Deriving precise orchard maps for unmanned ground vehicles from UAV images

Tjark Schütte[1,2], Volker Dworak[1] and Cornelia Weltzien[1,2]

Abstract: Mapping and environment representation are two of the main challenges in agricultural robotics and are vital to navigation tasks like localisation and path planning. In this work, we present a new method that enables the offline creation of orchard maps for unmanned ground vehicles based on unmanned aerial vehicle imagery. We employ photogrammetry to generate high-resolution 3D point clouds from aerial images. A cloth simulation filter is then used to classify ground and off-ground points. In order to obtain detailed probabilistic occupancy grid maps, per cell statistics are evaluated. First results show promising performance when compared to ground truth positions of orchard bushes and manual labelling.

Keywords: Precision Horticulture, photogrammetry, UAV, mapping, UGV navigation

1 Introduction

One of the main tasks in agricultural robotics and navigation of unmanned ground vehicles (UGV) is that of environment mapping. Generally, idiothetic and allothetic information collected by sensors on the robot is combined to generate a map of the robots environment [FM03]. This is often achieved by Simultaneous Localisation and Mapping (SLAM), combined with exploration algorithms, or by teaching the robot on the first visit by manually driving it through the new environment and generating a map. While fully autonomous SLAM depends on exploration strategies and is inefficient due to incomplete knowledge of the environment, teaching the robot is relatively labour-intensive and the results depend on the quality of the robot's localisation, just as in the automated case.

In order to improve the accuracy of the generated map, it is possible to generate it offline, in post-processing. This has the advantage that the whole set of data can be taken into account and a globally optimized map can be found by maximizing the posterior of the map over all given measurements and estimated robot poses. However, even if all available information is used, these maps may contain errors due to uncertainties in localisation and allothetic information. One way to tackle these challenges and reduce the effort of mapping larger orchards for agricultural robotics is to use remote sensing data that includes information on the complete area in a single measurement. As an example, Santos

[1] Leibniz-Institut für Agrartechnik und Bioökonomie e.V., Technik im Pflanzenbau, Max-Eyth-Allee 100, 14469 Potsdam-Bornim, tschuette@atb-potsdam.de; cweltzien@atb-potsdam.de; vdworak@atb-potsdam.de
[2] Technische Universität Berlin, FG Agromechatronik, Straße des 17. Juni 135, 10623 Berlin

et al. [Sa20] use satellite images to extract preliminary maps of vineyards in order to reduce the effort of mapping and help overcome local optima in the mapping process. This approach allows for 2D mapping and requires online refinement of the map. Furthermore, it uses either artificial neural networks or a support vector machine for crop row classification and segmentation. Both require manual labelling of images when training the classifiers for new data sets.

This paper presents a new method that enables the creation of maps from unmanned aerial vehicle (UAV) imagery by combining photogrammetry with probabilistic occupancy grid mapping. The latter is achieved by analysing and classifying point clouds using a cloth simulation filter for ground surface modelling.

2 Materials and Methods

2.1 Materials

To analyse the new method, a dataset was collected at the Fieldlab for Digital Agriculture of the Leibniz Institute for Agricultural Engineering and Bioeconomy. The dataset contains aerial images and GNSS position data of a blueberry field. The rows of blueberry bushes consist of more than 50 different blueberry varieties with smaller and larger gaps present in the rows. Fig. 1 shows an orthoimage of the orchard with ground truth positions of the blueberry bushes overlaid. The average height of the bushes is 127 cm (SD = 21cm). Mean and standard deviation of horizontal expansion are 110 cm and 23 cm, respectively. The row spacing between bush centres varies from 2.4 to 4.1 metres. In addition to this high heterogeneity of the orchard, due to an ongoing experiment investigating different mowing strategies, the grass was short (< 10cm) between the upper rows 1 to 6, while between rows 6 to 8 and below, the grass next to the bushes had a height of more than 40 cm.

