International Journal of Applied Earth Observations and Geoinformation 114 (2022) 103017

Contents lists available at ScienceDirect

International Journal of Applied Earth Observation and Geoinformation

journal homepage: www.elsevier.com/locate/jag

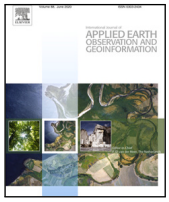

Occlusion detection of traffic signs by voxel-based ray tracing using highly detailed models and MLS point clouds of vegetation

Philipp-Roman Hirt [a,*], Jonathan Holtkamp [a], Ludwig Hoegner [b], Yusheng Xu [a], Uwe Stilla [a]

[a] Photogrammetry and Remote Sensing, TUM School of Engineering and Design, Technical University of Munich (TUM), Arcisstr. 21, 80333 Munich, Germany
[b] Department of Geoinformatics, Munich University of Applied Sciences, Karlstr. 6, 80333 Munich, Germany

ARTICLE INFO

Keywords:
Urban vegetation
3D city model
Mobile laser scanning
Voxel
Point cloud
Visibility analysis

ABSTRACT

Visibility analysis plays a vital role in the design and placing of traffic signs in the urban street environment. This work investigates the occlusion detection of traffic lights and traffic signs caused by vegetation. The presented analysis method is built upon the inputs from the expected situation reflected by a highly detailed 3D city model and the as-is situation captured by 3D Mobile Laser Scanning (MLS). The model contains the location and orientation of streets, traffic lights, and traffic signs; the measurements add detail on irregular-shaped and morphing objects such as vegetation, respectively. The analysis covers the visibility of traffic lights and traffic signs by ray-tracing in an occupancy grid that is generated by the voxelization of the space. The voxels facilitate the distinction between occupied and empty space. The identification of unknown volumes is added and considered in the decision process, to cope with the regions invisible to the sensor. As output, we provide a visibility metric and detailed 3D space descriptions on different levels of granularity, including the knowledge of the semantic classes of traversed voxels. During the whole process, the awareness of unknown volumes is added to an otherwise binary decision between visible and invisible targets. Experiments are conducted on the TUM-MLS-2016 dataset. Results demonstrate that the proposed method is feasible for the detection of occlusions by vegetation in the street scenario, and reveal that the identification of unknown volumes proves necessary for a reliable interpretation of the measurements.

1. Introduction

Urban planning and city management are undergoing rapid development in the frame of enhancing digitization and availability of information. The digital twin city is a key element in the frame of smart cities (Deng et al., 2021). Integrating multiple information sources into one common frame, it enables status assessment, functionality maintenance, and urban planning to be steered based on one system. Whereas classically, most planning information is based on two-dimensional maps and registers of the displayed objects, three-dimensional (3D) city models are gaining importance (Biljecki et al., 2015). Some objects, however, are irregularly shaped and not easily modeled—especially vegetation such as urban trees is challenging in terms of modeling and can profit from in-situ observations (Rutzinger et al., 2011; Weinmann et al., 2017; Wu et al., 2018). A way to capture the real world is by measurement of dense Light Detection and Ranging (LiDAR) point clouds (Beil et al., 2021). Such point clouds can be acquired by laser scanners mounted on a vehicle as part of a Mobile Mapping System (MMS) that captures the street environment. This way of data acquisition is linked to the traffic space. It therefore especially emphasizes applications in the context of urban traffic, consequently covering topics about mobility and safety.

In fact, urban traffic is a collection of a wide variety of participants. The clash of bikers, pedestrians, car and truck drivers, and other mobilized inhabitants in the narrow space of large cities involves dangers, especially for unprotected road users. In addition, smooth traffic flow can only be guaranteed by rational rules and appropriate traffic management (Karami and Taleai, 2020); for instance, by visual indicators like traffic signs and traffic lights. Both will be summarized by the term *traffic installations* in the following text.

The visibility of these installations is a key responsibility of municipalities (Ma et al., 2022; Wu et al., 2015; Soilán et al., 2018). Conflicts may arise especially due to urban trees that are slowly covering existing traffic installations (Huang et al., 2017). The example in Fig. 1 shows a driver approaching two traffic lights. In the beginning (position S_1), the upper traffic light T_{L2} is occluded by a nearby tree. At position S_2, at least the lower part is visible. S_3 suffers from the occlusion of traffic light T_{L1}, this time not by vegetation, but by a traffic sign in front of it. During the whole approach, both traffic lights are visible at position S_4

* Corresponding author.
E-mail address: philipp.hirt@tum.de (P.-R. Hirt).

https://doi.org/10.1016/j.jag.2022.103017
Received 31 March 2022; Received in revised form 5 September 2022; Accepted 6 September 2022
Available online 20 September 2022

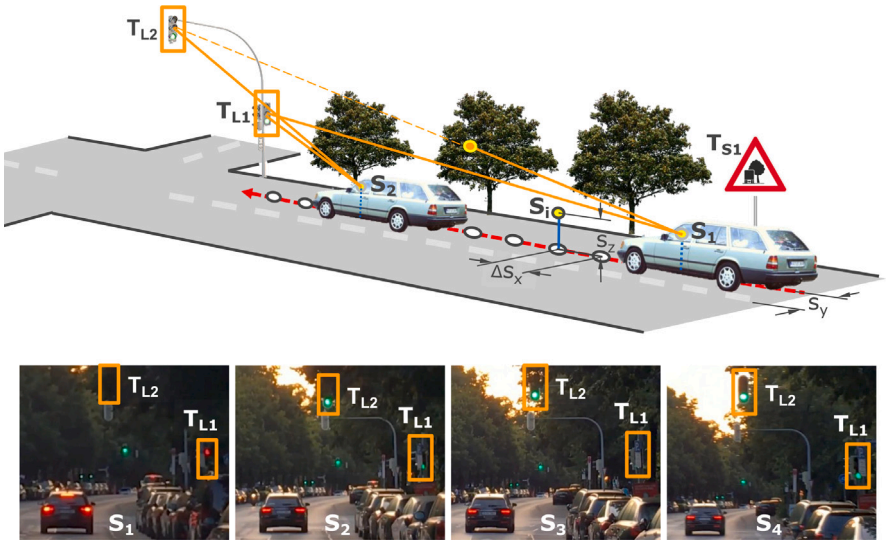

Fig. 1. Occlusion of traffic lights and traffic signs by vegetation. From the driver's sight positions S_i, the traffic lights T_{Li} in front are not always visible. Only S_4 provides full visibility. The offsets s_y, s_z, and the distance Δs_x are defined in a local coordinate system moving along the trajectory. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

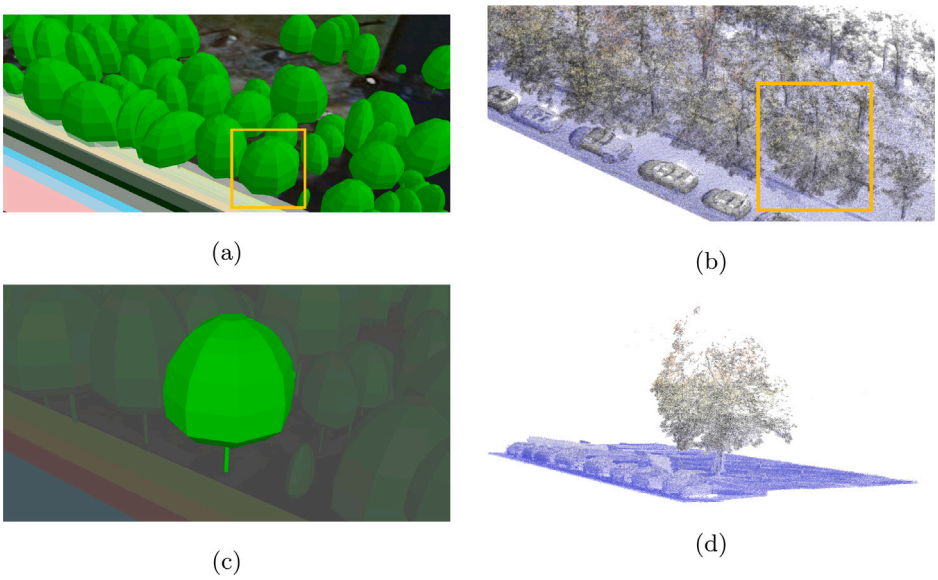

Fig. 2. Comparison of the detail in tree geometry between a parameterized tree in a 3D city model and a dense MLS point cloud. The views of the city model are from the visualization tool presented by Beil et al. (2020). (a, b) Multiple trees next to a road. (c, d) Same tree marked in the model and isolated in the point cloud.

only. In that context, in situ measurements can provide a current and detailed basis for necessary actions to be taken.

To this end, this article investigates the visibility of traffic signals and traffic signs employing 3D Mobile Laser Scanning (MLS) measurements. The analysis is built upon two kinds of inputs: One is the expected situation which is reflected by a 3D city model containing information on the location and orientation of streets and traffic installations. Another is the as-is situation captured by MLS. In particular, our focus is set on vegetation. For visualization purposes, 3D city models often contain coarse tree models. Yet in most cases, these tree models reflect only the shape of the trees, but could not sufficiently be used in the intended geometric analysis. As growth and change are to be expected, the visibility analysis has to be based on up-to-date dense data. Fig. 2 gives a comparison between parametric tree models often found in city models and actual point clouds from MLS. The detailed branch structure can only be depicted by the laser scanning data. This is also of importance in other terms of geometric evaluation, as our previous research on detailed change detection of urban trees has shown (Hirt et al., 2021). Hence combining city models with in situ measurements facilitates visibility analysis of existing structures as well as optimization of planning in cases where the best positions for new traffic signs have to be found.

Additionally, a thorough investigation has to go one step further: To capture obstacles in the line of sight, we use LiDAR as the data acquisition sensor. Once point clouds derived from that sensor are used, it is usually assumed that the point clouds reflect the entire scene, and visibility analysis is performed based on the 3D point data only. This neglects the fact that a measurement system never covers all the volume that is examined. Some parts of the environment are not reached by the sensor due to occlusion or simply because the sensor never points in that specific direction. Hence we have to distinguish between confirmed empty volumes (i.e., volumes without points that have been passed by a laser ray) and unknown volumes (i.e., volumes without points that never were reached by any measurement). In the latter case, it

cannot be decided whether the volume is indeed empty or whether a solid object was omitted due to the lack of measurements. This fact will be taken into account when performing the visibility analysis. The distinct evaluation based on occupied, empty, and unknown volumes can only be realized through a voxel grid, with each voxel containing information on its occupation (Hebel et al., 2013). This comprehensive status assessment of each voxel in the volume of interest is described by an occupancy grid.

