ISBN 978-82-326-3290-9 (printed version)
ISBN 978-82-326-3291-6 (electronic version)
ISSN 1503-8181

Doctoral Thesis

Trygve Utstumo

Doctoral theses at NTNU, 2018:250

Trygve Utstumo

Asterix

Robotic weed control in row crops

NTNU
Norwegian University of
Science and Technology
Faculty of Information Technology
and Electrical Engineering
Department of Engineering Cybernetics

NTNU
Norwegian University of
Science and Technology

Doctoral theses at NTNU, 2018:250

NTNU

NTNU
Norwegian University of
Science and Technology

Trygve Utstumo

Asterix

Robotic weed control in row crops

Thesis for the degree of Philosophiae Doctor

Trondheim, August 2018

Norwegian University of Science and Technology
Faculty of Information Technology
and Electrical Engineering
Department of Engineering Cybernetics

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology
and Electrical Engineering
Department of Engineering Cybernetics

ISBN 978-82-326-3290-9 (printed version)
ISBN 978-82-326-3291-6 (electronic version)
ISSN 1503-8181

Doctoral theses at NTNU, 2018:250

Printed by Skipnes Kommunikasjon as

Vitam navigemus amantes

Abstract

Vegetables and other row-crops represent a large share of the agricultural production. There is a large variation in crop species, and a limited availability in specialized herbicides. The robot presented here utilizes the systematic cultivation techniques of row crops to navigate and operate in the field. By the use of machine vision it separates seeded vegetable crops from weed, and treat each weed within the row with individual herbicide droplets, without affecting the crop. This Drop-on-Demand (DoD) method allow the use of non-selective herbicides and significant reductions in herbicide use.

This thesis presents six research papers concerning the development of the DoD system and the mobile robot. The robot is tailored to its purpose with cost, maintainability, efficient operation and robustness in mind. The three-wheeled design is unconventional, and the design maintains maneuverability and stability with the benefit of reduced weight, complexity and cost.

Topics within localization and navigation for agricultural robotics have been explored. Quaternion estimation by an Extended Kalman Filter and a Non-linear complimentary filter has been implemented on an ARM Cortex M3 microcontroller. A Bayesian framework for fusing delayed Visual Odometry measurements has been explored in simulations. A Non-linear Model Predictive Controller (NMPC) has been developed and explored in simulation to enable a controller guaranteed to not sway its wheels into the crop row and subsequently damage the crop. The framework is also suitable for implementing other constrains for operation in other environments, such as greenhouses or confined spaces. Path following by an adaptive controller and a Model Reference Adaptive Controller (MRAC) has been implemented and compared in indoor trials.

The DoD system for herbicide application has been developed within and in connection with this project. The influence of liquid properties viscosity and surface tension on the formation and stability of droplets has been tested in lab trials. A control circuit for synchronized control of solenoid valves was developed and tested.

Indoor pot trials with four weed species demonstrated that the Drop-on-Demand system (DoD) could control the weeds with as little as 7.6 µg glyphosate or 0.15 µg iodosulfuron per plant. The results also highlight the importance of liquid characteristics for leaf retention, as the common herbicide glyphosate had no effect unless mixed with suitable additives. The trials document the DoD effect on weed species not previously described in literature, and with an alternative herbicide to glyphosate, iodosulfuron. A field trial with the robot was performed in a carrot field, and all the weeds were effectively controlled with the DoD system.

The robot and DoD system represent an important contribution to the range of systems presented witin Precision Agriculture for in-row weed control - a movement which as a whole represent a paradigm shift to the environmental impact and health risks of weed control, while providing valuable new tools to the producers.

Contents

List of figures

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree of Philosophiae Doctor (PhD) at the Norwegian University of Science and Technology (NTNU).

My supervisors have been Professor Jan Tommy Gravdahl with the Department of Engineering Cybernetics, Norwegian University of Science and Technology, and Dr. Therese W. Berge with the Norwegian Institute of Bioeconomy Research (NIBIO). During the candidate period I have been an employee with Adigo AS, where I have performed this work as part of the Industrial PhD project "Asterix - Automatic weed detection and control in row crops", supported by the Research Council of Norway (ES218701). My internal supervisors with Adigo were Øyvind Overskeid and Steve Goldberg.

The pot trials and field trials described in this thesis have been a part of the project "Innovation for better weed control in vegetables" in cooperation with NIBIO and the Norwegian Extension Service (NLR), supported by the Research Council of Norway (234226). The project and development of the robotic system is also supported by Innovation Norway through "Miløteknologiordningen" (254238), The European Council - Horizon 2020 SME Instrument (736354) and "Grofondet" (180017). Drop-on-Demand weed control was initially explored in a predecessor to this project: "Changed cultivation methods for rutabaga - novel methods and strategies for cost efficient weed control" by NLR, NIBIO and Adigo, supported by the Research Council of Norway, (192311).

Acknowledgments

The work presented here would not have been possible without the cooperation and guidance of the cooperating vegetable producers who have hosted trials, and guided us on the development of the system. We extend our gratitude to all who have hosted us and helped us with data-collection, and we would like to extend our gratitude to Jon Frogner, Sverre Huseby, Roy Hasle and especially Jens Edvard Kaase.

Gerd Guren and Kari Aarekol with NLR have been important sources of knowledge and have always been ready to answer questions. Dr. Jan Netland and Kjell Wærnhus with NIBIO have been instrumental to the pot trials, and I am deeply thankful towards Dr. Therese W. Berge, who has been my advisor through this work. She has had an unmatched eye for detail and has shared her deep experience of statistical and experimental methods.

I have had the pleasure of involving six master students into this project. It may have been what I have enjoyed the most, as we have developed a companionship and explored the topics with the curiosity only the thick brick walls of an University allows us. I am proud to have guided and co-supervised Alexander Klungerbo, Torgrim Aalvik Lien, Frode Urdal, Øystein Grændsen, Jarle Dørum and Mathias Hauan Arbo through their work with their Masters thesis at NTNU. They have all made significant contributions to the project, and to the academic work within this thesis. Three of the papers presented here have been co-authored with one of them, and I have had the pleasure of working with Prof. Edmund Førland Brekke and Prof. Simen Ellingsen.

As an Industrial PhD candidate, I have had the chance to experience the best of two worlds. Adigo AS and its founders Anders Aker and Øyvind Overskeid have consistently ventured into new technologies through 22 years - and the Asterix project is so far the most ambitious project. They have been inspirational, and they have provide me with ample challenge within Adigo. Anders Brevik has been the Industrial Design group lead through these years, and my "partner in crime" and very good friend as I have led the Robotics group.

Steve Goldberg recruited me to Adigo in 2009 and has been my mentor and strong friend through the PhD project. He has been an inspiration, and he has strained himself to provide me with opportunity. I am very thankful to him for introducing me to Prof. Stergios Roumeliotis who hosted and tutored me at University of Minnesota in 2015, and Dr. Joel Hesch who hosted me at Google the summer of 2015. My mathematical and theoretical abilities took major leaps during this time as I tried to keep up with the pace of their research groups.

My advisor, Prof. Jan Tommy Gravdahl has the ability to delve deeply into a wide array of technical fields. He has been a solid and consistent support, which has been vital for me to bridge the divide between the industrial and academic worlds. I will strive for the opportunities to work more together.

Friendships have grown both at NTNU, UMN and Adigo which will outlast this thesis. I thank you, my family, colleagues and friends for the support to complete this journey. And not the least, I thank my wife Camilla.

Chapter 1

Introduction

The Ph.D. project has been part of a larger effort from Adigo AS to develop an autonomous robot for weed control in row crops, Figure 1.1. This involves research and development in several technical fields, where this thesis work has involved several of them.

Figure 1.1: The 2017 Asterix robot prototype in field trials in Central Norway.

1.1 Thesis Outline

Chapter 2 include conference papers and journal papers written and co-authored through the Ph.D. studies. The remaining parts of this thesis are intended to complement the papers and provide context, while avoiding repetition of information.

This introduction will describe the background for the project and its motivation, and the context and relationship between each paper is described in Section 1.4. The conclusion in Chapter 3 will tie together our findings and describe our perspective on the path forwards for Drop on Demand (DoD) herbicide application.

1.2 Background

The production of row crops represent a significant portion of the overall food production in the world. This production is composed of large variety of crops of which each individual crop has a smaller volume. In contrast to major crops such as corn, soy and cereal, the vegetable crops have a smaller selection of available herbicides. The research investments required to develop a new herbicide are significant, and with the smaller acreage of each vegetable crop, it is not likely that we will see the introduction of new herbicides to the market that will significantly improve the situation for vegetable producers, Gast (2008). In the past 20 years we have seen a significant increase in herbicide resistant weeds (Heap, 2014), while the availability of herbicides has been reduced by regulations due to health and environmental concern. The end result is an increasingly challenging situation for farmers who are left with fewer efficient herbicides.

Weed control is one of the most important factors in all agricultural production. Weeds compete with crop plants for moisture, nutrients and sunlight and will have a significant negative impact on yield without sufficient weed control. Typical weed control methods for row crops include a combination of pre-emergence herbicide application, pre-emergence tillage, mechanical row harrowing and post-emergence herbicide application - if a selective herbicide or crop resistance is available, (Fennimore et al., 2016; Slaughter et al., 2008).

In 2008, the European Commission withdrew the approval for several herbicides, among them herbicides with propachlor as the active ingredient (European Commission, 2008). The herbicide was a health risk and had been documented contaminating ground water and harmful to aquatic life. The consequence to farmers of some cabbages and rutabaga was that they lost access to their most effective herbicide. In Norway, this spurred a joint project with farmers, the Norwegian Extension Service, the Norwegian Institute of Bioeconomy Research and Adigo in the search for alternative weed control methods.

The project explored an array of possible technologies for weed control, including microwaves, lasers and DoD applications. The delivery of herbicide through individual controlled droplets was considered the most suitable, and the technologies was demonstrated in an early field trial, (Guren, 2010).

