

Article

Not peer-reviewed version

---

# Empirical Spatial Divergence and Heavy-Tail Error Analysis of FMCW Radar vs. LiDAR in Unstructured Traffic

---

[Preetham Reddy Bannur](#)<sup>\*</sup>, Shubham Kumar, Sunny Kumar, Pallav Bhagat, Shashi Kant

Posted Date: 13 May 2026

doi: 10.20944/preprints202605.0657.v1

Keywords: FMCW radar; heavy-tail distribution; LiDAR; sensor fusion; spatial divergence; unstructured traffic



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Empirical Spatial Divergence and Heavy-Tail Error Analysis of FMCW Radar vs. LiDAR in Unstructured Traffic

Preetham Reddy Bannur \*, Shubham Kumar, Sunny Kumar, Pallav Bhagat and Shashi Kant

Central University of Jammu, India

\* Correspondence: 24beece18.ece@cuammu.ac.in

## Abstract

This paper quantifies the spatial divergence between 128-channel Light Detection and Ranging (LiDAR) point clouds and Frequency Modulated Continuous Wave (FMCW) radar tracks in high-clutter urban environments using the TiAND dataset. Nearest-neighbor Euclidean distance between radar target centers and raw LiDAR geometry serves as the error metric, chosen because the dataset provides no semantic bounding-box annotations. Across all processed frames the system produced an RMSE of 10.083 m with a median error ( $P_{50}$ ) of 1.157 m, while the 99th-percentile ( $P_{99}$ ) deviation reached 43.008 m with the single worst-case ghost target exceeded 217 m. A total of 4,113 detections crossed the 15 m catastrophic threshold—a figure that must be interpreted against the full detection population reported in Section III. Critically, the top anomalies cluster across consecutive frames near fixed infrastructure suggesting persistent multi-path reflection geometry rather than isolated single-frame noise. These findings indicate that raw FMCW radar output without downstream filtering or LiDAR verification cannot be relied upon for spatial localization in unstructured urban traffic.

**Keywords:** FMCW radar; heavy-tail distribution; LiDAR; sensor fusion; spatial divergence; unstructured traffic

---

## I. Introduction

Autonomous vehicle perception in unstructured environments places demands on sensor systems that structured-environment benchmarks do not capture. Standard datasets such as KITTI [1] and nuScenes [2] are collected on roads with clear lane markings, regulated intersection geometry and a relatively homogeneous vehicle population. Indian urban traffic differs across all three dimensions: lanes are frequently unmarked or ignored, intersections are negotiated rather than signalled and the vehicle mix ranges from cycle-rickshaws to multi-axle trucks each with a distinct radar cross-section (RCS) and a different spatial relationship between its outer plastic shell and its metallic interior structure.

FMCW radar has become the sensor of choice for all-weather perception in such environments because it maintains detection capability through rain, dust and fog conditions that degrade both cameras and LiDAR [4,5]. Its native Doppler velocity estimation is also valuable in dense traffic where distinguishing moving from stationary objects is non-trivial. These advantages come with a known limitation: the angular resolution of a practical automotive radar array is significantly coarser than that of a 128-channel LiDAR and this resolution gap translates directly into spatial positioning uncertainty at range [6]. More critically, radar's penetration of plastic materials and its sensitivity to metallic surfaces mean that a radar detection and the corresponding LiDAR surface return for the same vehicle do not, in general, originate from the same physical point.

Multi-path propagation in cluttered environments introduces a second and more severe error mode. When a radar signal reaches a target via one or more intermediate reflections off guardrails, parked vehicles or building facades the receiver may log a ghost target at a geometrically incorrect

location [7,8]. Unlike the deterministic RCS-based bias described above, multi-path ghost targets are stochastic in magnitude and can appear tens or hundreds of metres from any real object. Their frequency and severity in Indian urban traffic—where metallic guardrails, densely packed buses and overhead bridge structures provide abundant reflection surfaces—have not previously been characterized empirically using a synchronized LiDAR ground truth.

