Tuning hierarchical spatial aggregation for three-dimensional perception under 29% point-cloud sparsity: a sensitivity sweep

ABSTRACT
We evaluated hierarchical spatial aggregation for three-dimensional perception under 29% point-cloud sparsity. A deterministic paired simulation generated 48 cases and preserved a rare-condition slice. Mean geometry recall changed from 0.523 to 0.568; the paired difference was +0.045 (95% interval +0.043 to +0.047). 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

Reproducible stress tests can expose weaknesses in three-dimensional perception before expensive field collection. The experiment focuses on Hierarchical Spatial Mamba Framework for Point Cloud Classification. 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 shifted-distribution; the stress level is fixed at 29% before analysis.

2. Methods

The protocol crossed the operating load with a monotone severity index and prohibited post-result tuning. We generated 48 cases with seed 50132. 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-R05-132.json`.

Reference [1] is used in Methods, not only as background: its work on Hierarchical Spatial Mamba Framework for Point Cloud Classification 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 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 error slice, [3] supplies abstract-backed evidence on lidar, depth, 3d through the study Real-Time Object Detection and Distance Measurement Enhanced with Semantic 3D Depth Sensing Using Camera-LiDAR Fusion. Its evidence ledger records the specific descriptors particularly, estimations, categories, evaluates, occlusion, preferred, requiring, assessed, bicycle, central, needing, faster, larger, offers, slower, trucks, cerde, corde; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [4] supplies abstract-backed evidence on point cloud, lidar, depth through the study Depth-prior-based lidar point cloud de-noising method leveraging range-gated imaging. Its evidence ledger records the specific descriptors penetrating, transparent, downstream, denoised, regarded, degrade, divides, noising, removed, cannot, caused, gated, kinds, noisy, inconsistent, leveraging, occlusions, consuming; 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 Calibration of an Outdoor Distributed Camera Network with a 3D Point Cloud. Its evidence ledger records the specific descriptors recalibration, homographies, localization, surveillance, distributed, extensively, universitat, exploiting, ubiquitous, assisting, barcelona, catalunya, assisted, becoming, covering, facultat, frequent, includes; we map those archived descriptors to the point-cloud sparsity control. For the measurement, [6] 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 stress range, [7] supplies abstract-backed evidence on point cloud, lidar, 3d through the study Deep learning for LiDAR-only and LiDAR-fusion 3D perception: a survey. Its evidence ledger records the specific descriptors collaboration, supplementary, complexity, increasing, overlooked, redundancy, understand, emerging, instance, property, starting, aspects, focuses, weather, handle, hoping, severe, draws; 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 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 reporting rule, [9] 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 baseline, [10] 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.

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 sensitivity sweep 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, shifted-distribution setting, and parameter table; it does not claim observed field measurements. This keeps the point-cloud sparsity experiment inexpensive while preserving a rerunnable sensitivity sweep.

4. Results

The baseline mean was 0.523; the intervention mean was 0.568. The paired change was +0.045, with a 95% interval from +0.043 to +0.047. In the prespecified adverse slice, the change was +0.039.

The machine-readable artifact `experiments/01-R05-132.json` contains all 48 paired records for three-dimensional perception. Consequently, the +0.045 result can be recomputed directly under the shifted-distribution setting rather than inferred from prose.

5. Discussion

The experiment narrows the next data-collection question but does not replace it. For three-dimensional perception under shifted-distribution, 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 sensitivity sweep 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 sensitivity sweep 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 shifted-distribution 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. Sun, Y., Zia, A., Long, Z., Qiu, Z., Xiang, W., & Zhou, J. (2026). Hierarchical Spatial Mamba Framework for Point Cloud Classification. Lecture Notes in Computer Science, 402-417. https://doi.org/10.1007/978-981-95-4395-3_28
2. 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
3. Yildiz, A. S., Meng, H., & Swash, M. R. (2025). Real-Time Object Detection and Distance Measurement Enhanced with Semantic 3D Depth Sensing Using Camera-LiDAR Fusion. Applied Sciences, 15(10), 5543. https://doi.org/10.3390/app15105543
4. Xiaoquan, L., & yang, Z. (2024). Depth-prior-based lidar point cloud de-noising method leveraging range-gated imaging. https://doi.org/10.1364/opticaopen.25909819.v1
5. Ortega, A., Silva, M., Teniente, E., Ferreira, R., Bernardino, A., Gaspar, J., & Andrade-Cetto, J. (2014). Calibration of an Outdoor Distributed Camera Network with a 3D Point Cloud. Sensors, 14(8), 13708-13729. https://doi.org/10.3390/s140813708
6. 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
7. Wu, D., Liang, Z., & Chen, G. (2022). Deep learning for LiDAR-only and LiDAR-fusion 3D perception: a survey. Intelligence & Robotics, 2(2), 105-129. https://doi.org/10.20517/ir.2021.20
8. 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
9. 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
10. 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