The weed that occur in between rows, inter-row weeds, can be controlled by row-harrowing, flaming or shielded spraying. Whereas the in-row weeds pose a greater challenge for the farmers. In lack of selective post-emergence herbicides

they are left with few other options than manual in-row hoeing by hand, which is much more expensive than conventional spraying.

In the past 10-20 years we have seen a significant push to bring new methods to the farmers to control in-row weeds. And for transplanted crops, there are methods available with vision-controlled in-row harrowing (Poulsen, 2018; Steketee, 2018; Tillett et al., 2008) and selective spraying (Blue River Technologies, 2018; Ecorobotix, 2018). These cultures are relatively sparse which allow for these methods. The Garford rotating disk cultivator was tested in celery, lettuce, and radicchio in trials presented by Fennimore et al. (2014). Economic analysis was performed on some of the yield data, which demonstrated that the rotating disc cultivator was generally more effective than the standard inter-row hoe cultivator combined with hand weeding for the transplanted crops. For direct seeded salad the system damaged 20 - 28 % of the salad stands, which incurred a net loss by using the system. The study demonstrates that there exists an economic rationale for robotic weed control, while the rotating disc cultivator that was tested was not suitable for the seeded cultures.

Seeded crops present a greater challenge than transplanted crops. The visual separation of crop and weed is harder as the weed and crop are similar in size, and for many crops there is not enough room in between crop plants for a mechanical hoe to pass in and out of the crop row. Herbicide application either requires a selective herbicide which does not harm the crop, or a better resolution application to not affect the crop.

The essence of DoD spraying is to detect the weeds within the plant row, and selectively shoot droplets of herbicide on those weed leaves. By targeting only the weed leaves, the crop and soil are left unaffected, which allows for the use of broad spectre herbicides that would normally harm the crop, Figure 1.2.

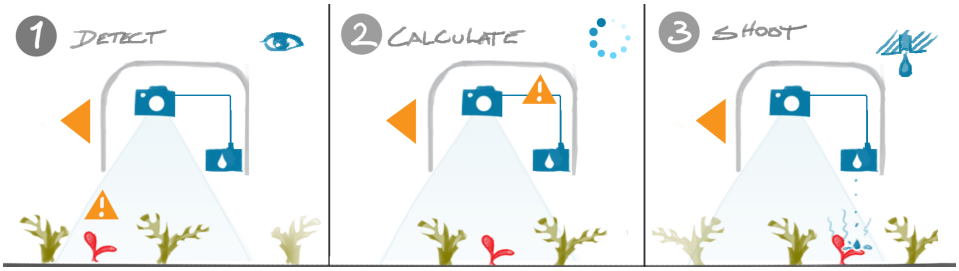

Figure 1.2: Concept illustration of the Drop-on-Demand herbicide application. Courtesy of Adigo AS.

1.2.1 Crops of focus

The DoD technology is suitable for a wide range of row crops, and an ambition for the project was to provide an alternative method for weed control also for seeded crops. To guide our selection of focus crops, we have considered the agricultural market in Germany, France, Netherlands and United Kingdom. The area harvested for a few selected crops is shown in Figure 1.3 using 2012 data from the Food

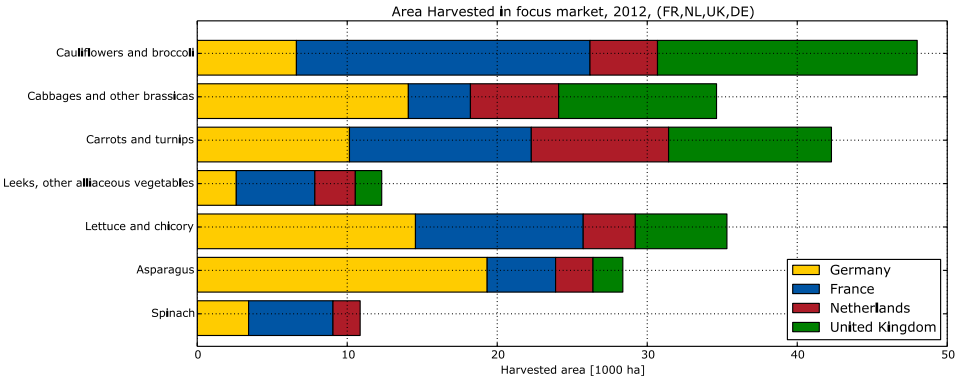

Figure 1.3: Harvested area for a few selected crops of interest in Germany, France, Netherlands and United Kingdom in 2012.

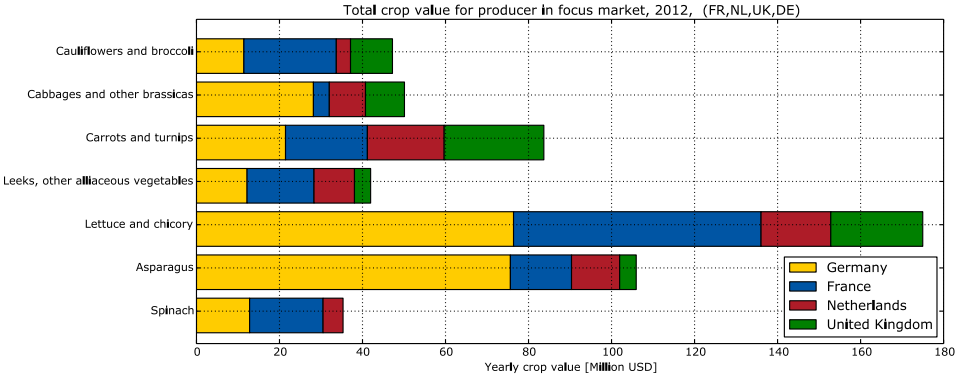

Figure 1.4: Total crop value for the farmers of the selected crops in Germany, France, Netherlands and United Kingdom in 2012.

and Agriculture Organisation of the United Nations, (FAO, 2014). Cauliflowers, broccoli, cabbages, other brassicas and lettuce is dominantly transplanted crops, for which mechanical in-row weeding solutions to some extent are commercially available.

Carrots and turnips are dominantly seeded crops, and we consider them a good example of challenging crops for robotic weed control. They are of high value as illustrated in Figure 1.4. In 2014 carrots and turnips accounted for 6.25 % of Europe's harvested area for vegetables, with 2.6 million Ha. The gross production value for Europe was above 3 billion USD in 2014 (FAO, 2014).

Carrot competes poorly with weeds especially in the early stages, as documented by Swanton et al. (2010) in a field trial in Ontario, Canada. The critical weed-free period for carrots was found to be 450 growing-degree-days (3 to 6 weeks at 10 to 20°C), or until the carrot plants have reached the six-leaf stage. The carrot crop row becomes increasingly dense in this period as illustrated by Figure 1.5 and 1.6. The crop will gradually out-compete any smaller weed by blocking for sunlight,

Figure 1.5: A carrot crop row at the one true leaf stage. The larger weeds in this image is *Chenopodium album* L. (Fat hen).

Figure 1.6: A carrot crop row at the 3 - 4 true leaves stage. Among other weeds *Tripleurospermum inodorum* (L.) Sch. Bip. (Scentless Mayweed) is present in this image, recognized by it tree-like leaf structure, which bear resemblance to the carrot leaves.

while it becomes increasingly hard to detect and target individual weeds within the row.

1.3 State of the Art

The available products for guided hoeing and selective thinning are paving the way for further advances in automatic weed control in speciality crops. Our attention will be focused on precision-spray application targeting individual weeds - a domain which is yet to see its first commercially available solution for seeded crops.

One of the first demonstrations of a Precision-Spray robot was by Lee, Slaughter, and Giles as early as 1999. They developed a robot for controlling weeds in tomato crops. The robot was equipped with an Cohu RGB camera which information was digitized to 256x240 pixels at 8 bit per channel. The processing was done by a 200 MHz Pentium Pro CPU running MSDOS. The system recognized 73 % of the tomato plants and 69 % of the weeds, and was able to treat 48 % of the weeds at a speed of 0.8 km/h.

Nearly 20 years have passed since then, and while the robots have become incrementally better, we are yet to see weeding robots make an impact on the use of herbicides in agriculture. A thorough overview of this field can be found in Fennimore et al. (2016) or Slaughter et al. (2008), while we here will focus on a few relevant technical aspects.

1.3.1 Drop-On-Demand herbicide application

A challenge presented by Lee et al. (1999) is to increase the accuracy, precision and efficacy of the herbicide application. This effort involves everything from the design of the droplet forming mechanism, the fluid dynamics of the droplets, the

droplets retention on the weed leaves, the choice of active ingredient, to the motion estimation and targeting algorithm.

Most of the previously presented systems for DoD herbicide application has either used adapted industrial print-heads (Lund and Mathiassen, 2010; Midtiby et al., 2011) or an array of solenoid valves and nozzles (Lee et al., 1999; Nieuwenhuizen, 2009; Søgaard and Lund, 2005) to form droplets. There is also a presented paper by Basi et al. (2012) where a pneumatic valve is presented for better dosing and formation of individual droplets. The fluid dynamics of the in-flight droplets has been investigated by Lund and Mathiassen (2010) and Lund and Olsen (2010). They described the disintegration of droplets and the effects of altering the viscosity and surface tension of the fluid. We expanded on this and also explored the effect of the electrical control signal to the solenoid valve on the droplet formation in our experiments presented in Urdal et al. (2014).

Lund and Mathiassen (2010) and Lund et al. (2006) demonstrated that herbicide droplets formulated with Glyphosate (27µg per plant) can effectively control *Solanum nigrum* L., (Black Nightshade) a weed which is resistant to most selective herbicides. Midtiby et al. (2011) presented a simulated row crop trial where plants passed under the system on a conveyor belt at 0.5 m/s. The system was able to effectively control weeds larger than 11x11 mm, which gave good results on *Brassica napus* L. (oilseed rape) and to some extent *Tripleurospermum inodorum*(L) Sch. Bip. (Scentless Mayweed). Koukiasas et al. (2016) demonstrated that *Galium aparine* L. is effectively controlled with 19.3 µg of glyphosate per plant.