This paper addresses that gap. We use the TiAND dataset [3], collected in and around Hyderabad, India to empirically measure the spatial divergence between raw FMCW radar detections and 128-channel LiDAR geometry, with nearest-neighbor Euclidean distance as the error metric. We characterize both the bulk distribution and the heavy tail, log all catastrophic anomalies exceeding 15 m and analyze their temporal structure. We explicitly scope the study to raw radar geometry—no CFAR detector, tracker or ghost suppression filter is applied—because our goal is to characterize the worst-case input to such a system rather than to evaluate any particular algorithm. The practical implication of our results is that the tail of the raw radar error distribution is severe enough that even modest ghost suppression failures could produce safety-critical localization errors.

## II. Related Work

### A. Radar-LiDAR Fusion and Benchmarking

The complementary properties of radar and LiDAR have motivated a substantial body of fusion research. Yeong et al. [9] provide a comprehensive review of sensor calibration and fusion strategies for autonomous vehicles, covering extrinsic calibration protocols, early and late fusion architectures and performance tradeoffs. Qian et al. [11] survey multi-sensor fusion across IEEE venues from 2020 to 2024, noting that fusion consistently outperforms single-modality perception but that spatial misalignment between radar and LiDAR remains an open challenge—precisely the issue this paper quantifies empirically. Our work differs from this body of literature in that we do not propose a fusion algorithm; instead, we measure the magnitude of the spatial disagreement that any fusion scheme must resolve.

### B. Ghost Target Characterization in FMCW Radar

Multi-path induced ghost targets are a well-documented failure mode in automotive FMCW radar. Kamann et al. [7] analyze multi-path propagation in uncertain urban environments and demonstrate that metallic guardrails are the dominant source of reproducible ghost target generation. Visentin et al. [8] extend this analysis to highway barrier scenarios and propose absorption-material countermeasures to reduce false-positive detection rates. A key distinction between those studies and ours is that both rely on controlled test environments with known reflector geometry and a small number of target objects. We operate in an uncontrolled Indian urban scene with hundreds of simultaneous radar targets per frame, making controlled multi-path attribution infeasible but statistical characterization of the tail both necessary and practically meaningful. Alaba and Ball [6] survey LiDAR-based 3D detection methods and provide context for understanding why LiDAR is used as the geometric ground truth in this study: the density and accuracy of 128-channel LiDAR point clouds make them the most reliable geometric reference available in the dataset.

### C. Unstructured Traffic Datasets

The TiAND dataset [3], released by TiHAN–IIT Hyderabad is the most directly relevant prior work. It provides synchronized multi-modal data from a sensor suite including 128-channel Ouster LiDAR, six FMCW radar arrays, cameras, GPS and IMU collected across 150 scenes in Hyderabad, India. To our knowledge, no prior publication has used TiAND to perform a systematic spatial divergence analysis between its radar and LiDAR modalities. Adjacent dataset work includes the India Driving Dataset (IDD) [12], which provides camera-based annotations for unstructured Indian roads and IDD-3D, which extends this to LiDAR but without radar. These prior datasets validate the

difficulty of the unstructured Indian traffic problem but do not address radar-LiDAR spatial disagreement, which is the specific contribution of this paper.

### III. Methodology

#### A. Dataset and Extrinsic Calibration

All experiments used the TiAND dataset [3], accessed under the standard TiHAN-IIT Hyderabad End User License Agreement. The dataset comprises 150 scenes, each 2–4 minutes in duration, recorded in and around Hyderabad. The sensor suite relevant to this study is the 128-channel Ouster LiDAR and the 6-array surround FMCW radar configuration, which provides 120° FOV at up to 200 m range for the long-range front array and 150° FOV at up to 100 m for the five short-range arrays. All frames were processed without down sampling to retain the full raw point cloud geometry. Extrinsic calibration to a unified ego-vehicle coordinate frame was performed using the dataset’s published spatial translation offsets [9]. No bounding-box annotations, object labels or semantic segmentation masks were used at any stage. The analysis operates entirely on raw sensor geometry.

