

# Sensor Visibility Estimation: Metrics and Methods for Systematic Performance Evaluation and Improvement

Joachim Börger<sup>1</sup>, Marc Patrick Zapf<sup>2</sup>, Marat Kopytjuk<sup>3</sup>, Xinrun Li<sup>2</sup>, and Claudius Gläser<sup>1</sup>

**Abstract**—Sensor visibility is crucial for safety-critical applications in automotive, robotics, smart infrastructure and others: In addition to object detection and occupancy mapping, visibility describes where a sensor can potentially measure or is blind. This knowledge can enhance functional safety and perception algorithms or optimize sensor topologies.

Despite its significance, to the best of our knowledge, neither a common definition of visibility nor performance metrics exist yet. We close this gap and provide a definition of visibility, derived from a use case review. We introduce metrics and a framework to assess the performance of visibility estimators.

Our metrics are verified with labeled real-world and simulation data from infrastructure radars and cameras: The framework easily identifies false visible or false invisible estimations which are safety-critical.

Applying our metrics, we enhance the radar and camera visibility estimators by modeling the 3D elevation of sensor and objects. This refinement outperforms the conventional planar 2D approach in trustfulness and thus safety.

## I. INTRODUCTION

The *visibility* of a sensor describes the area where it can or cannot collect data. Intelligent transport systems (ITS) and automated driving functions like trajectory planning, pedestrian protection and others need a dependable representation of the environment. The underlying perception algorithms rely on the visibility of the observed area for complete scene understanding: In particular, missing detections must reliably be traced back to true object absence (freespace) or a blind spot due to occlusion (potentially unsafe area).

In Fig. 1, a vehicle drives on a highway, with onboard object detection and long-range information from roadside sensors. The trucks obstruct all sensors, such that the accident at the on-ramp location cannot be detected. A *visibility map* estimating visible and occluded areas can be used to mark the merge-in as unsafe area. To increase safety, the vehicle speed can be reduced, and the driver alerted.

Despite the importance – any application using environment data should know where data can be received – there are way fewer works on visibility estimation than on classical challenges like object detection, tracking, segmentation etc.

Visibility was formalized for sensor placement optimization [1], terrestrial laser scanning [2], and autonomous agent navigation. However, these approaches and metrics are application-specific and hard to transfer to the automotive domain.

<sup>1</sup>Robert Bosch GmbH, Corporate Research, 71272 Renningen, Germany  
joachim.boerger@de.bosch.com

<sup>2</sup>Bosch (China) Investment Co. Ltd., Shanghai 200000, P.R. China  
marcpatrick.zapf@cn.bosch.com

<sup>3</sup>Robert Bosch GmbH, Cross-Domain Computing Solutions, Germany

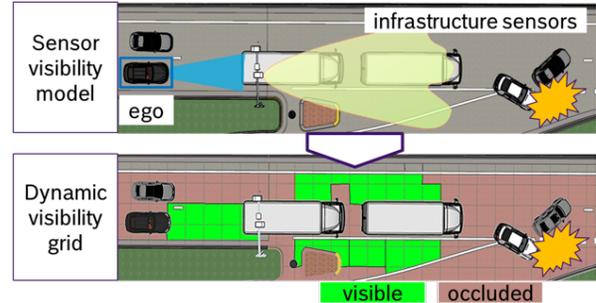


Fig. 1. Motivation for visibility estimation. Top: An ego-vehicle receives perception input (e.g., obstacle detection) from onboard and infrastructure sensors. The FoVs (Field of View) are depicted in blue and yellow. Due to partial occlusion by trucks, the crashed vehicles at the highway on-ramp point are not visible to any sensor. This could wrongly be interpreted as *freespace*. Bottom: The knowledge of visibility and occlusions serves as additional information for scene understanding and to warn of unsafe areas.

The main shortcomings of the state of the art are: Firstly, there is no common definition of visibility and its ground truth. Secondly, no universal metrics exist for a systematic evaluation or comparison of visibility estimators.