1.3.2 Leaf Classification

Weed and crop classification has largely followed the classical approach of segmenting plant material from the background soil, for subsequent classification based on shape, color and texture features, since Lee et al. (1999). Several systems have incorporated a Near-Infrared (NIR) channel to enhance the soil segmentation, e.g.: Nieuwenhuizen (2009). A mutual unsolved challenge has been overlapping leaves and higher density cultures - as the classification has mainly been reliant on segmentation and the shape of the leaves. This has been an important road-block for robotic weed control, as we have not had satisfactory performance of classification when weeds overlap. There has been much effort invested in improving these algorithms (Fennimore et al., 2016). One example is Haug et al. (2014), who was able to circumvent the reliance on segmenting individual plants by implementing a form of sliding-window classifier.

Computer Vision in Agriculture may have been slightly behind other technological domains. We have seen nearly all other image recognition and classification tasks being dominated by deep convolutional neural-networks (deep CNN) in recent years. We do see an important shift now towards artificial intelligence (AI) and CNN's making its way into weed detection. One of several compelling examples is presented by Milioto et al. (2018), where pixel-wise semantic segmentation into weed/crop is demonstrated.

(a) DeepField Robotics BoniRob, photo courtesy of Bosch AG

(b) Ecorobotix, photo courtesy of Ecorobotix Ltd

(c) Thorvald II platform, photo courtesy of Saga Robotics AS.

(d) AgBot II, photo courtesy of Queensland University of Technology

Figure 1.7: A selection of other robot platforms presented in litterature.

1.3.3 State-of-the art in Agricultural Robotic Platforms

There is a significant body of research and industrial push towards robotization in agriculture. There are philosophies towards automating tractors, building specialized robots for each task and towards making highly versatile and modular robots. A selection of comparable robots that have been presented for weed control is shown in Figure 1.7.

Modularity has been uphold as an important design criteria for the Armadillo (Nielsen et al., 2012), Naïo Dino (Naïo Technologies, 2018) and the Thorvald II platform (Grimstad and From, 2017) which can be customized to different configurations. Thorvald II, BoniRob, (Figure 1.7a and 1.7c), and Naïo Dino have drive and steering on all four wheels. This enables holonomic control of the robot: the robot can navigate freely in all directions, handle tight environments such as greenhouses and the front and rear wheels can follow the same tracks through a turn. This comes at the cost of having 8 motors for steering and drive.

The AgBot II shown in Figure 1.7d and presented by Bawden et al. (2017) is a robot platform for weed control, set up with differential drive front wheels and two rear castor wheels. The design emphasizes modularity and ease of on-site assembly of the system. A docking container covered with solar panels provide the power needs for charging, and the system has been tested with a range of chemical and mechanical weed control implements (McCool et al., 2018).

A more minimalistic approach has been taken by the Swiss company Ecorobotix (Ecorobotix, 2018) who are working on a fully solar powered robot, Figure 1.7b, which applies a micro-dose of herbicide by two robotic parallel arms.

The systems described above are intended to be a representative selection, and not an exhaustive review of the field, as there are several other systems that could have been mentioned.

1.4 System Development and Research

The development of an agricultural robot for weed control involves technology and knowledge that span across several fields. The author of this thesis has been deeply involved in all aspects of its development and testing. The robot platform that was used for the main part of the testing and development is shown in Figure 1.8.

Figure 1.8: The robot development platform built in 2012 and used with modifications through the following years to perform field trials and develop the components.

1.4.1 Localization and estimation

Localization is a fundamental part of robotics and a pre-requisite for accurate application of herbicide relative to the generated spray maps. The combination of GPS positioning with wheel encoders and inertial sensors (IMU) on the robot provide high accuracy localization.

The implementation of two attitude estimation filters on an IMU was presented in (Utstumo and Gravdahl, 2013) and is reproduced in Section 2.1. While the work only considers the three degrees of rotation for attitude estimation, it provides the fundamental understanding of quaternion implementations in an Extended Kalman Filter.

While working to expand the localization to include visual odometry, the question of including delayed measurements from a slower Visual Odometry algorithm

was explored. This work was presented in (Arbo et al., 2017), and is reproduced in Section 2.5.

1.4.2 Navigation and control

A robot working in a row crop will mostly have two navigational challenges at hand: Detecting and following the row, and upon reaching the headlands perform a turn and enter the subsequent row.

Operation in the crop row requires that the robot keeps its wheels within the existing wheel tracks. If the robot sways from its path, and the wheels enter the crop rows this will directly damage the crop and result in an economic loss for the producer. A Non-linear Model Predictive Crontroller (NMPC) was formulated to center the robot over the detected row crops while maintaining the constraints of the wheel tracks. The controller with simulations was presented in Utstumo et al. (2015), and is reproduced in Section 2.3.

The simulations with the NMPC controller assumed perfect knowledge of the model parameters. In a real world situation the robot response will vary with soil conditions, payload etc. Adaptive control is explored in Dørum et al. (2015), which is reproduced in Section 2.4. The paper presents two approaches to adaptive control: One is based on-line parameter estimation of the dynamic model, which is used to adjust the controller parameters. The second approach implements a direct Model Reference Adaptive Controller (MRAC), where instead of identifying the model parameters we directly identify the controller parameters so that the robot behaves as closely as possible to a simulated first-order reference model.

For a fleet of agricultural robots these methods may prove valuable in varying conditions. Where the operation in the crop row will not provide sufficient excitation for the adaptive controllers to estimate its parameters, the operations in the headlands may be used for system identification as well as any navigation to and from the field. The improved knowledge of the system response will enable a more precise control of the robot.

The NMPC controller was implemented using the Casadi open-source non-linear optimization framework, (Andersson, 2013). The intuitive structure of the control problem with a symbolic formulation of the cost function and constraints allows for adaptation to other scenarios and environments. For example if the robot was to operate within the constrained space of a greenhouse, additional constraints may be added to expand the controller.

1.4.3 Herbicide application

At the core of the system is the DoD module, which accurately images the ground, detects weeds and precisely shoot droplets of herbicide on weed leaves. The liquid properties are important through the three phases for the droplet. At the outset we seek to create one droplet without satellites and break-up. We require fast control of the solenoid valve at opening and closing, and as the droplet forms in air we seek liquid properties and an *Ohnesorge number*, so that the tail and filament of the droplet does not disintegrate into satellite droplets. In-flight, the droplet stability is governed by the *Weber number*, where a higher surface tension aids stability. In

the last phase, when the droplet lands on the leaf, we seek a good retention of the droplet and low surface tension so that it floats out and covers as much as possible of the leafs surface.

Together with the design of a H-bridge electronic controller for the solenoid valves, the droplet and liquid properties are further described in (Urdal et al., 2014), which is reproduced in Section 2.2.

1.4.4 The overall system and trials

The overall robot design is presented in (Utstumo et al., Submitted for publication 2018), which is reproduced in Section 2.6. The paper presents a liquid system for flushing the solenoid valves and pipes to avoid residues of herbicide and clogged valves, the row following method and results demonstrating the efficacy of the DoD method on weeds in a lab and a field trial.

Chapter 2

Conference and journal papers

This chapter presents 4 conference papers, and 2 journal papers, in chronological order.

Conference Paper - Attitude Estimation

T. Utstumo and J. T. Gravdahl. Implementation and comparison of attitude estimation methods for agricultural robotics. In *Proceedings of the 4th IFAC Conference on Modelling and Control in Agriculture, Horticulture and Post Harvest Industry, Espoo, Finland, 28-30 August 2013.*

Conference Paper - Drop on Demand Valve Control

F. Urdal, T. Utstumo, J. K. Vatne, S. A. A. Ellingsen, and J. T. Gravdahl. Design and control of precision drop-on-demand herbicide application in agricultural robotics. In *Proceedings of the 13th International Conference on Control Automation Robotics and Vision, ICARCV, Singapore, December 10-12, 2014.*

Conference Paper - Non-linear Model Predictive Control

T. Utstumo, T. Berge, and J. Gravdahl. Non-linear model predictive control for constrained robot navigation in row crops. In *Proceedings of the IEEE International Conference on Industrial Technology, Sevilla, Spain, March 17-19, 2015.*

Conference Paper - Adaptive Control

J. Dørum, T. Utstumo, and J. Gravdahl. Experimental comparison of adaptive controllers for trajectory tracking in agricultural robotics. In *Proceedings of the 19th International Conference on System Theory, Control and Computing, ICSTCC, Cheile Gradistei, Romania, October 14-16, 2015.*

Journal Paper - Delayed Measurements

M. H. Arbo, T. Utstumo, E. Brekke, and J. T. Gravdahl. Unscented multipoint smoother for fusion of delayed displacement measurements: Application to agricultural robots. *Modeling, Identication and Control*, 38(1):1, 2017.

Journal Paper - Robotic In-Row Weed Control

T. Utstumo, F. Urdal, A. Brevik, J. Dørum, J. Netland, Ø. Overskeid, T. W. Berge, and J. T. Gravdahl. Robotic in-row weed control for vegetables.

Computers and Electronics in Agriculture, Submitted for review, February 28, 2018.