For the aerial images, a consumer grade RGB-camera (α-6000, Sony, Tokyo, Japan) on a gimbal was attached to a quadcopter (HP-X4-E1200, Hexapilots, Germany). For ground truth position measurements of obstacles, in this case the blueberry bushes, as well as ground control points for georeferencing of the UAV data, an RTK-GNSS system (HIPer Pro, Topcon, Tokyo, Japan) was used. For data processing and analysis, Agisoft Metashape version 1.7 was used for point cloud creation. CloudCompare v 2.11.1, especially the cloth simulation filter (CSF) plugin [Zh16] was used for classification. QGIS Desktop 3.10 as well as python 3.8 with the packages rasterio version 1.2.10, numpy version 1.21.3 and scikit-learn version 1.0.1 were used for generation and analysis of the maps.

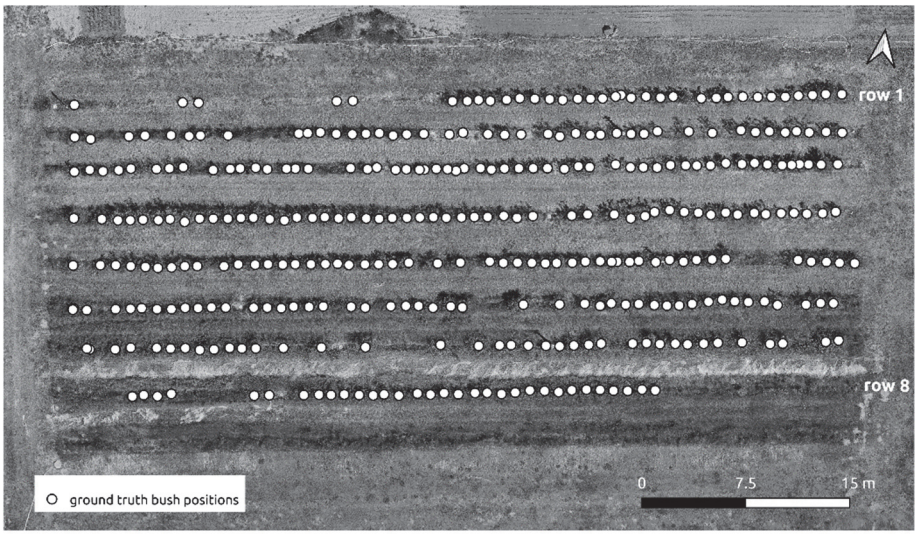

Fig. 1: Orthoimage of the analysed blueberry orchard, overlaid with ground truth positions (measured with an RTK GNSS systems) of the blueberry bushes (white markers)

2.2 Method

The presented method uses aerial images and RTK-GNSS positions of ground control points to first compute a georeferenced dense point cloud and then point cloud analysis to derive an occupation grid map from that data.

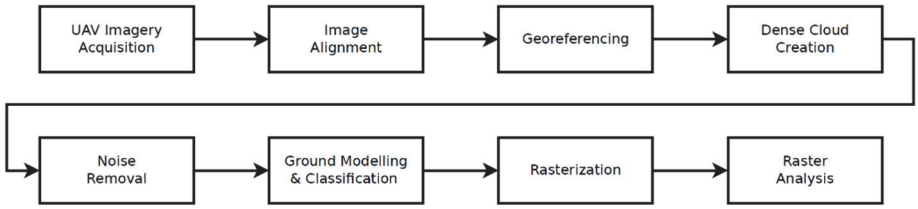

Fig. 2: High-level process of the map computation

Fig. 2 shows the high-level overview of the process. In the first step, the orchard is surveyed by the UAV system in a height of 16 m with 80 percent target overlap of the images. Three flight plans are flown to generate the images taken in nadir perspective and images taken in oblique view at a 45° angle to nadir, looking at the rows from two sides (North and South). This flight plan leads to a ground sampling distance (GSD) of 3.89 millimetres for the nadir images and 5.51 millimetres for the oblique view. The images are then aligned