In this paper, we contribute:

- 1) a formal definition of sensor visibility, derived from a review of relevant use cases
- 2) a metric and evaluation strategy for the performance assessment of generic visibility estimators
- 3) an example evaluation of our visibility estimators on real-world data from infrastructure sensors

## II. RELATED WORK

### A. Visibility estimation for cameras and sensor networks

Visibility modeling has been studied to optimize sensor networks, predominantly for multi-camera setups. *Coverage* describes the area visible from a sensor system. This is either the ideal field of view (FoV) in absence of occlusions [3], or *partial coverage*, which is impacted by static or dynamic obstacles and often computed via line-of-sight evaluation. We call the latter *visibility* for clear distinguishing.

Coverage is estimated with raytracing to optimize an infrastructure sensor setup for surveillance by Altahir et al. [4] and for vehicle-to-everything (V2X) by Vijay et al. [5]. They define coverage as the percentage of sensors that can see any given point on the surface. These works match our goal to define a sensor-agnostic visibility; however, the authors' objective is sensor placement at design-time, whereas our use cases require dynamic real-time data analysis.

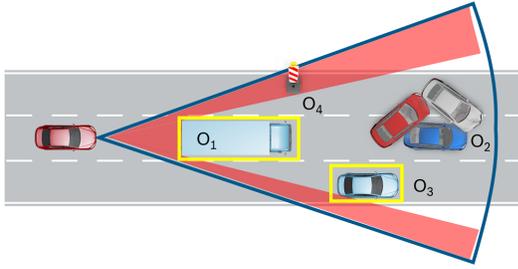


Fig. 2. Example of FoV, visibility and detectability. The *FoV* (blue) describes the maximum area where a sensor can detect objects given ideal conditions and free sight. *Visibility* (red) represents all points with free line-of-sight to the sensor:  $O_2$  is occluded by the truck  $O_1$ . *Detectability* (yellow) considers additional objects properties (Table I):  $O_4$  is visible, but nondetectable due to its material or orientation.  $O_3$  is occluded, but might still be detected by radar via beam reflection on the road-surface.

### B. Visibility estimation for radar sensors

Radar sensor modeling depends on numerous complex factors like object orientation [6] or radar cross-section [7] which are often unknown in real-world applications. Current approaches span from phenomenological [8] over deep learning-based [9], [10] to physical modeling [11].

Inverse sensor modeling is required for occupancy map estimation; most applications like ours use efficient approximations of the aforementioned methods [7], [12], [13].

Radar visibility estimation received little attention so far. Narksri et al. [14] estimate visibility for vehicle mounted sensors using a 3D map as external data source, focusing on hilly roads. Palffy et al. [15] include occlusion information for pedestrian detection. However, the occluded area itself is detected by camera sensors, not radar.

### C. Evaluation of visibility estimation

Porebski and Kogut [16] propose an evaluation method for occupancy grids by assessing the correctness of the shapes of landmarks such as road signs.

Schiegg et al. [17] summarize metrics related to collective perception and communication. The *environmental* or *spatial awareness ratios* describe the ratios of objects or area perceived vs. missed by a communication node.

Many visibility estimation approaches, ours included, are based on a preceding occupancy map. Weston et al. [10] evaluate occupancy map correctness using *intersection over union*. Collins et al. [18] propose a framework with three metrics: An image-based assessment of the map, a direct evaluation of obstacle detection based on it, and the usefulness of the map for path planning.

The diversity of these approaches points out the importance of a good metric and serves as basis for our proposed visibility evaluation framework.

## III. USE CASES

The following widely known use cases can benefit from visibility estimation, our metrics and evaluation framework:

- Object tracking: Temporary occlusions can be compensated by predicting objects through occluded areas

TABLE I

FoV, VISIBILITY AND DETECTABILITY: RELATION AND ASPECTS

Aspects considered	Detectability	Visibility	FoV
Constant physical sensor properties: Opening angle, max distances	x	x	x
Location of the object	x	x	x
Occlusion due to other objects	x	x	
Orientation of the object	x		
Characteristics of the object (height, material, motion, ...)	x		
Weather conditions: Rain, fog, temperature, ...	x		

instead of dropping them. Spawning new objects can be sped up at the edge of the visible area, where they typically appear for the first time. Both aspects improve the calculation of existence probability [19].

