False Confidence: Is Detection Reliable Enough for Autonomous Vehicles in LiDAR Point Cloud Corruption?

Grafika Jati[a,*], **Martin Molan**[a], **Francesco Barchi**[a], **Andrea Bartolini**[a], **Giuseppe Mercurio**[b] **and Andrea Acquaviva**[a]

[a]University of Bologna, Italy
[b]FEV Italia s.r.l., Italy

Abstract. LiDAR is a critical sensor for autonomous vehicles, helping them detect and understand their surroundings. But what happens when the sensor itself is compromised? In real-world driving, LiDARs can be affected by everyday contaminants like water, mud, dust, or engine oil. These substances can distort the data, causing the vehicle to miss objects—or worse, see things that aren't there—with dangerous levels of confidence. Most current AI models are trained in ideal, clean conditions and are not prepared for these real-world challenges.

Our research tackles this problem by creating the first dataset of real-world corrupted point cloud data affected by physical contamination. We also develop an intelligent safety layer that detects when LiDAR data is unreliable—before it affects critical decisions like braking or turning. Designed to run on lightweight edge devices, our system is fast, efficient, and adaptable. In future work, we aim to make the system smarter through few-shot learning and uncertainty-based decision support, helping autonomous vehicles become safer and more trustworthy in the messy, unpredictable real world. Data and code will be available at: https://gitlab.com/ecs-lab/lidaroc and https://gitlab.com/ecs-lab/anzil.

1 Resilience is an Overlook Problem

Sensor contamination presents a critical challenge distinct from environmental conditions, as it directly affects raw sensor data, leading to distorted or incomplete point clouds ([18]). Studies indicate that even synthetic contamination significantly degrades performance, with Gaussian noise reducing object detection accuracy from 80.57 to 61.20 ([5], [6, 24, 19] [8]) and LiDAR point reduction further impacting performance ([1]).

Not only does it degrade accuracy, but more critically, it leads to high-confidence false positives, false negatives, and ghost objects. Built-in object detection systems often fail to address contamination effectively, leading to false detections with high confidence scores (e.g., 0.9 out of 1.0) [10]. Real-world contaminants such as mud and lubricant degrade point cloud intensity and distort object geometry, causing detection failures—including the ego-vehicle being misclassified or pedestrians being missed (see Fig. 1). These issues pose critical safety risks, such as false emergency braking, yet are often unaddressed by built-in object detectors trained on clean data. Despite the severity, real-world contamination in LiDAR perception remains underexplored, underscoring the need for robust, contamination-aware detection systems [16], [2].

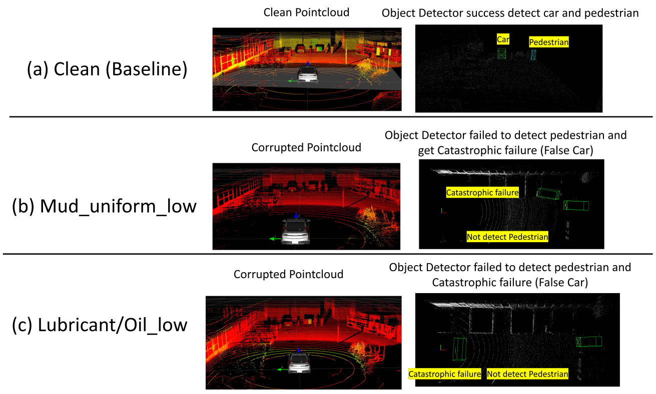

Figure 1. Example Pointcloud and 3D object detection result in data $20m$ for (a) clean LiDAR, (b)corrupted due to mud_uniform_low, (c)corrupted due to lubricant/oil_low.

2 Contributions

This research addresses a critical gap in autonomous vehicle perception: the lack of robustness under real-world LiDAR sensor contamination. Existing object detection models often fail silently when exposed to physical contaminants such as water, mud, or lubricant, producing high-confidence false detections that can lead to catastrophic decisions. To tackle this, we present a comprehensive, multi-level solution.

