

DETERMINATION OF PARKING SPACE AND ITS CONCURRENT USAGE OVER TIME USING SEMANTICALLY SEGMENTED MOBILE MAPPING DATA

A. Leichter^{1*}, U. Feuerhake¹, M. Sester¹

¹ Institute of Cartography and Geoinformatics, Leibniz Universität Hannover, Germany
(leichter, feuerhake, sester)@ikg.uni-hannover.de

WG I/7, II/10

KEY WORDS: Smart City, Mobile Mapping, Deep Learning, Data Fusion, Parking, Public Space.

ABSTRACT:

Public space is a scarce good in cities. There are many concurrent usages, which makes an adequate allocation of space both difficult and highly attractive. A lot of space is allocated by parking cars - even if the parking spaces are not occupied by cars all the time. In this work, we analyze space demand and usage by parking cars, in order to evaluate, when this space could be used for other purposes. The analysis is based on 3D point clouds acquired at several times during a day. We propose a processing pipeline to extract car bounding boxes from a given 3D point cloud. For the car extraction we utilize a label transfer technique for transfers from semantically segmented 2D RGB images to 3D point cloud data. This semantically segmented 3D data allows us to identify car instances. Subsequently, we aggregate and analyze information about parking cars. We present an exemplary analysis of the urban area where we extracted 15.000 cars at five different points in time. Based on this aggregated we present analytical results for time dependent parking behavior, parking space availability and utilization.

1. INTRODUCTION

Streets, sidewalks, roads or public spaces in general are places where advantages and disadvantages of urban life lead to overlapping challenges (de Magalhaes and Carmona, 2009). Public space is characterized by shared use through different actors. As the demands on cities intensify, shared use becomes a competition for this limited resource. To address this development, it is crucial to quantify public space itself and its usage. In this paper we propose to solve this task by exploiting the continuous acquisition of environmental data with vehicle sensors and the subsequent application of a deep learning (DL) model tailored to semantic segmentation of mobile mapping data (Peters and Brenner, 2019). The segmented data is aggregated to retrieve spatial and statistical usage information.

Public space is characterized by the accessibility for all users with as little restriction as possible. This open definition is necessary in densely populated urban spaces and has the advantage of the flexibility and the maximization of usage (Oranratmanee and Sachakul, 2014). Different types of users (private, commercial, governmental) use public spaces in the context of different applications (transportation, commercial activity, recreation, etc.). The disadvantages of shared use are competition for limited space and resulting poor planning ability (situational for users, global for administration). The quantification of the available space in the terms of mapping and dynamic recording of temporal usage opens the possibilities for the administration and users to compensate for the disadvantages described above, and increases the potential usage.

The problem of identifying on-street parking statistics has already been tackled by (Bock et al., 2015), using classical machine learning methods for the semantic segmentation of the points clouds. In our approach, we perform the semantic segmentation by means of a DL model. This makes the approach

more robust and flexible for investigation of different object types (semantic classes), which also allows us to perform further analyses, e.g., the pedestrians, cyclists. Identification of unique objects within a semantically segmented point cloud (or image) is the task of instance segmentation. Or, if organized in an end to end manner it is called panoptic segmentation (Kirillov et al., 2019). This kind of task can be tackled with machine (Deep) learning approaches, as proposed in (Liu et al., 2020), (Schlichting and Brenner, 2016) or (Hong et al., 2020). In our case the training data is missing for supervised machine learning approach. As a consequence we rely on a unsupervised approach to generate a sufficient amount of data for analysis. This study bases on car position information extracted from semantic segmentation of 3D point cloud with corresponding RGB images (Peters and Brenner, 2019). Such a semantic segmentation can be created independent of our approach with different models like the 3D point cloud driven model (Zhou et al., 2020) or by projecting semantic segmentation labels from RGB images to 3D point cloud (Kochanov et al., 2020). The point-wise semantic segmentation is then refined into the instance segmentation of the relevant objects. The lateral extents (bounding boxes) are transformed to global coordinates and accumulated over time in an occupancy grid. The acquisition of data for this study is done using a Mobile Mapping System (MMS). The MMS consists of two planar laser scanners and four RGB cameras coupled with global navigation satellite system (GNSS) receivers. The MMS allows recording of detailed spatial information. Use of MMS in smart city applications is common. In this study the temporal dimension of area usage is captured by acquiring several measurements in a defined area in a given temporal interval. This data set is used as a proof of concept of our approach. However, in the future, such a concept can potentially be carried out on large scale based on the sensor data of future (autonomous) vehicles.