In summary, we introduce two major innovations:

- Combining a street environment model with actual MLS measurements of vegetation, hence taking into account that realistic volumetric modeling of trees is often not present in city models
- Including information on volumes not covered by measurements, hence giving the analysis an additional level of confidence depending on lines of sight traveling through space of unknown occupation

For the remaining part of this article, after embedding our work thematically and methodologically (Section 2), Section 3 covers the novel approach for visibility analysis with consideration of unknown volumes. Using various ways of visibility representations, we examine exemplary traffic scenarios (Section 4). Section 5 then features the results of visibility analysis. Finally, we will discuss the quality (Section 6) and consider future extensions to a comprehensive analysis and planning framework (Section 7).

2. Related work

The proposed combination of models and measurements leads to the involvement of several research fields. From the view of applications, the visibility of traffic installations has been discussed for point cloud data or modeled environments (Section 2.1). In this context, the rasterization of 3D space into voxel grids is a primary aspect. Section 2.2 provides a basis for this occupancy grid-based approach. Since we combine real-world measurements with model data, Section 2.3 additionally summarizes ways of traffic environment modeling and introduces the featured format CityGML.

2.1. Visibility of traffic installations

Up to now and to the best of our knowledge, only visibility analyses from point cloud data alone or from a model perspective alone have been examined for this application. The first case includes work using point clouds to extract traffic signs from the measurements before analyzing the situation (Soilán et al., 2018; Huang et al., 2017). This captures the needed detail, yet does not offer the flexibility to change the position of traffic signs in case of planning new positions. The latter case, only using 3D city models, lacks the geometric detail and variety of actual 3D measurements in the context of vegetation.

Based on point clouds only, Soilán et al. (2018) evaluates the safety of pedestrian crossings in urban traffic examining the visibility of pedestrians to drivers. Occlusion in this case is defined as a point lying inside a cylinder of 0.5 m diameter around the line of sight from driver to target. Traffic lights are identified by classification. The line of sight operates on the point cloud data alone, neglecting the volume indeed covered by the sensor. The idea of ray tracing is also used by Voelsen et al. (2021) to identify static and dynamic objects, which will be discussed in the next section. More complex methods contain different approaches to the evaluation of visibility. Using the Hidden Point Removal (HPR) algorithm (Katz et al., 2007), Huang et al. (2017) analyzes occlusions based on visual perception: A field of view for an observer traveling on a trajectory is defined, similar to clipping approaches in computer graphics. By spherical inversion, all visible points are projected onto the surface of a sphere, while occluded points are identified in a separate point set. Comparison of these point sets leads to an occlusion degree and occlusion gradient. Simulating a human observer can be achieved similarly by projecting the points in a certain field of view onto an artificial retina (Zhang et al., 2019). Having generated a 2D visibility field, Zhang et al. (2019) propose a Traffic Sign Visual Recognizability Evaluation Model (TSVREM) consisting of several influencing factors like geometry, occlusion, and sight line deviation. Combined with the identification of lane markings and simulation of viewpoints, a so-called recognizability field can be computed for complete roads, also taking into account different lanes. Occluding points are identified in the MLS point cloud.

Similar to that, the general visibility of traffic signs can be evaluated by combining camera images with 3D point clouds (Wu et al., 2015). After detecting the traffic signs from the point cloud and—using 3D to 2D projection—in the image, a visibility value is derived from a vector of features in the 2D and 3D domains. This approach assesses the visibility of a traffic sign in a single scene captured by a mobile mapping system. The image data serves as validation for volumes not captured by the LiDAR scan.

Visibility evaluation is also possible based on a modeled environment. The advantages are the defined features of all used objects, the straightforward possibility of varying the environment model, and the absence of measurement noise or disturbing objects. Against this background, the combination of 3D and 2D space is proposed by Karami and Taleai (2020). In the context of traffic flow simulation, the visibility of highway traffic signs is evaluated based on models. The authors examine several indices for visibility, including their own index. It evaluates the projection of 3D obstacles along the line of sight onto the 2D plane of the traffic signs. Simulating different traffic participants, velocities, and sign positions, optimal locations for highway signs are identified. In terms of planning, this model-based approach is ideal for traffic environments that are to be constructed. For already existing scenarios, it is hard to model all real-world objects and to realistically analyze the placement of new objects. Hence planning the placement of new traffic signs in existing urban environments can prove difficult. Additionally, in the case of projecting occluding points onto a target surface, important knowledge is lost: where along the projection ray the occlusions occur and whether multiple occluding objects are present.

For the evaluation of general visibility in public spaces, the data structure of voxels is proposed by Aleksandrov et al. (2019). Voxels allow to not only represent existing objects but also to define empty space. Voxelization can be applied to predefined models or irregular real-world data like vegetation points, making it computationally more feasible while retaining realism. Aleksandrov et al. (2019) analyzes the voxel space seen by agents. An agent in that frame is considered an entity with a certain field of view. Raycasting is used to find all voxels inside the field of view not being occluded by objects. Agents are also used by Ma et al. (2022), this time in a dynamic real-time environment. The study deals with the inter-visibility between several agents to avoid crashes at intersections. The environment is given as a 3D point cloud, the traffic participants are simulated. For the line of sight analysis, the point cloud data is transferred to a voxel grid called occupancy matrix in that case. The paper provides extensive analyses with real-time trajectories, different traffic participants and use cases. The various simulations are depicted thoroughly and promise great potential in the optimization of traffic safety. However, it has to be noted that the entire analysis is based on the assumption that the point cloud completely reflects the real environment. Again, areas not covered by the sensor are not taken into account and there is no distinction between empty space and unknown space, which underlines the contribution of the work presented in this article. Finally, a summary of different techniques for occlusion detection, including Hidden Point Removal, occupancy grids, and ray tracing, is given by Alsadik et al. (2014).

There are also other fields for applying visibility analysis: Peng et al. (2022) from the domain of landscape architecture uses voxelized point clouds in combination with GIS information to analyze the design principles regarding visibility in a historic garden. In contrast to our

proposed work, no specific targets are evaluated, the focus lies on the distribution of view sheds over a certain larger area and the visual concept for a surrounding path. On the other hand, Barros-Ribademar et al. (2022) share our area of application—the public street space—yet concentrate on the occlusions in the point cloud data itself, targeting facades and street surfaces.

2.2. Description of space occupancy by voxel grids

One pivotal concept for occupancy analysis is the rasterization of 3D space by voxels, which also is the basis for the method introduced in this work. In that context, Aleksandrov et al. (2021) presents a review of the data structures and algorithms. Xu and Stilla (2021) and Xu et al. (2021) provide more general overviews on handling point clouds of urban infrastructure, discussing different sensors and data structures, including voxels. In the context of the visibility analysis, the approach of modeling three-dimensional space as a grid of voxels has been proven (Aleksandrov et al., 2019; Ma et al., 2022; Voelsen et al., 2021). In general, voxelization is useful when it comes to the analysis of complete volumes, be it in terms of change detection (Hebel et al., 2013), in other contexts like landscape architecture (Peng et al., 2022), or the modeling of empty spaces (Gorte et al., 2019). In the specific frame of mobile mapping, occupancy grids are used for applications like occlusion detection on facades and street surfaces (Barros-Ribademar et al., 2022), the detection of moving objects (Gehrung et al., 2017), change detection of urban environments in general (Gehrung et al., 2018), or with a focus on urban trees (Hirt et al., 2021). By including the position of the sensor, the path of the laser beams can be defined for each point measurement. Voxels passed by the beam are then regarded as empty. Based on that work, the occupancy grids in this article are constructed (Section 3.2), allowing distinction between empty space and space not covered by the sensor.

Voelsen et al. (2021) presents voxel grids in a different application: to classify a scene into static and non-static objects. Semantic segmentation followed by a change detection over several measurement epochs is combined with ray-tracing-based visibility analysis. The aim is to generate a static reference map while discarding moving objects like cars and pedestrians. Similarly in our case, dynamic objects have to be ignored too when checking the general visibility of traffic signs.

2.3. 3D modeling of urban space

Besides measuring the street environment, model information can be used as well, enabling extended ways of analysis. A variety of data standards and formats exist for traffic space modeling (Schwab and Kolbe, 2019). In the frame of autonomous driving, High Definition (HD) maps containing all traffic elements at a high level of detail are gaining importance (Seif and Hu, 2016), for instance, based on the OpenDRIVE standard. High accuracy and exact depiction of lanes and traffic signals can be expected when using HD maps; however, their public availability is still limited. On the other hand, 3D city models aim at constructing all elements of an urban environment, including the traffic space. In that context, the Geographic Markup Language (GML) is expanded to the CityGML standard (Biljecki et al., 2015). With its current development, CityGML is also able to include detailed lane information, different types of traffic areas, and geometrically modeled traffic installations. As introduced by Beil et al. (2020), the latest version of CityGML 3.0 (Kolbe et al., 2021) will enable volumetric traffic space modeling. On that basis, clearance space above roads can be defined and included in a clearance check.

3. Methodology

The proposed method consists of two major stages: static scene simulation and dynamic agent-based visibility analysis. Fig. 3 illustrates the relations between these elements. Concerning the scene simulation, we distinguish between the expected traffic environment from a 3D city model (Section 3.1) and the actual, sensed environment from 3D measurements (Section 3.2). The occupancy grid derived from MLS is constructed with specific consideration of volumes not covered by the sensor, hence remaining unknown in the constructed scene. Into the scene, an agent is placed moving along a user-definable trajectory (Section 3.3). Finally, Section 3.4 elaborates on the visibility analysis based on the line-of-sight from the agent to traffic installations.

3.1. Static traffic environment

For the selected scene, we use a 3D city model encoded in CityGML. CityGML provides all the necessary information, including objects for each lane, sidewalk, etc., and objects for traffic installations and their fixtures. Especially, in the latest version of CityGML (i.e., V3.0), this will also include clearance volumes (Beil et al., 2020). To correctly define the target areas for the line of sight analysis, the model must contain objects of the type *City Furniture* with associated geometry descriptions. This description can be a detailed, realistic 3D model of elements like the traffic light body, or a bounding box representation. The indispensable features needed for analysis are spatial extent and knowledge about which face is the front. The presented approach will derive a spatial rectangle from the bounding box information, which is defined by three points. As the order of those points is given—upper left, lower left, lower right corner—it defines the orientation of the front.

For the possible trajectories, separate lanes need to be defined. This includes sidewalks and bike lanes. As these are vector trajectories, a reference line can describe them. Additional attributes like width or curvature radius can be added, as well as the clearance space above, making them volumetric traffic spaces.