2.1 Conference Paper - Attitude Estimation

T. Utstumo and J. T. Gravdahl. Implementation and comparison of attitude estimation methods for agricultural robotics. In *Proceedings of the 4th IFAC Conference on Modelling and Control in Agriculture, Horticulture and Post Harvest Industry, Espoo, Finland, 28-30 August 2013.*

Implementation and Comparison of Attitude Estimation Methods for Agricultural Robotics

Trygve Utstumo [*,**] **Jan Tommy Gravdahl** [*]

[*] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway.*
jan.tommy.gravdahl@itk.ntnu.no, trygve.utstumo@itk.ntnu.no
[**] *Adigo AS, NO-1415 Oppegård, Norway.*

Abstract: The field of precision agriculture increasingly utilize and develop robotics for various applications, many of which are dependent on high accuracy localization and attitude estimation. Special attention has been put towards full attitude estimation by low-cost sensors, in relation to the development of an autonomous field robot. Quaternions have been chosen due to its continuous nature, and with respect to applications in the pipeline with on other platforms.
The performance and complexity of two approaches to attitude estimation has been investigated: One Multiplicative Extended Kalman Filter (MEKF) and one non-linear observer. Both were implemented on an ARM Cortex M3 microcontroller with sensors for a Attitude Heading Reference System (AHRS), and benchmarked towards a relative high grade commercial AHRS device.
The relative computational burden of the MEKF have been underlined, by execution times more than 10 times those of the non-linear estimator. The implementation complexity is also significantly lower for the non-linear observer, which facilitate test and verification through more transparent software.

Keywords: Instrumentation and Sensing; Robotics and Mechatronics for Agricultural Automation; Precision Agriculture

1. INTRODUCTION

Modern agriculture is increasingly utilizing advanced technology to automate and better manage its production processes. The use of autonomous systems for weed control is a research field with growing interest, and several autonomous systems have been demonstrated, where some are presented in the review by Slaughter et. al. Slaughter et al. (2008).

Adigo is developing a mobile robot, illustrated in Figure 1, for research on precision agriculture[1]. Building on previous experience with autonomous robots, the attitude estimation is given special attention as part of the localization and navigation for the robot. Attitude heading reference systems (AHRS) are widely utilized in other applications of autonomous ground vehicles, and provide important input to the localization sensor fusion.

The work presented can facilitate customized and better integrated solutions with attitude estimation and enable the use of low cost sensors.

1.1 Multiplicative Extended Kalman Filter

The survey by Crassidis, Markley and Cheng Crassidis et al. (2007) provides a good background and review of various attitude filters, observers and smoothers. For a singularity free representation of attitude we have considered quaternion estimators.

The Extended Kalman filter (EKF) has become the workhorse of attitude estimation, largely through the research effort and numerous applications in space exploration. There are numerous variations on how to implement an EKF for quaternion estimation, especially in the update step where both quaternion addition and multipli-

Fig. 1. A robot developed for autonomous N_2O measurements on cereal fields. The robots autonomy is currently under development where efficient attitude estimation is a focus.

[1] Consortium research program "Multisensory Precision Agriculture - Improving yields and reducing environmental impact" sponsored by the Norwegian Research Council [207829].

cation can be utilized. Multiplication should however be preferred Shuster (1993).

Implementing multiple vector measurements directly in the MEKF is not trivial, and is accompanied with a complex set of tuning parameters Markley et al. (2003). It is however quite straightforward to implement a quaternion measurement in the MEKF. Using the QUEST Shuster and Oh (1981), or an equivalent algorithm, for preprocessing measurement vectors greatly simplify the filter interface. The QUEST algorithm provides support for an arbitrary number of measurement vectors, and the QUEST covariance matrix can directly feed the Kalman $\mathbf{R}$ matrix Shuster (1990),Crassidis et al. (2007).

1.2 Non-linear observer: Explicit Complementary Filter

The EKF does have a number of drawbacks: Implementation is not straight-forward, the numerous parameters require tuning, it is computationally expensive, and it is usually difficult to prove its convergence Martin and Salaun (2010).

Many of these issues can be addressed by nonlinear observers, and especially the stability properties can be proven, using Lyapunov-based methods. A significant step was taken with the *Explicit Complementary Filter* proposed by Hamel and Mahony (2006) and refined in Mahony et al. (2008). It makes use of the vector measurements directly in body frame, and also includes gyro bias estimates. The filter provides near global stability. The concept has later been extended to include time-varying reference vectors by Hua (2010) and Grip et al. (2011). For systems with low accelerations the filter performs as well as these later designs (Hua, 2010), and is thus suitable for wheeled robotic applications with slow dynamics.

2. MODELLING

We operate with three coordinate frames, where our reference frame is North-East-Down (NED), the robot body frame is defined as forward-right-down (BODY) and the instrument frame depending on how the sensor is mounted in the robot (I). The frames are indicated by subscripts n, b and i respectively.

2.1 Unit Quaternions

Rotations and attitude are represented by unit quaternions, where the scalar is defined as the first element in

$$\mathbf{q} = \eta + i\epsilon_2 + j\epsilon_2 + k\epsilon_3 = [\eta, \boldsymbol{\epsilon}]^T \tag{1}$$

Quaternion multiplication is expressed by:

$$\mathbf{p} \otimes \mathbf{q} = \begin{bmatrix}\eta_p \\ \boldsymbol{\epsilon_p}\end{bmatrix} \otimes \begin{bmatrix}\eta_q \\ \boldsymbol{\epsilon_q}\end{bmatrix} = \begin{bmatrix}\eta_p\,\eta_q - \boldsymbol{\epsilon_p^T}\,\boldsymbol{\epsilon_q} \\ \eta_p\,\boldsymbol{\epsilon_q} + \eta_q\,\boldsymbol{\epsilon_p} + \boldsymbol{\epsilon_p} \times \boldsymbol{\epsilon_q}\end{bmatrix}. \tag{2}$$

Consecutive rotations by quaternions are done by post-multiplication, in contrast to rotation matrices. Thus, the rotation from NED to I, can be composed by rotations from NED to BODY to I as $\mathbf{q}_n^i = \mathbf{q}_n^b \otimes \mathbf{q}_b^i$

2.2 Angular velocity

For a rotation from NED to BODY, the kinematic differential equation is given by Egeland and Gravdahl (2002)

$$\dot{\mathbf{q}} = \tfrac{1}{2}\,[0, \boldsymbol{\omega}_n]^T \otimes \mathbf{q} = \tfrac{1}{2}\mathbf{q} \otimes [0, \boldsymbol{\omega}_b]^T \tag{3}$$

where $\boldsymbol{\omega}_n$ and $\boldsymbol{\omega}_b$ are rotational velocities.

2.3 Measurements

The AHRS receives measurement vectors from a MEMS accelerometer, $\mathbf{a}_b$, and magnetometer, $\mathbf{m}_b$, in BODY frame. They are normalized and compared with their reference vectors in NED frame, $\mathbf{a}_n$ and $\mathbf{m}_n$. The MEMS gyro is modelled with a bias in body frame as $\boldsymbol{\omega}_b = \boldsymbol{\omega}_{actual} + \mathbf{b}$

3. MULTIPLICATIVE EXTENDED KALMAN FILTER

Representing the full quaternion in the filter would lead to a singularity in the co-variance matrix $\mathbf{P}$, which results in numerical errors and possibly negative eigenvalues in $\mathbf{P}$, Shuster (1993).

An intuitive solution is to leave out the scalar element, η, of the quaternion. Since the quaternion is of unit length, it can be reconstructed by

$$\mathbf{q}(\boldsymbol{\epsilon}) = \begin{bmatrix}\sqrt{1 - \|\boldsymbol{\epsilon}\|^2} \\ \boldsymbol{\epsilon}\end{bmatrix} \tag{4}$$

With this modification singularities arise at multiples of π. By only representing the rotation error in the filter, $\delta\boldsymbol{\epsilon}$, the singularities are less likely to occur. Error representations with better margins are described by Markley et al. (2003), which should be considered for a robust implementation. The resulting state vector, $\mathbf{x} = [\delta\boldsymbol{\epsilon}\,;\,\mathbf{b}]$ maintains the error of each update, and the gyro bias estimate.

Update	$\bar{\mathbf{q}}(k) \rightarrow$	$\hat{\mathbf{q}}(k)$
Error est.	$\delta\mathbf{q}(k) =$	$\begin{bmatrix}\delta\eta \\ \delta\boldsymbol{\epsilon}\end{bmatrix} = \hat{\mathbf{q}}^{-1}(k) \otimes \mathbf{q}_{ref}(k)$
Innovation	$\Delta\mathbf{x} =$	$\mathbf{K}\,\delta\boldsymbol{\epsilon}$
State update	$\hat{\mathbf{q}}(k) =$	$\mathbf{q}\left(x_{[1:3]}(k)\right) \otimes \bar{\mathbf{q}}(k)$
Reset	$\mathbf{x}_{[1:3]} =$	$\mathbf{0}$

Table 1. Details of the MEKF State update, $\bar{\mathbf{q}}(k)$ is the prediction and $\hat{\mathbf{q}}(k)$ is the posterior.

The essence of the MEKF is how the state update is performed by quaternion multiplication. The error is calculated relative to the QUEST position estimate, and is represented by its vectorial part $\boldsymbol{\epsilon}$. This allows the three dimensional error representation to construct the Kalman matrices, $\mathbf{P}$, $\mathbf{K}$ and $\mathbf{H}$ without singularities. The details of the MEKF update in discrete time is shown in Table 1.

The full singularity free unit quaternion represent the attitude estimate and is used in the nonlinear state propagation, through the kinematic equation (3), discretized by Euler's method Crassidis et al. (2007); Markley et al. (2003); Lefferts et al. (1982); Shuster (2009).