The remainder of this paper is structured as follows: Section 2 presents the approach to determine and map the concurrent us-

* Corresponding author

ages of the space. In section 3 we conduct an experiment using real data and present its results, which are discussed in the subsequent section 4. With the temporal evaluation of the parking space we show different applications which are based on our mapping of the MPS. We conclude this paper by giving a summary and different aspects, which will be addressed in future work.

2. APPROACH

Although on-street parking spaces usage or availability is a relevant information, they are hardly mapped in today's navigation systems. Therefore, the idea of this paper is to automatically extract this information by road users and their sensors (Figure 1 left). The basic assumption is that the available space can be automatically determined by observing the whole space's usage over time. With the focus on public spaces in the vicinity of roads, Mobile Mapping Systems are a suitable measurement tool, as they scan the environment around the vehicle with LiDAR and/or cameras. As it is rather difficult to automatically detect a parking space, our approach follows the idea of (Bock et al., 2015) and observes parked cars, as a proxy. Those objects can be easily and reliably detected with today's Machine Learning approaches (Tao et al., 2020). However, a parked car does not necessarily mean that there is a (legal) parking place. Thus, the temporal behavior of parking also has to be analyzed.

The approach to implement this idea consists of the three consecutive steps "car extraction", "information aggregation" and "information analysis and visualization". In the following, these steps are described in detail.




Figure 1. Overview of the processing pipeline.

2.1 Car extraction

Given a point cloud capturing a street scene our goal is to assign all unique car instances a unique ID and map this to all 3D points belonging to the respective car. In our workflow we extract the car bounding boxes by means of six serially applied processing steps (Figure 2). Since we are interested in the spatial coverage of vehicles in on-street parking, the process is designed to extract parked vehicles on the street covered by the MMS data captures. The vehicles in the crossed side streets are not relevant in this context. In this way, it is not necessary to interpret the less dense, more distant parts of the point cloud. The processing workflow is designed in such a way that vehicles are extracted without outliers as far as possible, so that the spatial information about the occupancy is not distorted. This means in the trade-off between precision and recall, precision was prioritized. The resulting lower number of extracted vehicles is compensated for by the number of repeated recordings and thus still a complete picture of usage is generated.

Besides cars, also other objects e.g. pedestrian, cyclist and parked bicycles in the road environment can be identified using the same workflow, which, however, is not considered in this paper.




Figure 2. Overview of the workflow for bounding box extraction.

(1) The point-wise semantic segmentation of the point clouds is performed using the label transfer framework by (Peters and Brenner, 2019), which avoids the time consuming and error prone labeling of 3D point clouds. The input data is composed of LiDAR data and RGB images. In the first step, the RGB images are pre-segmented by the pre-trained (dataset: (Cordts et al., 2016)) DeepLabV3 (Chen et al., 2017) model. In the next step, the pixel-wise predictions from the RGB images are projected to the LiDAR data (Peters and Brenner, 2019). The output contains predictions of each 3D point into 14 classes, namely building, bicycle, fence, wall, traffic sign, person, pole, vegetation, car, sidewalk, rider, road. Our car class definition differs from the original one in (Cordts et al., 2016); we consider points labeled as bus and truck also as cars.

(2) In order to reduce the influence of the projection errors and labeling errors in general, we apply a second processing step which utilizes information about homogeneous regions (Brenner, 2016) to filter falsely classified points (Felzenszwalb and Huttenlocher, 2004).




Figure 3. Example homogeneous regions.

(3) In the third step of the process we remove not only the points that are classified as ground points but also such that have similar height values. In this way the objects are separated into isolated clusters and an additional result is that falsely classified ground points are also removed from the scene. To identify ground points we investigate the heights distribution of the points labeled as ground in earlier steps. We ignore 5% of the highest points as outliers and use the maximal height of the remaining 95% we use to crop of all ground points.

(4) The points labeled as car from the remaining point cloud are clustered by means of the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm (McInnes and Healy, 2017) (Campello et al., 2013).