3.2. Mobile LiDAR for voxel-based occupancy grid generation

For describing a complete volume of interest, an occupancy grid is employed to assign each volume element a certain status of occupancy. This is derived from the MLS data and includes the knowledge of empty and unknown space. To achieve this, we revert to the voxel grid concept as presented by Hirt et al. (2021). Each voxel holds the attribute variables *occupied* and *empty*. At initialization, both values are set to 0, and the occupancy of the voxel is *unknown*. To determine these values for each voxel, we need the sensor position for each point in the point cloud. Thus reconstructing the measurement ray from start to end point, each voxel traversed by the beam increases its *empty* attribute by 1. Each voxel containing a measured point increments the *occupied* attribute by 1. Regarding voxels containing a sensor position, an additional attribute *sensor* is introduced. It can later be used to recover the sensor's positions in the voxel space. Fig. 4 sketches an example of the voxel occupancy determination.

This approach leads to an occupancy grid with the status of each voxel not being fixed, but being represented by the tuple:

$$V = (empty, occupied), \text{ with } empty, occupied \in \mathbb{N}, \tag{1}$$

and thus being open for interpretation in the following steps. Here, we identify two ideal cases: a voxel $V_e = (n, 0)$ with $n > 0$ representing empty space, and a voxel $V_o = (0, n)$ with $n > 0$ representing occupied space. All voxels $V_u = (0, 0)$ neither contain measured points nor were traversed by measurement rays. Due to the lack of information, the status of the occupancy cell is hence called *unknown*. The conflict case $V_c = (n, m)$ with $n > 0, m > 0$ is caused by the borders of objects or noise. Mostly, it can be expected that either $n \gg m$ or $m \gg n$, yet more problematic proportions have to be taken into account as well. The handling of these cases during the evaluation will be discussed in Section 3.4.

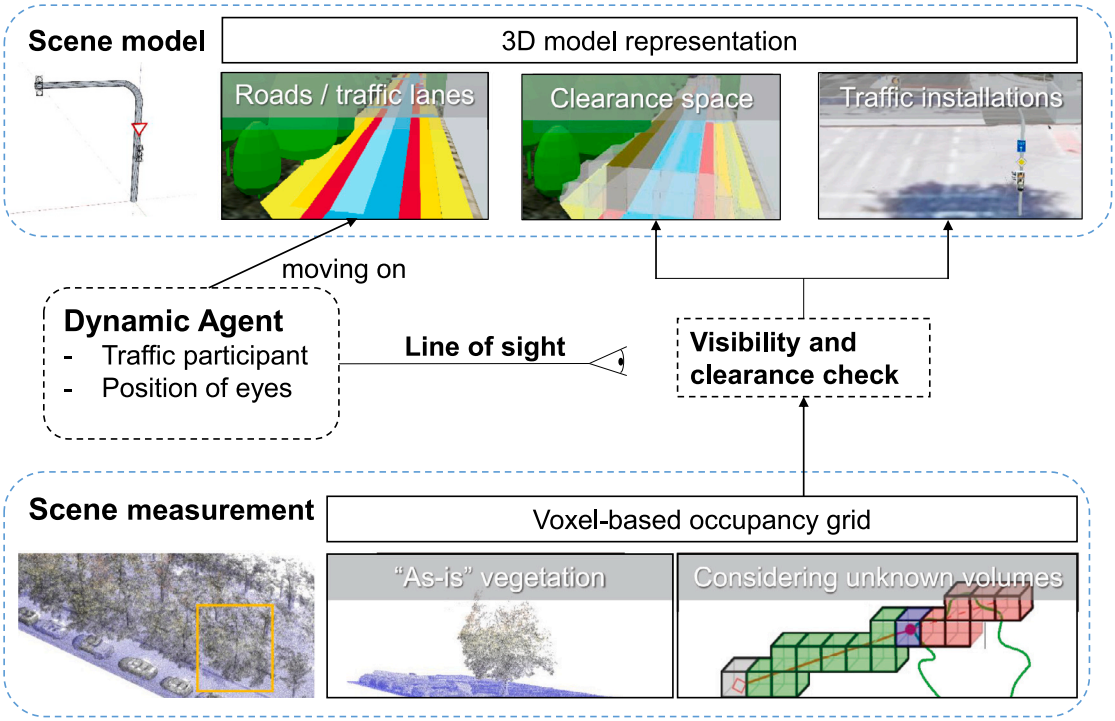

Fig. 3. Relationship between static scene and dynamic visibility analysis. Images of the model environment were taken from the online viewer presented by Beil et al. (2020).

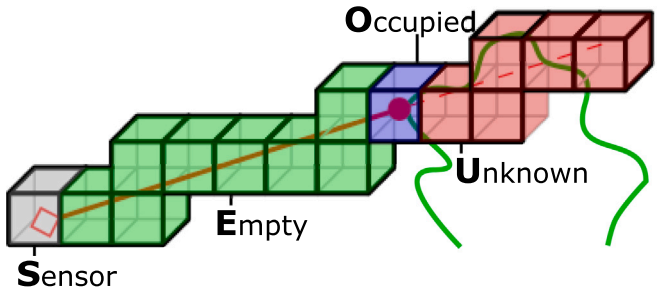

Fig. 4. Voxel status determination by ray-casting from the sensor (red square) to the measured point (red dot) (Hirt et al., 2021). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.3. Agent definition

The static scene as described in Section 3.1 is populated by traffic participants which are modeled by agents. Their ability to observe certain important traffic installations shall be evaluated. Possible agent types are car, truck, or bus drivers, cyclists, or pedestrians. To keep the definition as flexible as possible, an agent is defined by three inputs: (i) The main reference is a complete or partial path element from the city model, for example, a lane on a street, a bike lane, or a sidewalk. (ii) The offset to the left or right in the direction of travel is defined. In the trajectory coordinate system, this is denoted by s_y (cf. Fig. 1). This reflects a driver's position not being exactly in the middle of a lane. (iii) At last, the eye-height s_z is added as a distance above the height level on which the lane is defined. With these three parameters and the direction of travel, a trajectory for the driver's head can be derived. The points on the trajectory are the origins of the lines of sight that will be checked for occlusions.

3.4. Line of sight analysis

We define test cases to perform the analysis. A test case consists of an agent who is traveling along a predefined trajectory, and a set of traffic control elements to be checked for visibility. These elements are drawn from the static model and can be traffic lights, traffic signs, or other objects of interest. They must have one main face that can be described as a spatial plane, with its normal vector being perpendicular to the side that contains the information: the text or symbols on traffic signs, the color on a traffic light, etc. The line of sight to a raster of points on that surface can then be checked for occlusions (Sections 3.4.1 & 3.4.2). Section 4.4 introduces the different levels of output information.

3.4.1. Line of sight and occupancy decision

Before evaluating the Line of Sight (LoS), we need to define the resolution of the computations during analysis. The traffic sign is described by a spatial plane with its corners. Within that area, we define a grid of points on which the analysis will be executed. Furthermore, the agent's trajectory is split into several viewpoints, with the distance Δs_x between each other. By adjusting the grid width and the viewpoints separation, we define the complexity of the analysis steps.

As a result, we obtain the set of viewpoints P_V and target points per target T, respectively. Now the visibility of each target point $P_{T,i}$ from each viewpoint P_V can be calculated by tracing the line of sight. Fig. 5 illustrates the decision process. If the angle between the LoS and the normal vector of the spatial plane describing the target is greater than 80°, the target point is assumed as invisible and the visibility indicator $\xi_{\text{vis},i} \in \{0, 1\}$ for that point is set to 0. If the angle is $\leq$80°, the LoS is analyzed.

By tracing the LoS from P_V to $P_{T,i}$, we obtain the amount of occupied and unknown voxels intersected by this line. Based on their value, the decision on the visibility of $P_{T,i}$ is made. If the amount of occupied voxels is greater than 0, the target point is assumed as invisible and the visibility indicator $\xi_{\text{vis},i}$ is set to 0. If there are no occupied voxels within the line of sight, $\xi_{\text{vis},i}$ is set to 1. Furthermore, we define an *unknown*-indicator $\xi_{\text{u}} \in \{-1, 1\}$ which is initialized as 1 and changed to -1 if any clear LoS from one viewpoint to any of the points $P_{T,i}$ within the target T includes an unknown voxel. Finally, the result is a vector containing two coefficient sets for each viewpoint:

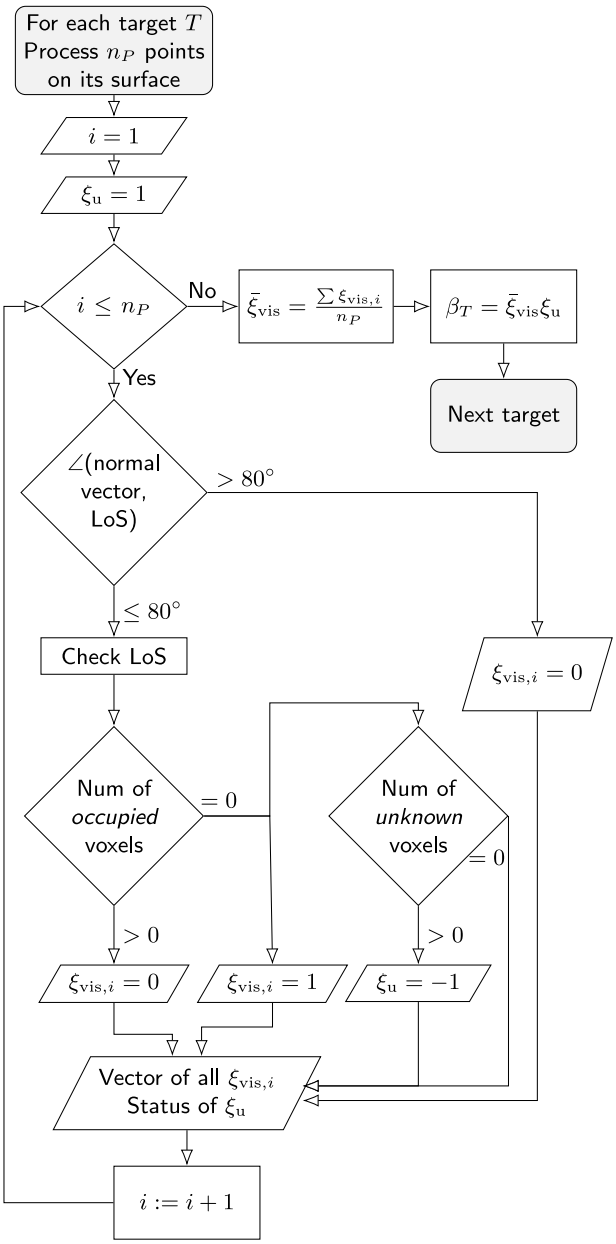

Fig. 5. Delineation of the visibility degree β_T. Inputs are the Line of Sight (LoS) from viewpoint to the target point and the normal vector of the target. This process is executed for the lines of sight from each viewpoint P_V to each target point $P_{T,i}$.

the visibility indicators $\xi_{vis,i}$ for all target points $P_{T,i}$ and one *unknown*-coefficient ξ_u indicating whether any of the clear lines of sight from P_V to $P_{T,i}$ contains unknown voxels. The visibility degree β_T for one target is a percentage of the visible target area, computed via the mean of all visibility indicators $\bar{\xi}_{vis}$:

$$\beta_T = \xi_u \, \bar{\xi}_{vis} \, . \tag{2}$$

Multiplying by ξ_u sets the visibility degree negative in case unknown voxels are involved, indicating unreliable data for that target.

