Joint Camera and LiDAR Risk Analysis

Oliver Zendel*, Johannes Huemer, Markus Murschitz, Gustavo Fernandez Dominguez, Amadeus Lobe
AIT Austrian Institute of Technology Gmbh, Austria

Abstract

Sensor fusion of camera-based sensors and LiDAR is widely used in autonomous systems with the premise of increased robustness. Contrarily, due to their overlapping functional principles, they also share many risk factors that may result in degraded operation. This work applies the risk analysis method Hazard and Operability study (HAZOP) to LiDAR sensors and connects it with an existing camera-based HAZOP. This systematic approach leads to a structured listing of potential sources of data quality reduction in LiDAR data. Many risk factors identified for camera-based systems (e.g. transparency or reflections) can be correlated to degradation in the corresponding LiDAR data. To validate our findings, the public dataset A2D2 is analyzed for such co-occurring camera-LiDAR risk factors. Additionally, experiments under controlled laboratory conditions are performed to quantify the impact of various identified risks. Our HAZOP results are released publicly and are intended to improve the design and usage of sensor systems as well as training and test datasets for safer autonomous systems.

1. Introduction

Autonomous driving systems rely on camera data for detection, navigation, and other safety-relevant tasks. An ongoing discussion is the benefit of additional sensor modalities to increase robustness and prevent dangerous situations due to corrupted camera data. LiDAR sensors are gaining popularity as they can provide reliable depth data in harsh conditions. These systems measure the time of flight by emitting a laser pulse and detecting its reflection. The correspondence problem is solved by emitter beam encoders. Distances can be calculated from the measured time which is then aggregated into 3D point clouds. Both cameras and LiDAR sensors share many similarities as both rely on light and interactions of objects with light. This poses a difficult question: are LiDAR sensors susceptible to the same influences which may already corrupt camera data? If so, then simply adding an additional LiDAR sensor to increase robustness of a computer vision task might not work as planned.

While several models for data reconstruction, object detection, object tracking and precision positioning using LiDAR data were presented in last years [7, 21, 27, 35, 37–39, 41], little work has been dedicated to analyzing the risk factors and their impact on the acquired data.

This work shows a systematic approach toward identifying potential performance-reducing conditions (called hazards) which may impact both camera and LiDAR output quality at the same time. Risk factors for LiDAR sensors used for autonomous driving are analyzed by applying the well-established risk analysis method HAZOP [20] to LiDAR sensors and connecting it with an existing camera-based CV-HAZOP [44]. In the context of this document, risks and hazards are to be understood as in ISO 21448:2022 Safety of intended functionality [17], as insufficiencies of the performance of electronic elements in the system.

Section 2 revisits previous works regarding LiDAR data quality and joint autonomous driving LiDAR/Camera datasets. The main contributions of this work (summarized in Section 7) are:

- A novel model of a generic LiDAR system (Section 3),

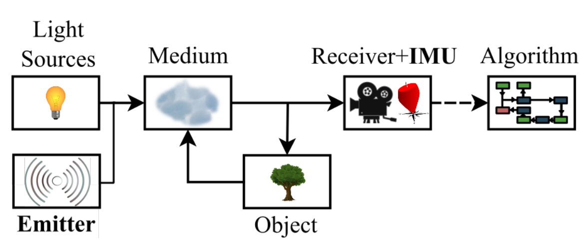

Figure 1. LiDAR-HAZOP generic system model. Boxes represent locations, solid lines signal/noise transfer by light, and dashed lines transfer by digital data. Bold font indicates change compared to the existing CV-HAZOP model.

*Thanks to Philipp Glira and Christian Zinner for contributing to the HAZOP list. This project received funding from EU JU grant No 692480 www.iosense.eu. Corresponding Author: oliver.zendel@ait.ac.at

- The first published HAZOP analysis including LiDAR sensor (Section 4) resulting in a list of 136 combined camera and LiDAR hazards.

- Two independent LiDAR hazards evaluation methods: First, by analyzing the A2D2 dataset for instances where both camera and LIDAR data fail (Section 5) and second, by conducting controlled laboratory experiments testing for specific hazards with quantifiable impact (Section 6).

Figure 2. Examples from A2D2 data [12]. Top: frame 20181016082154 with strong overexposure on left building, Middle: frame 20181204135952 where fog blocks most of the point cloud, Bottom: frame 20180823750 with low albedo traffic signs. Each example shows a camera image with LiDAR depth superimposed and the top view of the corresponding point cloud.

2. State-of-the-Art

LiDAR sensors have seen a rise in popularity in the field of autonomous driving, due to lower prices, power consumption, compact designs, and improved data quality.

Amann et al. [1] pointed out that the main sources of inaccuracy in laser range finders are noise-generated timing jitter, speckle noise, drift, non-linearity, and timing errors. Lichti et al. [23, 24] list typical mechanical, electrical, and calibration errors for terrestrial laser scanners. Yang and Wang [42] addressed the problem of mirror reflection using LiDAR information and they summarised how surface properties affect the amount of light reflected, absorbed, and transmitted. Petit and Shladover [31] analyzed the threats on autonomous vehicles and described possible attacks as well as relevant security hazards for autonomous fleets. Using 3D LiDAR data, Rachman presented a framework for multi-object detection and multi-object tracking on urban environments [34]. Goodin et al. [14] developed a model for the performance degradation of LiDAR sensors under rain conditions in ADAS applications. Li and Ibanez-Guzman [22] reviewed the SotA of automotive LiDAR technologies and related algorithms. Henley et al. [15] discuss multi-beam LiDAR to handle specular surfaces.