(5) The cluster from the previous step contain outliers, namely points from neighboring clusters which have been wrongly assigned in the previous step. Furthermore, not all the points depicting a car are assigned to the car clusters. In this step the the outliers are removed and the clusters are maximized to the all points depicting a car. This step is executed isolated per single identified cluster C as set of 3D points. We calculate centers of mass (com) for all cluster points C . We select within a radius r_s to com as car labeled points which are most probable to be part

of a car and call this points $seed \subseteq C$. Then we select points within a radius r_c to com from the result from the step (3) (complete point cloud without the ground points). These points contain all potential car points. We call them *candidates*. Subsequently, we voxelize both seed and candidates and use the region growing algorithm to identify all car points within the *candidate* grid based on the *seed* input.

(6) Finally, the instances represented by different clusters are filtered based on a set of rules. A car instances is assumed to be valid if all following criteria are met.

- The car instance consists of more than $minp$ minimal number of points.
- The length of the bounding box is less than $maxl$ the maximal length.
- The width of the bounding box is more than $minw$ the minimal width.

2.2 Information aggregation

Subsequently, similar to the work of (Schlichting and Brenner, 2016), the footprints of the extracted cars are put into relationship with the street geometry to identify potential parking space. For this purpose, prior environmental information like street and building geometries and the previously extracted car observations at certain locations are taken into account. In order to represent the temporal dynamics of the car occupancies, the space is discretized by an occupancy grid, in which each grid cell holds the information about the car observations and their observation times. The observation data model for each cell is given by

$$observations = \{o_1, o_2, \dots, o_n\}. \quad (1)$$

$$o_i = \{x1, y1, x2, y2, x3, y3, x4, y4, t, type\} \quad (2)$$

and

where o_i = the i -th observation
 x_i, y_i = bounding box of the object's footprint
 t = the observation time
 $type$ = the object's type.

The prior information about the environment is used to exclude certain regions which usually cannot be used for parking or other purposes. To this end, building footprints and the lanes of the street are left out during the discretization. The information about the footprints and the street lanes are obtained from OpenStreetMap (OSM) (contributors, 2017). The latter are approximated by applying a buffer operation to the line geometries of the roads, because OSM does not provide any areal information for roads. Further, we also omit spaces, which are too far from any street. The resolution of the grid, i.e. the cell dimensions, is adjustable to the later use case. In general, with increasing resolution, the accuracy of the approximation of the space but also the processing effort will increase, too.

2.3 Information analysis and visualization

The above-described pipeline allows us to present an easy updateable map of parking space. Furthermore, for this map the usage over time is available. In addition, the generated data

allows us to compile an overview of possible alternating usages. For instance, free parking slots can be temporarily used for other purposes such as mobile electrical charging stations or mobile logistics delivery hubs. Such instances can be determined using GIS-analyses by searching for locations of a certain size, which are available for a certain duration in time.

3. EXPERIMENT AND RESULTS

In order to evaluate the proposed approach, we conducted an experiment based on data obtained from an already existing mapping campaign. In the following, the application and the results of the individual intermediate steps will be described in detail.

3.1 Mapping campaign and data

The campaign provides data from a defined round course (cf. Figure 4) on different days and day times. This round leads to the northern part of the city of Hanover, Germany, and covers 26 km. During the campaign, it has been completed 5 times in total. Certain locations have been visited twice in one round.




Figure 4. The route of the mapping campaign which provides the data is in the northern part of Hanover, Germany (basemap: Stamen Maps).

For the acquisition of the required data the mobile mapping system Riegl VMX-250 (cf. Figure 5) has been used. This system uses two 2D lasers VQ-250 with a sampling rate of 300k measurements per second, a precision of 5 mm and an accuracy of 10 mm. The lasers are mounted on the roof of a Volkswagen T4 and are directed backwards. For the positioning of the system, position high-precision global navigation satellite system (GNSS) is used with an Inertial Measurement Unit (IMU) (Applanix POS-LV 510) with a position accuracy of 15-30cm in height and 20cm in lateral plane. The RGB images are captured by 4 rear- and side-facing cameras.

3.2 Car Extraction

Applying the car extraction process (2.1) to the experiment provides the following results.

The result of the first processing step (1) for the car extraction is presented in the upper part of the Figure 6. The visualization of the 3D point cloud colored by the assigned semantic label is mostly valid. The trees are classified as vegetation. Wall, fence, road and sidewalk are also mostly classified correctly.



Figure 5. Riegl VMX-250 mounted on the top of the volkswagen T4.

