GPS-free Localisation and Navigation of an Unmanned Ground Vehicle for Yield Forecasting in a Vineyard

SAM MARDEN, MARK WHITTY

UNSW Australia
School of Mechanical and Manufacturing Engineering

Abstract. Yield forecasting is an important practice for vineyards to ensure that seasonal yield targets are met. Currently this is performed by manually collecting samples from around the vineyard that are measured and weighed. There has been increased research into automating this process, particularly through computer vision techniques.
This paper presents a localisation and navigation system for an autonomous ground vehicle that is capable of traversing a vineyard and acquiring geo-referenced images for use in yield forecasting, without relying on high- precision GPS. Localisation and navigation is achieved using laser scanners to detect the rows in the vineyard, which are used as features in line-based SLAM.
Successful autonomous navigation of a block, including travelling along the length of the rows and transitioning between rows, in combination with geo-referencing images acquired during this time, has been demonstrated using this approach.

1 Introduction

Major growers and purchasers of grapes are stipulating that particular yield targets should be met, in the belief that this will improve and maintain wine quality. Consequently, there is a strong demand for improved systems to forecast yield.

The current industry standard for yield forecasting is to manually sample and weigh bunches taken from various sections of the vineyard, and use this information to infer the yield of the entire vineyard [6]. Due to the inefficiency and inaccuracy of these approaches[4], there has been increased research into methods of automating this yield estimation.

Much of this work has focused on using computer vision techniques [12][14], and these approaches have demonstrated significant improvements over the current industry practices in terms of their ability to forecast the yield of a vineyard. However, there still exists the problem of how to efficiently acquire the images necessary for this processing, since manually collecting photos of grapes from across an entire vineyard is not feasible. Research has shown [21] that while vineyard managers strive for a forecasting accuracy of 5%, many have to tolerate errors as large as 10% in order to cover multiple vineyard blocks while keeping costs within a prescribed budget.

Autonomous systems capable of navigating a vineyard or orchard have been demonstrated using high-precision RTK-GPS [17][18] that are capable of operating safely in agricultural environments. The lack of GPS availability due to environmental conditions such as large canopies, the need for prior surveying of the area, and unreliable connectivity in certain scenarios (see Figure 5) make GPS-free approaches desirable.

Laser scanners have been used in agricultural robotics in order to create 3D maps of plants, using GPS for localisation [13][20][22]. Localisation using only lasers has been performed using artificial reflective markers [1][11] to aid in navigation, as well as in combination with GPS [9].

Recently, localisation of a robot in a maize field has been performed using only 2D LiDAR [10], however in this case there were specific requirements on the height and overhang of leaves that

Fig. 1: *Robot autonomously travelling between the rows of the vineyard.*

are observed. Additionally, the localisation was not performed simultaneously with autonomous navigation and row following.

Row-following has been performed by autonomous agricultural robots by using 2D LiDAR to detect the rows in an orchard setting [8][19][2], however in these instances global localisation has not been performed.

Table 1 below compares these existing approaches to agricultural localisation and navigation.

Reference	Localisation	Mapping	Navigation
[17][18]	GPS	No	Yes
[13][20][22]	GPS	Yes	No
[1][11]	LiDAR with reflective landmarks	Yes	No
[9]	LiDAR with line features	Yes	No
[8][19][2]	Dead reckoning	No	Yes
[16]	GPS	Yes	Yes

Table 1: *Comparison of existing approaches to localisation and navigation in a vineyard.*

The clear limitation of existing approaches is that they do not provide simultaneous localisation, mapping, and navigation without relying on GPS, which are necessary for the task performed in this paper. In certain applications, such as slashing and spraying of crops, it is not necessary for the autonomous vehicles to be able to build a map of their environment and localise themselves within it, and in others this problem has been solved through the use of high-precision GPS.

In the context of creating a map of geo-referenced images taken by the autonomous vehicle, it is necessary that the robot is capable of creating an accurate map of the vineyard, in addition to being able to localise itself within this map. In terms of navigation, the robot must be able to stay close to the centre of the row. From this position, it is able to take pictures of both sides of the row that it is currently in, and it will also avoid colliding with the rows themselves. It is also desirable that this approach does not rely on high-accuracy GPS [2]. In order to be able to determine exactly which vine the robot is capturing images of at a given time, the localisation should be accurate to within 1 metre of ground truth parallel to the rows of the vineyard, and 2.5 metres perpendicular to the rows.