Yulong et al. [5] specialize in adversarial sensor attacks with specialized signal spoofing. Their recent work [4] includes an overview of potentially malicious attacks against LiDAR sensors. They used a custom-made photodiode receiver with a delay circuit and attack laser to replay the original emitter signal and thus confuse LiDAR systems resulting in artificial points in the 3D data. In a second setup, signals with an intensity below the minimum operational threshold are spoofed, so that they get filtered out and create deliberate gaps in the point cloud. Their work is investigating potential security risks deliberately created by malicious attackers, while the hazard analysis in this work focuses on safety hazards during regular operation.

CV-HAZOP [44, 45] applies HAZOP (Hazard and Operability Analysis) [20], a procedure devised by the safety community to validate complex systems, to camera-based sensors. It is designed to systematically search and identify difficult, performance-decreasing situations and aspects. The main steps of the risk analysis are: (i) model the system, (ii) partition the model into subcomponents (i.e. locations), (iii) find appropriate parameters for each location which describe its configuration, (iv) define useful guide words, (v) try to find meanings for each guide word/parameter combination, and derive consequences as well as hazards from each meaning. It provides a public referenceable list [10] of entries representing visually challenging hazards applicable to computer vision tasks.

Multiple datasets were released to the scientific community in last years containing both camera and LiDAR data intended for autonomous vehicles [3, 6, 8, 11, 12, 16, 19, 28–30, 32, 33, 36, 40][1]. Section 5 is evaluating potential performance degradation affecting both camera and LiDAR data from the Audi Autonomous Driving Dataset (A2D2) [12]. This was chosen as it includes visually challenging situations with diverse scenery and weather conditions (examples in Figure 2).

This paper differs from the aforementioned work by analyzing safety risk factors of LiDAR sensors, applying the risk analysis HAZOP [20] to LiDAR sensors, and connecting it with an existing camera-based HAZOP [44]. To the best knowledge of the authors, there is no published work considering a LiDAR system as a whole identifying the risks on a generic level.

[1] See supplemental for a detailed specification survey on these datasets

3. Model

The foundation for the following risk analysis is a model of the system to be analyzed. The generic model introduced in [44] is used as a template for the LiDAR-HAZOP. Two changes transform the generic CV-HAZOP model into a generic LiDAR model (depicted in Figure 1):

- New location Emitter: The laser emitter is an important component in every LiDAR sensor. The CV-HAZOP's *Light Sources* (*L.S.*) location is split into two locations for the LiDAR-HAZOP: *Light Sources* and *Emitter*. CV-HAZOP *L.S.* entries refer to both passive (e.g. sun, street lights) and active light sources (e.g. projectors, laser lines). Thus a mapping of either *L.S.* or *Emitter* can be easily achieved: passive *L.S.* entries map to the LiDAR-HAZOP *L.S.* location, active *L.S.* entries to the *Emitter* location. Both new locations are described using the parameters of the same original CV-HAZOP *L.S.* location.

- Registration parameters: Registration of single distance measurements into a coherent coordinate system forms an important aspect of all LiDAR systems. Two new parameters are added to the receiver location[2]: local and global registration.

  The local registration joins individual distance measurements of one laser sweep resulting in a sensor coordinate system. This is typically done using internal signals (e.g. time-based and angle encoders) and is shared between encoders on the emitter and receiver end. Emitter parts that help local registration are still represented at the *Emitter* location to improve functional hierarchy within the risk analysis.

  The global registration joins multiple sweeps into one global coordinate system. It typically uses inputs from external forces and signals (e.g. IMU, GNSS).

All remaining locations and parameters of the original CV-HAZOP are reused in the LiDAR-HAZOP. All existing CV-HAZOP entries can be mapped to the LiDAR-HAZOP hierarchy and vice-versa (excluding *Registration* parameters). This parity is a crucial feature to allow reuse and cross-references between both risk analyses.

4. Risk Analysis

The execution of the LiDAR HAZOP risk analysis follows the steps of CV-HAZOP: experts interpret each parameter/guide-word combination to assign meaning in the context of LiDAR data generation.

[2]The CV-HAZOP location *Observer* is named *Receiver* for easier readability.

The existing stereo vision HAZOP result list [43] is used as a baseline due to the likeliness of their output data: stereo vision creates depth maps and LiDAR point clouds. The camera-based HAZOP entries are used as an additional inspiration to identify potentially connected hazards for LiDAR data.

LiDAR hazards identified by individual experts are collected and discussed to create a joint uniform result. In total five experts created a result list with 136 unique entries for LIDAR hazards, each linked to a pre-existing stereo vision hazard. Table 1 shows an exemplary reduced excerpt of the result list[3].

5. Dataset Evaluation

This section showcases a practical application using LiDAR-HAZOP entries to find critical cases. The existing dataset A2D2 containing automotive scenes with both camera and LiDAR data is examined in a qualitative analysis. Frames containing a visual hazard (i.e. safety-relevant degradation of camera images) are detected using classifier networks. For each identified frame the corresponding point cloud is manually checked for signs of degradation in the LiDAR data. This should highlight cases where both sensor modalities fail. The same procedure can be applied to other existing datasets or during the creation of new datasets.