The classification of the fence at the left side of the upper Figure 6 includes several artifacts. Several blocks of points are falsely classified as sidewalk, vegetation and car. The three cars visible (in blue) at this scene are also mostly classified correctly. At the bottom part of the scene in the foreground on the left several misclassifications as road are visible. After applying the second processing step (2) from the car extraction process several misclassification are fixed. Three prominent examples of such an improvement are marked on the left side of the bottom part of Figure 6.




Figure 6. Visualization of the 3D labeled data. At the top we see the results of the projected semantic segmentation from the processing step (1). Bottom part shows the improved segmentation from the processing step (2). The color map: Blue - Car, Purple - Road, Pink - Sidewalk, Green - Vegetation, Light Green - Terrain, Yellow - Traffic Sign, Grey - (Wall, Pole, Fence), Light Blue - Accent Marker

In the third step of the car extraction process (3) the ground points are removed. In Figure 7 an exemplary result is shown. Single objects like trees and cars on the street are separated into isolated clusters. Furthermore this step removes some false positive car classifications. In Figure 8 points classified as car are shown before the subtraction of the ground points (top) and after the processing step (3). The light blue circle shows an

area where ground points falsely classified as car have been removed.




Figure 7. Visualization of the 3D labeled data with removed ground points. The result of the third car extraction processing step (3). The color map: Same as in Figure 6




Figure 8. Visualization of the points labeled as car. Top part shows point cloud before removing the ground points and at the bottom part after removing the ground points.

Following the removing of the ground points, the remaining points labeled as car are clustered in the fourth processing step (4). The result is shown in the upper part of the Figure 9. The three cars in the scene are assigned to unique clusters. The cluster of the car on the left contains a large number of outliers. In addition, some of the points falsely classified as car have been also assigned to unique clusters, they can be seen in Figure 9 marked with green and yellow color. This assignments to the clusters are used in the fifth processing step (5) - the region growing - to generate *seed* points and the data shown in Figure 7 is used to generate *candidate* points for the region growing algorithm. Here we set r_s and r_c to $0.5m$ and $3m$ respectively. The result of the region growing step is shown in Figure 9, bottom part. The shown point clouds follow a homogeneous grid as they show centroids of the voxel generated within the region growing step. All three cars present in the scene are identified as unique cluster with no visible outliers included. The only remaining outliers are two identified clusters based falsely as car classified fence points. This false classified clusters can be seen on the left side of the Figure 9 marked with green and yellow color.

In the final step we apply rules to filter not reasonable car clusters. Nonparallel lateral bounding boxes are estimated for each cluster. Such cluster with width less than $0.5m$ and length larger than $6m$ are removed as they don't fit the typical extent of a car. Also sparse and/or small clusters with less than 500 points are ignored. The result is shown in Figure 10. All three cars in scene are identified as unique, separate instances.

After processing 183GB of data and 5.2 billion points in total 15.914 cars are extracted. Their footprints are depicted in Fig-



Figure 9. Visualization of the points labeled as car clustered and colored in single random color per cluster: before (upper part) and after region growing (lower part).




Figure 10. Visualization of the final result of the car extraction. The points are colored in single random color per car instance.

Figure 11 with red color. In this example, besides the cars located along the streets also those cars on a dedicated parking place (center left) are detected. To verify the accuracy, We investigated manually three randomly selected sections of data. Out of 100 cars present in the data 72 have been instantiated correctly and 7 objects have been falsely identified as cars.

3.3 Information aggregation

According to the aggregation step (subsection 2.2) the occupancy grid (Figure 13) is generated. To have a good trade-off between approximation accuracy and computational effort, we use a cell size of 1 m x 1 m. This leads to about 11M grid cells in total.

In order to calculate the occupancies, the extracted cars are projected into the grid. This results in 37235 cells, which own at least one car observation. The maximum count of observations on the grid per cell corresponds to the number of visits during the campaign. This means that this location has been occupied each time it has been visiting during the experiment. In contrast to that, we can also observe cells holding only a single observation. Figure 12 shows the percentage of cells according to the count of observations. Approximately 25 % of the cells are visited at least 5 times.

The results are visualized in Figure 13. Each grid cell holds the information about the different objects and their observation



Figure 11. The extracted cars (red boxes) on top of a basemap (Stamen maps).