This paper presents a system implemented on an autonomous ground vehicle that acquires images referenced to their location within the vineyard, which is necessary in order to accurately determine the distribution of yield. This is performed without relying on GPS.

The remainder of the paper is structured as follows: Section 2 provides a brief overview of the system itself, Section 3 presents the method of localisation using line-based Simultaneous Localisation and Mapping (SLAM), Section 4 presents the approach to navigation used, and Section 5 presents the results of the performance of the system in simulation and in a vineyard.

2 System Overview

The system presented in this work autonomously traverses the vineyard acquiring images as it travels. In order to achieve this, the robot performs two primary tasks; localisation and mapping, which will provide accurate references for the locations at which the images are taken, and navigation, will allow the robot to safely traverse the vineyard and ensure that photos are captured from the correct positions. A brief overview of the approaches to these two tasks is discussed below, with the following sections going into further detail regarding how these are achieved.

- Localisation and mapping: localisation and mapping are performed using lines that are extracted from a 2D LiDAR mounted at the front of the vehicle.
- Navigation: navigation is achieved by using the currently estimated rows of the vineyard in order to ensure the safe and precise control of the robot. At any given time, the robot aims to navigate along the centre of the row that it is currently in. Additionally, the robot uses the end points of the extracted rows to navigate around the ends of the rows in order to transition between them.

3 Localisation and Mapping

3.1 Line-Based SLAM

Localisation and mapping is performed in this work using the rows of the vineyard as features in an EKF-SLAM framework. The use of lines as features in an EKF-SLAM framework has been shown previously to provide highly accurate localisation in structured environments, even in the presence of poor quality sensors [3].

In the case of the robot travelling through the vineyard, we have to consider the pose of the robot in three, rather than two, dimensions, due to the fact that the terrain in the vineyard varies considerably in height and slope. The rows of the vineyard, which are used as features in line-based SLAM, are assumed to be perfectly straight. This results in the following parametrisation of the robot's state:

$$\left(\underbrace{x \quad y \quad z}_{\text{robot translation}} \quad \underbrace{q_w \quad q_x \quad q_y \quad q_z}_{\text{robot rotation}} \quad \underbrace{\rho_1 \quad \theta_i \quad \ldots \quad \rho_N \quad \theta_N}_{\text{line features}} \right)$$

The process model is implemented using the linear velocity estimated from the motor encoders and the angular velocities estimated from the gyroscope, while observations are performed using the lines that are extracted from the laser scans, with the observed lines being parametrised in polar coordinates.

The measured lines are those that are extracted from the current accumulation of laser scans, using the method discussed below. The expected measurements of the line features are the lines currently in the state vector parametrised in the robot-local frame. This conversion from global to local coordinates is performed as follows:

$$\begin{pmatrix} \rho_{\text{local}} \\ \theta_{\text{local}} \end{pmatrix} = \begin{pmatrix} \rho_{\text{global}} - x_{\text{robot}} \cos \theta_{\text{global}} - y_{\text{robot}} \sin \theta_{\text{global}} \\ \theta_{\text{global}} \end{pmatrix}$$

3.2 Line Extraction

RANSAC [7] is used to extract lines from the laser scans obtained from the 2D LiDAR. Once the lines are extracted, they are parametrised in polar coordinates r, θ, and are assumed to be infinitely long.

Laser scans are accumulated into 3D point clouds over a period of several seconds due to the fact that single scans are individually of very poor quality, which is caused by a number of factors, including uneven terrain, overhanging branches, and long grass. Note that when the lines are extracted, the point clouds are projected to 2D. This 2D projection of the accumulated laser scans is shown in Figure 2 below.