5.1. Visual Hazard Detection

Automatic visual hazard detection based on image classifiers allows for a quick reduction of potential hazard frames [46]. The existing classifiers are based on the Wilddash 2 dataset hazard labels and are identical to the ones presented in the dataset paper. From these ten hazard detectors in [46], only some are applicable for the mixed-data analysis on A2D2. The following hazards have been disregarded:

- image blur, compression artifacts, lens distortion, and interior reflections (*screen*) are characteristics of the camera sensor itself (or the mounting position thereof).

- motion blur is characteristic of moving objects and artifacts associated with registration. Both factors are relevant and significantly reduce the quality and usefulness of the point cloud data in the analyzed datasets in basically every frame (see subsection 5.2). A frame-by-frame analysis is thus not necessary.

- road coverage and intra-class variations are relevant for specific semantic scene understanding. A similar connected case for point cloud data is possible but requires out-of-distribution examples both in image and 3D data at the same time (which are missing in A2D2).

[3]See supplemental for the full list

Table 1. Examples from LiDAR HAZOP shortened for brevity. HID: hazard identifier, GW: guide word

HID	Location	GW	Parameter	Stereo Entry	LiDAR Entry
125	Light Sources	More	Intensity	Directly lit object is overexposed	Bright light source saturates sensor, signal is lost.
e141	Emitter	Less	Beam	Only a small part of the scene is sufficiently lit.	Emitter fingers are too far apart and therefore missing crucial details.
449	Object	Less	Texture	Two objects on epipolar lines have very little texture thus allowing a mismatch	Signal is canceled by an object's surface (e.g. Helmholtz resonator)
476	Object	No	Reflectance	Object appears to have neither texture nor shading due to its low albedo	Object with very low albedo is not returning signal becoming invisible
478	Object	More	Reflectance	Object has shiny material that creates mirror-like reflections	Mirroring parts of an object creates distorted response.
489	Object	Where else	Reflectance	Not applicable	Retro-reflectors create glare effect which surrounds object with copies.
502	Object	More	Transparency	Highly transparent object is invisible	Scenery behind transparent object is measured
729	Objects	Part of	Transparency	Object has transparent parts showing a different object	Transparent parts of object is mixed with background object
883	Reciever Optics	Part of	Viewing position	Near objects are out of focus	Object is missed due to neglecting the minimal working distance

This leaves three potential combined hazard categories: *overexposure*, *particles*, and *underexposure*. The hazard detector is applied on the $41k$ A2D2 frames with semantic ground truth and reduces the manual search space drastically: only 3455 frames have to be checked (see Table 2).

5.2. Typical Degradation of LiDAR Data

The A2D2 dataset supplies camera images and a consolidated and cleaned point cloud for the respective moment the camera image was taken. Sensor sweeps have to be integrated over time. Therefore, ego-motion as well as movement of traffic participants potentially introduce errors. Before publishing A2D2, its creators performed filter operations and pre-processing steps to the raw data. During annotation, these additional sources of error had to be accounted for by checking previous and next frames and observing the quality of data surrounding the potential hazard location. Inconclusive frames were discussed by multiple annotators and among disagreement, the severity of the potential hazard was lowered.

5.3. Evaluation

A simple GUI is used to show a camera image with superimposed LiDAR data and additionally a 3D point cloud viewer of the same scene. Each of the 3455 frames that contain automatically detected potential hazards is then manually checked for degradation visible in the point cloud (either missing data or false phantom data). This results in one of four possible labels: no visual hazard present in camera data (detector failed); LiDAR severity none (no degradation); low (potential degradation but small scope/impact); high (strong degradation). Figure 2 shows examples of frames annotated with high severity for each hazard. Table 2 summarizes the results for all three hazards. In general, the visual hazard detection worked very well creating only minimal additional overhead due to false positives. A large portion of camera overexposure hazards are caused by bright (overexposed) sky without the sun in direct view. The affected overexposed pixels can thus not contribute to the point cloud. This explains the large number of overexposure frames with no LiDAR data corruption (severity *none*). Many examples for particles manifested due to water spray kicked up by cars driving on wet highway roads. The point clouds are missing data for both the water, the road, and vehicles. Thus reducing the quality of the full scene and resulting in a fair number of *high* severity annotations. A2D2 only contains well-lit daylight scenes so underexposure camera hazards are only caused by dark objects. Most are moving cars where data quality is hard to attribute (see 5.2). Such unclear entries were annotated with *low* severity. However, multiple static dark objects (e.g. black signs) creating gaps in the LiDAR data could also be identified.

Table 2. Results of joint hazard analysis for A2D2. Frames preselected using automatic hazard detection in the camera image. Severity is manually annotated per frame and hazard.

Hazard	Total	No Vis.	None	Low	High
overexp.	2691	1	2462	224	4
particles	639	16	197	245	181
underexp.	125	31	42	38	14

6. Experiments

In addition to the LiDAR-HAZOP and its evaluation on A2D2, a subset of hazards is selected for replication in a small-scale, indoor laboratory setup. Subsequently, a quantitative analysis is performed, with the aim of determining the effect of the selected hazards on the LiDAR measurement quality. To this end, the corresponding point cloud recordings are evaluated using four different performance metrics.

6.1. Selection of hazards

In a laboratory environment, several test objects were selected (see Figure 3), each designed to induce at least one of the selected hazards of Table 1:

- Reflective material: HID 478 and 489 are replicated using several highly reflective materials including retro-reflective foils (b), cat-eye reflectors (b) and a reflective white box made of PVC (c).