Real-World Point Cloud Corruption Data We present LIDAROC [10], a real-world LiDAR corruption dataset collected using a 128-channel RS-Ruby automotive-grade LiDAR. The dataset includes six physical contaminants—water, mud, dust, salt water, lubricant, and foam—applied at three intensity levels via spraying, dripping, and drying (3). LIDAROC comprises three subsets: tunnel-5m, outdoor-10m, and outdoor-20m, covering different sensor-to-object distances. Target objects include cars, motorcycles, pedestrians, and reflectors (figure 4). An additional road-driving subset with in-situ contamination is also included (5). LIDAROC addresses the gap in existing large-scale datasets (e.g., KITTI [7], NuScenes [3], Waymo [22]), which are limited to clean, ideal conditions ([20]). LIDAROC is available here [11], [12], [13].

* Corresponding Author. Email: grafika.jati2@unibo.it

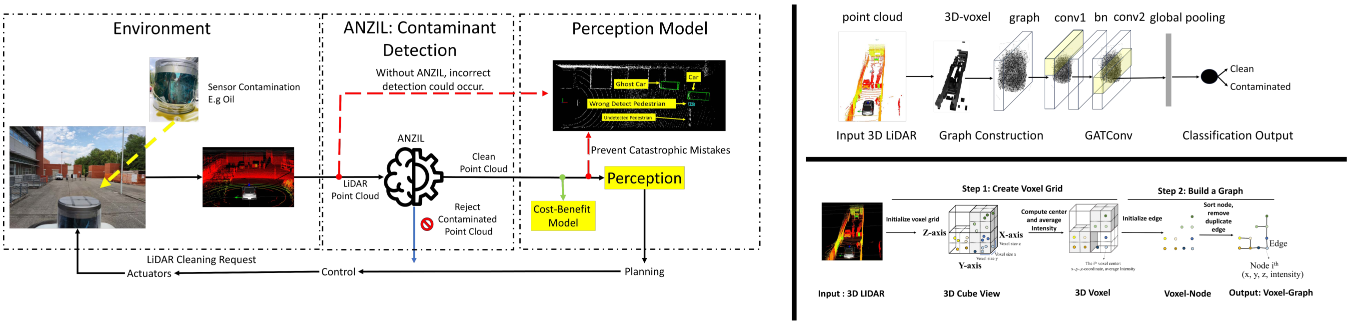

Figure 2. (Left) Proposed integration of contaminant detection as a safety layer in autonomous perception. (Right) Proposed Graph-based LiDAR representation and contaminant detection via Graph Attention Networks.

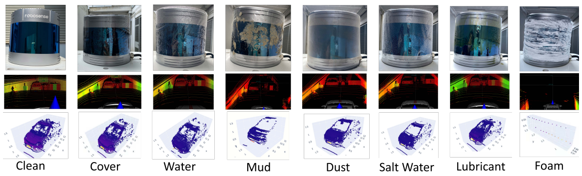

Figure 3. Top: LiDAR contamination procedures in 5m setup: clean, (2) cover-clean, (3–8) water, mud, dust, salt, lubricant, foam. Middle: Scene-level corruption. Bottom: Object-level corruption.

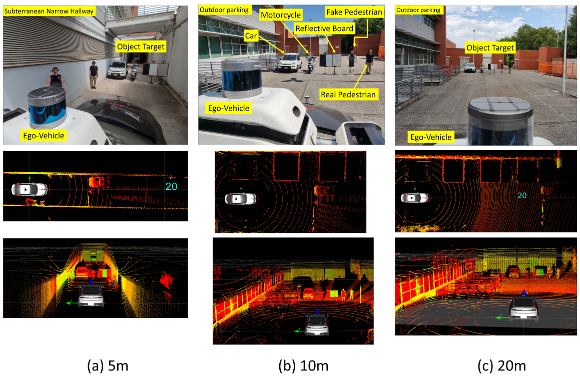

Figure 4. Top: Environment setup. Middle & Bottom: Clean LiDAR point cloud (top view). Target objects show varying point intensity based on material and distance.