using low accuracy in Metashape. After a first sparse point cloud has been created, ground control points are marked in all images they are visible in, in order to update the referencing of the model as well as to optimize the image alignment. The sparse cloud is then thinned based on reprojection error and reconstruction uncertainty, which are set to 0.2 and a level of 20 respectively. In the next step, the dense cloud is computed using ultra high quality settings and disabled depth filtering. This process takes several hours, even for smaller areas. After exporting, the dense point cloud is further processed using the software CloudCompare. In a first step, an initial ground model is computed, using the CSF plugin. The settings used are "Relief" and a grid spacing of 2.0 meters. Classification results are disregarded. Subsequently, signed distances from this model are computed for each point. All points that are below ground are removed, allowing for a tolerance of -0.2 meters. Points above a manually set threshold, which can be chosen as the height of the UGV, are also removed. In the data that is presented in this paper, a maximum height of 6 meters was used. In the next step, the CSF is applied again, this time using a finer grid spacing of 0.5 meters and a classification threshold for ground and off-ground points of 0.2 meters. The point cloud including the classification results is then rasterized using a cell size of 5 x 5 centimetres in x and y dimensions. For each raster cell, the expected value of the posterior occupation probability using a non-informative beta prior is computed based on the cell population and the classification results. The posterior can be calculated according to formula one, where $P(\theta|X)$ is the posterior occupation probability given the classification results X of the points in the raster cell. η is the normalizing constant, k the number of off-ground points in the cell and n is the total number of points in the cell. α and β are the parameters of the beta prior.

$$P(\theta|X) = \frac{1}{\eta} \text{Beta}(\alpha + k, \beta+n-k) \tag{1}$$

The prior reflects the binomial distribution of the average points per cell in the categories "off-ground" and "ground", weighted by the inverse average surface densities of the respective partial clouds. The latter is done to account for lower surface densities of bushes compared to flat ground. Using the beta prior reduces the problem of probabilities trending towards one or zero. The expected value of the posterior is finally calculated according to formula two, where p and q are the parameters of the posterior beta distribution:

$$E(X) = \frac{p}{p+q} = \frac{\alpha + k}{(\alpha + k)+ (\beta+n-k)} \tag{2}$$

3 Results

Fig. 3 shows a map that was generated using the previously described method. The ground truth positions of the bush centres are overlaid as white markers. It can be seen that for a large part the map seems to match the topography of the orchard well. However, in the lower rows, noise is present. Fig. 4 shows a detailed view of this part as well as of the upper right part of the map, again with GPS measurements overlaid.

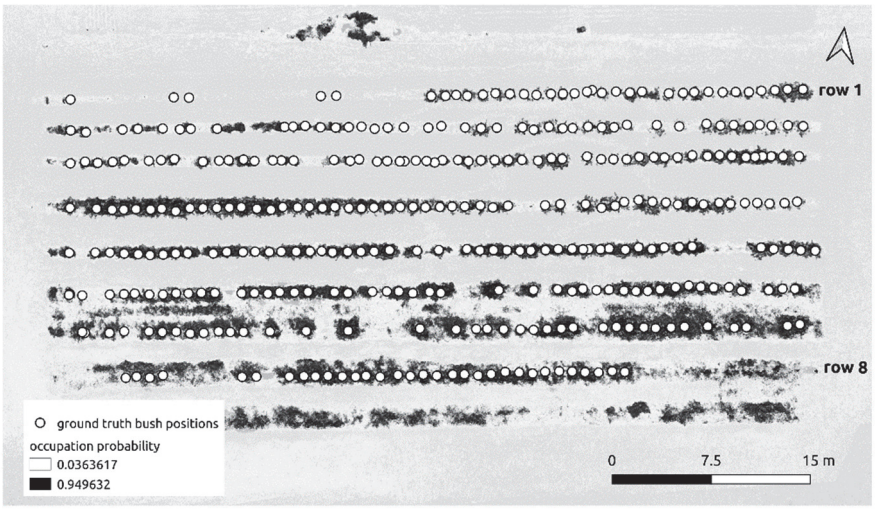

Fig. 3: Resulting map with overlaid RTK-GNSS positions of the blueberry bush centres. Depicted is the occupation probability (white=0 to black=1)