- Low albedo material: Surfaces with very low albedo, i.e. black, matte objects, have little to no reflectance. To assess this hazardous effect (HID 476) on LiDAR measurements, several low albedo surfaces were prepared: Black 3.0 paint [2] (d), regular black acrylic spray (d), and also a foil (b).

- Transparent material: A variety of documented hazards is related to transparent objects (e.g.: HID 502, HID 729) which are considered to cause confusion in distance measurements. To this end, a cylindrical glass carafe (e), as well as an acrylic glass object (f), are used.

- Overexposure: The last investigated hazard (HID 125) concerns the overexposure of the scene by an additional light source. Thus, an array of IR light sources (g) were added to the scene. Commercially only 850nm and 940nm lamps are available so a mixture was used to potentially affect the LiDAR's 905nm signal (see Figure 8).

6.2. Experimental setup

The hardware setup for the laboratory experiments is depicted in Figure 6. It comprises a low-cost automotive LiDAR Livox Mid-100 (released 2019) which is specifically designed for applications like autonomous driving and robot navigation with a maximum range of 260m. The Mid-100 consists of three Mid-40 elements, but for the experiments, only the central sensor is used. The detailed technical specifications can be found in [25]. Additionally, an industrial ethernet camera is mounted on the test rig to color the LiDAR point cloud with RGB data. Both sensors are extrinsically calibrated using the approach presented in [9]. The setup is completed by a high-precision linear stage, which is used for accurate movement of the sensor setup in the direction of the test object. The LiDAR-to-object distance was empirically set around 2.5m, after an initial investigation of the minimum working distance of the sensor. The specified distance according to the manual is 1m, however, our experiments show significant artifacts up to a distance of 2.2m (see Figure 5). To avoid background thermal or optical noise caused by artificial light or incident sunlight, the measurements were carried out in a dark room, additionally covered with thermal curtains. Moreover, the test objects were placed on a table wrapped with a matte tablecloth, to reduce disturbances caused by reflections.

6.3. Methodology

The Livox Mid-40 has already been investigated in recent research with a focus on temporal stability in [18] and accuracy in general in [13]. In comparison, our work is not a designated analysis of the limits of a specific LiDAR model but aims to identify hazards to LiDAR systems in general. The term *system* hereby refers to a grey box model of all the hardware as well as the software components, as it is not possible, with reasonable effort, to exactly determine the influence of every sub-component for itself.
Each HAZOP scenario was recorded as a point cloud, utilizing the Livox ROS driver from [26]. The integration time for each measurement was set to 8 seconds. Thereby, it was ensured that despite the Lissajous-like scanning pattern of the Mid-40 (described in [25]), every region of the test object was represented by a sufficient amount of points for the following processing steps. Moreover, the resulting point cloud was uniformly resampled using a voxel grid representation with a voxel size of 2mm, with the aim of eliminating the angular dependency of the point density caused by the scanning pattern.
The concept of the analysis is based on three regions of interest (ROI). Figure 7 shows the selected ROIs for an exemplary scenario. The red area represents the actual object surface, represented by the plane P, where the points

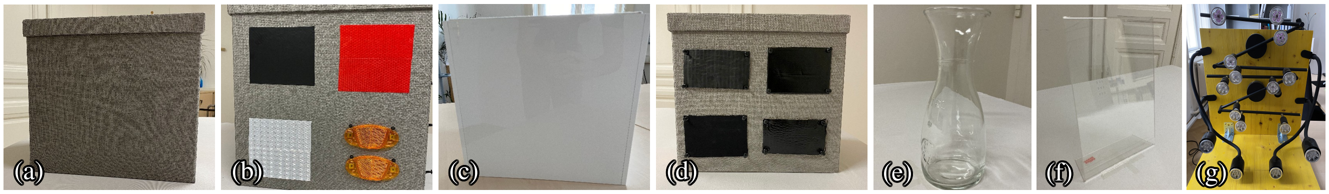

Figure 3. Selection of test objects prepared for LiDAR-HAZOP.

