Evaluating LiDAR Compression for 3D Semantic Segmentation in Diverse Off-Road Environments

Adam Niemczura
aniemcz@clemson.edu
Clemson University
Clemson, South Carolina, USA

Max Faykus (Advisor)
Clemson University
Clemson, South Carolina, USA
mfaykus@clemson.edu

Oyinlolu Odetoye (Advisor)
Clemson University
Clemson, South Carolina, USA
oodetoy@clemson.edu

Melissa Smith (Advisor)
Clemson University
Clemson, South Carolina, USA
smithmc@clemson.edu

Jon Calhoun (Advisor)
Clemson University
Clemson, South Carolina, USA
jonccal@clemson.edu

Scott Groel (Advisor)
Clemson University
Clemson, South Carolina, USA
sgroel@clemson.edu

CCS Concepts

• **Information systems → Data streaming**; *Sensor networks.*

Keywords

Compression, Point Clouds, LiDAR, 3D Semantic Segmentation

ACM Reference Format:
Adam Niemczura, Max Faykus (Advisor), Oyinlolu Odetoye (Advisor), Melissa Smith (Advisor), Jon Calhoun (Advisor), and Scott Groel (Advisor). 2025. Evaluating LiDAR Compression for 3D Semantic Segmentation in Diverse Off-Road Environments . In *Proceedings of ACM SRC Undergraduate Competition (SC'25).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Transmitting point cloud data is vital for applications like autonomous vehicle navigation, especially when vehicles are compute limited and unable to run large 3d semantic segmentation models locally. LiDAR data filmed over several hours or days easily grows to gigabytes or terabytes uncompressed, making cloud-based model inference or peer to peer vehicle networks costly, especially on low-power devices like drones. While recent research has advanced point cloud compression, most work evaluates performance using object detection and on urban datasets like SemanticKITTI [1] or nuScenes [2]. These do not exactly reflect performance on off-road outdoor data, which is typically noisier and less structured. We benchmark 3 LiDAR compressors, RENO (neural-based) [8], TMC13 (rules-based baseline) [5], and LCP [9] (scientific particle compressor untested in this domain) on the GOOSE dataset [6]. We trained 2 3d semantic segmentation models on this decompressed LiDAR data to observe their downstream segmentation performance.

Code is available at https://github.com/AdamNiem/goose-reno-compression-suite.

SC'25, St Louis, MO

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

2 Methodology

2.1 Preprocessing and Overall Pipeline

To conduct this experiment, we followed these general steps:

(1) Convert files to ply with only xyz data
(2) Quantization (if necessary)
(3) Compression + Decompression (including dequantization)
(4) Due to point loss, reconnect labels by connecting each point to the closest label from original via K-Nearest Neighbors using a neighborhood of 1
(5) Restore intensity values via the same method as step 4 into file to turn it from xyz only to xyz and intensity. Also converts them back into binary files
(6) Sphereformer (25 epochs) [4] and PTV3 (50 epochs) [7] model is trained on this decompressed data with configuration based on GOOSE PTV3 example config from scratch for each compression level [6]. Both were allotted 30 hours and 2 A100s

Additionally, for RENO, since it is a machine learning approach, it requires training the compressor. Therefore the compressor was trained on the training data of goose. This was done to prevent the compressor from leaking information about the test data to the downstream semantic segmentation models.

2.2 Timing

We also benchmarked the speeds of the commpressors to serve as a rough comparison only since RENO utilizes GPUs [8] while LCP and TMC13 are currently CPU only [9] [5]. The timings were collected from the printed times provided by each compressor

Table 1: Resources allocated for compressing the entire GOOSE dataset with different compressors. All CPU cores were Intel® Xeon® Gold 6258R @ 2.70GHz (x86_64).