- Multi-sensor fusion: Imagine an object is detected by one, but not by another sensor. If this can be explained by occlusion, trust to the sensors is adapted according to their visibility to solve the contradiction.
- Sensor Monitoring: When monitoring false positive (FP) and false negative (FN) detections of one sensor compared to another, visibility can prevent wrong alerts.
- Safe trajectory planning: Freespace information can be used redundantly to object detection to verify that a planned trajectory is not blocked [20], [21]. Naturally, occluded areas must be marked *unknown* (Fig. 1).
- Predicted visibility. Trajectory planning avoids driving into a blind area: Maneuvers which increase the predicted visibility can be prioritized, e.g., moving slightly towards the center line in a traffic jam situation.
- Sensor topology evaluation: A visibility estimator can predict the perception performance expected by a sensor setup [22]. The *coverage rate* (Section VII-B) measures system robustness or sensor redundancy.

The visibility information for these use cases can be summarized: *Given location  $(x, y)$  and assumed there is an unknown object  $O$  – can it be detected by the sensor or not?*

## IV. DEFINITIONS

The above use cases need knowledge of *Detectability*: the probability that an object  $O$  is detected by the sensor. This not only requires a complex physical sensor model, but also information of the object's orientation, material etc., which usually is unknown in real-world applications.

The term *visibility* originates from Latin "*visibilis*" – *that may be seen*. We use this meaning to approximate detectability: It considers only the line-of-sight visibility of a spatial location by checking for occlusions. See Fig. 2 and Table I for a comprehensive summary.

### A. Definitions

Let  $m = (x, y)$  represent a cell of a Cartesian grid, indexed by its upper left corner, and  $O = (x, y, \phi, v_x, v_y, \phi')$  an

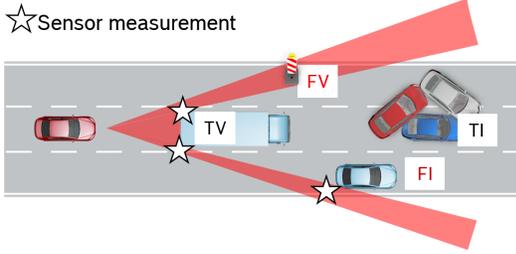


Fig. 3. Performance metrics for visibility estimators: Analogue to binary classifier evaluation, the metrics measure the frequency of the four situations true visible (TV), false visible (FV), true invisible (TI), false invisible (FI).

object with position, orientation, velocity and yaw rate. We introduce the main definitions:

- *Static Field of View* (FoV): Maximum area where a sensor can detect objects in ideal conditions, usually given by maximum detection distance, azimuth and elevation angle, w.r.t some reference object of a minimum size.
- *Visibility of a cell*:  $V_m = 1$  if some point  $p$  within cell  $m$  is visible from the sensor, i.e. the direct line of sight is not blocked.
- *Detection status* of an object:  $D_O = 1 \Leftrightarrow$  at least one sensor measurement  $z$  can be associated with  $O$ . Radar: The position of  $z$  and Doppler velocity  $v_D$  match  $O$  within some tolerance. Lidar: Only the position matches. Camera: At least one pixel at the projected position of  $O$  differs from the background.

### B. Definitions of basic terms

Let  $p = (p_x, p_y, p_z) \in \mathbb{R}^3$  be a point. For the definition of the features mentioned before, we also rely on:

- Occupancy  $Occ_m \in [0..1]$  describes the occupied ratio of a cell  $m$ . This is the inverse of freespace and required for the use case safe trajectory planning.
- *Sensor measurement*  $z$ : A low level measurement, e.g.,  $z = (z_x, z_y, z_z, v_D)$ . Usually, radars deliver 3D positions with Doppler speed  $v_D$ . Lidars or stereo cameras provide 3D positions with intensities, and mono cameras measure 2D pixel colors in the image plane.
- *Visibility of an object*:  $V_O = 1$  if any cell  $m$  overlapping with  $O$  is visible, i.e.  $\exists m : m \cap O \neq \emptyset, V_m = 1$

## V. METHODS

The metrics defined and computed by our evaluation framework adhere to binary classifier evaluation such as *true visible* or *false invisible rate* and thus are generic and scalable. We apply our framework to various visibility estimators for radar and camera presented here.