As discussed in Section 3.2, the occupation status is ambiguous in the conflicting case that both n and m are larger than 0 for the voxel status $V = (n, m)$. To avoid false negatives—occupied voxels marked as empty—and hence increasing the sensitivity of visibility analysis, we define a voxel as occupied as soon as n is greater than 0, regardless of the empty value m.

3.4.2. Including semantic knowledge

Relying on measurement data always includes static and dynamic objects. For analyzing the street environment, only static objects like road furniture or urban trees are of interest, while cars and pedestrians should not be considered in terms of visibility evaluation (Voelsen et al., 2021). Semantic segmentation of the measured data can be achieved by baseline labeling methods as presented by Zhu et al. (2020) and Xu et al. (2020). It results in point clouds like the one shown in Fig. 8 where a class label like street, building, vegetation, etc. is assigned to each point. With semantic knowledge available for the area of interest, we can consider it during ray tracing. If a traversed voxel is identified as occupied, its nearest neighbor in the labeled point cloud will then be checked. If the distance between the voxel center and the corresponding labeled point is inside realistic limits (e.g., 1.5 times the voxel size), the voxel is regarded to be of that class. When the class of the occupied voxel is a dynamic one like car or measurement noise, the status of the voxel is set back to empty, as we assume it to be empty in general. Moreover, the class of each occupied voxel is recorded and a histogram visualizes the classes causing occlusion.

4. Experiments

The application of the different evaluation levels is shown by several test scenarios on a publicly available benchmark data set. Besides the pure visibility analysis, the optional inclusion of semantic knowledge is also evaluated in Section 4.3.

4.1. Data set TUM-MLS-2016

The experiments are performed on the benchmark data set TUM-MLS-2016 (Gehrung et al., 2017; Zhu et al., 2020). This data set fulfills the requirements needed to evaluate our method: Scanner data with the viewpoints for each recorded point, and coverage of an area where model data of traffic installations is available. TUM-MLS-2016 contains data from April 2016, recorded by an MLS system. Two Velodyne HLD-64E scanners are mounted at the front corners of a vehicle, slightly tilted outwards. The sensor configuration allows for dense point cloud acquisition of the street environment. The exact geo-reference is provided by an Applanix POS-LV integrated navigation system. The scanners are rotating 360°, with 64 range measuring units acquiring points simultaneously. In addition, the Ladybug 360°-camera provides visual information on the surroundings. The forward-facing images from the Ladybug include parts of the sensors and the vehicle roof. The examined test area *TUM City Campus* is located in the downtown area of Munich, Germany (48.1493° N, 11.5685° E), and mainly contains four streets around the quarter of the Technical University of Munich. For the evaluation of our approach, we chose two junctions at the *Arcisstraße*. On one side of that street, the park-like area around the *Alte Pinakothek* museum contains many tall trees that partly can cause occlusions on the nearby street in leaf-on season.

As discussed in previous work (Hirt et al., 2021), the absolute point accuracy as the combination of measurement and geo-reference errors reaches several centimeters. In addition, especially the volume filled by branches will have different appearances throughout repeated measurements caused by leaves and wind. A sampling in cubic voxels of 10 cm is considered sufficient for capturing significant structures while maintaining efficient processing. A test for edge cases in voxel size is performed in Section 5.3 to validate that assumption.

4.2. Test scenarios

As discussed in Section 3.4, a test scenario is defined by an agent traveling on a trajectory and a set of traffic installations. To evaluate the whole functionality, we define four test scenarios. Fig. 6 gives an overview of the examined junctions and roads.

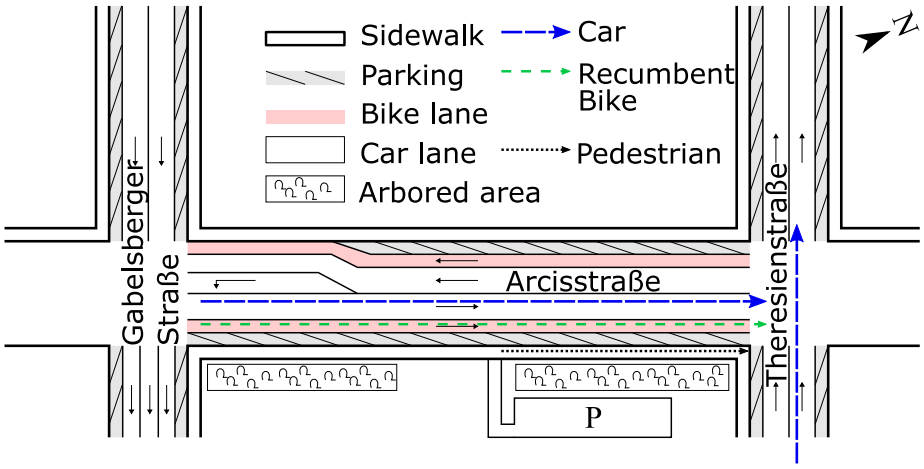

Fig. 6. Streets and junctions in the test area (not true to scale). The small arrows indicate the driving direction. Contains the trajectories of the featured agents.

We structure the scenarios by junction, a set of traffic installation objects, and agents. We chose the rather exotic agent riding a recumbent bike as from the lying position of the rider and the resulting unusually low eye height, interesting viewing configurations can be expected.

(1) Arcisstraße to junction with Theresienstraße
Target set: {Traffic light; no parking sign on opposite side} (Figs. 7(a), (b))

(a) Car driver, on right lane, eye height $s_z = 1.50\,\mathrm{m}$
(b) Recumbent bike rider, on bike lane, $s_z = 0.50\,\mathrm{m}$

(2) T-junction from parking lot into Arcisstraße
Target set: {No entrance sign} (Figs. 7(c), (d))

(a) Car driver, on right lane, $s_z = 1.50\,\mathrm{m}$
(b) Recumbent bike rider, on bike lane, $s_z = 0.50\,\mathrm{m}$

(3) Arcisstraße to junction with Theresienstraße
Target set: {Pedestrian light ahead, pedestrian light to the left} (Figs. 7(e), (f))

(a) Pedestrian, on sidewalk, $s_z = 1.80\,\mathrm{m}$

(4) Theresienstraße to junction with Arcisstraße
Target set: {Priority sign, One-way sign on opposite side} (Figs. 7(g), (h))

(a) Car driver, on right lane, $s_z = 1.50\,\mathrm{m}$

The traffic installation objects are given in the CityGML format based on data from the city administration of Munich, containing the 3D outline of the installations. Up to now, this data is not openly available.

4.3. Inclusion of semantic knowledge

For consideration of static and dynamic occluding objects (Section 3.4.2), class labels are introduced. A manually labeled ground truth is given for parts of the test data set (Zhu et al., 2020). The available classes are: Unlabeled (usually measurement noise), Artificial Terrain, Natural Terrain, High Vegetation, Low Vegetation, Building, Hardscape, Artifact, and Vehicle. We use the available labels at the resolution of a 20 cm voxel grid. Association of an occupied voxel with a class from ground truth is performed for the nearest point inside a radius of 30 cm. As not the complete area of interest is labeled, the case of unknown classification has to be considered as well. If no classification is available, we assume a static object and hence an occlusion.

Fig. 8 gives a bird's eye view of the area of our examinations. It is generated from the labeled point cloud. The visible classes are depicted in different colors as shown by the figure's legend.

4.4. Decision output and evaluation metric

The result of the visibility analysis can be delivered in different levels of granularity. A single test case checks the visibility of one sign for one agent on a predefined trajectory. Table 1 gives an overview of the possible outputs.

The basic level of output is derived by a simple yes/no question, namely whether occlusions have been detected at any of the viewpoints along the trajectory or not. Moreover, another indicator states whether any voxels of unknown occupation have been traversed during the analysis. To derive one single metric for comparing different settings, we compute the mean visibility degree $\bar{\beta}_T$ of the target along the trajectory. In a similar way, the percentage of viewpoints where unknown voxels are in the lines of sight (unreliable viewpoints) can be used for comparison. If any conflict is detected, then more detailed spatial information is given as 3D point clouds.

As a basis for any decision on further measures, the output of Level 1 provides information on the compromised viewpoints and occluding objects. The agent's trajectory is given as discrete viewpoints along a line, with each viewpoint holding the visibility degree of the target. The visibility degree β_T is defined as the percentage of the target surface visible from that viewpoint (Eq. (2)). If that value comes to be 0, the traffic sign or signal is not visible to the agent. Negative values indicate unknown volumes along the line of sight, hence a reliable visibility degree cannot be given. If semantic knowledge is available, occlusions by car or noise points—dynamic occlusions—do not negatively influence the visibility degree.

With those trajectory points indicating problematic positions, a point cloud of the occluding voxels shows where the occlusions can be resolved. The points are labeled according to their semantic classes, so it can easily be seen where for instance trees have to be pruned to ensure good visibility.