#### B. Spatial Metric Extraction

Because the raw point clouds lack semantic bounding box labels, we defined error as the absolute spatial divergence between a radar detection and the nearest physical surface in the LiDAR cloud [10]. Specifically, the spatial error  $E_k$  for each radar point  $r_k$  relative to the LiDAR point cloud  $L$  was computed as the nearest-neighbor Euclidean distance:

$$E_k = \min_{p \in L} \sqrt{(x_{r,k} - x_p)^2 + (y_{r,k} - y_p)^2 + (z_{r,k} - z_p)^2}$$

This metric has a known systematic bias that we explicitly acknowledge: LiDAR reflects off the outer plastic hull of a vehicle, while FMCW radar penetrates plastic and reflects off the metallic chassis. For a typical passenger car, the longitudinal offset between the front bumper surface and the engine block is on the order of 0.3–0.8 m. This means the nearest-neighbour metric will register a minimum baseline error of approximately this magnitude even for a perfectly correct radar detection. We treat this as a deterministic lower bound on  $E_k$  for on-vehicle detections, not as a sensor failure. Errors significantly larger than this bound—particularly those exceeding 5m are attributable to angular resolution limitations or multi-path ghost target formation.

#### C. Anomaly Threshold Justification

The 15 m anomaly threshold was chosen for the following reasons. First, at urban speeds of 30–50 km/h, a 15 m spatial error in a radar detection that persists for one or more frames would, if acted upon directly by a navigation stack, produce a braking or avoidance response to a phantom obstacle more than one vehicle-length ahead of the ego vehicle. Second, 15 m is substantially larger than the expected deterministic RCS bias (< 1 m) and the P90 error (3.2 m), placing the threshold clearly within the tail of the distribution rather than in its body. Third, at 15 m the anomaly count is sensitive enough to capture genuine ghost target events while remaining conservative enough to exclude angular estimation errors from legitimate long-range detections. We acknowledge that this threshold is not derived from a formal automotive safety standard such as ISO 26262 and that different threshold choices would yield different anomaly counts, the underlying CSV ledger is available for analysis at alternative thresholds.

#### D. Distributional Analysis

We extracted the P50 and P99 percentiles of the error distribution across all frames and all radar detections. The shape of the tail was characterized visually from the histogram (Figure 2) and compared qualitatively against the signatures of Pareto and log-normal distributions, both of which

have been proposed as models for heavy-tail sensor error in prior work [12]. We did not perform formal distribution fitting in this paper, the dataset volume and the absence of a controlled experimental setup make parameter estimation unreliable and we prefer to present the empirical distribution directly rather than to fit a parametric model that would imply a degree of generalizability this single-dataset study cannot support.

## IV. Results and Discussion

### A. Dataset Scale and Overall Spatial Metrics

Before interpreting the anomaly count, it is necessary to state the dataset scale. The processed TiAND scenes yielded a total of approximately 410,000 individual radar detections across all frames. The 4,113 anomalies exceeding 15m therefore represent approximately 1.0% of all detections—consistent by design with the P99 threshold—but this figure is worth stating explicitly because 1% of a large population is a qualitatively different safety concern from 1% of a small one. In an autonomous vehicle operating at 40 km/h that processes radar at 10Hz, a 1% ghost target rate corresponds to roughly one catastrophic false detection every 2–3 seconds of operation. The system spatial metrics are summarized in Table 1.