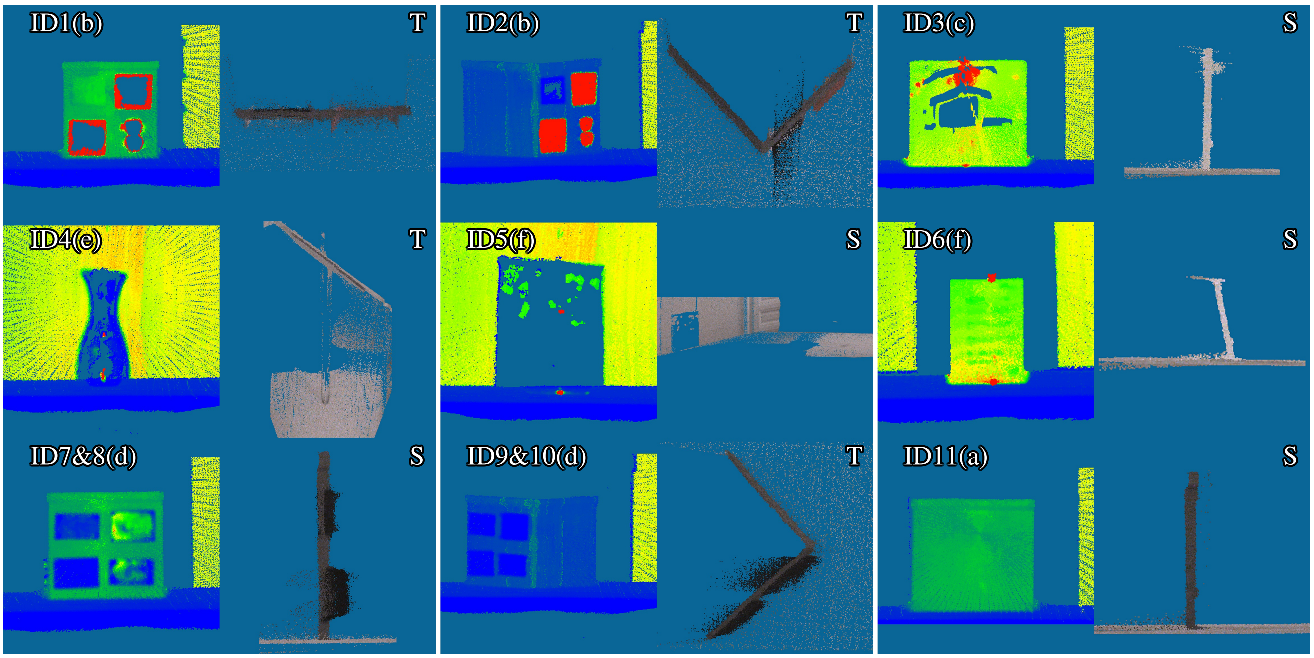

Figure 4. Visualization of results with visible degradation of point cloud data. Two images per visualization: front view LiDAR reflectance, camera data mapped to point cloud (T)op or (S)ide view. Upper corner shows experiment ID & test object (see Table 3 and Figure 3).

would lie in an ideal scanning scenario. The green cuboid V_t is formed to include all points that can be related to a measurement in the object's direction. These include e.g.: points that lie behind the measured surface in the case of transparent objects. Finally, the blue volume V_i contains the surface inlier points and is defined as

$$V_i = \{p | p \in V_t \wedge d_p \leq 4 * std_{ref}\}, \tag{1}$$

where d_p is the distance of point p to plane P distances and std_{ref} represents their standard deviation.

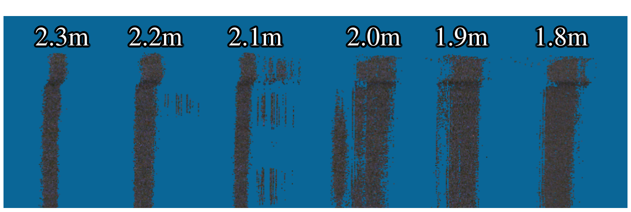

Figure 5. Analysis of Mid-40 reasonable minimum working distance. Artifacts on the flat surface occur for distances $\leq$ 2.2m.

6.4. Selection of performance metrics

To enable both qualitative and quantitative evaluation of the investigated hazards, four performance metrics are selected. Absolute metrics are not expressive for comparing the results among the test objects due to differences in size or shape. To this end, a surface with supposedly no haz-

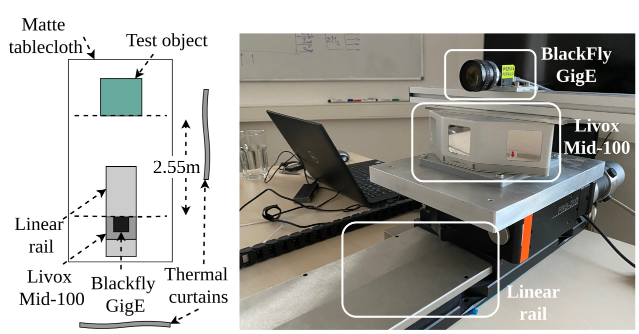

Figure 6. Laboratory setup for LiDAR-HAZOP experiments.

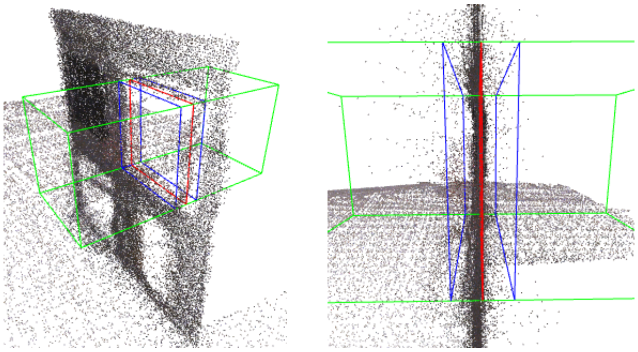

Figure 7. ROIs for the test scenarios: The object surface plane P (red), the inlier bounding box V_i (blue) and a volume for points that lie in the direction of the test object V_t (green).

ardous effects (see Figure 3 (a)) is used as a reference for normalizing the results. The following indicators are selected for the analysis:

- Normalized Point Density: To assess the influence of a certain hazard on the point density, the number of points n within a specified volume V_t (green in Figure 7) in the test scenario is related to a volume of the same size in the reference scenario n_{ref}. It is calculated as

$$\hat{\rho} = \frac{n/V_t}{n_{ref}/V_t} = \frac{n}{n_{ref}}. \quad (2)$$

- Outlier Ratio: is constituted by the number of points that lie in V_t but not in V_i, divided by the total amount of points in V_t.

- Offset: the average distance of points inside V_t with respect to the surface plane P.