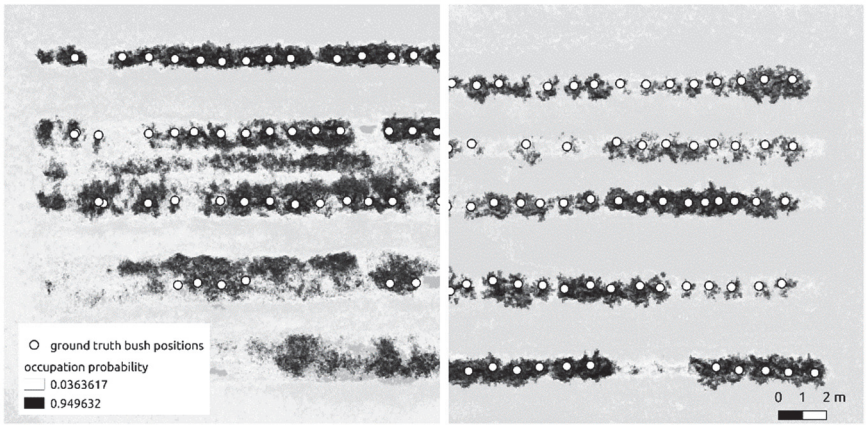

Fig. 4: Detailed view of the lower left (left image) and upper right (right image) part of the map. White markers: centre points of the blueberry bushes. In the left image, high grass that is classified as obstacle is visible between the rows and close to the bushes.

In order to analyse the map further, a dataset with ground truth labels was created, using an orthorectified image of the orchard, by selecting areas that could clearly be identified as driveable surface or obstacles/bushes in QGIS. The results were analysed in python. Using a threshold of 0.4013, the two classes are almost linearly separable, with precision

and recall both over 0.97. It is important to mention that especially at the edges of vegetation even with the high-resolution orthoimage (GSD=4.99), an accurate manual labelling of the data was not always possible. Therefore, only 87 percent of the map that could be identified as clearly belonging to either the ground or the obstacle class were labelled. This leaves out specifically hard-to-decide cases and can therefore be expected to have significantly increased these performance metrics.

4 Discussion

The approach of using point cloud data from UAV photogrammetry offers great potential for high-precision mapping of semi-structured outdoor areas. In particular, it shows potential for high-precision mapping of larger orchards with highly automated processes. While this work shows some of its drawbacks, mainly due to the classification by protrusion from the ground surface model, this also offers potential for increased generalisability due to the independence of object classes. To further evaluate the method, more datasets with more orchards need to be analysed. In addition to the presented results, 3D-mapping like OctoMaps [Ho13] and further plant analysis like growth height determination [Ho20] could be performed using the same data basis. Usage of airborne LiDAR devices on UAVs with high precision positioning could drastically reduce the effort of the general method since it will make the steps of photogrammetry and geo-referencing of the point cloud obsolete.

References

[FM03] Filliat, D.; Meyer, J.-A.: Map-based navigation in mobile robots. II. A review of map-learning and path-planning strategies. Cognitive Systems Research 4, pp. 283-317, 2003.

[Ho13] Hornung, A. et al.: OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees. Autonomous Robots, 2013.

[Ho20] Hobart, M. et al.: Growth Height Determination of Tree Walls for Precise Monitoring in Apple Fruit Production Using UAV Photogrammetry. Remote Sensing 10/12, p. 1656, 2020.

[Sa20] Santos, L. C. et al.: Occupancy Grid and Topological Maps Extraction from Satellite Images for Path Planning in Agricultural Robots. Robotics 4/9, p. 77, 2020.

[Zh16] Zhang, W. et al.: An Easy-to-Use Airborne LiDAR Data Filtering Method Based on Cloth Simulation. Remote Sensing 6/8, p. 501, 2016.