Compressor	CPUs	GPUs	Memory
RENO	32 × Xeon Gold 6258R	1 × V100	64 GB
LCP	4 × Xeon Gold 6258R	None	16 GB
TMC13	4 × Xeon Gold 6258R	None	16 GB

2.3 Evaluation Metrics

For evaluation 4 key metrics were used. Those being mIoU, Compression Ratio, PSNR D1 (Point to Point), and PSNR D2 (Point to Plane). Here mIoU serves as a metric to highlight the downstream semantic segmentation model's accuracy on the validation data where higher mIoU is better. Compression Ratio is used here to measure the effectiveness of each compressor in compression and in this case is calculated as the file size in bytes of the original LiDAR point cloud as an ASCII PLY file divided by the size of the compressed file in bytes. This is done so as to keep all compression ratios between compressors consistent And for measuring the quality of the decompressed point clouds we use PSNR D1 and D2.

3 Results

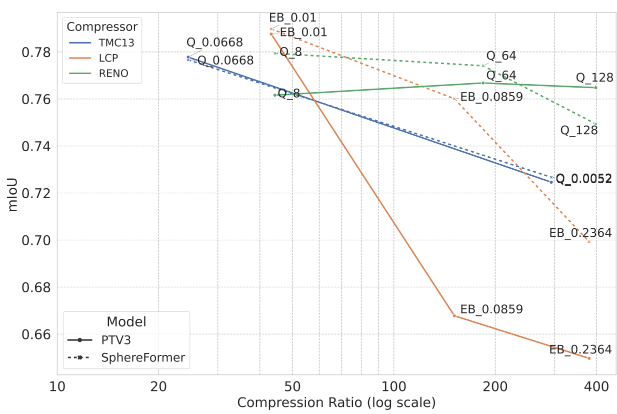

Figure 1

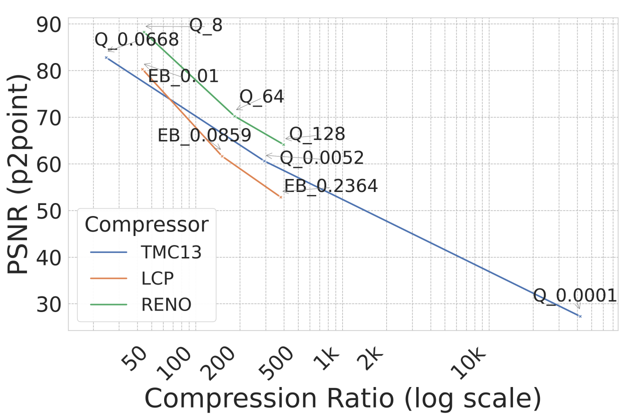

Figure 2

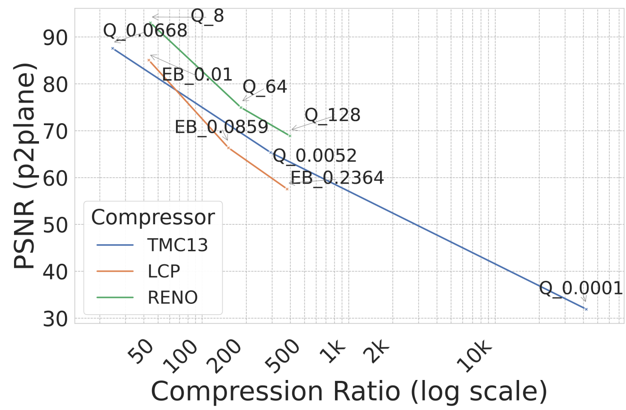

Figure 3

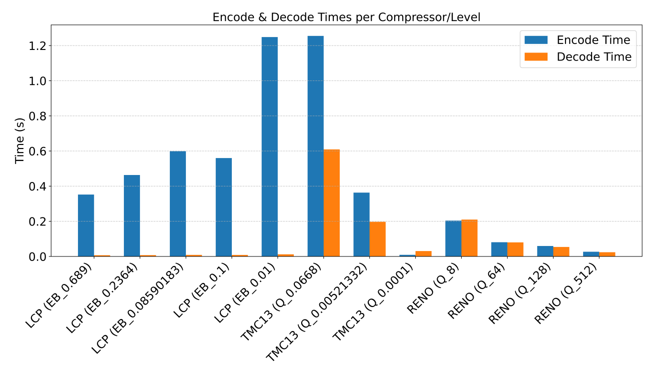

Figure 4

3.1 LCP

LCP achieves competitive results to RENO and TMC13 despite never having been used in this domain before. Although it also has wide disparities across models, with a 9% difference in mIoU at error bound 0.0859 between the 2 models tested. It also has the fastest decoding time of those tested although with a similar encoding time to TMC13.