Figure 12. Proportion of the cells according to the count of stored observations.

times as shown in the top part of the figure. The color intensity encodes the observation frequency. In this example, a row of cars (black footprints) has been observed in several rounds of the campaign. This can be a hint that there is a potential space for parking. However, the observation frequency is not the same for all the locations. For instance, the red marked cell is occupied more often than the blue cell. The obvious reason for this is that the underlying parking possibility seems not to be used that regularly. By having a closer look at the occupancy times, it seems that blue location is usually used during lunch time, whereas the red one is also used in the morning.

3.4 Information analysis and visualization

Based on the generated occupancy grid different analyses can be performed. The first and directly derivable evaluation is the characterization of the popularity of the parking spaces based on their occupancy rates. The higher it is, the higher the popularity is. On the opposite, low occupancy rates indicate unpopular or atypical parking possibilities. Please note, that we considered cells with at least one parking car as a potential parking space - of course, before doing so, the distinction into legal and non-legal has to be made. Cells with no observation at all are

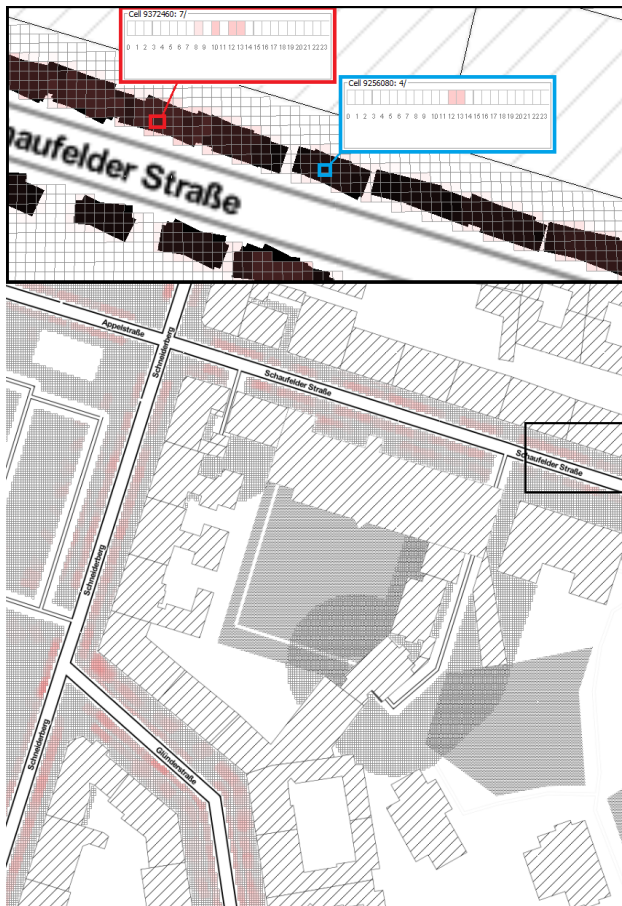


Figure 13. The number of occurrences of extracted cars are encoded by the color intensity of the grid cells (basemap: Stamen Maps).

not considered for this analysis as there is no evidence that there is parking space in general.

Another analysis is performed by evaluating the usual occupancy times. To this end, the time of the day is split up into different intervals. Subsequently, the cells are classified by the occupancy during those time intervals. If there is at least one observation in an time interval, the latter will be treated as occupied. Otherwise it will be treated as free. In the example given by Figure 14, there are two intervals: 1 - 12 am and 13 - 24 pm. While the green and blue cells indicate spaces which are usually occupied during the first and second interval, the red cells show regions which are occupied in both intervals. In this example, the framed street shows a different behavior than the neighboring streets. The occupancy is limited to the first interval. From this it can be inferred that it is a residential street, where most cars have left during the day.

The latter analysis can be generalized to enable a search for free parking spaces, which can be used for other usages, in a certain time interval. For this purpose, the cells' occupancy time-series have to be evaluated. It has to be checked whether the cells are free of observations during this interval. An example is illustrated by Figure 15. There, the result of a search for the time period between 12 am and 2 pm is visualized on the map. Similar to this, the search for places which offer maximum time periods can be determined accordingly.