The output of Level 2 reflects the information basis available to the analysis algorithm. From each viewpoint, the lines of sight to discrete points on the target's surface are encoded as a point cloud. Each point is a voxel traversed by the ray. It holds an intensity value. For one output case, this intensity value is the index of the viewpoint, if the voxel is empty—so at perfect visibility conditions, these values will increase from one viewpoint to the other. Only once a voxel is identified as occlusion, this value becomes 0. This description allows for easy identification of how many and which occlusions are present from an individual viewpoint. In the second output configuration, this intensity

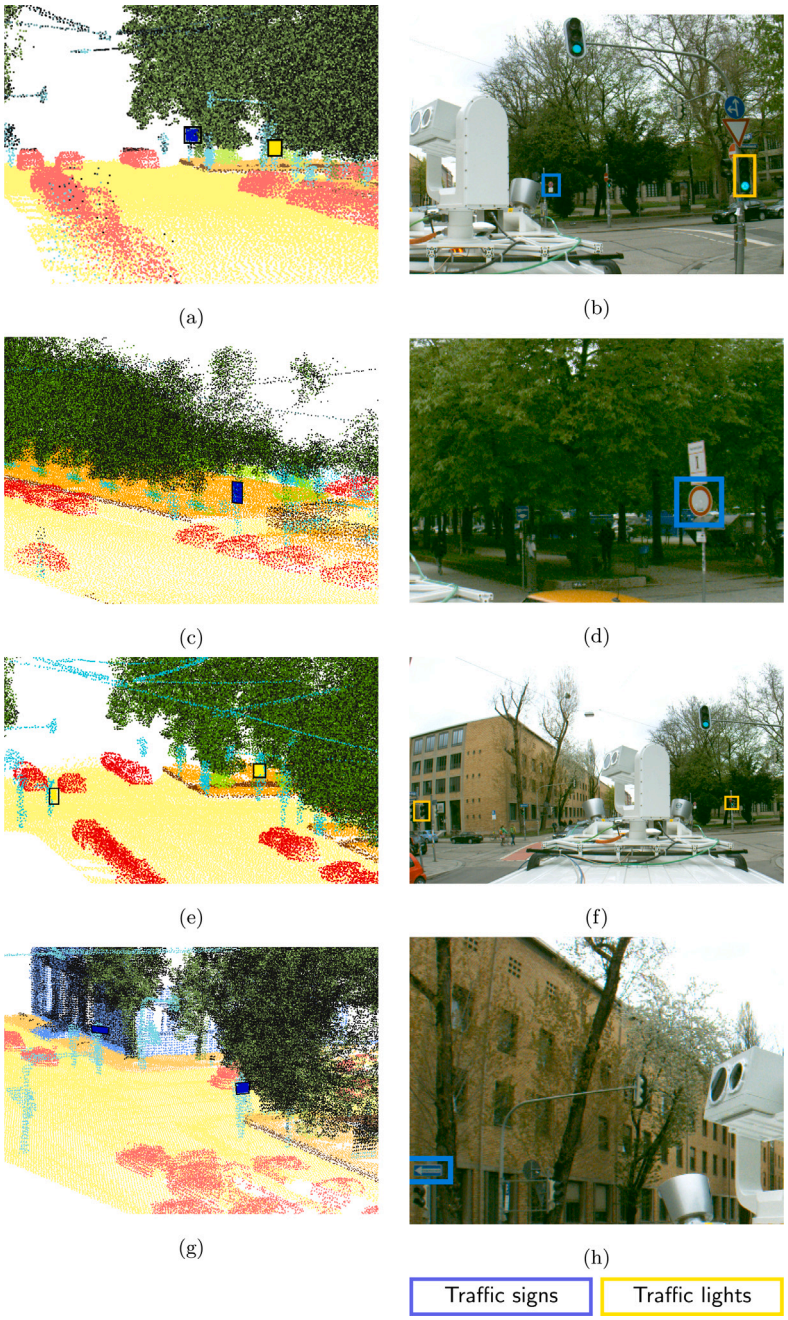

Fig. 7. Locations of the target traffic installations. (a, c, e, g) In the point clouds, a labeled scene is given for orientation (color coding as in Fig. 8). The traffic signs are framed in blue, traffic lights in yellow. (b, d, f, h) For validation, the front-facing images from the panoramic camera are included, too. (h) only contains the one-way traffic sign. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Outputs of visibility analysis at different levels of detail.

Level	Output format	Conveyed information	Encoding
0	Boolean	Occlusion detected	Yes/No
0	Boolean	Unknown voxels included	Yes/No
0	Percent	Mean visibility degree	0 – 100%
0	Percent	Amount of unreliable viewpoints along trajectory	0 – 100%
1	3D trajectory	Visibility degree from viewpoint	Percentage of visible object area (−1-0)/(0-1)
1	3D point cloud	Occluding voxels	3D points with class labels of disturbing objects
2	3D ray bundle	Lines of sight (LoS) from each viewpoint	3D points, intensity 0 if occlusion detected
2	3D ray bundle	LoS with occupancy and uncertainty information	3D points, intensity is occupancy value (0-255), −1 if unknown

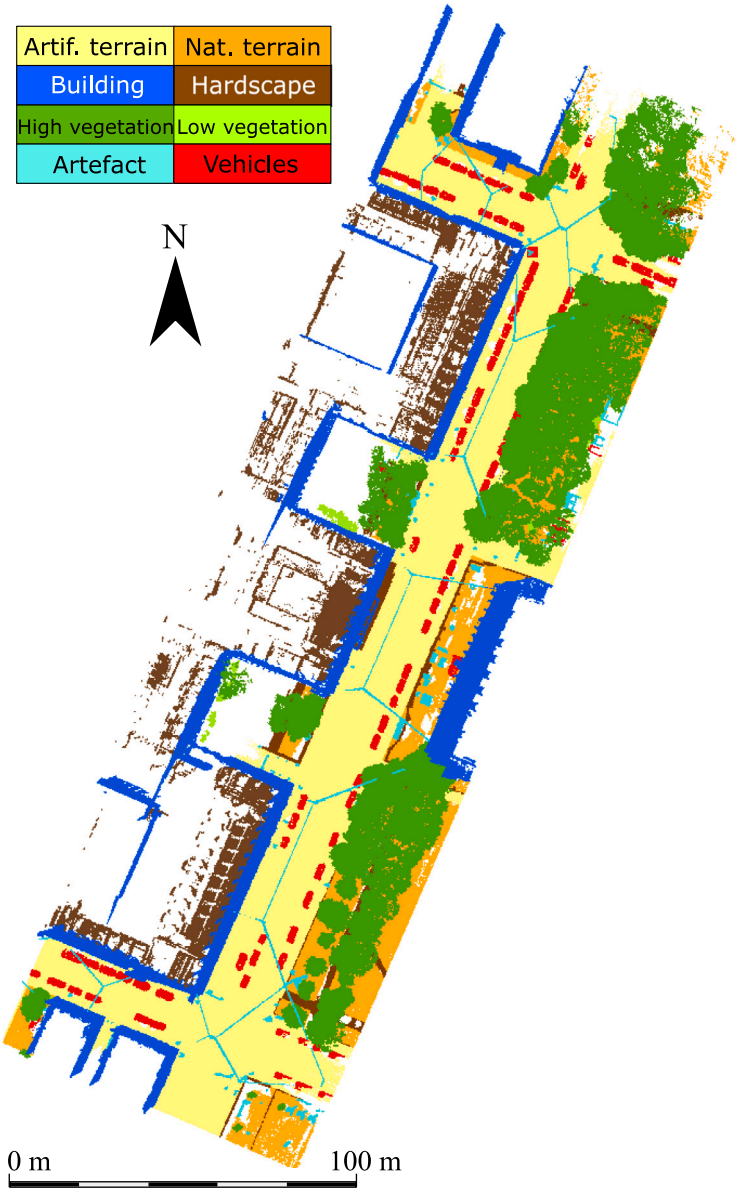

Fig. 8. View on the labeled point cloud from above. Colors encode the semantic classes. The main street is Arcisstraße, with the junctions to Gabelsberger Straße (South) and Theresienstraße (North, cf. Fig. 6). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

value ranges from 0 to 255, reflecting the occupancy value of that voxel. An intensity of −1 identifies voxels of unknown status. This output hence shows the exact decision basis of occlusion and reliability analysis.

5. Results

The previously presented output formats offer different possibilities for analysis. To focus on the most interesting findings, some exemplary results are shown in this section. The outline follows the scenarios described in Section 4.2 and the output categories (Table 1). As the main result, Table 2 gives for each scenario the most definite information: whether a visibility conflict has been detected and whether unknown voxels were involved.

5.1. 3D output for detailed analysis

In the spatial domain, the trajectories colored by visibility degree are the next stage of analysis. In the figures, the background scene is always a labeled point cloud, using the color scheme as presented in Fig. 8. We computed the normals for high-rising objects for visualization so that the trees get a more plastic appearance. Fig. 9 shows these trajectories for the first scenario, with the agents being *car driver* and *recumbent bike rider*. The car driver (Fig. 9(a)) gets clear visibility on the traffic light, except for positions distant to the junction or behind the traffic light (red color). Only the trajectory points far away from the junctions suffer from bad visibility with occlusions caused by tree branches (Fig. 10).

As for the recumbent bike rider (Fig. 9(b)), unknown voxels were detected: some visibility degrees are negative. To analyze the reason for that, Fig. 11 gives further details. In Fig. 11(a), the voxels identified as unknown are colored in blue. In combination with the lines of sight (Fig. 11(b)), this gives a clear picture of which viewpoints are affected by unknown voxels.

The treatment of the incidence angle is visualized in Scenario 2, with a no entrance sign parallel to the street and a car driving by (Fig. 12). The visibility degrees indicate good visibility in general (Fig. 12(a)), yet at a certain distance to the sign, the viewing angle

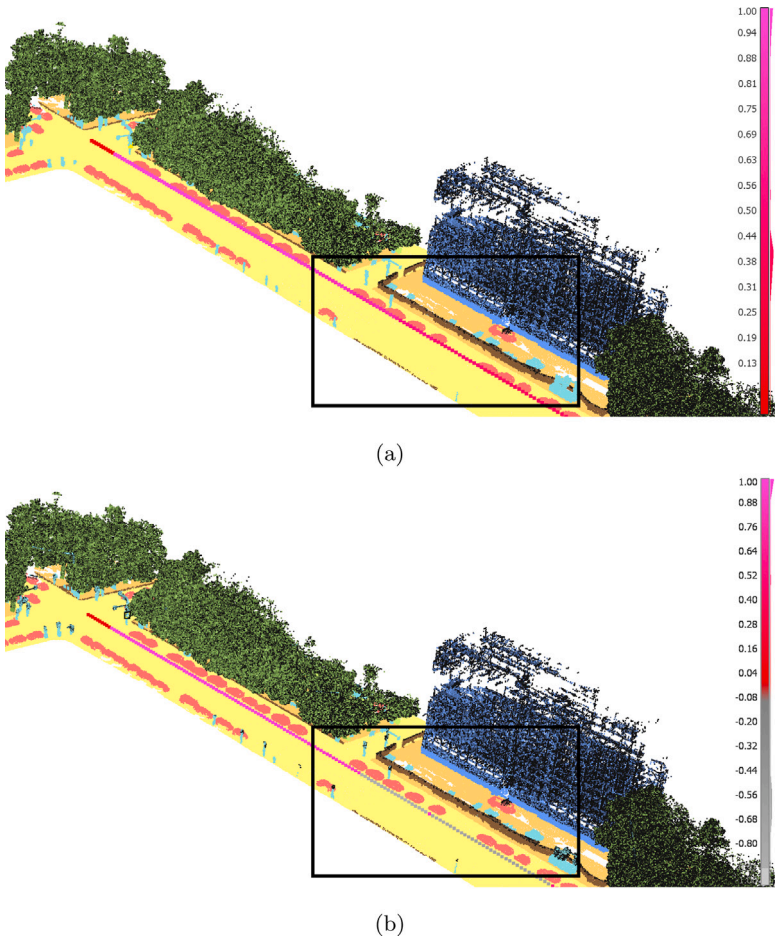

Fig. 9. Agent trajectory colored by visibility degrees for Scenario 1. (a) Agent *car driver*. (b) Agent *recumbent bike rider*. As (b) contains points from which unknown voxels are involved, the negative visibility degrees are depicted by grey values. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