3.2 RENO

RENO compressor performs the best on downstream semantic segmentation with minimal decay in performance over the levels of compression tested. Across the compression ratios tested, RENO is able to maintain a downstream mIoU above 74% for all compression ratios on PTV3 and Sphereformer, ranging from a compression ratio of 50 all the way to 400.

3.3 TMC13

TMC13 achieves median to low mIoU compared to LCP and RENO, performing worse or similarly to LCP at similar compression ratios. Specifically, it initially gets an mIoU of 77% for both models at a CR of 25 and gradually degrades consistently for both models

4 Conclusion

Overall we find RENO to perform the best of the compressors evaluated, bolstering neural-based compressors as viable for diverse off road navigation. We also find LCP to be competitive in this domain although the data suggests LCP introduces distortion to the point clouds that can greatly vary the performance of downstream semantic segmentation models. Meanwhile TMC13 does moderately in comparison to both LCP and RENO overall. We ultimately find RENO's compression capabilities transfers well to outdoor off-road scenarios as well as highlighting LCP's potential usefulness in forming future point cloud compressors.

5 Future Work

Future work would include exploring the reason for LCP's discrepancy between models as well as looking into further exploring how sensitive RENO is to out of training distribution point clouds. Another promising direction would be investigating building upon RENO to go from a single frame compressor to a multi-frame compressor. For example, one paper, when testing on the KITTI LiDAR dataset, "...find that on average 99% of each point cloud is geometrically overlapped with the previous point cloud." [3]. They also claim to have been able to have a "40x to 90x compression rate" via exploiting this temporal relationship. Therefore expanding RENO to take advantage of this temporal redundancy could see major improvements in compression performance.

Acknowledgments

This work was supported in part by the Clemson University Creative Inquiry + Undergraduate Research Program. Clemson University is also acknowledged for their generous allotment of compute time on the Palmetto Cluster.

References

[1] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. 2019. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

[2] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, and O. Beijbom. 2020. nuScenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[3] Yu Feng, Shaoshan Liu, and Yuhao Zhu. 2020. Real-Time Spatio-Temporal LiDAR Point Cloud Compression. arXiv:2008.06972 [eess.IV] https://arxiv.org/abs/2008.06972

[4] X. Lai, Y. Chen, F. Lu, J. Liu, and J. Jia. 2023. Spherical Transformer for LiDAR-based 3D Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[5] G. Li, W. Gao, and W. Gao. 2024. MPEG Geometry-based Point Cloud Compression (G-PCC) Standard. In *Point Cloud Compression: Technologies and Standardization*.

[6] P. Mortimer, R. Hagmanns, M. Granero, T. Luettel, J. Petereit, and H. J. Wuensche. 2024. The GOOSE dataset for perception in unstructured environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[7] X. Wu, L. Jiang, P. S. Wang, Z. Liu, X. Liu, Y. Qiao, and H. Zhao. 2024. Point Transformer v3: Simpler Faster Stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[8] K. You, T. Chen, D. Ding, M. S. Asif, and Z. Ma. 2025. Reno: Real-time neural compression for 3D LiDAR point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[9] L. Zhang, R. Li, C. Ren, S. Di, J. Liu, J. Huang, and K. Zhao. 2025. LCP: Enhancing Scientific Data Management with Lossy Compression for Particles. In *Proceedings of the ACM on Management of Data*.

Received 25 August 2025