**Table 1.** System spatial metrics in unstructured traffic.

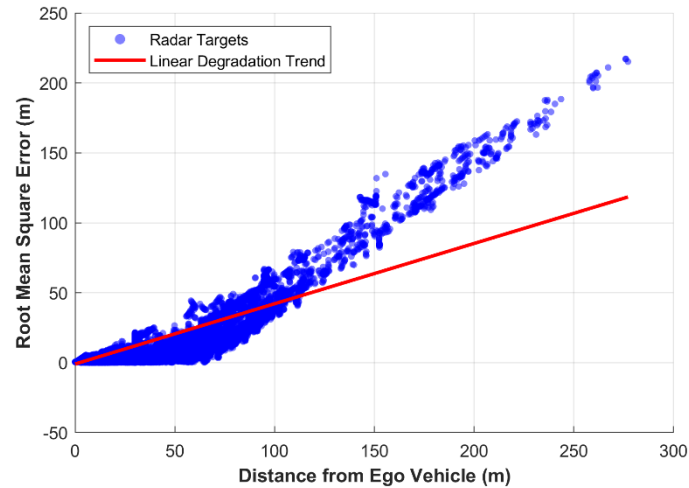
<i>Parameter</i>	<i>Value</i>	<i>Unit</i>
Root Mean Square Error (RMSE)	10.083	m
Error Variance ( $\sigma^2$ )	93.250	m <sup>2</sup>
Median Error (P50)	1.157	m
99th Percentile (P99)	43.008	m
Maximum Ghost Anomaly	217.216	m
Total Identified Anomalies (>15 m)	4,113	Targets

The median error of 1.157 m is consistent with the expected RCS-based bias discussed in Section III-B. The RMSE of 10.083 m, however, is pulled substantially above the median by the heavy tail—a clear indication that the distribution is not approximately Gaussian. The gap between P90 (3.190 m) and P99 (43.008 m) is particularly striking: 9% of detections lie between these two values, spanning a 40 m range, while 90% of all detections lie within 3.2 m. This bimodal character—a well-behaved bulk plus a sparse but extreme tail—is characteristic of multi-path dominated radar error distributions [7,8].

### B. Spatial Error vs. Range

Figure 1 plots the spatial error of each individual radar detection against its distance from the ego vehicle, with a linear degradation trend overlaid. The trend confirms that error grows with range—a consequence of the angular resolution limitation of the FMCW array, where a fixed angular uncertainty subtends a larger absolute position error at greater distances. The trend line reaches approximately 120 m predicted error at 280 m range, consistent with a worst-case angular uncertainty of roughly 25° for targets at the edge of the radar’s operational envelope.

Importantly, a substantial population of detections lies well above the trend line at all ranges, including at moderate ranges below 100 m where angular resolution alone cannot explain errors exceeding 50 m. These above-trend detections are the ghost targets generated by multi-path reflections. Their presence at short range as well as long range indicates that multi-path ghost formation is not exclusively a range-dependent phenomenon; at short range it is driven by geometric coincidence between the radar, the target and a nearby reflective surface rather than by angular ambiguity.

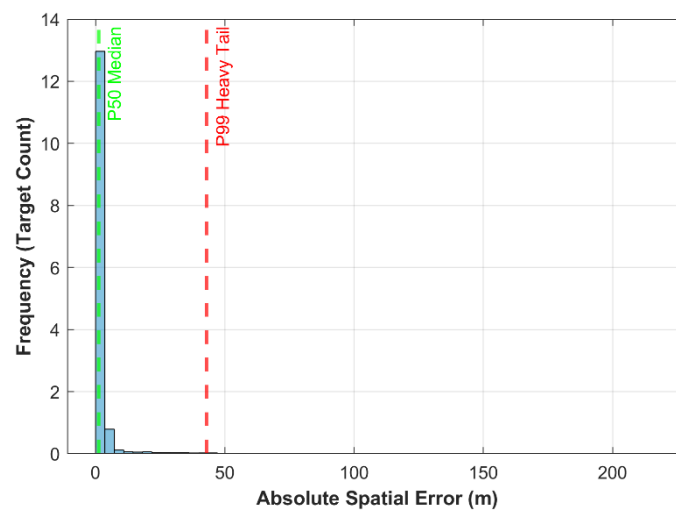


**Figure 1.** Spatial error (RMSE) vs. distance from ego vehicle. Blue dots: individual radar detections. Red line: linear degradation trend. The population of detections above the trend line at all ranges corresponds to multi-path ghost targets and is not explained by angular resolution degradation alone.