### A. Metrics for the performance of estimated visibility

We consider algorithms that estimate the visibility  $V_m$  for cells  $m$ . Recap that visibility approximates detectability, which is hard to compute and requires sensor and object knowledge. Still, the true *detection status* of an object can be checked in measured data and will serve as *gold standard*.

TABLE II

METRICS TO ASSESS THE PERFORMANCE OF VISIBILITY ESTIMATORS

	Prediction $V_O$ : Visible	Prediction $V_O$ : Invisible
Ground truth $D_O$ : Object detected by the sensor	True visible rate (TVR)	False invisible rate (FIR): Decreased function availability
Ground truth $D_O$ : Object not detected	False visible rate (FVR): Unknown blind spots	True invisible rate (TIR): Correct detection of occluded areas

This idea implies that the evaluation becomes object-based rather than spatially-based, which fits perfectly to the use cases described before. Two criteria are relevant:

- **Robustness**: For safe trajectory planning and object tracking, it is essential to detect invisible areas.
- **Availability**: At the same time, fearing too many blind spots reduces function availability; estimating too many occlusions therefore shall be avoided.

Table II illustrates the metrics to evaluate these criteria quantitatively. For each ground truth object  $O$  and time step  $t$ , the comparison of the estimated object visibility  $V_O$  to the true detection status  $D_O$  yields one of the four situations TV, TI, FV, FI (Fig. 3). With this terminology, we achieve analogy to the evaluation of binary classifier predictions. The frequencies are calculated as usual, e.g., false visible rate

$$FVR = \frac{FV}{I} = \frac{FV}{FV + TI}$$

### B. Conventional 2D radar visibility grid

A 2D radar visibility estimator is chosen as baseline. It computes the visibility  $V_m^{\text{Radar2D}}$  for all cells  $m$  in these steps:

- **Preprocessing**: The raw data are filtered based on quality, elevation angle, RCS and radial velocity.
- **Occupancy**: A 2D Cartesian occupancy grid with resolution 1m x 1m is processed with the filtered data: The inverse sensor model  $p(Occ_m|z)$  estimates the occupancy of  $m$  given measurement  $z$  by a scaled multivariate Gaussian probability density function (PDF) as in Elfes [23]. A *decay function* accounts for dynamic objects as in Porębski and Kogut [16].
- **Visibility**:  $V_m$  is calculated based on  $Occ_m$  by the line-of-sight approach similar to Adarve et al. [24]: Cells between the sensor origin and last occupied cell are marked visible, cells behind invisible. For directions without any occupied cells, the area up to the maximum FoV is assumed visible.

### C. Improved 3D radar visibility grid

With elevated sensor mounting positions, we now add height information in 3D to avoid false visible and invisible estimations (Section VII-A):

- **3D spherical occupancy grid**: Matching the nature of radar measurements, which provide a radial distance  $r$ , azimuth  $\varphi$  and elevation angle  $\alpha$ , the internal Cartesian

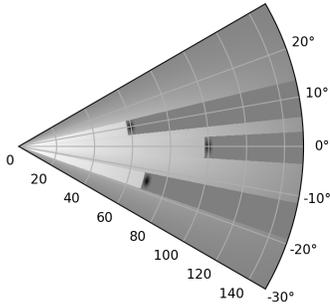


Fig. 4. Spherical 3D radar visibility grid in polar coordinates (distance and azimuth angle), sliced to 2D as defined. Black shows occupied cells, white visible cells, and the intensity of gray depicts the grade of occlusion.

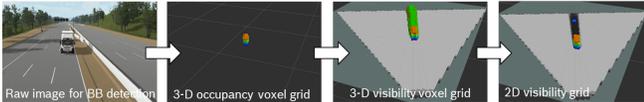


Fig. 5. Creation of a camera visibility grid. 3D vehicle bounding boxes are estimated from raw camera images and voted into a 3D occupancy grid via homography projection. Using raytracing, a 3D visibility grid with visible (not shown) and occluded (green) voxels is created. For the final 2D output, the 3D visibility grid is squashed to 2D using z-averaging.

grid is replaced by a polar grid; for 3D modeling, it is enriched to a sphere by elevation information. Raytracing thus becomes very efficient, in particular for stationary sensors.

