Stress-testing hierarchical spatial aggregation for three-dimensional perception under 48% point-cloud sparsity: a sensitivity sweep

ABSTRACT

We evaluated hierarchical spatial aggregation for three-dimensional perception under 48% point-cloud sparsity. A deterministic paired simulation generated 64 cases and preserved a late-arriving block. Mean geometry recall changed from 0.521 to 0.570; 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

This study isolates one operational question in three-dimensional perception: whether a transparent control survives noisy-input inputs. 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 late-arriving block. The operating setting is noisy-input; the stress level is fixed at 48% before analysis.

2. Methods

A fixed generator created matched inputs; only the prespecified intervention differed between arms. We generated 64 cases with seed 50018. 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-R05-018.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, 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 error slice, [3] supplies abstract-backed evidence on point cloud, lidar, 3d through the study F-PointNet multi-modal 3D object detection based on RGB images and LiDAR point cloud. Its evidence ledger records the specific descriptors generalisation, perturbations, localisation, perturbation, adjustments, contributes, emphasising, capability, designated, suggesting, advancing, conducted, employing, scenarios, explores, included, optimise, pointnet; we map those archived descriptors to the point-cloud sparsity control. For the reporting rule, [4] 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 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 Real Pseudo-Lidar Point Cloud Fusion for 3D Object Detection. Its evidence ledger records the specific descriptors accuracies, impressive, addressed, emergence, heavily, pseudo, crude, fuses, reach, rely, generating, essential, resulting, cyclists, matching, recently, popular, costs; 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 Up-Sampling Method for Low-Resolution LiDAR Point Cloud to Enhance 3D Object Detection in an Autonomous Driving Environment. Its evidence ledger records the specific descriptors interpolation, conventional, automobile, channels, conserve, increase, neighbor, selected, budget, pillar, sample, simple, solves, apart, empty, faced, super, reconstructing; we map those archived descriptors to the point-cloud sparsity control. For the error slice, [8] supplies abstract-backed evidence on point cloud, lidar, camera through the study Adaptive point cloud acquisition and upsampling for automotive lidar. Its evidence ledger records the specific descriptors participants, registering, automotive, popularity, ultrasonic, upsampling, acquiring, proposals, retaining, upsampled, autonomy, detector, infrared, numerous, paradigm, spectrum, decades, gaining; 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 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 baseline, [10] 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.

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 late-arriving block. 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, noisy-input 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 0.521; the intervention mean was 0.570. The paired change was +0.048, with a 95% interval from +0.046 to +0.051. In the prespecified adverse slice, the change was +0.041.

The machine-readable artifact `experiments/01-R05-018.json` contains all 64 paired records for three-dimensional perception. Consequently, the +0.048 result can be recomputed directly under the noisy-input 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 noisy-input, the sign of the geometry recall change remained positive in both the full set and the late-arriving block. 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 late-arriving block, 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 noisy-input generator. A real-data study should retain point-cloud sparsity and the late-arriving block 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. 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

3. Luo, Y. (2024). F-PointNet multi-modal 3D object detection based on RGB images and LiDAR point cloud. Advances in Engineering Innovation, 12(1), 77-85. https://doi.org/10.54254/2977-3903/12/2024138

4. 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

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. Fan, X., Xiao, D., Cai, D., & Ding, W. (2023). Real Pseudo-Lidar Point Cloud Fusion for 3D Object Detection. Electronics, 12(18), 3920. https://doi.org/10.3390/electronics12183920

7. You, J., & Kim, Y. K. (2022). Up-Sampling Method for Low-Resolution LiDAR Point Cloud to Enhance 3D Object Detection in an Autonomous Driving Environment. Sensors, 23(1), 322. https://doi.org/10.3390/s23010322

8. Jovanov, L., Lee, W. Y., & Philips, W. (2023). Adaptive point cloud acquisition and upsampling for automotive lidar. Applied Optics, 62(17), F8. https://doi.org/10.1364/ao.482535

9. 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

10. 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