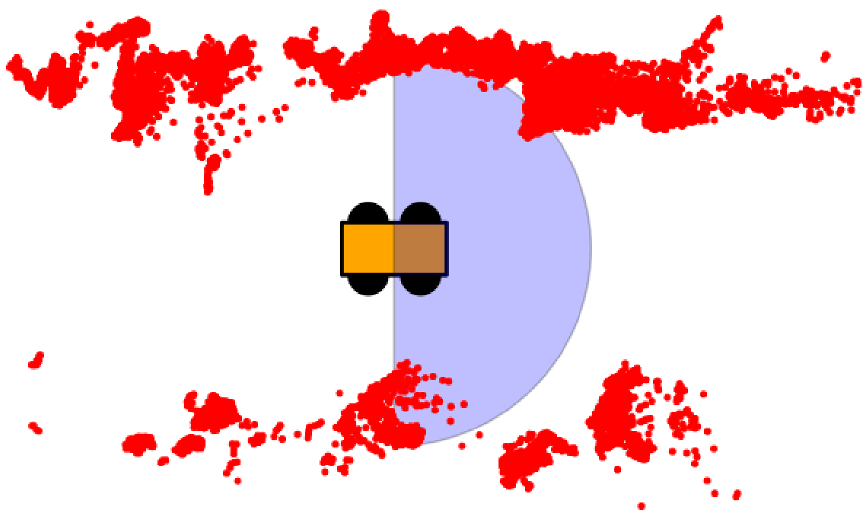

Fig. 2: *A top-down view of the robot located within the rows of the vineyard, in addition to the accumulated laser scans, illustrating the challenging operating conditions brought about by the unstructured environment.*

Note the readings from the laser that are greater than a certain distance away from the robot are ignored in order to prevent issues where the ground is observed.

3.3 Data Association and Feature Management

In a similar approach to that of [3], lines that are extracted from the accumulated laser scans are associated to existing lines according to the following metric:

$$d_{\text{line}} = \sqrt{w_\rho \Delta\rho^2 + w_\theta \Delta\theta^2}$$

Here $\Delta\rho$ and $\Delta\theta$ are the differences, in local coordinates, between the observed and expected radii and angles of the lines respectively, while w_ρ and w_θ are the weights for the radii and angles. Provided that this metric is below a certain threshold, the two lines are considered to have been matched, whilst lines that are extracted erroneously are rejected.

The feature management policy detailed in Table 2 below outlines the approach used in this work to manage existing features and handle observations of new features.

Event	Action
New line observed	Add line to candidate list, set score to S_{new}
Candidate line observed	Increment candidate score
Candidate line not observed	Decrement candidate score
Candidate score above S_{max}	Add line to EKF state
Candidate score below S_{min}	Remove line from candidate list

Table 2: *Feature management policy.*

Here S_{new}, S_{max} and S_{min} are adjustable thresholds used to control how quickly lines are removed from the candidate list. This approach, although being relatively simplistic, is sufficiently robust that it is able to account for erroneous detections of lines. It is critical that these lines are estimated and managed accurately and reliably, as they are used to guide the robot in its navigation. Note the erroneous observations do not occur so frequently that memory becomes a concern.

4 Navigation

4.1 Row Following

Two existing approaches have been combined to produce the approach to navigation presented in this work. The centre lines of the currently extracted row are used as a target lines for the vehicle to follow [8], with the control law being adapted from an Ackermann-steered tractor [5] in order to accommodate the differentially-driven vehicle here.

Navigation is achieved using the control law discussed below to follow the centre line of the currently estimated rows of the vineyard. We assume that the robot is driving at a constant linear velocity, and as such the control input is parametrised in terms of the angular velocity of the robot, the control law for which is given by:

$$\omega = v \cos^3(\theta)\left(-K_d \tan(\theta) - K_p x\right)$$

Here ω and v are the linear and angular velocities of the robot respectively, with v assumed to be constant. K_p and K_d are the control constants. The linear (x) and angular (θ) errors are calculated based on the current estimate of the robot's pose and the lines of the vineyard rows.

4.2 Row Transition

In order to cover the area of an entire block in the vineyard, the robot must be able to not only travel along the lengths of the rows, but also transition between rows. This task of exiting a row and reliably entering the next one has recently been identified as the most challenging aspect of navigating in a vineyard-like environment [8].

In this work, the robot continuously attempts to navigate around the end of the currently observed rows, as estimated in the EKF's state. However, since the robot is going to be continuously observing the rows, it will not actually transition into the next row until it is has reached the end of the current row, as illustrated in Figure 3 on the following page.