- To estimate a cell's occupancy  $Occ_m$  given a measurement in polar coordinates  $z = (r, \varphi, \alpha)$ , the dual inverse sensor model [12] is applied.  $p(Occ_m|z)$  is simplified to a scaled multivariate Gaussian PDF.

Since only the area of potential vehicle presence is relevant, the 3D grid is sliced at height  $h = 1$  meter above ground (Fig. 4). The final output  $V_m^{\text{Radar}3D}$  on a 2D Cartesian grid, as defined in Section IV-A, is achieved by resampling.

The parameters (resolution, decay factor, positive and negative sensor models) were found using Bayesian optimization [25] with training data recorded on a different day.

Note that this visibility estimator is still very basic: The underlying occupancy grid is static, using a decay mechanism for dynamics rather than an advanced technique like particle filters [26]. Still, the 3D enhancement achieves big improvements and is a good showcase for our visibility metrics.

#### D. 3D camera visibility grid

A 3D visibility estimator based on raytracing is applied to infrastructure camera data (Fig. 5). It follows our basic logic (Section V-B) in occupancy grid updating, decaying and visibility grid creation via line-of sight.

- Vehicle bounding box estimation: Detection of 2D boxes via a CNN (YOLOv3) and subsequent 3D box estimation using vehicle size constraints and a homography projection from image space to world coordinates.
- Occupancy: A 3D Cartesian occupancy voxel grid is filled with the 3D data. Voxels covered by bounding boxes are updated as occupied.



Fig. 6. Real-world evaluation data is collected on a highway (depicted here as digital twin for privacy reasons). Radar and camera sensors are mounted on the gantry above the road, facing into the direction of travel.

- Visibility: 3D visibility voxel grids are computed using raytracing similar to radar (Section V-C).

We squash the 3D voxel grid to achieve the 2D output: For each cell  $m$ , the required  $V_m^{\text{Camera}}$  is calculated as the average visibility value of all voxels above  $m$  with height  $h \in [0, 4]$  meters. This method slightly differs from radar due to better height measurements of the camera.

#### E. Reference estimator with perfect occupancy grid

To identify root causes of performance deficits of the visibility estimators, we recap the two main steps:

- 1) Computation of the occupancy grid  $Occ_m$
- 2) Estimation of visibility  $V_m$

For an isolated analysis, step 1) is replaced by a perfect occupancy grid created from ground truth object positions (Section VI-B). This can be interpreted as upper bound for the maximum achievable performance of a occupancy estimator. The evaluation then addresses only the visibility estimation (step 2) and reveals insights to the complexity and limitations of the underlying sensor model.

The reference estimator is a generic concept. It was applied to the radar data in this investigation.

## VI. EXPERIMENTS

We use labeled real world and simulation data to assess and improve the performance of the visibility estimators.

#### A. Data description and sensor setup

The evaluation scene is a public motorway with three lanes in each driving direction (Fig. 6). Sensor data was collected from one radar and one camera mounted on a gantry above the road in 6 m height, facing into the direction of travel.

60 seconds of data with dense traffic and many occlusions from trucks were recorded at 18:45 evening time and labeled. The radar data contained 41 vehicle trajectories and is sufficiently representative to prove our evaluation method. For camera, due to privacy constraints, an equivalent 60-seconds timeseries from digital twin simulation was used.

#### B. Labeling objects

Thanks to our evaluation method, expensive highly precise labels can be skipped: The ground truth detection status  $D_O$  only requires knowledge of object presence close to the grid cells rather than exact object positions. This was achieved

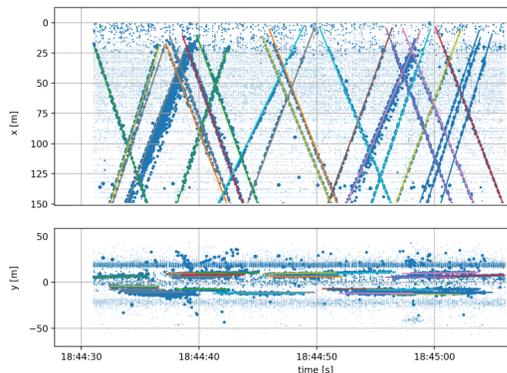


Fig. 7. Labeling objects. Top: Radar reflections (dots) and labeled object trajectories (lines) with IDs are displayed with longitudinal position vs. time. Bottom: Lateral position vs. time. The raw radar data measure the object positions, so we just have to select those reflections corresponding to a vehicle and attribute an ID. To that aim, reflections with high Doppler velocities are plotted bold. Knowing the typical driving behavior of traffic participants in the surveilled area, we can easily select those points corresponding to one vehicle by drawing lines in a simple GUI.

efficiently with the procedure described in Fig. 7: Labeling of 1 minute data took roughly only 1 minute working time.