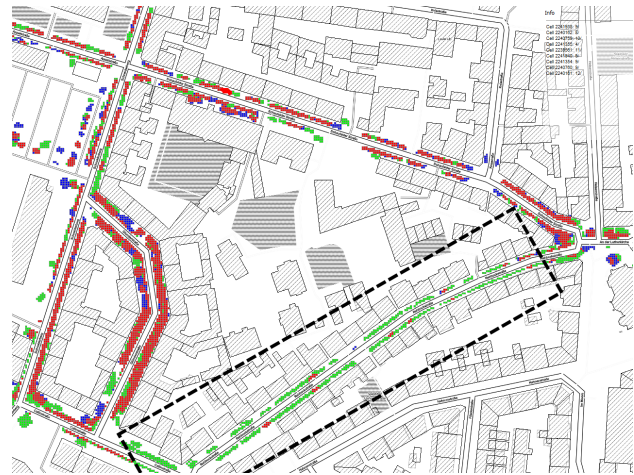


Figure 14. The usage before and after 12 am is encoded by different colors. Green: occupied before 12 am, blue occupied after 12 am, red occupied all the time. (basemap: Stamen Maps).



Figure 15. The availability of the space during the requested time period (12 am to 2 pm). Green cells encode free, red cells occupied space (basemap: Stamen Maps).

4. DISCUSSION

The **car extraction step** of our approach provided 15.000 unique car bounding boxes. This number made a sufficient analysis of the space usage possible. Nevertheless the manual validation of the results showed that 28 out of 100 cars have not been identified correctly and 7 objects has been misclassified. This error has been compensated by the large amount of data processed and analyzed and should not have influenced the analytical results. Nevertheless improvements in this step would allow analyzing smaller data sets and allow complementary assumption, if no vehicles are sighted. For example, if there is no vehicle presence at a location, we would be able to assume that there were actually no vehicles there. At the moment we can only assume that if several cars are observed at certain positions the false positive error is not decisive.

The reasons for the errors in identification of car instances can be suspected in single processing steps as follows: (1) The projection of the semantic segmentation from RGB images in to 3D point cloud allows making use of larger available reference data to combine with effective models tailored for RGB images. The results shown here suffer from projection errors. Some of

the errors introduced in this step errors cannot be corrected by the subsequent step. A possible strategy to tackle this problem is to switch the paradigm and to use 3D point cloud oriented models. (2) Segmentation improvement step is capable of reducing number of segmentation errors which are few in number and located within a homogeneous region. Obviously it is not sufficient as the larger erroneously classified areas have a higher potential to become source for instantiation failure. Although this homogeneous region based error filtering is not suitable for the instantiation of objects with small homogeneous areas or no homogeneous areas at all like pedestrians or bicycles. A possible alternative could be to apply a more elaborate framework for label transfer like (Kochanov et al., 2020) or (Boulch et al., 2018). (3) The ground removal approach used in this study is straightforward and simple yet effective for small scenes with no or little ground height variation. For a more robust behavior this ground removal step should be applied to tiles of limited size. (4) The used clustering approach is suitable for the earlier defined requirements of identifying on-street parking of unique cars. The drawback is the limited robustness regarding the variable density of point clouds. In this study this problem is not crucial due to its focus on on street parking within the homogeneously and densely scanned space. For future studies with other objects classes, alternative approaches like in (Hong et al., 2020) can be considered. Generally, this methodology can be similarly applied to analyze space usage by other traffic participants like pedestrians or cyclists. The prerequisite for this is the availability of corresponding semantically segmented data.

The **information aggregation step** discretizes the MPS into a grid and takes into account underlying environmental information. During that process the street areas are only approximated by a buffer operation to the OSM street center line. A more accurate way can be the integration of official cadastral map data provided by the municipalities.

Further, as this work is based on a limited amount of data, a temporally more dense and better distributed data acquisition will provide even more reliable results. Especially, the temporal analyses will significantly benefit by this.

5. CONCLUSION AND OUTLOOK

In this work we present an approach to determine parking space and its concurrent usage possibilities. The proposed method processes Mobile Mapping Data and consists of three consecutive steps. During the *car extraction step* a semantic segmentation is applied to the LiDAR data and RGB image data to extract the required car objects. The following *information aggregation step* generates an occupancy grid, which also stores information about the observation counts and times. In the last *Information analysis and visualization step* the occupancy grids are evaluated to obtain the results for different proposed analyses. Finally, this approach is evaluated by conducting an experiment, which uses a real data set with 5 acquisitions of the same area.