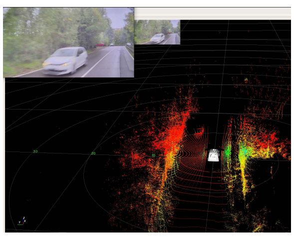

Figure 5. Point cloud of LIDAROC_road_mud in sparse traffic.

Perception Model Robustness Benchmarking We evaluate the detection performance of state-of-the-art 3D object detectors—PointRCNN, SECOND, Part-A$\hat{2}$, and PointPillars—trained on the KITTI dataset [21], [23] using only clean-labeled samples. Following [6], these models are assessed on physically corrupted point cloud, simulating realistic deployment under sensor degradation. Results show a notable decrease in accuracy and frequent high-confidence misdetections, revealing the vulnerability of current detectors to real-world LiDAR corruption [15]. Its open new problem to develop robust object under real-world corrupted data.

Model-Agnostic, Generalizable Safety Layer for AI Pipelines We propose ANZIL, a model-agnostic contaminant detection framework based on graph representations and Graph Attention Networks (GAT) [15] (figure 2). This method acts as a safety layer, filtering out contaminated frames before they are processed by object detection or segmentation modules, thereby enhancing system robustness without requiring modifications to existing models. The use of GATs provides robustness to variations in point cloud density, resolution, and intensity, outperforming traditional CNNs that struggle with sparse spatial representations [14]. Unlike CNNs, which process fixed grids—including empty space—graph-based methods selectively focus on topological relationships between points, resulting in more efficient and reliable feature extraction across diverse scenes and LiDAR type [17]. The interaction between the safety system and downstream perception is illustrated in Figure 2. Finally, we introduce a cost-benefit evaluation methodology that quantifies the trade-off between safety gains (e.g., reduction of high-confidence false detections) and potential data loss, offering a practical metric for real-world deployment.

Edge Deployment and Near-Sensor Classification Fourth, we explore the edge deployment of our method on NVIDIA Jetson AGX Xavier for real-time performance [15], and additionally design TinyLid, a lightweight 2D CNN for contaminant classification implemented on RISC-V-based hardware for near-sensor, low-power inference [9].

3 Future Directions

Binary labeling of LiDAR point clouds is insufficient; severity quantification is required for quality-aware perception. A system must be capable of deciding whether a point cloud should be accepted or discarded based on its quality. As ongoing research, we propose a few-shot LiDAR point cloud processing framework that enables both classification and severity quantification of corruption. This framework will also support real-time adaptation to novel types of contamination through few-shot training. Our work contributes to building robust AI perception models, particularly for object detection. Future work includes investigating point cloud corruption as a physical adversarial attack and its impact on perception reliability, bridging the gap between simulated and real-world data [4], and using generative AI to synthesize realistic corrupted point clouds—advancing robustness from theory to deployment in safety-critical autonomous systems.

References

[1] T. Beemelmanns, Q. Zhang, C. Geller, and L. Eckstein. Multicorrupt: A multi-modal robustness dataset and benchmark of lidar-camera fusion for 3d object detection. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 3255–3261, 2024. doi: 10.1109/IV55156.2024.10588664.

[2] D. Bogdoll, M. Nitsche, and J. M. Zöllner. Anomaly detection in autonomous driving: A survey. In *the IEEE/CVF conference on computer vision and pattern recognition*, pages 4488–4499, 2022.

[3] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.

[4] Y. Chen, Z. Wang, L. Zou, K. Chen, and K. Jia. Quasi-balanced self-training on noise-aware synthesis of object point clouds for closing domain gap. In *European Conference on Computer Vision*, pages 728–745. Springer, 2022.

[5] Y. Dong, C. Kang, J. Zhang, Z. Zhu, Y. Wang, X. Yang, H. Su, X. Wei, and J. Zhu. Benchmarking robustness of 3d object detection to common corruptions. In *the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1022–1032, June 2023.

[6] Y. Dong, C. Kang, J. Zhang, Z. Zhu, Y. Wang, X. Yang, H. Su, X. Wei, and J. Zhu. Benchmarking robustness of 3d object detection to common corruptions. In *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1022–1032, 2023.