- Standard Deviation of the point-to-plane distance of object points around the object surface P

6.5. Results

Hazards selected in Section 6.1 are now evaluated using the proposed metrics. Figure 4 visualizes the results summarized in Table 3. Multiple scenarios with reflective material are evaluated for HID 478/489. Retro-reflectors show problematic cavities at an incident angle of 90° (ID1) for all three segments and a low point density value of 0.1. The 45° setup (ID2) results in a much more dense point cloud but pose a significant offset of 1.5cm and an outlier ratio of 75%. The reflective PVC box (ID3) also shows distinct holes in the measured surface, resulting in a normalized point density of 0.7.
Scenarios ID4/ID5 represent a composition of both reflective and transparent surfaces (discussed in HID 729). Both glass and acrylic glass have optical characteristics that lead to light passing through as well as being reflected at a certain angle. The results for the glass carafe (ID4) are comprised of a combination of wrong distance measurements especially at the round edges as well as measurements of the wall that lies behind the actual test object, as the offset of 1.7m indicates. The reflected intensity is also noticeably lower compared to the neighboring regions, which can be attributed to the refractive characteristics of the glass. The same applies to the results of the empty acrylic glass stand (ID5). Compared to the accurately detected acrylic glass stand containing a sheet of paper (ID6), most of the surface is not detected and only some parts of the wall behind are registered. This suggests that the combination of multiple layers of acrylic glass and the corresponding sequence of refraction and reflection of light causes the problematic measurements.
The underexposure scenarios comprising low albedo surfaces, described in HID 476, are investigated in scenarios with ID7-10. By mere visual analysis, the surfaces appear very similar, however, the differences in reflected intensity is quite distinct. In addition, the measured distances vary between the surfaces. In the case of the black paint (ID7), which is the bottom left surface, the offset is around 2.2cm, which exceeds the specified range precision of the sensor model. In comparison, the results for the black spray (ID8) at an orthogonal incident angle are very similar to the reference surface. Nevertheless, at an incident angle of 45 degrees, the outlier ratio of 53% and standard deviation of 2.2cm imply rather problematic measurements, which indicates that the surface (ID10) has a specular component. This is not the case for the black paint (ID9), which underlines the diffuse properties of the material.
The last investigated hazard is overexposure caused by background/additional light sources (HID 125). In order to verify our test setup, the power spectral density of the LiDAR's emitter and the IR light source array (f) are measured using a spectrometer (see Figure 8). This shows that the array can deliver more than four times the power of the LiDAR emitter at 905nm. Despite that, in the conducted experiments (ID11), the point cloud was not noticeably influenced by lighting the test object (a) with the array. As a follow-up experiment (ID12), the array was positioned directly opposite the LiDAR. Here, gaps appear in the point cloud data and reflectivity is lowered as soon as the array is activated (see Figure 9). For evaluations, two reference planes are chosen: one plane for the wooden back-plate (12A) and one plane encompassing the infrared LED elements (12B). For both, (12A) and (12B) the regions in close proximity to the IR lamps of the array are barely reconstructed, whereas the background gets influenced to a smaller extent.

Table 3. Results for laboratory LiDAR-HAZOP experiments. Bold data indicate insufficient or hazardous performance. The corresponding thresholds for acceptable performance, shown in the table header, are either empirically set or derived from the sensor specifications.

ID	**HID**	**Description**	**Incident angle [deg]**	**Norm. point density (>0.9)**	**Outlier ratio (<0.1)**	**Offset [m] (<0.01)**	**Std. dev. [m] (<0.01)**
1	489	Retro-reflectors	90.0	**0.10**	**0.12**	-0.003	0.007
2	489	Retro-reflectors	45.0	**0.89**	**0.75**	**-0.015**	0.005
3	478	Reflective box	90.0	**0.70**	0.01	-0.002	0.004
4	502	Round glass pitcher	90.0	**0.23**	**0.99**	**1.699**	**0.639**
5	502	Empty Plexiglass stand	74.5	**0.05**	**1.00**	**1.849**	**0.455**
6	729	Plexiglass stand	74.5	1.00	0.00	0.001	**0.012**
7	476	Black paint	90.0	1.00	**0.91**	**-0.022**	0.007
8	476	Black spray	90.0	1.00	0.00	0.002	0.003
9	476	Black paint	45.0	1.00	**0.43**	-0.010	0.007
10	476	Black spray	45.0	1.00	**0.53**	0.005	**0.022**
11	125	Overexposure	90.0	1.00	0.00	0.000	0.003
12A	125	Overexposure	direct	0.91	**0.11**	0.001	0.007
12B	125	Overexposure	direct	**0.57**	**0.47**	-0.005	**0.022**

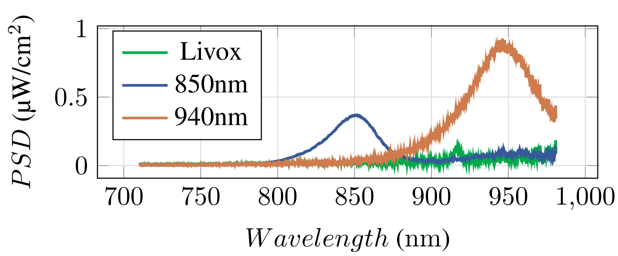

Figure 8. Spectrometer data for overexposure experiments showing *Power Spectral Density* (PSD) per wavelength.

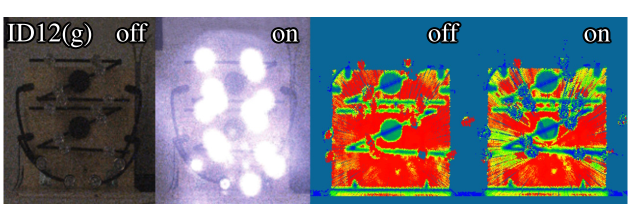

Figure 9. Visualization of results for overexposure experiment (ID12A&B). Camera image with infrared light sources (test object (f)) turned off/on and corresponding point cloud data. Red indicates high- whereas blue indicates low reflectivity. The leftmost visualization shows clear degradation of 3D data emanating outwards from the individual infrared light sources.

