Stabilizing hierarchical spatial aggregation for three-dimensional perception under 19% point-cloud sparsity: a blocked comparison

ABSTRACT

We evaluated hierarchical spatial aggregation for three-dimensional perception under 19% point-cloud sparsity. A deterministic paired simulation generated 72 cases and preserved a median slice. Mean geometry recall changed from 0.582 to 0.630; the paired difference was +0.048 (95% interval +0.046 to +0.051). 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

The central concern is whether hierarchical spatial aggregation remains measurable when point-cloud sparsity increases. 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 sparse-observation; the stress level is fixed at 19% before analysis.

2. Methods

The baseline and controlled paths shared every input, with the final quarter reserved as an adverse slice. We generated 72 cases with seed 50063. 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-R05-063.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 Robust Human Tracking Using a 3D LiDAR and Point Cloud Projection for Human-Following Robots. Its evidence ledger records the specific descriptors surroundings, differences, individuals, potentially, projecting, proportion, vulnerable, following, failures, revealed, devices, finders, observe, outlier, reduced, shorter, typical, humans; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [3] supplies abstract-backed evidence on point cloud, lidar, depth through the study Geometric calibration for LiDAR-camera system fusing 3D-2D and 3D-3D point correspondences. Its evidence ledger records the specific descriptors correspondences, contribute, establish, utilize, boards, verify, extra, hard, improvements, constraints, chessboard, auxiliary, computing, monocular, because, besides, planar, stable; 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 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. For the baseline, [5] 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 measurement, [6] supplies abstract-backed evidence on point cloud, lidar, 3d through the study 3D CITYGML BUILDING MODELLING FROM LIDAR POINT CLOUD DATA. Its evidence ledger records the specific descriptors revolutionized, diversified, initiatives, comparable, government, arcgispro, citygml4j, initially, exchange, isolated, managing, packages, workflow, citygml, storage, create, demand, mapped; we map those archived descriptors to the point-cloud sparsity control. For the stress range, [7] supplies abstract-backed evidence on lidar, depth, 3d through the study View Synthesis: LiDAR Camera versus Depth Estimation. Its evidence ledger records the specific descriptors configured, synthesizer, adequately, inspection, plications, synthesize, multiview, realsense, recommend, rendering, synthesis, delicate, neverthe, slightly, tational, borders, exposed, minutes; 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 Real-time 3D multi-pedestrian detection and tracking using 3D LiDAR point cloud for mobile robot. Its evidence ledger records the specific descriptors convolutional, heterogeneous, insufficient, autoencoder, continuing, identifies, pedestrian, recognizes, movements, recognize, evaluate, movement, navigate, segments, crowded, realize, frames, mixing; 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 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. For the baseline, [10] 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.

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, 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 blocked comparison.

4. Results

The baseline mean was 0.582; the intervention mean was 0.630. The paired change was +0.048, with a 95% interval from +0.046 to +0.051. In the prespecified adverse slice, the change was +0.042.

The machine-readable artifact `experiments/01-R05-063.json` contains all 72 paired records for three-dimensional perception. Consequently, the +0.048 result can be recomputed directly under the sparse-observation setting rather than inferred from prose.

5. Discussion

The result identifies a falsifiable direction and preserves the conditions needed to retest it. For three-dimensional perception under sparse-observation, 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 sparse-observation 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. Kitamoto, S., Hiroi, Y., Miyawaki, K., & Ito, A. (2025). Robust Human Tracking Using a 3D LiDAR and Point Cloud Projection for Human-Following Robots. Sensors, 25(6), 1754. https://doi.org/10.3390/s25061754

3. An, P., Ma, T., Yu, K., Fang, B., Zhang, J., Fu, W., & Ma, J. (2020). Geometric calibration for LiDAR-camera system fusing 3D-2D and 3D-3D point correspondences. Optics Express, 28(2), 2122. https://doi.org/10.1364/oe.381176

4. 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

5. 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

6. Jayaraj, P., & Ramiya, A. M. (2018). 3D CITYGML BUILDING MODELLING FROM LIDAR POINT CLOUD DATA. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-5, 175-180. https://doi.org/10.5194/isprs-archives-xlii-5-175-2018

7. Xie, Y., Fachada, S., Bonatto, D., Teratani, M., & Lafruit, G. (2021). View Synthesis: LiDAR Camera versus Depth Estimation. CSRN. https://doi.org/10.24132/csrn.2021.3101.35

8. Na, K. I., & Park, B. (2023). Real-time 3D multi-pedestrian detection and tracking using 3D LiDAR point cloud for mobile robot. ETRI Journal, 45(5), 836-846. https://doi.org/10.4218/etrij.2023-0116

9. 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

10. 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