4. THE EXPLICIT COMPLEMENTARY FILTER

The Explicit Complementary Filter described by Mahony et al. (2008), for magnetometer and accelerometer input can be presented as:

$$\boldsymbol{\sigma} = k_a \left(\mathbf{a}_b \times \hat{\mathbf{a}}_b\right) + k_m \left(\mathbf{m}_b \times \hat{\mathbf{m}}_b\right) \tag{5a}$$

$$\dot{\hat{\mathbf{q}}} = \frac{1}{2}\hat{\mathbf{q}} \otimes \begin{bmatrix} 0 \\ \boldsymbol{\omega}_{gyro} - \hat{\mathbf{b}} + k_p \boldsymbol{\sigma} \end{bmatrix} \tag{5b}$$

$$\dot{\hat{\mathbf{b}}} = -k_I \boldsymbol{\sigma} \tag{5c}$$

Where $\boldsymbol{\sigma}$ is the filter correction term, and $k_{[p,I,a,m]}$ are the gains for correction, bias integration and weights on accelerometer and magnetometer measurements. The reference vectors $\hat{\mathbf{a}}_b$ and $\hat{\mathbf{m}}_b$ are found by rotating the reference in NED frame by the transposed rotation estimate, $\mathbf{R}(\mathbf{q})^T$.

$$\hat{\mathbf{a}}_b = \mathbf{R}(\hat{\mathbf{q}})^T \mathbf{a}_i / |\mathbf{a}_i|, \quad \hat{\mathbf{m}}_b = \mathbf{R}(\hat{\mathbf{q}})^T \mathbf{m}_i / |\mathbf{m}_i| \tag{6}$$

The magnetometer can have large influx of noise, especially in vehicles with electrical motors. This problem is well known and different solutions are proposed in literature.

To minimize the impact of this it is possible to reduce the weighting of the magnetometer in periods with high noise on the magnetometer,(Mahony et al., 2008). With more permanent noise on the magnetometer one can limit its effect to only the yaw rotation, (Martin and Salaun, 2010). This can be done by aligning the magnetometer cross product with the measured accelerometer vector, changing equation (5a) to:

$$\sigma_m = k_a \left(\mathbf{a}_b \times \hat{\mathbf{a}}_b\right) + k_m \left(\left(\mathbf{m}_b \times \hat{\mathbf{m}}_b\right)^T \mathbf{a}_b\right) \mathbf{a}_b \tag{7}$$

Local exponential stability can be shown with this modification, but it complicates the analysis for region of attraction, (Martin and Salaun, 2010). Hua (2010) confirms the insulation of magnetic perturbations from roll and pitch in simulations, and do comparisons with the update in equation (5a). A price to pay for this modification, is increased error amplitude from accelerations.

5. HARDWARE

The implementation of the filter algorithms investigated has been done on the AHRS CHR-6dm by CH Robotics, and benchmarked towards the Microstrain 3DM-GX3-25 ARHS.

The two sensors have been aligned and mounted to an aluminum bar, with double-sided tape, to minimize magnetic disturbances.

The CHR-6dm AHRS was chosen because of its Open Source firmware, potent ARM Cortex-M3 processor and its low cost. The individual sensors are surface mounted to the PCB, as seen in Figure 2. The accelerometers and gyros are mounted in agreement with the BODY-frame, whereas the magnetometer is constructed with the z-axis pointing up, which result in the following rotation from the magnetometer instrument frame:

$$\mathbf{m}_b = \mathbf{R}_b^i \mathbf{m}_i = \begin{bmatrix} \mathbf{m_y} & \mathbf{m_x} & -\mathbf{m_z} \end{bmatrix}_i^T \tag{8}$$

The gyro and accelerometers are analog devices, sampled by an AD-converter at 400Hz. The magnetometer is connected over the I2C bus, and reports it's measurements at 87Hz.

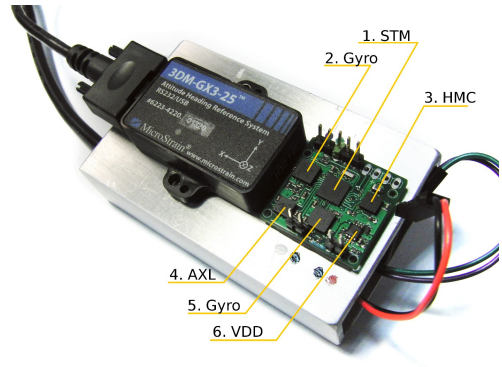

Fig. 2. The components of the CHR-6dm are: **1.** Microcontroller, STM32F103T8 **2.** Gyro Pitch-Roll, LPR510AHL **3.** Magnetometer, HMC5843 **4.** Accelerometer, ADXL335 **5.** Gyro Yaw, LY510AHL **6.** 3.3V Voltage regulator

The gyro bias, $\mathbf{b}$, was measured on several devices, and in various temperatures. The bias values varied in tests in the range of 0.004 to 0.120 rad/s.

The magnetometer measurements are corrected both by a scaling factor and a constant bias, calculated from an initial calibration routine.

6. FIRMWARE DEVELOPMENT

The algorithms for attitude estimation have been developed by first prototyping the algorithms in Matlab with recorded data, then implemented in C both on the PC and then on the microcontroller. The filter parameters were tuned in the Matlab filter prototypes using recorded sensor data.

The QUEST algorithm was implemented on the basis of Shuster (2006) and the flow chart in Takahashi et al. (2009), for two input vectors. Further details on this can be found in Utstumo (2011).

The MEKF was implemented with QUEST as a preprocessor as described in Mahony et al. (2008). The detailed steps of the algorithm are shown in Table 2. The covariance matrix, $\mathbf{P}$, is treated by its block-diagonal elements, $\mathbf{P}_a$, $\mathbf{P}_b$, and the correlation blocks $\mathbf{P}_c$. This further enables us to exploit the simple form of $\mathbf{H}$ to simplify the covariance propagation and reduce the number of matrix operations, as described by Markley et al. (2003).

The covariance matrix returned from the QUEST-algorithm is presented directly to the Kalman filter as the measurement covariance matrix, $\mathbf{R}$. The $\mathbf{P}$ matrix is set initially large on the error estimate, $\mathbf{P}_a(0) = \text{diag}(100)$, to quickly converge to the correct attitude, while the initial bias estimate dynamics are limited by setting, $\mathbf{P}_b(0) = \text{diag}(0.1)$.

The time step, Δt, is calculated between each step by using the internal timer `TIM3` on the microcontroller.

The Explicit Complementary filter is expressed in three equations (5), and those three equations are implemented

Step	*Equations*
Update step	(87 Hz)
Quest	$[\mathbf{q}_q, \mathbf{R}_q](k) = \text{QUEST}(\mathbf{a_n}, \mathbf{m_n}, \mathbf{a_b}, \mathbf{m_b}, \sigma_\mathbf{a}, \sigma_\mathbf{m})$
Kalman gain	$\mathbf{K}(k) = \begin{bmatrix}\bar{\mathbf{P}}_a \\ \bar{\mathbf{P}}_c\end{bmatrix}^T [\bar{\mathbf{P}}_a + \mathbf{R}_q]^{-1}$
P update	$\hat{\mathbf{P}}(k) = \bar{\mathbf{P}}(k) - \mathbf{K}(k)[\bar{\mathbf{P}}_a\ \bar{\mathbf{P}}_c]$
Error estimate	$\delta\mathbf{q}(k) = \bar{\mathbf{q}}'(k) \otimes \mathbf{q}_q(k)$
Innovation	$\hat{\mathbf{x}}(k) = \bar{\mathbf{x}}(k) + \mathbf{K}(k)\,\delta\boldsymbol{\epsilon}_b$
State update	$\hat{\mathbf{q}}(k) = \bar{\mathbf{q}}(k) \otimes \begin{bmatrix}\sqrt{1 - \lvert\hat{\mathbf{x}}_{[1:3]}(k)\rvert} \\ \hat{\mathbf{x}}_{[1:3]}(k)\end{bmatrix}$
Normalize	$\hat{\mathbf{q}}(k) = \hat{\mathbf{q}}(k) \,/\, \lvert\hat{\mathbf{q}}(k)\rvert$
Reset	$\hat{\mathbf{x}}_{[1:3]}(k) = [0\ \ 0\ \ 0]^T$
Propagate	(400 Hz)
Propagate state	$\bar{\mathbf{q}}(k+1) = \hat{\mathbf{q}}(k) + \frac{\Delta t}{2}\left(\hat{\mathbf{q}}(k) \otimes \begin{bmatrix}0 \\ \boldsymbol{\omega}_b - \hat{\mathbf{x}}_{[4:6]}(k)\end{bmatrix}\right)$
Linearization	$\mathbf{F} = \begin{bmatrix}-[\boldsymbol{\omega}_b\times] & -\mathbf{I} \\ \mathbf{0} & \mathbf{0}\end{bmatrix},\ \mathbf{G} = \begin{bmatrix}-\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}\end{bmatrix}$
Propagate P	$\bar{\mathbf{P}}(k+1) = \bar{\mathbf{P}}(k) + \Delta t\,(\mathbf{FP} + \mathbf{PF}' + \mathbf{GQG}')$

Table 2. The implementation of the MEKF

Step	*Equations*
Update	(87Hz) If new magnetometer data
Rotate reference	$\hat{\mathbf{m}}_b = \mathbf{R}(\hat{\mathbf{q}})^T\mathbf{m}_n$ $\hat{\mathbf{a}}_b = \mathbf{R}(\hat{\mathbf{q}})^T\mathbf{a}_n$
Correction	$\sigma_m = k_a(\mathbf{a}_b \times \hat{\mathbf{a}}_b) + k_m\left((\mathbf{m}_b \times \hat{\mathbf{m}}_b)^T\mathbf{a}_b\right)\mathbf{a}_b$
Update	$\delta\hat{\mathbf{q}} = \frac{1}{2}\hat{\mathbf{q}} \otimes \begin{bmatrix}0 \\ \boldsymbol{\omega} - \hat{\mathbf{b}} + k_p + \sigma_m\end{bmatrix}$ $\hat{\mathbf{q}}(+) = \hat{\mathbf{q}}(-) + \Delta t\,\delta\hat{\mathbf{q}}$ $\hat{\mathbf{b}}(+) = \hat{\mathbf{b}}(-) + \Delta t\,(-k_I\sigma_m)$
Normalize	$\hat{\mathbf{q}}(+) = \hat{\mathbf{q}}(+)/\lvert\hat{\mathbf{q}}(+)\rvert$
Propagate	(400Hz) If no new magnetometer data
Propagate	$\dot{\hat{\mathbf{q}}} = \frac{1}{2}\hat{\mathbf{q}} \otimes \begin{bmatrix}0 \\ \boldsymbol{\omega} - \hat{\mathbf{b}}\end{bmatrix}$ $\hat{\mathbf{q}}(+) = \hat{\mathbf{q}}(-) + \Delta t\,\dot{\hat{\mathbf{q}}}$

Table 3. The Mahony implementation

directly, with some surrounding logic to handle asynchronous updates, and correction of numerical drift on the quaternion. With the magnetic noise in mind, the implementation has been adapted to use the update (7) from Martin and Salaun (2010). The detailed steps of the algorithm are shown in Table 3.