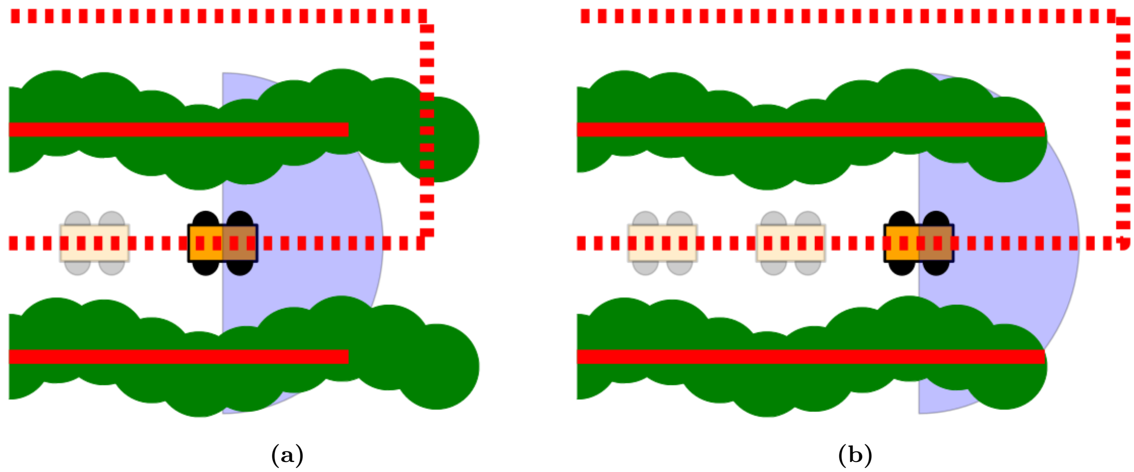

Fig. 3: *The proposed approach to row transitioning utilises the current estimates of the endpoints of the rows of the vineyard to provide reliable and safe navigation behaviour. While the robot is inside the rows of the vineyard (A), it continuously attempts to drive around the currently observed endpoints of the row. Since it will continue to see more of this row, it will not actually transition to the next row until it reaches the end of the current one (B).*

5 Experiments and Results

The system has been tested extensively in a variety of simulated environments, in addition to an actual vineyard.

The platform used for testing the real-world performance of the navigation is shown in Figure 1, and has been equipped with several sensors, including:

– SICK LMS 111: LiDAR acquiring laser scans at a frequency of 50Hz. The laser has been mounted high on the robot, approximately one metre off the ground.
– Microstrain IMU: includes an accelerometer, gyroscope, and magnetometer.
– Encoders: both of the motors driving two motors on either side of the Husky are equipped encoders.
– GPS: a high-precision Trimble GPS has been mounted on the Husky in order to provide a ground truth reference.
– Cameras: the robot has been equipped with a Nokia D80 DSLR.

The software has been implemented using ROS as the framework for inter-process communication [15], and is capable of operating in real-time.

The system has been tested in the Jarrett's of Orange vineyard, and has successfully demonstrated the ability to travelling the length of a block in the vineyard and transition between rows, as well as being able to start navigating from any point in the block, whether inside or outside the rows. Figure 1 illustrates the robot autonomously travelling along the length of the block.

5.1 Localisation and Mapping

The system has been tested extensively in simulation under a variety of conditions including varying noise in odometry, laser scans, and quality of the observable rows. The following table presents an evaluation of the performance of the algorithm with changes in these conditions. Figure 4 below illustrates the performance of the system in a simulated vineyard.

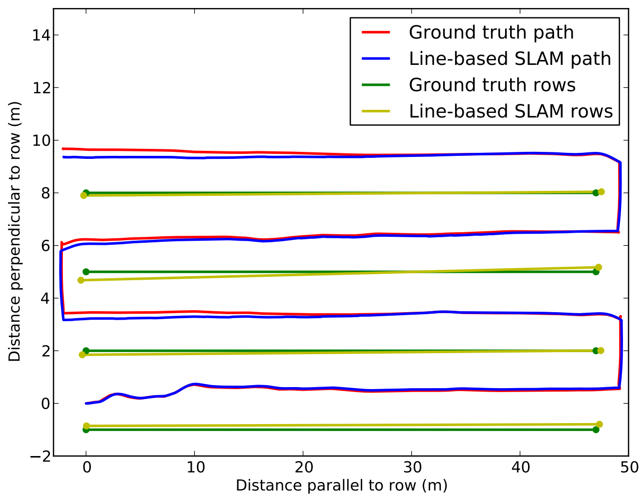

Fig. 4: *Comparison between the line-based SLAM and ground truth paths and rows in simulation.*

The RMS errors of the estimated path compared with the ground truth reference and the estimated map compared with the known map (quantified as the distance between the endpoints of the rows) are shown in Table 3 below.

Noise in observations of vineyard rows	Low	Medium	High
Localisation parallel to row RMSE (metres)	0.04	0.04	0.04
Localisation perpendicular to row RMSE (metres)	0.20	0.49	0.44
Mapping RMSE (metres)	0.39	0.70	0.79

Table 3: *Analysis of the performance of localisation and mapping in different operating conditions.*

In this testing, it was found that the most significant factor in determining the performance of the localisation and mapping was the quality of the rows of the vineyard. As aforementioned, the required level of accuracy for the localisation of the robot is 2.5 metres parallel to the rows, and 1 metre perpendicular to the rows. These results clearly demonstrate the achievement of the target levels of accuracy for localisation and mapping using the approach presented in this paper. Nevertheless, we note that the localisation accuracy slowly drifts over time, and further work is being conducted into the use of loop closure to solve this issue.