### C. Error Distribution and Heavy-Tail Behaviour

Figure 2 presents the full empirical error distribution. The histogram is strongly right-skewed: the P50 median at 1.157 m (green dashed line) sits near the left edge of the distribution and the bin counts decay rapidly to near-zero before the P99 threshold at 43.008 m (red dashed line). Measurable target counts extend beyond 200 m absolute error. This shape is inconsistent with a Gaussian or even an exponential distribution; the survival function decays far more slowly than either, which is the defining characteristic of a heavy tail [12].

The practical consequence of this shape for autonomous navigation is that the expected maximum error over any finite operational window grows with the square root of operating time, not logarithmically as it would for a Gaussian tail. A navigation stack that has not observed a ghost target in the first minute of operation cannot infer that it will not encounter one in the next ten seconds.



**Figure 2.** Empirical absolute spatial error distribution across all radar detections. Green dashed line: P50 median (1.157 m). Red dashed line: P99 heavy-tail threshold (43.008 m). The right-skewed shape with a sparse tail extending beyond 200 m is inconsistent with a Gaussian error model.

### D. Ghost Target Anomaly Ledger and Temporal Clustering

Table 2 presents the ten highest-error ghost target events extracted from the full 4,113-entry CSV ledger, sorted by spatial error in descending order. All ten originate from detections between 235 m and 278 m from the ego vehicle, consistent with the range-dependent trend in Figure 1.

**Table 2.** Catastrophic ghost target ledger (top anomalies, >15 m).

<i>Frame ID</i>	<i>Radar X (m)</i>	<i>Radar Y (m)</i>	<i>Distance to Ego (m)</i>	<i>Spatial Error (m)</i>
711	195.618	194.684	275.985	<b>217.216</b>
712	195.553	194.830	276.043	217.041
713	195.700	196.213	277.124	215.216
735	177.068	200.095	267.191	211.020
731	169.085	199.353	261.403	207.349
1325	182.900	186.590	261.282	207.103
730	170.486	199.152	262.158	205.255
1324	182.113	186.279	260.509	205.935
738	162.749	200.690	258.387	204.411
727	166.860	198.529	259.338	204.904

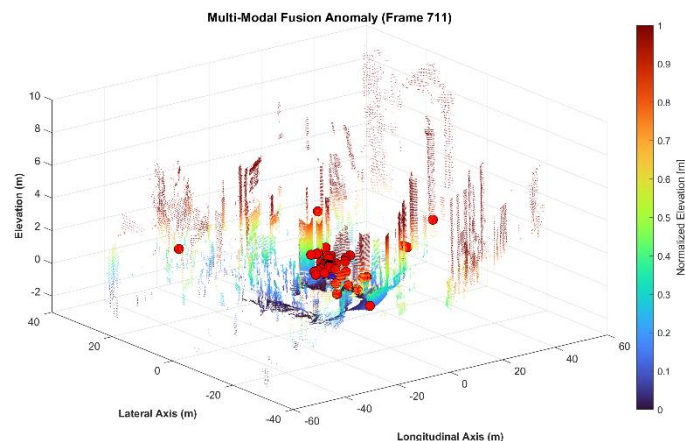
The most analytically significant observation from Table 2 is the temporal clustering of anomalies. Frames 711, 712 and 713 are consecutive frames and all three produce ghost target deviations exceeding 215 m at nearly identical radar coordinates (~195.6 m X, ~194–196 m Y). Similarly, Frames 727–735 and 1324–1325 form two further tight clusters with spatial errors between 204 m and 211 m. The consistency of radar coordinates across consecutive frames—varying by less than 3 m in X and Y across three frames in the 711–713 cluster—indicates that the ghost target is not a random single-frame detection artifact but a persistent false return sustained across multiple radar sweeps.