A significant part of the filter implementation is the supporting libraries to handle matrix, vector and quaternion operations in 32 bit floating point precision. We have chosen to write the methods specific for each matrix dimension, eliminating the overhead accompanied with generic functions supporting arbitrary length vectors.

7. EVALUATION

The hardware implementation has been tested with the QUEST algorithm, MEKF and The Explicit Complementary Filter. And the run-time of each algoritm has been recorded. The timer `TIM4` was set up specifically to time the algorithms, with a clock resolution of 1 μs. The results are shown in Table 4. Note that this is only timing the algorithm run time, excluding the time spent fetching the measurement from the sensors.

Algorithm	Propagate [μs]	Update [μs]
MEKF	1 337	5 580
QUEST, (part of MEKF)		571
Explicit Complimentary Filter	92	376

Table 4. Execution time of the algorithms

To evaluate the accuracy and estimation performance of the sensor, it has been manually aligned and mounted together with the Microstrain 3DM-GX3, see Figure 3.

7.1 Runtime

Our QUEST implementation runs at 571 μs. Directly comparable results have been published earlier by Takahashi et al. (2009), where two processors are compared at 24MHz. Execution time for the ARM Cortex M3 at 72MHz in 32 bit precision is reported to 2444 μs. The large discrepancy in execution time may be due to the custom 3x3 matrix math library, differences in timing, including, or not, the time taken to read and decimate sensors, or communicate data.

The execution time of the Multiplicative Extended Kalman Filter is slow in comparison, and we cannot uphold a constant output rate of 400 Hz through the update.

The Mahony non-linear observer, on the other hand is more than 10 times faster, which leave runtime for auxiliary tasks such as communication etc.

7.2 Benchmarking towards the 3DM-GX3

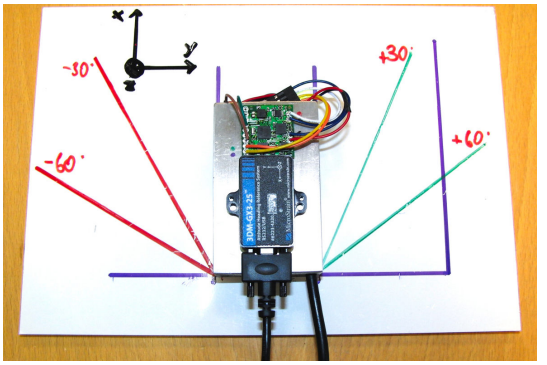

Fig. 3. This setup is used to test the step-response and accuracy of the yaw estimate. The filter is benchmarked towards the known movement, as well as the reference Microstrain 3DM-GX3 sensor.

The relative high-grade MEMS sensor by Microstrain was used as a benchmark in this project. To be able to compare the MEKF and Mahony filters on the same data-sets, the AHRS was set up to output raw sensor measurements. The algorithms have then been run in Matlab on the logged data. This also facilitates the process of tuning the filters.

To test the accuracy of the algorithms, a level plate was attached by double-sided tape to a table, Figure 3. The plate was assured level by using a hand level tool, and directed north as reported by the Microstrain sensor.

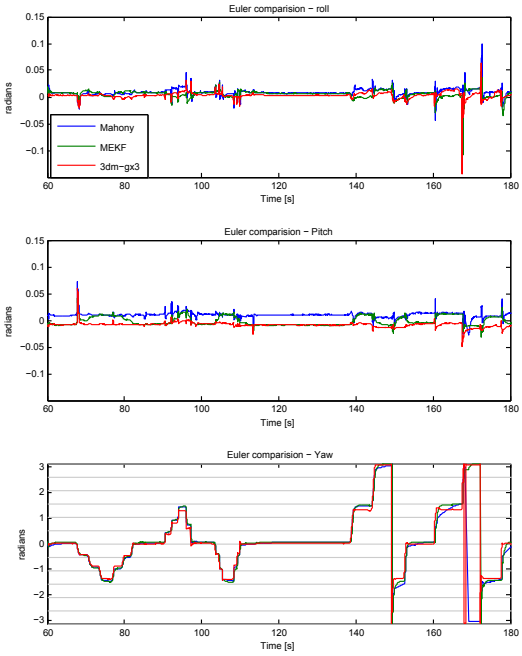

Fig. 4. Accuracy test for Yaw-angle, the lines in the Yaw plot are drawn at $n\pi/6$.

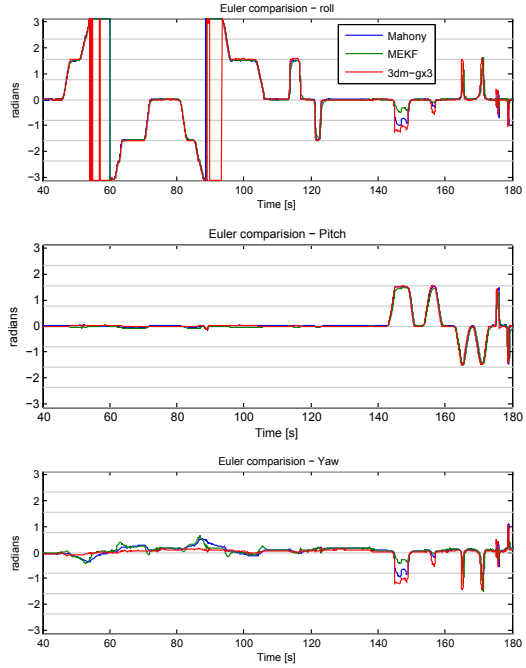

Fig. 5. Accuracy test for Roll and Pitch angle.

Three test sequences were performed: One for yaw accuracy and one for roll and pitch, where the sensor was moved in steps to test transient response and absolute accuracy. The last test was in free hand motion with smoother motions.

7.3 Yaw accuracy

In Figure 4 the filters output is shown. There are several interesting features in these plots. Please note that the y-axis range on Roll and Pitch is $\pm 8.6°$.

The pitch estimate shows how a small positive bias on the accelerometer y-axis directly affects the Explicit Complementary filter, with the modified update (7). The magnetometer readings are discarded from the pitch and roll updates, and the filter then solely rely on the accelerometer.

The MEKF on the other hand, uses the QUEST algorithm with near equal weights on the two measurements. As the two vector measurements conflict, QUEST alternates between trusting the magnetometer, or the accelerometer the most.

This is a desirable trait to the QUEST algorithm, and the pure Mahony update would rather have output a weighted average between the two.

Further, notice that the transients caused by the abrupt stop and go motion affect the two sensor nearly the same. This is an effect of the accelerometer measurement swinging out when rotating the sensor in a stop-and-go fashion.

In the first part of the Yaw plot, the three lines follow closely, and the Microstrain slightly underestimate the rotation which was performed in three steps to $\pm\pi/2$. In the latter part of the graph, the Microstrain display better performance under fast dynamics, where the two filters on the AHRS data first underestimate the rotation, but slowly swing in to near correct levels.

7.4 Roll and pitch accuracy

The deviation in Figure 5 is mostly due to a combination of non-perfect alignment of the two sensors, and differences in how the magnetic field is measured.

The filters also show excellent accuracy in estimating the roll and pitch angles when the gravitation vector is aligned with a measurement vector. At 170s the Euler conversion experience near-gimbal-lock conditions, and the Roll and Yaw measurements does not represent the physical motion.

7.5 Hand held

The dataset shown in Figure 6 was recorded by freely holding the sensor-block, and moving it around. In contrast to the Yaw Accuracy test, the Microstrain now show an higher amplitude in yaw rotation. Otherwise the filters display very similar dynamics, where the MEKF is the most conservative one.

At approximately 52 seconds, the Microstrain Euler conversion experience a gimbal lock, and spin both the Roll and Pitch angle 360°, a mathematical artifact in the conversion to Euler angles.

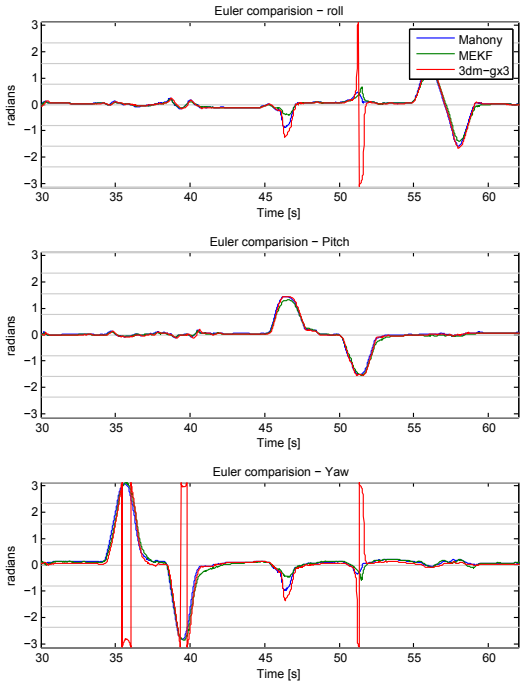

Fig. 6. The filters under controlled hand motion.