Note that only moving objects can be labeled this way: The process relies on the visualization of the Doppler velocity; thus, stationary objects are hard to distinguish from clutter.

## VII. RESULTS

For a demonstration of our metric, we present the performance assessed for the radar and camera visibility estimator.

### A. Examples scenes

With our evaluation framework, correct and wrong visibility estimations could easily be identified (Fig. 8 and 9). For both radar and camera, the height over- or underestimation of vehicles often leads to FV or FI.

Some radar FV stem from the absence of a reflection due to the nondeterministic nature of the sensor (Section VII-B). Not depicted, also the underlying static occupancy grid sometimes reacts too slow, causing similar errors.

At higher distance, the camera detector missed some vehicles. Thus, occluded space behind these vehicles could not be identified correctly, causing FV errors.

### B. Statistical evaluation of radar visibility

The resulting performance of the visibility estimators is summarized in Table III.

Radar achieves a total detection ratio for  $D_O$  of 92.6%, leaving 7.4% undetected due to occlusion, averaged over all objects and time steps within the maximum static FoV. This number is interpreted as *coverage rate* and can be regarded for optimization of the sensor position layout.

The improved 3D radar grid outperforms the 2D version by 36% vs. 61% FVR and thus achieves a higher trustfulness. The improvements stem from the height modeling of sensor and objects by considering the limited elevation FoV in near range, as well as the improved inverse sensor model.

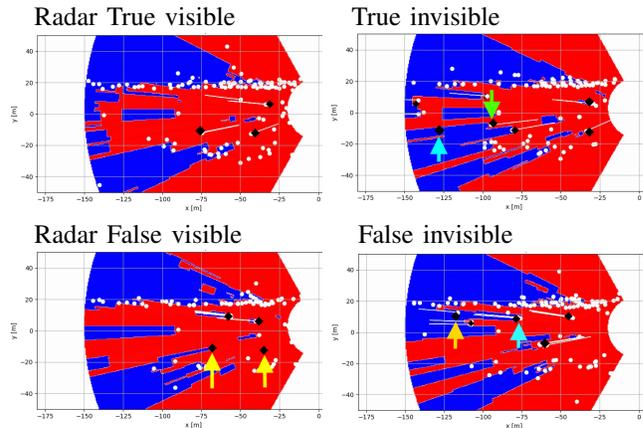


Fig. 8. Estimated radar visibility (red), unknown (white) and invisible (blue). The radar is located at (0, 6) facing along the negative x-axis. Reflections  $z$  are shown as white dots with Doppler velocity as lines. True visible: All objects (black diamonds) return a reflection and are correctly estimated visible. True invisible: The cyan object is occluded by the green one - it returns no reflection and is correctly invisible. False visible: The yellow objects return no reflection although in free line of sight to the sensor, which can have multiple root causes in the nondeterministic nature of radar (Section VII-B). False invisible: The yellow object is detected (white dot), but estimated invisible because the cyan objects height is overestimated.

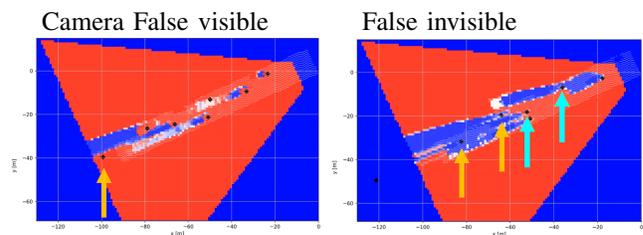


Fig. 9. Estimated camera visibility (red), actual detectability area (faint white dotted area), invisible areas (blue) and ground truth objects (black diamonds). The camera is located at (0, 6) facing along the negative x-axis and slightly negative y-axis. False visible: The yellow object is outside the detection range (no faint white dot) but wrongly estimated visible. False invisible: The yellow objects are detected, yet wrongly estimated invisible because the cyan objects' heights are over-estimated, which results in too large invisible areas.