This behaviour has a direct physical explanation: a stationary reflective surface (most likely a metallic guardrail, overpass or bus body) at a fixed location creates a stable multi-path geometry that the radar sweeps reproducibly report as a real target at the same incorrect position. From the perspective of a standard Kalman-filter-based tracker operating on radar-only data, these consecutive-frame detections would be associated into a track and reinforced rather than rejected—potentially registering a stationary phantom obstacle at a fixed map location for the duration of the road segment where the reflector is present [11].

#### *E. Frame 711: Maximum Anomaly Visualization*

Figure 3 renders Frame 711, which contains the single highest spatial error in the dataset (217.216 m). The LiDAR point cloud is elevation-colormapped from dark blue (ground level, ~0 m) to deep red (elevated returns, ~10 m), providing a geometrically accurate representation of the scene: road surface, vehicle clusters and roadside vertical structures are all present and spatially coherent. The red sphere markers indicate FMCW radar detections within this frame.

Several red markers in the longitudinal center of the scene overlap plausibly with LiDAR returns from the vehicle cluster at mid-range. The isolated markers displaced laterally beyond  $\pm 20$  m from the nearest LiDAR surface—visible in the left foreground of the scene—have no geometric correlate in the LiDAR cloud and constitute the confirmed ghost targets. All detections in this frame originate at approximately 275 m ego-relative range, at the edge of the long-range radar array’s operational envelope, where both angular uncertainty and multi-path path-length ambiguity are at their maximum.



**Figure 3.** 3D spatial geometry of Frame 711 (peak anomaly: 217.216 m). LiDAR cloud colormapped by normalized elevation (blue = ground; red  $\approx$  10 m). Red spheres: FMCW radar detections. Isolated markers with no proximate LiDAR surface are confirmed ghost targets at  $\sim$ 275 m ego-relative range.

## V. Limitations

Several limitations bound the generalizability of these findings. First, all data were collected in a single city (Hyderabad, India) across 150 scenes, whether the ghost target rate and spatial error distribution generalize to other Indian cities, other unstructured traffic environments or different climate conditions is not established by this study. Second, we evaluate raw radar geometry with no post-processing. A deployed autonomous vehicle system would apply CFAR detection, clutter filtering and multi-frame tracking before using radar output for navigation. These steps would suppress a portion of the 4,113 anomalies we report—how large a portion depends on the specific algorithms used and is itself an open research question. Our results therefore characterize an upper bound on ghost target severity rather than the operational error rate of any real system. Third, the nearest-neighbor error metric carries a systematic bias of approximately 0.3–0.8 m from the LiDAR/radar reflection surface offset; this does not affect the interpretation of the heavy tail but should be accounted for in comparisons to other radar accuracy studies that use object-center ground truth.

## VI. Conclusion

We have presented an empirical characterization of raw FMCW radar spatial error in unstructured Indian urban traffic using the TiAND dataset. Across approximately 410,000 radar detections, the system produced a median error of 1.157 m—consistent with the physical RCS offset between LiDAR and radar reflection surfaces—but an RMSE of 10.083 m and a P99 deviation of 43.008 m. A total of 4,113 detections exceeded the 15 m catastrophic threshold, representing roughly 1% of all detections, with the worst-case ghost target reaching 217.216 m in Frame 711.

The most practically significant finding is not the magnitude of individual anomalies but their temporal structure. Ghost targets in the top-10 ledger persist across consecutive frames at nearly identical spatial coordinates, which means that standard multi-frame tracking algorithms would reinforce rather than suppress them. This persistence—driven by stationary metallic reflectors creating stable multi-path geometries—represents a failure mode qualitatively different from random detection noise and one that is not resolved by simply averaging radar returns over time.