8. RESULTS

We have demonstrated the implementation of the QUEST algorithm, the much implemented Multiplicative EKF and a non-linear observer, the Explicit Complimentary Filter.

The algorithms have been implemented on the CHR-6dm, where most of the firmware has been re-engineered to better accompany the filters. The implemented algorithms have been analysed by run-time on the microprocessor. The QUEST algorithm runs at 571 μs, which is four times faster than a previously published implementation Takahashi et al. (2009), on nearly the same processor.

The Mahony non-linear estimator demonstrate its computational advantage, by running more than ten times faster than the MEKF.

By comparing the algorithms with a relative high-grade sensor, the usability and accuracy of the filters have been positively indicated.

9. CONCLUSION

The implementation, and run-time results of the MEKF and the Mahony non-linear observer illustrate clearly the relative computational cost of Kalman filters to non-linear observer designs. Whereas relatively cheap micro-controllers are fully capable of running a MEKF filter, the chosen non-linear observer run ten times faster.

The Extended Kalman filter design provides few obvious advantages, besides being the industry standard for decades. Non-linear observers may be more demanding in design, but provide attractive stability properties and the implementation in code is significantly more lightweight and transparent. This in turn leave less room for software bugs and facilitate test and verification.

For our agricultural robotic attitude estimation, the explicit complementary filter is preferred over a Kalman based design. And we focus on applying observers which also incorporate heading estimates from the forward motion measured by GPS to improve absolute accuracy.

REFERENCES

Crassidis, J., Markley, F., and Cheng, Y. (2007). Survey of nonlinear attitude estimation methods. *Journal of Guidance Control and Dynamics*, 30(1), 12.

Egeland, O. and Gravdahl, J. (2002). *Modeling and Simulation for Automatic Control*.

Grip, H., Fossen, T., Johansen, T., and Saberi, A. (2011). Attitude estimation based on time-varying reference vectors with biased gyro and vector measurements.

Hamel, T. and Mahony, R. (2006). Attitude estimation on so [3] based on direct inertial measurements. In *Robotics and Automation, 2006. ICRA 2006. Proc. 2006 IEEE Int. Conference*, 2170–2175. IEEE.

Hua, M. (2010). Attitude estimation for accelerated vehicles using gps/ins measurements. *Control Engineering Practice*, 18(7), 723–732.

Lefferts, E., Markley, F., and Shuster, M. (1982). Kalman filtering for spacecraft attitude estimation. *Journal of Guidance, Control, and Dynamics*, 5(5), 417–429.

Mahony, R., Hamel, T., and Pflimlin, J. (2008). Nonlinear complementary filters on the special orthogonal group. *Automatic Control, IEEE Transactions on*, 53(5), 1203–1218.

Markley, F. et al. (2003). Attitude error representations for Kalman filtering. *J. of Guidance Control and Dynamics*, 26(2), 311–317.

Martin, P. and Salaun, E. (2010). Design and implementation of a low-cost observer-based attitude and heading reference system. *Control Engineering Practice*, 18(7), 712–722.

Shuster, M. (1990). Kalman filtering of spacecraft attitude and the QUEST model. *J. of the Astronautical Sciences*, 38, 377–393.

Shuster, M. (1993). The Quaternion in Kalman Filtering. *Advances in the Astronautical Sciences*, 85, 25–37.

Shuster, M. (2006). The quest for better attitudes. *Journal of the Astronautical Sciences*, 54(3/4), 657.

Shuster, M. (2009). Filter QUEST or REQUEST. *Journal of Guidance, Control, and Dynamics*, 32, 643–645.

Shuster, M. and Oh, S. (1981). Three-axis attitude determination from vector observations. *J. of Guidance and C.*, 4(1), 70–77.

Slaughter, D., Giles, D., and Downey, D. (2008). Autonomous robotic weed control systems: a review. *Computers and Electronics in Agriculture*, 61(1), 63–78.

Takahashi, S.N., de Souza, A., and Tosin, M. (2009). Execution time of QUEST algorithm in embedded processors. In *Symposium on Aerospace Eng. and Applications*.

Utstumo, T. (2011). *Attitude Estimation in Agricultural Robotics*. Master's thesis, Department of Engineering Cybernetics, Norwegian University of Science and Technology.

2.2 Conference Paper - Drop on Demand Valve Control

F. Urdal, T. Utstumo, J. K. Vatne, S. A. A. Ellingsen, and J. T. Gravdahl. Design and control of precision drop-on-demand herbicide application in agricultural robotics. In *Proceedings of the 13th International Conference on Control Automation Robotics and Vision, ICARCV, Singapore, December 10-12, 2014.*

Design and Control of Precision Drop-on-demand Herbicide Application in Agricultural Robotics

Frode Urdal[†*], Trygve Utstumo[†*], Jan Kåre Vatne[*], Simen Andreas Ådnøy Ellingsen[‡], Jan Tommy Gravdahl[†]

[†] Department of Engineering Cybernetics, Norwegian University of Science and Technology, NO-7491 Trondheim
[‡] Department of Energy and Process Engineering, Norwegian University of Science and Technology, NO-7491 Trondheim
[*] Adigo AS, NO-1405 Langhus, frode@adigo.no

***Abstract*—Drop-on-demand weed control is a field of research within Precision Agriculture, where the herbicide application is controlled down to individual droplets. This paper focuses on the fluid dynamics and electronics design of the droplet dispensing. The droplets are formed through an array of nozzles, controlled by two-way solenoid valves.**

A much used control circuit for opening and closing a solenoid valve is a spike and hold circuit, where the solenoid current finally is discharged over a Schottky diode on closing. This paper presents a PWM design, where the discharge is done by reversing the polarity of the voltage. This demands an accurate timing of the reverse spike not to recharge and reopen the valve. The PWM design gives flexibility in choosing the spike and hold voltage arbitrarily, and may use fewer components. Calculations combined with laboratory experiments verify this valve control strategy.

In early flight the stability of the tail, or filament, is described in theory by the Ohnesorge number. In later flight, when a droplet shape has formed, the droplet stability is governed by the Weber number. These two considerations have opposite implications on the desired surface tension of the fluid. The Weber number is more important for longer distances, as the filament satelites normally catch up and join the main droplet in flight.

I. Introduction

In this paper we study the use of a H-bridge, PWM, as the valve control strategy for a drop-on-demand(DOD) herbicide application in precision agricultural robotics. The design and control strategy has been guided by experiments with droplet dynamics, and the effect of reverse voltage overshoot has been illustrated.

Weed control is a vital part of agriculture, and herbicide application is the most efficient and common control strategy. Environmental and health concerns lead to restrictions and regulations on the use of herbicides, which stimulate initiatives for other weed control strategies [1]. Precision agriculture is an active area of research and methods in agriculture which focuses on adapting the field treatment to the spatial and temporal heterogenity of a field. Weed control in row crops, such as carrots, can be seperated into controlling weeds within, and in between the crop rows: Intra- and inter-row weed control.

DOD herbicide application for intra-row weeding has been investigated by several research groups: [2] designed a robotic weed control system for tomatoes, [3] developed an automated detection and control system for volunteer potatoes in sugar beet fields and [4] created a crop/weed discriminating microsprayer. Common for all tree applications is the use of a valve array to only target the weeds, thus avoiding crop and soil. The literature displays promising results, and experiments indicate that the herbicide usage can be reduced by more than 95 % [1]. The literature also illustrates that there are remaining challenges with precisly targeting droplets, classifying weeds by machine vision and maintaining a precise motion estimate for the robotic platform and nozzle array. The review article [5] presents a good overview of the field.

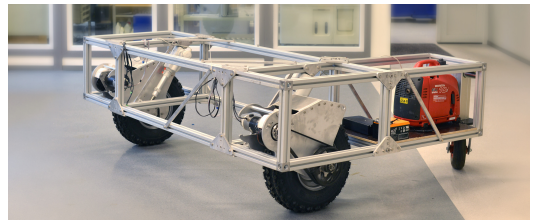

Fig. 1. The Asterix robot platform design for operation in row crops. The platform has two driven wheels and a passive caster wheel. The Asterix modules with the DOD system and machine vision will be mounted in the open area between the two wheels.

The work presented in this paper has been done in the framework of the Asterix project, which works towards a functional robot for DOD intra-weed control in carrots and other row crops. The robotic prototype platform for Asterix is shown in Figure 1 and the localization and attitude estimation for this robot has been described in [6]. In the following sections we will focus on the design and control of the DOD array of nozzles and the control strategy, while we also present our experimental results accompanied with some of the fluid dynamic theory of droplet stability. Droplet stability for at least 15 cm is necessary for this application.

A. Valve and nozzle limitations

The valve and nozzle used are of type INKX0514300A and INZA4710975H respectively, from The Lee Company, as illustrated in Figure 2. The requirement on resolution of control decides what time of the season a system is effective. The control resolution will have a practical lower limit depending

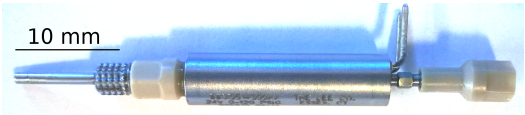

Fig. 2. A VHS valve, INKX0514300A, with minstack mountings and nozzle, INZA4710975H, from The Lee Company.

on the droplet accuracy. If the droplets have an accuracy of ± 2.5 mm there is no need to have finer resolution than 5 mm$\times$5 mm as it would result in many droplets missing the target. The sideways resolution is only a function of nozzle placement, ref. [4] used one row with a spacing of 10.5 mm. Their results and calculations showed that the system was not suitable for targeting weeds smaller than 11 mm$\times$11 mm. Ref. [3] used a similar system with resolution of control about 100 mm^2. Thus neither system will be efficient in the early stages of the season when the weeds are still smaller than 100 mm^2.

The resolution in driving direction can be controlled by the frequency of the valves and the velocity of the vehicle. For instance, ref. [3] used a valve limited to a maximum frequency of 80 Hz, and the demand for control resolution was 100 mm^2, thus limiting the velocity of the vehicle to 0.8 m/s.