The system has also been tested extensively in an actual vineyard in order to further demonstrate the robustness of the vehicle's performance and to verify that the system performs as robustly in the real-world as it does in simulation. Figure 5 below illustrates the path taken by the robot as it autonomously navigates the length of a block and transitions to the next row, as compared with the high-precision RTK-GPS. In addition to this, we also present a comparison between the lines extracted and filtered by the robot and lines extracted from aerial LiDAR acquired at the vineyard.

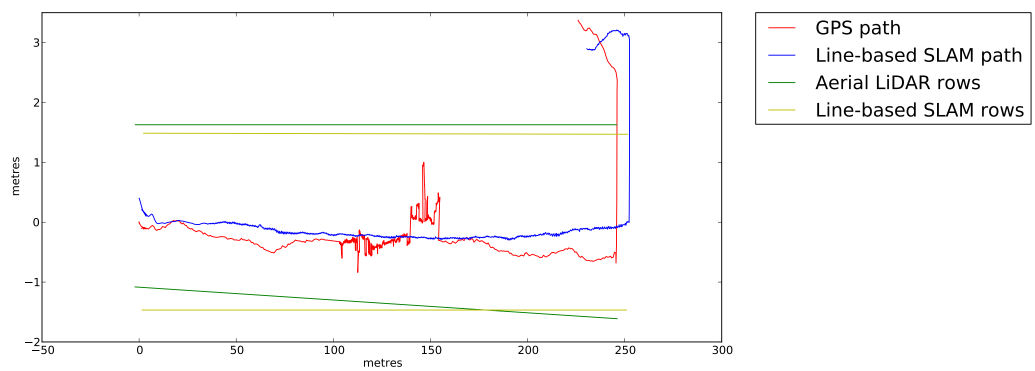

Fig. 5: *Comparison between the line-based SLAM and GPS paths in a vineyard. This clearly illustrates the shortcomings of the use of GPS in such an application, as we can clearly see that the the quality of GPS readings deteriorates for a substantial period of time, which is most likely due to loss of contact with the base station.*

The estimation of the rows of the vineyard is shown here to be accurate in terms of both the length and the spacing of the vines, and we can also observe that the localisation is more accurate when the line-based SLAM is used. It is clear that even the high-precision GPS is inaccurate enough that the use of GPS alone would not be sufficient.

5.2 Navigation

The RMS errors of the linear and angular offsets of the robot from the centre of the row as it traversed the length of the block were 0.03 metres and 1.35 degrees respectively. These low errors clearly demonstrate the high performance of the system is accurately navigating along the centre lines of the rows of the vineyard.

The use of global localisation and mapping in facilitating navigation is a key component of the system. These results demonstrate that even in situations where, in general, individual observations are poor, the combination of the information that is fused over time in the EKF allows the robot to accurately navigate along the lengths of the rows of the vineyard, and to safely transition between them.

6 Conclusion and Future Work

This paper has presented autonomous localisation and navigation in a vineyard for the purpose of geo-referencing images, which uses lines extracted from 2D LiDAR to avoid relying on GPS. The system has been demonstrated to be able to autonomously traverse and transition between the rows of a vineyard with a sufficient level of accuracy that it is possible to determine which vine the vehicle is currently taking pictures of.

The images acquired during this process will be used for automated yield forecasting. The ability to efficiently and autonomously obtain these images will provide vineyard operators with significantly more information regarding the distribution of yield within their vineyards, thereby allowing them to make more informed decisions regarding the management and maintenance of their vineyards.

7 Acknowledgements

The authors would like to acknowledge Justin and Pip Jarrett for providing the use of their vineyard for testing the system. The authors would also like to thank Jose Guivant for the use of the Husky base platform.

