LIDAR and Depth Camera Data for Robust Perception of Obstacle Distance. Volume 5: Dynamics, Vibration, and Control, V005T07A069. https://doi.org/10.1115/imece2022-95281

5. CANDAN, L., & KAÇAR, E. (2023). Methodology of real-time 3D point cloud mapping with UAV lidar. International Journal of Engineering and Geosciences, 8(3), 301-309. https://doi.org/10.26833/ijeg.1178260

6. 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

7. He, C., Gong, J., Yang, Y., Bi, D., Lan, J., & Qie, L. (2021). Real-time Track Obstacle Detection from 3D LIDAR Point Cloud. Journal of Physics: Conference Series, 1910(1), 012002. https://doi.org/10.1088/1742-6596/1910/1/012002

8. Zhang, W., Qiu, W., Song, D., & Xie, B. (2019). Automatic Tunnel Steel Arches Extraction Algorithm Based on 3D LiDAR Point Cloud. Sensors, 19(18), 3972. https://doi.org/10.3390/s19183972

9. Mi, G., Zhang, J., Li, J., & Zhuo, L. (2026). Hyper-source Enhanced Retrospective Attention Fusion for Camera-LiDAR 3D Object Detection. https://doi.org/10.2139/ssrn.6689267

10. Deng, W., Chen, X., & Jiang, J. (2024). A Staged Real-Time Ground Segmentation Algorithm of 3D LiDAR Point Cloud. Electronics, 13(5), 841. https://doi.org/10.3390/electronics13050841