7. Conclusion

This work represents the first joint LiDAR and camera risk analysis identifying situations and aspects capable of degrading data from both sensors at the same time. The existing CV-HAZOP is used as a blueprint for the novel LiDAR HAZOP model and inspiration for the systematic risk analysis. This process results in 136 novel hazard entries. The A2D2 dataset is used to evaluate the identified risks in an autonomous driving scenario. First, the camera images are classified by a hazard detector to identify potential hazard frames. Then, experts categorize associated LiDAR data quality based on the presence of hazards. This process identified multiple scenes for each of the three identified hazards in A2D2 data. Additionally, laboratory experiments are designed to replicate hazards out of LiDAR HAZOP entries. The measured point cloud data clearly shows degradation of data quality in presence of the described hazards. Both evaluations show that the theoretic hazards identified during the risk analysis indeed reduce data quality in real-world situations. Autonomous systems that rely on camera data cannot expect to solve robustness issues completely just by adding LiDAR sensors to the mix. The publicly available LiDAR HAZOP list allows for quick checking of common sources affecting the two sensor modalities, both during system design, data-based training, and evaluations.

References

[1] Markus-Christian Amann, Thierry Bosch, Marc Lescure, Risto Myllylä, and Marc Rioux. Laser ranging: a critical review of usual techniques for distance measurement. *Society of Photo-Optical Instrumentation Engineers*, 40(1):10–19, 2001. 2

[2] The blackest acrylic paint in the world - black 3.0 - culture hustle. https://culturehustle.com/products/black-3-0-the-worlds-blackest-black-acrylic-paint-150ml. Accessed: 2023-03-02. 5

[3] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2

[4] Yulong Cao, S Hrushikesh Bhupathiraju, Pirouz Naghavi, Takeshi Sugawara, Z Morley Mao, and Sara Rampazzi. You can't see me: Physical removal attacks on lidar-based autonomous vehicles driving frameworks. *arXiv preprint arXiv:2210.09482*, 2022. 2

[5] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z Morley Mao. Adversarial sensor attack on lidar-based perception in autonomous driving. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 2267–2281, 2019. 2

[6] Ming-Fang Chang, John W Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, DeWang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8748–8757, 2019. 2

[7] T. Chen, R. Wang, B. Dai, D. Liu, and J. Song. Likelihood-field-model-based dynamic vehicle detection and tracking for self-driving. *IEEE Transactions on Intelligent Transportation Systems*, 17(11):3142–3158, 2016. 1

[8] Yukyung Choi, Namil Kim, Soonmin Hwang, Kibaek Park, Jae Shin Yoon, Kyounghwan An, and In So Kweon. KAIST multi-spectral day/night data set for autonomous and assisted driving. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):934–948, 2017. 2

[9] Jiahe Cui, Jianwei Niu, Zhenchao Ouyang, Yun Qin He, and Dian Liu. Acsc: Automatic calibration for non-repetitive scanning solid-state lidar and camera systems. *ArXiv*, abs/2011.08516, 2020. 5

[10] CV-HAZOP - VITRO. https://vitro-testing.com/cv-hazop/. Accessed: 2023-03-02. 2

[11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 2

[12] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, et al. A2D2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320*, 2020. 2

[13] Craig Glennie and Preston Hartzell. Accuracy assessment and calibration of low-cost autonomous lidar sensors. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIII-B1-2020:371–376, 08 2020. 5

[14] Christopher Goodin, Daniel Carruth, Matthew Doude, and Christopher Hudson. Predicting the Influence of Rain on LIDAR in ADAS. *Machine Learning and Embedded Computing in Advanced Driver Assistance Systems (ADAS)*, page 204, 2019. 2

[15] Connor Henley, Siddharth Somasundaram, Joseph Hollmann, and Ramesh Raskar. Detection and mapping of specular surfaces using multibounce lidar returns. *Optics Express*, 31(4):6370–6388, 2023. 2

[16] Xinyu Huang, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. The apolloscape open dataset for autonomous driving and its application. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2702–2719, 2019. 2

[17] ISO/TC 22/SC 32 Electrical and electronic components and general system aspects. ISO 21448:2022 Road vehicles — Safety of the intended functionality. Published https://www.iso.org/standard/77490.html and under review, 2022. Accessed: 2023-03-03. 1

[18] Carter Kelly, Benjamin Wilkinson, Amr Abd-Elrahman, Orlando Cordero, and H Andrew Lassiter. Accuracy Assessment of Low-Cost Lidar Scanners: An Analysis of the Velodyne HDL–32E and Livox Mid–40's Temporal Stability. *Remote Sensing*, 14(17):4220, 2022. 5

[19] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Lyft Level 5 AV Dataset 2019. https://level5.lyft.com/dataset/, 2019. Accessed: 2020-08-17. 2

[20] Trevor A Kletz. *Hazop & Hazan: Hazard Workshop Modules: Notes on the Identification and Assessment of Hazard*. Institution of Chemical Engineers, 1983. 1, 2

[21] S. Kraemer, C. Stiller, and M. E. Bouzouraa. Lidar based object tracking and shape estimation using polylines and free-space information. In *Proceedings 2018 IEEE RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, pages 4515–4522, 2018. 1