[7] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.

[8] S. Huch, L. Scalerandi, E. Rivera, and M. Lienkamp. Quantifying the lidar sim-to-real domain shift: A detailed investigation using object detectors and analyzing point clouds at target-level. *IEEE Transactions on Intelligent Vehicles*, 8(4):2970–2982, 2023. doi: 10.1109/TIV.2023.3251650.

[9] G. Jati, M. Molan, F. Barchi, A. Bartolini, and A. Acquaviva. Tinylid: a risc-v accelerated neural network for lidar contaminant classification in autonomous vehicle. In *the 21st ACM International Conference on Computing Frontiers*, CF '24, page 249–257, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705977. doi: 10.1145/3649153.3649201.

[10] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Lidaroc: Realistic lidar cover contamination dataset for enhancing autonomous vehicle perception reliability. *IEEE Sensors Letters*, 8(9): 1–4, 2024. doi: 10.1109/LSENS.2024.3434624.

[11] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Lidaroc dataset 5m: Realistic lidar cover contamination dataset for enhancing autonomous vehicle perception reliability., July 2024. URL https://doi.org/10.5281/zenodo.12800039.

[12] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Lidaroc dataset 10m: Realistic lidar cover contamination dataset for enhancing autonomous vehicle perception reliability., July 2024. URL https://doi.org/10.5281/zenodo.12800559.

[13] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Lidaroc dataset 20m: Realistic lidar cover contamination dataset for enhancing autonomous vehicle perception reliability., July 2024. URL https://doi.org/10.5281/zenodo.12800632.

[14] G. Jati, M. Molan, J. A. Khan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Autogran: Autonomous vehicle lidar contaminant detection using graph attention networks. In *Companion of the 15th ACM/SPEC International Conference on Performance Engineering*, ICPE '24 Companion, page 112–119, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704451. doi: 10.1145/3629527.3652896.

[15] G. Jati, M. Molan, F. Barchi, A. Bartolini, G. Mercurio, and A. Acquaviva. Anzil: Attention-based network for zero-risk inspection of lidar point cloud in self-driving cars. *Expert Systems with Applications*, 292: 128412, 2025. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2025.128412. URL https://www.sciencedirect.com/science/article/pii/S0957417425020317.

[16] M. A. Khan, H. Menouar, M. Abdallah, and A. Abu-Dayya. Lidar in connected and autonomous vehicles - perception, threat model, and defense. *IEEE Transactions on Intelligent Vehicles*, pages 1–19, 2024. doi: 10.1109/TIV.2024.3510787.

[17] D. Li, C. Lu, Z. Chen, J. Guan, J. Zhao, and J. Du. Graph neural networks in point clouds: A survey. *Remote Sensing*, 16(14), 2024. ISSN 2072-4292. doi: 10.3390/rs16142518.

[18] F. Matos, J. Bernardino, J. Durães, and J. Cunha. A survey on sensor failures in autonomous vehicles: Challenges and solutions. *Sensors*, 24 (16), 2024. ISSN 1424-8220. doi: 10.3390/s24165108.

[19] M. J. Mirza, C. Buerkle, J. Jarquin, M. Opitz, F. Oboril, K.-U. Scholl, and H. Bischof. Robustness of object detectors in degrading weather conditions. In *IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 2719–2724, 2021. doi: 10.1109/ITSC48978.2021.9564505.

[20] J. Ren, L. Pan, and Z. Liu. Benchmarking and analyzing point cloud classification under corruptions. In *International Conference on Machine Learning*, pages 18559–18575. PMLR, 2022.

[21] S. Shi, X. Wang, and H. Li. Pointrcnn: 3d object proposal generation and detection from point cloud. In *the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[22] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[23] O. D. Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. https://github.com/open-mmlab/OpenPCDet, 2020.

[24] X. Zhao, Y. Fang, H. Min, X. Wu, W. Wang, and R. Teixeira. Potential sources of sensor data anomalies for autonomous vehicles: An overview from road vehicle safety perspective. *Expert Systems with Applications*, 236:121358, 2024. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2023.121358.