However, the remaining FVR of 36% still indicates too many undetected occlusions. Even the reference estimator with perfect occupancy information yields a FVR of 31%. Thus, the major error source must be estimation of visibility rather than the preceding occupancy grid: By construction, the underlying sensor models regard only line-of-sight visibility but no other aspects that influence the detectability (Table I).

The error of 31% thus quantifies the gap between the simplified line-of-sight visibility and true detectability (Section IV). It also ignores that the radar principle is nondeterministic by just evaluating binary visibility estimations versus binary detection results, which yields many FV without direct cause. We will comment on that in chapter VIII.

The FIR improves from 23% to 17%: The 3D estimator models that elevated sensors can overlook obstacles after some distance, whereas the 2D estimator predicts an unlimited occlusion area behind the first obstacle.

TABLE III  
PERFORMANCE OF THE VISIBILITY ESTIMATORS

	TVR	FVR	FIR
Conventional 2D visibility grid	77%	61%	23%
3D radar visibility grid	83%	36%	17%
3D radar reference estimator with perfect occupancy grid	90%	31%	10%
3D camera visibility grid	93%	2%	7%

### C. Statistical evaluation of camera visibility

The metrics of the camera visibility estimator were assessed on simulation data and are summarized in Table III. The estimator achieves a good FVR of only 2% due to overall dependable object bounding box detection. The FIR of 7% mostly stems from over-estimation of bounding boxes, as well as voxel sizes in the 3D occupancy grid overestimating the actual size of the vehicles.

Overall, the camera visibility estimator achieves a much better performance than radar. This comes from the deterministic nature of the camera to reliably detect objects. Additionally, the simulation data contain fewer trucks and occlusions, which explains the low FIR rate. This is subject to change in future investigations.

## VIII. SUMMARY

We contributed definitions of visibility applicable to various sensor modalities used in smart transportation systems. Our evaluation framework systematically assesses the performance and trustfulness of visibility estimators with *run-time* data, which is important for safety-related applications.

The metrics were applied to radar and camera visibility estimators; the former could be augmented by 3D object and sensor height information and thus lower the FVR by 25%.

We shed light on the importance of the underlying occupancy grids [27] for the resulting visibility estimators: Their influence with an improvement potential of only 5-7% is smaller than expected, which in turn emphasizes the importance of the sensor modelling part.

To avoid overestimating the FVR, the nondeterministic nature of sensors like radar should be considered: We propose to compare the estimated visibility probability against the empiric detection distribution rather than the binary detection event, e.g., via a Kolmogorow-Smirnow-Test.

## ACKNOWLEDGMENT

The authors would like to thank Thomas Nürnberg and Ronja König for thorough reviews and Peter Baumann, Nils Uhlemann and Adwait Kale for technical contributions.

## REFERENCES

- [1] J. Liu, S. Sridharan, and C. Fookes, "Recent advances in camera planning for large area surveillance: A comprehensive review," *ACM Comput. Surv.*, vol. 49, no. 1, may 2016.
- [2] B. Lecigne, J. U. H. Eitel, and J. L. Rachlow, "viewshed3d: An R package for quantifying 3D visibility using terrestrial lidar data," *Methods in Ecology and Evolution*, vol. 11, no. 6, pp. 733–738, 2020.
- [3] B. Wang, "Coverage problems in sensor networks: A survey," *ACM Comput. Surv.*, vol. 43, no. 4, oct 2011.