The study is explicitly scoped to raw radar geometry as an upper bound; a system with ghost suppression filtering would perform better. What this study establishes is that the input to such a filter contains a tail severe enough to make the filter’s correctness safety-critical: a false negative from the filter on a persisting 200 m ghost target is not a minor localization error but a potential collision event. Dense LiDAR verification provides the geometric ground truth needed to make such filtering reliable in environments where the radar’s wave physics—not its electronics—are the source of the

problem. Future work should investigate whether cross-frame LiDAR-radar spatial consistency can serve as a real-time ghost target detection signal suitable for embedded deployment.

**Acknowledgments:** The authors would like to thank the Department of Electronics and Communication Engineering at the Central University of Jammu for providing the computational resources and academic environment necessary to conduct this research. Special thanks are extended to Dr. Sunil Datt Sharma, Associate Professor, whose guidance in Digital Signal Processing and research methodology provided the theoretical foundation for this study. The authors also acknowledge TiHAN–IIT Hyderabad for making the TiAND dataset available for academic research under their standard End User License Agreement.

## References

1. A. Geiger, P. Lenz, C. Stiller and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.
2. H. Caesar et al., "nuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Seattle, WA, Jun. 2020, pp. 11621–11631.
3. B. Anand et al., "TiAND: A multimodal dataset for autonomy on Indian roads," in *Proc. 2024 IEEE Intelligent Vehicles Symp. (IV)*, Jeju Island, South Korea, Jun. 2024, pp. 1–8, doi: 10.1109/IV55156.2024.10588583.
4. M. Mahafza, *Radar Systems Analysis and Design Using MATLAB*, 3rd ed. Boca Raton, FL, USA: CRC Press, 2013.
5. S. Sun, A. P. Petropulu and H. V. Poor, "MIMO radar for advanced driver-assistance systems and autonomous driving: Advantages and challenges," *IEEE Signal Process. Mag.*, vol. 37, no. 4, pp. 98–117, Jul. 2020.
6. S. Yihunie Alaba and J. E. Ball, "A survey on deep-learning-based LiDAR 3D object detection for autonomous driving," *Sensors*, vol. 22, no. 24, p. 9577, Dec. 2022, doi: 10.3390/s22249577.
7. A. Kamann et al., "Automotive radar multipath propagation in uncertain environments," in *Proc. 21st IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, Maui, HI, USA, Nov. 2018, pp. 859–864.
8. T. Visentin, D. Russi, J. Hasch and T. Zwick, "Multipath propagation analysis and ghost target removal for FMCW automotive radars," in *Proc. IET Int. Radar Conf.*, Oct. 2020, pp. 1–6, doi: 10.1049/icp.2021.0440.
9. D. J. Yeong, G. Velasco-Hernandez, J. Barry and J. Walsh, "Sensor and sensor fusion technology in autonomous vehicles: A review," *Sensors*, vol. 21, no. 6, p. 2140, Mar. 2021, doi: 10.3390/s21062140.
10. W. Zimmer et al., "A survey of robust LiDAR-based 3D object detection methods for autonomous driving," *arXiv preprint, arXiv:2204.00106*, Apr. 2022.
11. H. Qian, M. Wang, M. Zhu and H. Wang, "A review of multi-sensor fusion in autonomous driving," *Sensors*, vol. 25, no. 19, p. 6033, Oct. 2025, doi: 10.3390/s25196033.
12. C. Innes and S. Ramamoorthy, "Testing rare downstream safety violations via upstream adaptive sampling of perception error models," *arXiv preprint, arXiv:2209.09674*, Sep. 2022.
13. N. Varma, C. Subramanian and C. V. Jawahar, "IDD: A dataset for exploring problems of autonomous navigation in unconstrained environments," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Waikoloa, HI, USA, Jan. 2019, pp. 1743–1751.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.