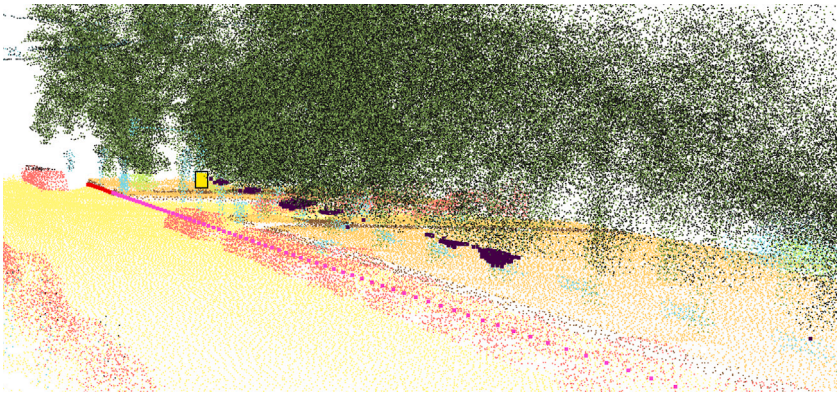

Fig. 10. Occluding voxels for the car driver approaching the traffic light (marked yellow). The occluding voxels (dark purple) belong to the class high vegetation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

becomes too shallow, rendering the sign invisible. This explains the red-colored viewpoints in the upper left part of Fig. 12(a). For the short range of viewpoints with bad visibility before the sign is passed, Fig. 12(b) gives the reason. The occluding element is depicted by the turquoise voxels. Labeled as an artifact, the object is indeed a ticket machine, which in that case occludes the target sign on a small interval of the trajectory.

The former cases dealt with analyses of agents traveling along a similar trajectory as the measurement system. Scenario 4 describes an agent coming from another street into the acquired area. The visibility of a one-way sign is examined in Fig. 13. We obtain negative visibility degrees along the first part of the trajectory (grey trajectory voxels in Fig. 13(a)) while entering the area of interest. As the measurements did not cover the complete volume around the trajectory, a definite statement of visibility cannot be obtained on the first part of the trajectory.

Besides that, visibility is good except for one conflicting case. A voxel is identified as occlusion directly in front of the sign (Fig. 13(b)). In the resolution of the voxel grid, it is associated with an artifact label (indicating that it is part of another traffic sign or traffic light).

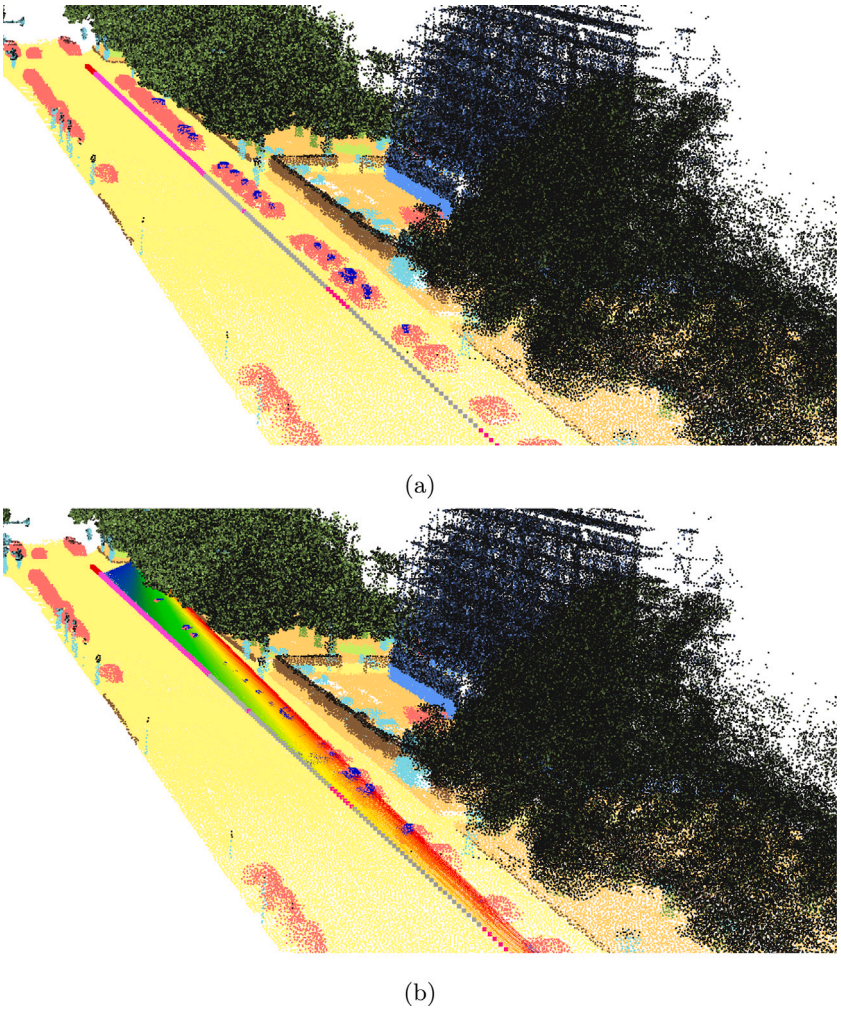

Fig. 11. Visibility of the traffic light (Scenario 1) for a recumbent bike rider. (a) The unknown voxels are marked in blue. (b) Additional lines of sight from each viewpoint to the target. The LoS are colored by their originating viewpoint, from red following the rainbow to blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
Level 0 results: Indicators on visibility conflicts and the existence of unknown voxels in the lines of sight.

Scenario	Visibility conflict	Unknown voxels
1: Traffic light		
(a)	Yes	No
(b)	Yes	Yes
1: No parking sign		
(a)	Yes	No
(b)	Yes	No
2: No entrance sign		
(a)	Yes	No
(b)	Yes	No
3: Pedestrian light ahead		
(a)	Yes	No
3: Pedestrian light to the left		
(a)	Yes	No
4: Priority sign		
(a)	Yes	Yes
4: One-way sign		
(a)	Yes	Yes

Comparison to the image data in Fig. 13(c) proves that, although an occlusion by the pedestrian light seems rather unlikely. We will give a possible explanation in the discussion section (Section 6).

5.2. Distribution of occluding semantic classes

The spatial output formats in Levels 1–3 include semantic knowledge of occluding voxels. In the following, the types of occlusion voxels are depicted by a class histogram to illustrate the working principle. Each histogram evaluates the trajectory of one agent in terms of the visibility of different targets. In that way, characteristic occlusion candidates can be identified for the individual traffic participant. Fig. 14 exemplarily shows the class histograms for a car and a recumbent bike driving along Arcisstraße. With the targets being a traffic light and traffic sign at the junction ahead (Scenario 1) and a traffic sign beside the street (Scenario 2), different viewing geometries are included. For the two installations in front of the car, high vegetation is the main occluding factor. This has also been shown by the occluding voxels in Fig. 10. In addition, the low eye height of the recumbent bike rider causes the lines of sight to traverse through parked cars (Fig. 14(b)). As described in Section 3.4.2, dynamic classes such as car and noise are ignored in the following decision process.

5.3. Variation of voxel sizes

One key parameter is the size of the voxels forming the underlying occupancy grid. Our choice for 10 cm was first motivated by the expected absolute point accuracy in combination with reasonably expected variations of branch positions during the season. We validate

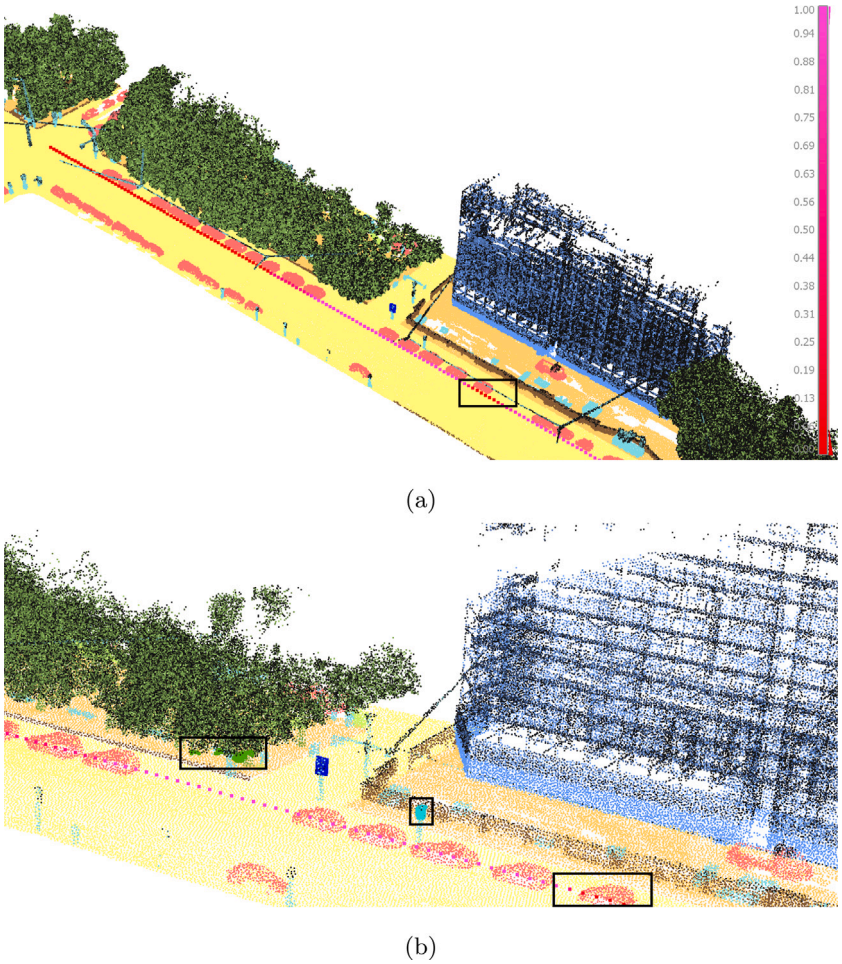

Fig. 12. Scenario 2: A car driving by a sign parallel to the street. (a) Visibility of the sign (colored blue) described by the visibility degrees. (b) Occluding voxels labeled according to their semantic class (green: high vegetation; turquoise: artifact). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Mean visibility value and rate of viewpoints containing unknown voxels for each trajectory, depending on the chosen voxel size for the occupancy grid.