References

1. Ahamed T, Kulmutiwat S, Thanpattranon P, Tunti S, Ryozo N, Takigawa T. Monitoring of plant growth using laser range finder. *2011 ASABE Annual International Meeting, Louisville, Kentucky*, 2011.
2. Oscar C Barawid Jr, Akira Mizushima, Kazunobu Ishii, and Noboru Noguchi. Development of an autonomous navigation system using a two-dimensional laser scanner in an orchard application. *Biosystems Engineering*, 96(2):139–149, 2007.
3. Young-Ho Choi, Tae-Kyeong Lee, and Se-Young Oh. A line feature based slam with low grade range sensors using geometric constraints and active exploration for mobile robot. *Autonomous Robots*, 24(1):13–27, 2008.
4. Peter R Clingeleffer, Steve Martin, Mark Krstic, and Gregory M Dunn. *Crop Development, Crop Estimation and Crop Control to Secure Quality and Production of Major Wine Grape Varieties: A National Approach: Final Report to Grape and Wine Research & Development Corporation*. Grape and Wine Research & Development Corporation, 2001.
5. L Cordesses, P Martinet, B Thuilot, and M Berducat. GPS-based control of a land vehicle. In *Proc. IAARC/IFAC/IEEE Int. Symp. Automation Robotics Construction (ISARC)*, pages 22–24, 1999.
6. Gregory M Dunn and Stephen R Martin. The current status of crop forecasting in the Australian wine industry. In *ASVO Seminar Series: Grapegrowing at the Edge*, pages 4–8, 2003.
7. Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
8. Bradley Hamner, Sanjiv Singh, Marcel Bergerman, et al. Improving orchard efficiency with autonomous utility vehicles. In *American Society of Agricultural and Biological Engineering Annual Convention Proceedings Abstract and Paper*, 2010.
9. Søren Hansen, Enis Bayramoglu, Jens Christian Andersen, Ole Ravn, Nils Andersen, and Niels Kjølstad Poulsen. Orchard navigation using derivative free Kalman filtering. In *American Control Conference (ACC), 2011*, pages 4679–4684. IEEE, 2011.
10. Santosh A Hiremath, Gerie WAM Van Der Heijden, Frits K Van Evert, Alfred Stein, and Cajo JF Ter Braak. Laser range finder model for autonomous navigation of a robot in a maize field using a particle filter. *Computers and Electronics in Agriculture*, 100:41–50, 2014.
11. Libby J, Kantor G. Deployment of a point and line feature localization system for an outdoor agricultural vehicle. *2011 IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China*, pages 1565–1570, May 2011.

12. Scarlett Liu, Samuel Marden, and Mark Whitty. Towards automated yield estimation in viticulture. *Proceedings of Australasian Conference on Robotics and Automation (ACRA), Sydney, Australia*, 2-4 Dec 2013.
13. Fadhillah Norzahari, Kate Fairlie, Adrian White, Mitchell Leach, Mark Whitty, Stephen Cossell, Jose Guivant, and Jayantha Katupitiya. Spatially Smart Wine: Testing Geospatial Technologies for Sustainable Wine Production. *Civil Engineering Surveyor*, pages 12–18, 2011.
14. Stephen T. Nuske, Supreeth Achar, Terry Bates, Srinivasa G Narasimhan, and Sanjiv Singh. Yield estimation in vineyards by visual grape detection. In *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '11)*, September 2011.
15. Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. ROS: an open-source robot operating system. In *ICRA workshop on open source software, Kobe, Japan*, volume 3, page 5, 2009.
16. A Ruckelshausen, P Biber, M Dorna, H Gremmes, R Klose, A Linz, F Rahe, R Resch, M Thiel, D Trautz, et al. Bonirob–an autonomous field robot platform for individual plant phenotyping. *Precision agriculture*, 9:841, 2009.
17. Albert Stoll and Heinz Dieter Kutzbach. Guidance of a forage harvester with GPS. *Precision Agriculture*, 2(3):281–291, 2000.
18. B Thuilot, C Cariou, L Cordesses, and P Martinet. Automatic guidance of a farm tractor along curved paths, using a unique CP-DGPS. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 2, pages 674–679. IEEE, 2001.
19. R Tsubota, N Noguchi, and AKIRA Mizushima. Automatic guidance with a laser scanner for a robot tractor in an orchard. In *Proceedings of the automation technology for off-road equipment conference, Kyoto, Japan*, 2004.
20. Ulrich Weiss and Peter Biber. Plant detection and mapping for agricultural robots using a 3d lidar sensor. *Robotics and autonomous systems*, 59(5):265–273, 2011.
21. JA Wolpert and EP Vilas. Estimating vineyard yields: Introduction to a simple, two-step method. *American Journal of Enology and Viticulture*, 43(4):384–388, 1992.
22. Yokota M, Mizushima A, Ishii K, Noguchi N. 3-D GIS map generation using a robot tractor with a laser scanner. *2005 ASAE Annual International Meeting, Tampa, Florida.*