- [4] A. A. Altahir, V. S. Asirvadam, N. H. Hamid, P. Sebastian, N. Saad, R. Ibrahim, and S. C. Dass, "Modeling multicamera coverage for placement optimization," *IEEE Sensors Letters*, vol. 1, no. 6, pp. 1–4, 2017.
- [5] R. Vijay, J. Cherian, R. Riah, N. De Boer, and A. Choudhury, "Optimal placement of roadside infrastructure sensors towards safer autonomous vehicle deployments," in *IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 2589–2595.
- [6] Tom Schipper, Lars Reichardt, and Thomas Zwick, "RCS Measurement Results for Automotive Related Objects at 23-27 GHz," *Proceedings of the 5th European Conference on Antennas and Propagation (EUCAP)*, 2011.
- [7] K. Werber, M. Rapp, J. Klappstein, M. Hahn, J. Dickmann, K. Dietmayer, and C. Waldschmidt, "Automotive radar gridmap representations," jun 2015.
- [8] S. Bernsteiner, Z. Magosi, D. Lindvai-Soos, and A. Eichberger, "Radar Sensor Model for the Virtual Development Process," *ATZelektronik worldwide*, vol. 10, no. 2, pp. 46–52, 2015.
- [9] T. A. Wheeler, M. Holder, H. Winner, and M. Kochenderfer, "Deep stochastic radar models," 2017.
- [10] R. Weston, S. Cen, P. Newman, and I. Posner, "Probably unknown: Deep inverse sensor modelling in radar," 2019.
- [11] S. M. Patole, M. Torlak, D. Wang, and M. Ali, "Automotive radars: A review of signal processing techniques," *IEEE Signal Processing Magazine*, vol. 34, no. 2, pp. 22–35, 2017.
- [12] M. Slutsky and D. Dobkin, "Dual Inverse Sensor Model for Radar Occupancy Grids," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1760–1767.
- [13] M. Li, Z. Feng, M. Stolz, M. Kunert, R. Henze, and F. Küçükay, "High resolution radar-based occupancy grid mapping and free space detection," mar 2018, pp. 70–81.
- [14] P. Narksri, H. Darweesh, E. Takeuchi, Y. Ninomiya, and K. Takeda, "Visibility estimation in complex, real-world driving environments using high definition maps," in *IEEE International Intelligent Transportation Systems Conference*, 2021, pp. 2847–2854.
- [15] A. Palffy, J. F. P. Kooij, and D. M. Gavrilu, "Occlusion aware sensor fusion for early crossing pedestrian detection," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1768–1774.
- [16] J. Pořębski and K. Kogut, "Performance Evaluation of the Highway Radar Occupancy Grid," *Sensors (Basel, Switzerland)*, vol. 21, no. 6, 2021.
- [17] F. A. Schiegg, I. Llatser, D. Bischoff, and G. Volk, "Collective perception: A safety perspective," *Sensors*, vol. 21, no. 1, 2021.
- [18] T. Collins, J. Collins, and D. Ryan, "Occupancy grid mapping: An empirical evaluation," in *2007 Mediterranean Conference on Control Automation*, 2007, pp. 1–6.
- [19] M. Aeberhard, "Object-level fusion for surround environment perception in automated driving applications," Ph.D. dissertation, 2017.
- [20] T. Michalke, C. Wüst, D. Feng, C. Gläser, M. Dolgov, and F. Timm, "Where can I drive? Deep Ego-Corridor Estimation for Robust Automated Driving," apr 2020.
- [21] F. A. Schiegg, I. Llatser, and T. P. Michalke, "Object detection probability for highly automated vehicles: An analytical sensor model," in *VEHITS*, 2019.
- [22] M. Gabb, H. Digel, T. Müller, and R.-W. Henn, "Infrastructure-supported perception and track-level fusion using edge computing," *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1739–1745, 2019.
- [23] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [24] J. D. Adarve, M. Perrollaz, A. Makris, and C. Laugier, "Computing occupancy grids from multiple sensors using linear opinion pools," in *IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 4074–4079.
- [25] F. Madrigal, C. Maurice, and F. Lerasle, "Hyper-parameter optimization tools comparison for multiple object tracking applications," *Machine Vision and Applications*, vol. 30, no. 2, pp. 269–289, 2019.
- [26] R. Danescu, F. Oniga, and S. Nedevschi, "Modeling and Tracking the Driving Environment With a Particle-Based Occupancy Grid," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1331–1342, 2011.
- [27] S. Thrun, "Learning occupancy grid maps with forward sensor models," *Autonomous Robots*, vol. 15, no. 2, pp. 111–127, 2003.