Scn.	Agent	Target	Unknown @ 7.5 cm	Unknown @ 10 cm	Unknown @ 20 cm	MeanVis @ 7.5 cm	MeanVis @ 10 cm	MeanVis @ 20 cm
1a.1	Car	Traffic light ahead	0.0%	0.0%	0.0%	67.7%	60.7%	69.0%
1a.2	Car	No parking far ahead	0.0%	0.0%	0.0%	35.9%	35.6%	42.5%
1b.1	Rec. bike	Traffic light ahead	48.1%	44.4%	0.5%	74.9%	73.2%	74.2%
1b.2	Rec. bike	No parking far ahead	11.1%	9.1%	0.5%	37.2%	36.8%	43.6%
2a	Car	No entrance drive-by	0.0%	0.0%	0.0%	28.8%	30.9%	76.9%
2b	Rec. bike	No entrance drive-by	0.0%	0.0%	0.0%	34.0%	33.8%	76.9%
3a.1	Pedestrian	Ped. light ahead	0.0%	0.0%	0.0%	12.8%	13.1%	91.1%
3a.2	Pedestrian	Ped. light left	0.0%	0.0%	0.0%	72.9%	72.5%	53.6%
4a.1	Car	Priority sign	20.0%	20.0%	0.0%	56.7%	56.0%	100.0%
4a.2	Car	One-way sign	100.0%	15.0%	0.0%	98.6%	99.0%	100.0%

that choice by additionally processing all scenarios at a voxel size of 7.5 cm and 20 cm.

Table 3 compares the main metrics of mean visibility value and rate of unreliable viewpoints for each trajectory. Between the voxel sizes 7.5 cm and 10 cm, the largest difference in visibility are 7 percentage points (Scenario 1a.1). The largest percentage difference of viewpoints with unknown voxels in their lines of sight is 85 percentage points (Scenario 4a.2, caused by the sparsity of the point cloud), whereas in the other scenarios that number is always less than 4 percentage points.

Compared to these differences we now look at the metrics for voxel size 20 cm. The mean visibility value can vary up to a difference of 88 percentage points (Scenario 3a.1), completely negating the conclusion from the finer voxel sizes: The pedestrian light now seems to be almost entirely visible. The same effect can be observed with the unknown voxels, being almost always not detected.

From these comparisons, we deduct that at 20 cm too much granularity of information is lost. On the other hand, increasing the sampling to 7.5 cm only leads to small differences in the metrics. Without sufficient ground truth of space occupation, it cannot be stated whether the values at 7.5 cm are more realistic; let alone the small variation in comparison to 10 cm validates our assumption that the chosen voxel size is adequate for this task.

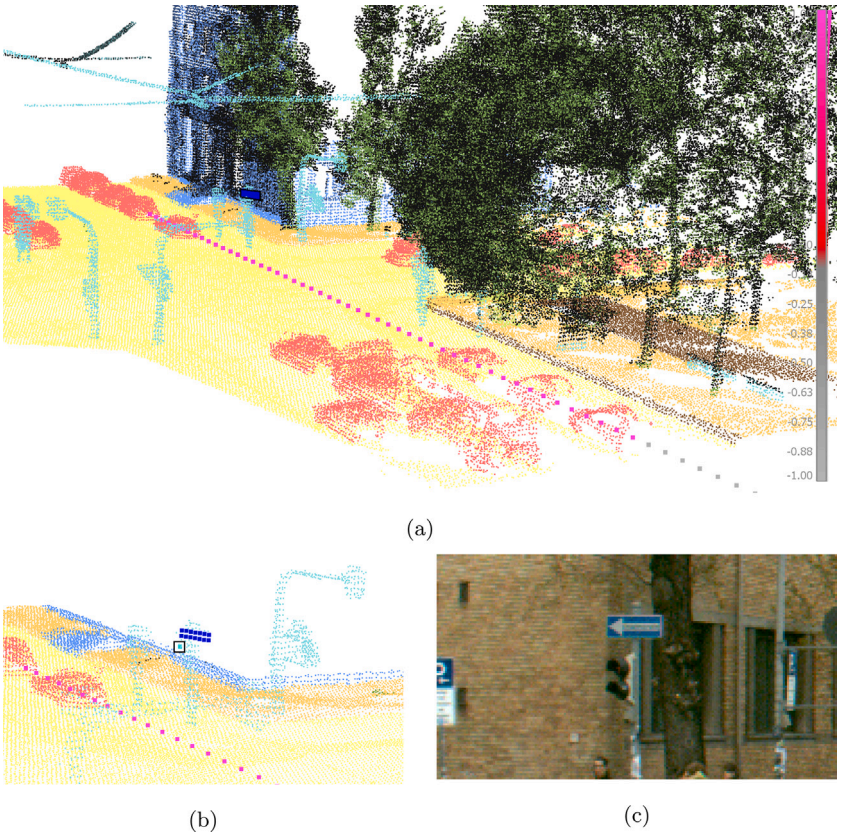

Fig. 13. Car passing a one-way sign along Theresienstraße. (a) Visibility degrees along the trajectory. (b) Single occluding point in turquoise in front of the sign. (c) Detail view on the situation from camera footage.

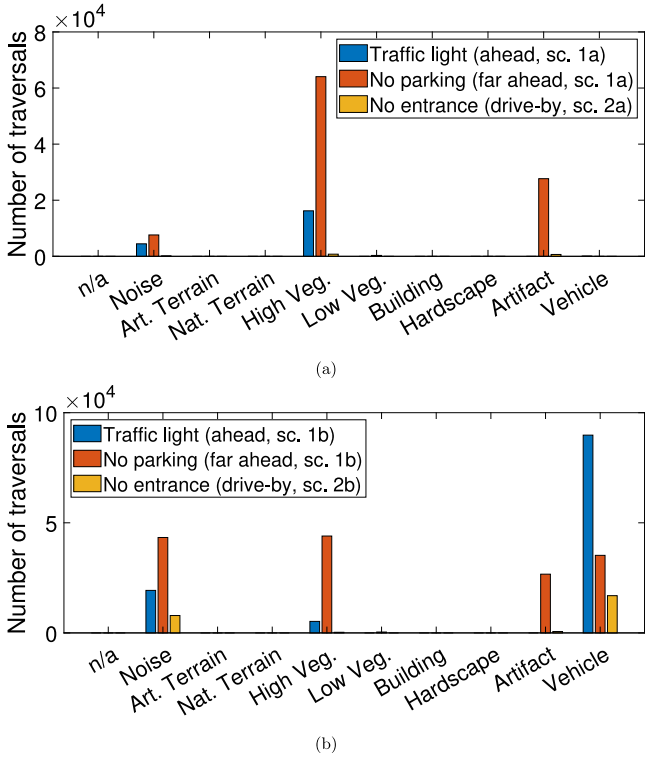

Fig. 14. Distribution of semantic classes traversed by the lines of sight. (a) Agent *car driver*. (b) Agent *recumbent bike rider*, as defined in Scenarios 1 and 2. One voxel can be traversed multiple times from different viewpoints.

6. Discussion

This section discusses the outcomes based on visual interpretation, the performance of technical procedures, and the phenomenons observed during our analysis.

The first discussion point is about the unknown car interior. The examination of Scenario 1 for the recumbent bike riders resulted in many unknown voxels. Fig. 11 reveals all of these voxels being on or inside cars. This can be explained by the laser beam being reflected away by car windows depending on incident angles, or the interior of cars only being partially traversed by measurements. Hence, although the occlusion effect of parked cars is filtered out by semantic information, they are still responsible for unknown volumes. A possible solution could be that we should use the labeled point cloud to further exactly identify the whole volumes covered by parked cars. Then, these volumes could be marked as dynamically occupied in the occupancy grid.

The second discussion point is regarding the occlusion by one point. Scenario 4 demonstrated a visibility conflict caused by one voxel, which was classified as a pedestrian light (Fig. 13). From the perception of the image data, an actual occlusion seems unlikely. The fact that we detect a visibility conflict here might be caused by the effect of voxelization. In other words, due to the process of mapping unordered points into a grid structure, a point belonging to the pedestrian light might have been associated with a voxel with its main part further above the actual object. Thus, the extent of the pedestrian light was extended due to rasterization, and a presumably false occlusion was caused.

At last, regarding the approach in general, the definition of voxel occupation and the treatment of conflict cases needs to be discussed. As explained in Section 3.2, each voxel holds a tuple $V(n, m)$, where n indicates the empty value and m indicates the occupation value. The cases $n = 0$ or $m = 0$ are clear, as is the decision where $n \gg m$ or $m \gg n$.

Nevertheless, there is also a possible case that both n and m could be larger than 0. An option to address this issue could be to count the voxel as occupied only if the occupation value is greater than the empty value, namely in the situation that $n > m$. However, in our test cases this approach resulted in several voxels marked as empty even though they should be occupied. Therefore, we assume a voxel is occupied as soon as $n > 0$ (see Section 3.4.1). This might result in more false positives (empty voxels marked as occupied), but by filtering noise artifacts and introducing semantic knowledge, these influences can be minimized.

7. Conclusion

In this work, we carried out an analysis of the visibility of traffic signs and traffic lights by ray-tracing in an occupancy grid in the urban street environment. Both city models and MLS point clouds were utilized as datasets for conducting this work, followed by voxelizing the data into a 3D grid for structured and quantitative analysis. The generated voxels facilitated the distinction between occupied and empty space. During the occlusion detection, the identification of unknown volumes was added and considered in the decision process. For further analysis, we provided detailed three-dimensional results on different levels of granularity, including the knowledge of the semantic classes of traversed voxels. Ultimately, the awareness of unknown volumes was added to an otherwise binary decision between visible and invisible targets.

In conclusion, our proposed method is feasible for occlusion detection in the street scenario, achieving promising results. Based on the experiments, we particularly discussed difficulties with parked cars, realizing that the idea of ignoring occlusions by dynamic objects was contradicted by the inner space of cars not being covered by the sensor and hence causing unknown voxels. A further step would be to include parking spaces provided by the 3D city model, acknowledging that if the model contains a parking space for cars, the whole volume above—up to an average car height—can be expected to be occupied, whereas driving cars on lanes are still ignored in visibility analysis. This opens further possibilities for combining city models and mobile measurements. Our experiments also reveal that the identification of unknown volumes proves necessary for reliable interpretation of the measurements, especially if the agents differ from the characteristics of the MLS vehicle or its trajectory. However, identifying unknown volumes can only be seen as an intermediate step. Following the identification of problematic viewpoints, a final decision must be reached otherwise: for instance by human interaction based on different output formats, further in situ observations for selected cases, or by an extension of the current analysis.