Although the used methods provide promising results, there are different open aspects to be tackled in future work. As we only focus on cars in this work, our approach can be extended to also consider other object types like pedestrians or bicycles, which share the available MPS. To this end, only the extraction process has to be adjusted to provide the footprints of those objects. The occupancy grid generation is already designed to handle different object classes.

In future work, we will also investigate and analyze possible measurement schemes in order to acquire the temporal information in a necessary resolution and quality. To this end, a sensitivity study should be performed.

Finally, analyses, which also take into account cell neighborhood relations, can be applied to the occupancy grid. In this way, it is possible to search for certain constellations by adding neighborhood constraints. For instance, in order to place mobile charging stations, one has to look usually free spaces, which are adjacent to popular and highly used parking spaces.

ACKNOWLEDGEMENTS

This work is funded by the Federal Ministry for Economic Affairs and Energy (BMWi), Germany under the project "d-E-mand" (grant ID 01ME19009B).

REFERENCES

- Bock, F., Eggert, D., Sester, M., 2015. On-street parking statistics using lidar mobile mapping. *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, IEEE, 2812–2818.
- Boulch, A., Guerry, J., Le Saux, B., Audebert, N., 2018. SnapNet: 3D point cloud semantic labeling with 2D deep segmentation networks. *Computers & Graphics*, 71, 189–198. <https://www.sciencedirect.com/science/article/pii/S0097849317301942>.
- Brenner, C., 2016. Scalable estimation of precision maps in a mapreduce framework. *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 1–10.
- Campello, R. J. G. B., Moulavi, D., Sander, J., 2013. Density-based clustering based on hierarchical density estimates. J. Pei, V. S. Tseng, L. Cao, H. Motoda, G. Xu (eds), *Advances in Knowledge Discovery and Data Mining*, Springer Berlin Heidelberg, Berlin, Heidelberg, 160–172.
- Chen, L.-C., Papandreou, G., Schroff, F., Adam, H., 2017. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- contributors, O., 2017. Planet dump retrieved from <https://planet.osm.org>.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B., 2016. The cityscapes dataset for semantic urban scene understanding. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- de Magalhaes, C., Carmona, M., 2009. Dimensions and models of contemporary public space management in England. *Journal of Environmental Planning and Management*, 52(1), 111–129. <https://doi.org/10.1080/09640560802504704>.
- Felzenszwalb, P. F., Huttenlocher, D. P., 2004. Efficient graph-based image segmentation. *International journal of computer vision*, 59(2), 167–181.
- Hong, F., Zhou, H., Zhu, X., Li, H., Liu, Z., 2020. LiDAR-based Panoptic Segmentation via Dynamic Shifting Network. *arXiv preprint arXiv:2011.11964*.

Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P., 2019. Panoptic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9404–9413.

Kochanov, D., Nejadasl, F. K., Booij, O., 2020. KPRNet: Improving projection-based LiDAR semantic segmentation. *arXiv preprint arXiv:2007.12668*.

Liu, S.-H., Yu, S.-Y., Wu, S.-C., Chen, H.-T., Liu, T.-L., 2020. Learning Gaussian Instance Segmentation in Point Clouds. *arXiv preprint arXiv:2007.09860*.

McInnes, L., Healy, J., 2017. Accelerated hierarchical density based clustering. *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, IEEE, 33–42.

Oranratmanee, R., Sachakul, V., 2014. Streets as public spaces in Southeast Asia: Case studies of Thai pedestrian streets. *Journal of Urban Design*, 19(2), 211–229. <https://doi.org/10.1080/13574809.2013.870465>.

Peters, T., Brenner, C., 2019. Automatic generation of large point cloud training datasets using label transfer. 39. *Wissenschaftlich-Technische Jahrestagung der DGPF e.V.*, 20. – 22. Februar 2019 in Wien.

Schlichting, A., Brenner, C., 2016. Generating a hazard map of dynamic objects using LiDAR Mobile Mapping. *Photogrammetric Engineering & Remote Sensing*, 82(12), 967–972.

Tao, A., Sapra, K., Catanzaro, B., 2020. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*.

Zhou, H., Zhu, X., Song, X., Ma, Y., Wang, Z., Li, H., Lin, D., 2020. Cylinder3d: An effective 3d framework for driving-scene lidar semantic segmentation. *arXiv preprint arXiv:2008.01550*.