[22] You Li and Javier Ibanez-Guzman. Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Processing Magazine*, 37(4):50–61, 2020. 2

[23] Derek D Lichti. Error modelling, calibration and analysis of an am–cw terrestrial laser scanner system. *ISPRS journal of photogrammetry and remote sensing*, 61(5):307–324, 2007. 2

[24] Derek D Lichti and Maria Gabriele Licht. Experiences with terrestrial laser scanner modelling and accuracy assessment. *Proceedings of the ISPRS Commission V Symposium Image Engineering and Vision Metrology*, 2006. 2

[25] Livox Technology Company Limited. Livox Mid Series User Manual v1.2. https://www.livoxtech.com/mid-40-and-mid-100, 2020. Accessed: 2023-03-03. 5

[26] Livox Technology Company Limited. Livox ros driver. https://github.com/Livox-SDK/livox_ros_driver, 2021. Accessed: 2023-03-03. 5

[27] Z. Luo, S. Habibi, and M. Mohrenschildt. Lidar based real time multiple vehicle detection and tracking. *World Academy of Science, Engineering and Technology, International Journal of Computer, Electrical, Automation, Control and Information Engineering*, 10(6):1125–1132, 2016. 1

[28] Yuexin Ma, Xinge Zhu, Sibo Zhang, Ruigang Yang, Wenping Wang, and Dinesh Manocha. Trafficpredict: Trajectory prediction for heterogeneous traffic-agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6120–6127, 2019. 2

[29] Michael Meyer and Georg Kuschk. Automotive radar dataset for deep learning based 3d object detection. In *Proceedings of the 16th European Radar Conference 2019*, 2019. 2

[30] Abhishek Patil, Srikanth Malla, Haiming Gang, and Yi-Ting Chen. The h3d dataset for full-surround 3d multi-object detection and tracking in crowded urban scenes, 2019. 2

[31] Jonathan Petit and Steven Shladover. Potential cyberattacks on automated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 16(2):546–556, 2015. 2

[32] Quang-Hieu Pham, Pierre Sevestre, Ramanpreet Singh Pahwa, Huijing Zhan, Chun Ho Pang, Yuda Chen, Armin Mustafa, Vijay Chandrasekhar, and Jie Lin. A*3D dataset: Towards autonomous driving in challenging environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2267–2273. IEEE, 2020. 2

[33] Matthew Pitropov, Danson Garcia, Jason Rebello, Michael Smart, Carlos Wang, Krzysztof Czarnecki, and Steven Waslander. Canadian adverse driving conditions dataset, 2020. 2

[34] Abdul Rachman. 3D-LIDAR Multi Object Tracking for Autonomous Driving. Master's thesis, Delft University of Technology, The Netherlands, 2017. 2

[35] M. Shirazi and B. Morris. Looking at intersections: a survey of intersection monitoring, behavior and safety analysis of recent studies. *IEEE Transactions on Intelligent Transportation Systems*, 18(1):4–24, 2017. 1

[36] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset, 2019. 2

[37] Y. Sun, H. Xu, J . Wu, J. Zheng, and K. Dietrich. 3-D data processing to extract vehicle trajectories from roadside lidar data. *Transportation research record*, 2672(45):14–22, 2018. 1

[38] M. Velas, M. Spanel, M. Hradis, and A. Herout. Cnn for very fast ground segmentation in velodyne lidar data. In *Proceedings 2018 IEEE Int. Conf. Autonomous Robot Systems and Competitions (ICARSC)*, pages 97–103, 2018. 1

[39] J. Wu. An automatic procedure for vehicle tracking with a roadside lidar sensor. *Institute of Transportation Engineers, ITE Journal*, 88(11):32–37, 2018. 1

[40] Pengchuan Xiao, Zhenlei Shao, Steven Hao, Zishuo Zhang, Xiaolin Chai, Judy Jiao, Zesong Li, Jian Wu, Kai Sun, Kun Jiang, et al. Pandaset: Advanced sensor suite dataset for autonomous driving. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 3095–3101. IEEE, 2021. 2

[41] W. Xiao, B. Vallet, Y. Xiao, J. Mills, and N. Paparoditis. Occupancy modelling for moving object detection from lidar point clouds: A comparative study. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, IV-2(W4):171–178, 2017. 1

[42] Shao-Wen Yang and Chieh-Chih Wang. On solving mirror reflection in lidar sensing. *IEEE/ASME Transactions on Mechatronics*, 16(2):255–265, 2011. 2

[43] Oliver Zendel, Kathrin Honauer, Markus Murschitz, Martin Humenberger, and Gustavo Fernández Domínguez. Analyzing computer vision data - the good, the bad and the ugly. In *Computer Vision and Pattern Recognition, CVPR 2017*. CVPR, 2017. 3

[44] Oliver Zendel, Markus Murschitz, Martin Humenberger, and Wolfgang Herzner. CV-HAZOP: Introducing test data validation for computer vision. In *ICCV*, 2015. 1, 2, 3

[45] Oliver Zendel, Markus Murschitz, Martin Humenberger, and Wolfgang Herzner. How good is my test data? introducing safety analysis for computer vision. *International Journal of Computer Vision*, 125:95–109, 2017. 2

[46] Oliver Zendel, Matthias Schörghuber, Bernhard Rainer, Markus Murschitz, and Csaba Beleznai. Unifying panoptic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21351–21360, June 2022. 3