To deal with the above-mentioned issues, in the future, the used dataset would have to go from semantically labeled point clouds to actual volumetric car objects created from the labeled point cloud. This would require segmentation and modeling of individual car instances. Their complete volume could then be marked and ignored in visibility analysis. In case traffic spaces do also model the clearance volume, this volume can be included in the analysis. Vegetation protruding into that clearance space can then be identified without causing actual visibility conflicts. A general assessment of visibility can then be performed by clearance checks, with the agent-based solution being a refinement at especially dangerous points, or if changes to the existing environment are planned. Thus, further research will offer different applications by an extended environment model.

CRediT authorship contribution statement

Philipp-Roman Hirt: Conceptualization, Methodology, Software, Writing – original draft, Visualization. **Jonathan Holtkamp:** Software, Investigation, Writing – original draft. **Ludwig Hoegner:** Conceptualization, Writing – review & editing. **Yusheng Xu:** Writing – review & editing. **Uwe Stilla:** Supervision, Visualization, Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Co-author is guest editor of this Special Issue – Y.X.

Data availability

The MLS benchmark dataset TUM-MLS-2016 is openly available on the web. The authors do not have permission to publish the CityGML models.

Acknowledgments

We thank Dr. Marcus Hebel from Fraunhofer IOSB, Ettlingen, for providing the additional camera footage from the 2016 measurement campaign. The work was carried out in the frame of the Leonhard Obermeyer Center (LOC) at the Technical University of Munich.

References

Aleksandrov, M., Zlatanova, S., Heslop, D.J., 2021. Voxelisation algorithms and data structures: A review. Sensors 21 (24), 8241. http://dx.doi.org/10.3390/s21248241.

Aleksandrov, M., Zlatanova, S., Kimmel, L., Barton, J., Gorte, B., 2019. Voxel-based visibility analysis for safety assessment of urban environments. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. IV-4W8, 11–17. http://dx.doi.org/10.5194/isprs-annals-IV-4-W8-11-2019.

Alsadik, B., Gerke, M., Vosselman, G., 2014. Visibility analysis of point cloud in close range photogrammetry. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. II-5, 9–16. http://dx.doi.org/10.5194/isprsannals-II-5-9-2014.

Barros-Ribademar, J., Balado, J., Arias, P., González-Collazo, S.M., 2022. Visibility analysis for the occlusion detection and characterisation in street point clouds acquired with mobile laser scanning. Geocarto Int. 1–18. http://dx.doi.org/10.1080/10106049.2022.2032392.

Beil, C., Kutzner, T., Schwab, B., Willenborg, B., Gawronski, A., Kolbe, T.H., 2021. Integration of 3D point clouds with semantic 3D city models – providing semantic information beyond classification. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. VIII-4/W2-2021, 105–112. http://dx.doi.org/10.5194/isprs-annals-VIII-4-W2-2021-105-2021.

Beil, C., Ruhdorfer, R., Coduro, T., Kolbe, T.H., 2020. Detailed streetspace modelling for multiple applications: Discussions on the proposed CityGML 3.0 transportation model. ISPRS Int. J. Geo-Inf. 9 (10), 603. http://dx.doi.org/10.3390/ijgi9100603.

Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., Öltekin, A., 2015. Applications of 3D city models: State of the art review. ISPRS Int. J. Geo-Inf. 4 (4), 2842–2889. http://dx.doi.org/10.3390/ijgi4042842.

Deng, T., Zhang, K., Shen, Z.-J., 2021. A systematic review of a digital twin city: A new pattern of urban governance toward smart cities. J. Manag. Sci. Eng. 6 (2), 125–134. http://dx.doi.org/10.1016/j.jmse.2021.03.003.

Gehrung, J., Hebel, M., Arens, M., Stilla, U., 2017. An approach to extract moving objects from MLS data using a volumetric background representation. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. IV-1/W1, 107–114. http://dx.doi.org/10.5194/isprs-annals-IV-1-W1-107-2017.

Gehrung, J., Hebel, M., Arens, M., Stilla, U., 2018. A voxel-based metadata structure for change detection in point clouds of large-scale urban areas. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. IV-2, 97–104. http://dx.doi.org/10.5194/isprs-annals-iv-2-97-2018.

Gorte, B., Zlatanova, S., Fadli, F., 2019. Navigation in indoor voxel models. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. IV-2/W5, 279–283. http://dx.doi.org/10.5194/isprs-annals-IV-2-W5-279-2019.

Hebel, M., Arens, M., Stilla, U., 2013. Change detection in urban areas by object-based analysis and on-the-fly comparison of multi-view ALS data. ISPRS J. Photogramm. Remote Sens. 86, 52–64. http://dx.doi.org/10.1016/j.isprsjprs.2013.09.005.

Hirt, P.-R., Xu, Y., Hoegner, L., Stilla, U., 2021. Change detection of urban trees in MLS point clouds using occupancy grids. PFG – J. Photogram. Remote Sens. Geoinform. Sci. 89 (4), 301–318. http://dx.doi.org/10.1007/s41064-021-00179-4.

Huang, P., Cheng, M., Chen, Y., Luo, H., Wang, C., Li, J., 2017. Traffic sign occlusion detection using mobile laser scanning point clouds. IEEE Trans. Intell. Transp. Syst. 18 (9), 2364–2376. http://dx.doi.org/10.1109/TITS.2016.2639582.

Karami, S., Taleai, M., 2020. An innovative three-dimensional approach for visibility assessment of highway signs based on the simulation of traffic flow. J. Spat. Sci. 1–15. http://dx.doi.org/10.1080/14498596.2020.1787253.

Katz, S., Tal, A., Basri, R., 2007. Direct visibility of point sets. ACM Trans. Graph. 26 (3), 24–es. http://dx.doi.org/10.1145/1276377.1276407.

Kolbe, T.H., Kutzner, T., Smyth, C.S., Nagel, C., Roensdorf, C., Heazel, C., 2021. OGC City Geography Markup Language (CityGML) Version 3.0 Part 1: Conceptual Model Standard. Open Geospatial Consortium, International Standard.

Ma, Y., Zheng, Y., Wong, Y.D., Easa, S., Cheng, J., 2022. A virtual procedure for real-time monitoring of intervisibility between conflicting agents at intersections using point cloud and trajectory data. Transp. Res. C 134, 103486. http://dx.doi.org/10.1016/j.trc.2021.103486.

Peng, Y., Nijhuis, S., Zhang, G., Stoter, J.E., Agugiaro, G., 2022. Towards a practical method for voxel-based visibility analysis with point cloud data for landscape architects: Jichang garden (Wuxi, China) as an example. J. Digital Landsc. Archit. 7–2022 (7), 682–691. http://dx.doi.org/10.14627/537724063.

Rutzinger, M., Pratihast, A., Oude Elberink, S., Vosselman, G., 2011. Tree modelling from mobile laser scanning data-sets. Photogramm. Rec. 26 (135), 361–372. http://dx.doi.org/10.1111/j.1477-9730.2011.00635.x.

Schwab, B., Kolbe, T.H., 2019. Requirement analysis of 3D road space models for automated driving. ISPRS Ann. Photogram. Remote Sens. Spat. Inf. Sci. IV-4/W8, 99–105. http://dx.doi.org/10.5194/isprs-annals-IV-4-W8-99-2019.

Seif, H.G., Hu, X., 2016. Autonomous driving in the iCity—HD maps as a key challenge of the automotive industry. Engineering 2 (2), 159–162. http://dx.doi.org/10.1016/J.ENG.2016.02.010.

Soilán, M., Riveiro, B., Sánchez-Rodríguez, A., Arias, P., 2018. Safety assessment on pedestrian crossing environments using MLS data. Accid. Anal. Prev. 111, 328–337. http://dx.doi.org/10.1016/j.aap.2017.12.009.

Voelsen, M., Schachtschneider, J., Brenner, C., 2021. Classification and change detection in mobile mapping LiDAR pointclouds. PFG – J. Photogram. Remote Sens. Geoinform. Sci. 89 (3), 1–13. http://dx.doi.org/10.1007/s41064-021-00148-x.

Weinmann, M., Weinmann, M., Mallet, C., Brédif, M., 2017. A classification-segmentation framework for the detection of individual trees in dense MMS point cloud data acquired in urban areas. Remote Sens. 9 (3), 277:1–277:28. http://dx.doi.org/10.3390/rs9030277.

Wu, S., Wen, C., Luo, H., Chen, Y., Wang, C., Li, J., 2015. Using mobile LiDAR point clouds for traffic sign detection and sign visibility estimation. In: 2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS). pp. 565–568. http://dx.doi.org/10.1109/IGARSS.2015.7325826.

Wu, J., Yao, W., Polewski, P., 2018. Mapping individual tree species and vitality along urban road corridors with LiDAR and imaging sensors: point density versus view perspective. Remote Sens. 10 (9), 1403. http://dx.doi.org/10.3390/rs10091403.

Xu, Y., Stilla, U., 2021. Toward building and civil infrastructure reconstruction from point clouds: A review on data and key techniques. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens. 14, 2857–2885. http://dx.doi.org/10.1109/JSTARS.2021.3060568.

Xu, Y., Tong, X., Stilla, U., 2021. Voxel-based representation of 3D point clouds: Methods, applications, and its potential use in the construction industry. Autom. Constr. 126, 103675. http://dx.doi.org/10.1016/j.autcon.2021.103675.

Xu, Y., Ye, Z., Yao, W., Huang, R., Tong, X., Hoegner, L., Stilla, U., 2020. Classification of LiDAR point clouds using supervoxel-based detrended feature and perception-weighted graphical model. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens. 13, 72–88. http://dx.doi.org/10.1109/JSTARS.2019.2951293.

Zhang, S., Wang, C., Lin, L., Wen, C., Yang, C., Zhang, Z., Li, J., 2019. Automated visual recognizability evaluation of traffic sign based on 3D LiDAR point clouds. Remote Sens. 11 (12), 1453. http://dx.doi.org/10.3390/rs11121453.

Zhu, J., Gehrung, J., Huang, R., Borgmann, B., Sun, Z., Hoegner, L., Hebel, M., Xu, Y., Stilla, U., 2020. TUM-MLS-2016: An annotated mobile LiDAR dataset of the TUM city campus for semantic point cloud interpretation in urban areas. Remote Sens. 12 (11), 1875. http://dx.doi.org/10.3390/rs12111875.