Flat fan nozzles are an alternative that allows for smaller weeds to be targeted by spraying a small patch. Recent work has investigated the efficiency of patch spraying with flat fan nozzles [7]. These tests showed promising results for spraying of 100 mm$\times$100 mm patches. When working with row crops, especially carrots, a DOD application with finer resolution is interesting, as the seeds are placed close to each other and weed in between should be controlled. The use of flat fan nozzles in row crops was also examined in ref. [3], where DOD was found beneficial.

Solenoid valves have an upper limit for droplet frequency, and for some microdispensing valves this limit may be hundreds of hertz. However, due to the required droplet volume, the real upper limit may end up around 100 Hz, as a higher frequency would further reduce the volume. Relevant volumes per droplet for a DOD herbicide application lies between 1 μL and 5μ L, and on-times of about 10 ms.

One aspect that needs to be considered when dealing with valve opening time intervals of a few milliseconds, is the fluid dynamics. The fluid in a straight tube can be modelled as an equivalent electrical circuit [8]. This can then be applied to simulate the fluid response in the nozzle under ideal conditions. Increasing the diameter or decreasing the length of the nozzle will result in increased volume rate deposition, but may alter the properties for the droplet in flight.

B. Droplet formation

A droplet produced by a DOD system consists of two or three sections, the main droplet, the filament and a tail. The filament is a cylindrical stream of flow following the main droplet, while the tail is a thin flow behind the filament. The different parts are illustrated in Figure 3. For more information consult [9].

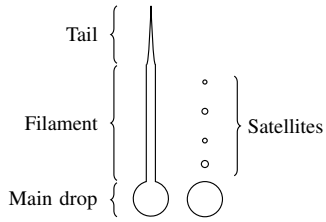

Fig. 3. Droplet definitions

The relative importance to the filament stability from surface friction and viscosity can be expressed through the Ohnesorge number:

$$Oh = \frac{\eta}{\sqrt{\rho \sigma R}} \tag{1}$$

where η, ρ and σ are the viscosity, density and surface tension of the liquid, respectively, while R denotes the radius of the cylindrical filament [10]. Furthermore, the initial filament aspect ratio, $\Lambda = L/2R$, will decide if the filament breaks up or not, L is the length of the filament. The critical value for filament breakup, Λ_c, increases with Oh [10].

When the droplet is falling, a number of scenarios may occur: the filament may be absorbed into the main droplet, it may break at the main droplet thus creating a single satellite droplet or a Rayleigh-Platou instability may occur, creating multiple satellite droplets [9].

Satellite droplets are small droplets lagging behind the main droplet, often caused by the disintegration of the tail or filament. Without wind and other disturbances that could be present for a DOD application in movement, the satellite droplets will typically catch up with the main droplet and coalesce with it, and the dispensed fluid volume reaches the target as a single droplet [11]. An image sequence illustrating this is presented in Figure 8. This is a result of less drag on the satellites as they are smaller and travel in the wake of the main droplet. This will happen under ideal conditions, but how the satellites will behave in the field is not certain. Most of the research described above are results from ink jet printers with droplets much smaller than what is needed for herbicide applications. However, ref. [12] verifies that the theory applies for larger droplets as well, which is more relevant for this project.

The Weber number is of importance when studying the droplets in air, and is defined as, [13]:

$$We = \frac{\rho u^2 d}{\sigma} \tag{2}$$

Where ρ is the density of air, u the droplet velocity, d the droplet diameter and σ the surface tension. With the assumption of spherical droplets, the droplet is stable if its Weber number is below the critical Weber number, which lies between 10 and 40 [14].

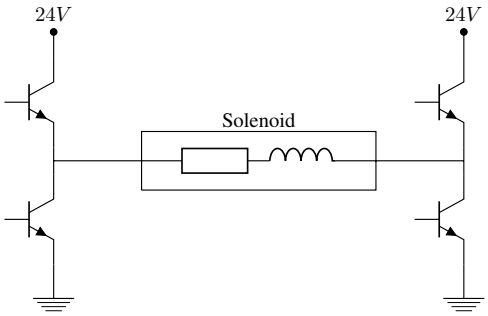

Fig. 4. PWM H-bridge valve driver for a single solenoid

C. Herbicide efficiency with DOD

By far the most common herbicide in use today is glyphosate. It has been widely used through the past 40 years [15], and a water solution of glyphosate is a natural and common choice for DOD weed control [5].

Tests on the efficiency of single droplets of herbicide is presented in [16]. The tests were done with seeds of *Solanum nigrum* planted in pots under outdoor conditions. Results showed that approximately 0.8 μg of glyphosate per plant reduced the biomass by 95 % when applied by hand.

In field trials with a DOD system, the microspray system was set to dispense droplets of 2.5 μL with 5 μg glyphosate each. The system achieved 82 % efficiency when the average dose per plant was 22.6 μg. This is only about 4 % of the recommended application [16].

II. Valve Control

In DOD applications the ideal solenoid valve would open and close instantaneously, and the droplet size would be directly proportional with the opening time of the valve. Any physical solenoid valve has a response time τ, which allows for the solenoid coil to charge and the plunger to open. In selecting a valve for DOD applications, one should focus on achieving a response time significantly smaller than the open time, $\tau < T_{open}$.

Several methods are in use for valve control. Typical configurations are: **Single voltage source** controlled by a transistor. This is a simple driver, but it takes longer to open the valve, as the voltage cannot be higher than the hold voltage as it may burn off the coil. Thus charging the coil takes longer than using a higher voltage source. **Spike and hold** drivers with two different voltage sources, one for the spike and another for the hold voltage. They are more complex, but achieve a much faster response. Common for both configurations when closing the valve is that the energy in the coil is burnt off over two diodes in reverse series parallel to the valve. Another solution is to use **PWM** control to create a spike and hold driver equivalent, with diodes to discharge the coil. However, if the PWM is extended to a full H-bridge, it can be used to discharge the energy in the coil.

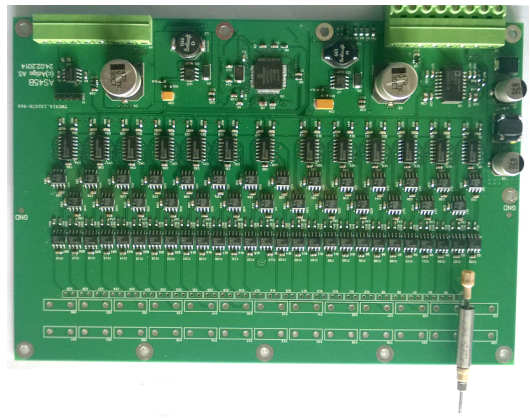

Fig. 5. The DoD demand control unit with one valve and nozzle mounted for an experimental setup.

A. Comparison of PWM and Schottky for solenoid discharge

The main idea for controlling the valve with PWM is that only one voltage source is needed, in a traditional spike and hold driver, two sources are needed, as the spike voltage will overheat the valve if applied for too long. When using PWM the voltage source can be adjusted to fit the spike voltage, that way a large spike followed by a PWM signal to reduce the voltage to the hold value will simulate a spike and hold driver circuit. The PWM control can discharge the solenoid by reversing the voltage over the diode for a significant time, so the current in the coil reaches zero. It is important that the current avoids excessive undershoot as this may open the valve for a short duration before it is closed. This solution will be detailed below. The schematic principle for one single valve driver is presented in Figure 4. When closing the valve, the voltage is reversed over the valve, thus discharging the energy in the coil. The discharge time is reduced with increased voltage, just as the opening time is reduced by increased spike voltage. The Schottky diode solution discharges the coil by burning off the energy in the coil over two schottky diodes in reverse series.

When using the PWM method, the voltage across the valve is limited to the spike voltage, but when using diodes the voltage can be increased further. 50 V reverse voltage is quite common for the schottky diodes for small solenoids with a hold value of about 3.5-4.5 V. The time to close the valve with an internal resistance of 40 Ω, inductance of 12 mH, and hold voltage of 4 V can be calculated for the different solutions. The current response of a resistor in series with an inductor follows the first order response:

$$I(t) = I_0 + (I_1 - I_0)(1 - e^{-t/\tau}) \quad (3)$$

where $I(t)$ is the current at time t, I_0 is the initial current, I_1 is the steady state current for the final solution and τ is

TABLE I
ELECTRICAL CHARACTERISTICS FOR THE LEE INKX0514300A VALVE FROM DATASHEET

Description	Value
Resistance	40 Ω
Inductance	12 mH
Hold voltage	4 V

TABLE II
EXPERIMENTAL AND CALCULATED REVERSE VOLTAGE SPIKE TIMES, WHERE THE FINAL CURRENT IS THE EXPECTED OVERSHOOT OR RESIDUE CURRENT IN THE SOLENOID COIL.

Description	Negative spike duration (ms)		Final current
	Theoretical	Experimental	(mA)
Scottky	0.0231	-	0
Ideal PWM 24V	0.0463	0.05	-8.31
Control PWM 24V	-	0.10	-99.15

the time constant. For the PWM solution with 24 V the time to reduce the current level to zero is 0.0463 ms while using 50 V diodes results in a time of 0.0231 ms. This is about half the time, but represent a very small portion of the time which the valve is open. A typical open time interval for the solenoid is $T_{open} = 8$ ms. The response time of the valve is about $\tau \approx 0.3$ ms. The effect on the tail will be examined by experiments to ascertain whether this control strategy works or not.

The complexity of the control configurations is another aspect that needs to be inspected. For a microdosing system it is important to have a fast response circuit as a spike and hold circuit. There are many solutions for such a driver, but the main difference discussed here is how to discharge the energy in the coil. Regardless of the solution chosen the PWM approach will result in fewer components than the schottky diode solution for a valve matrix. This is achieved by using a half H-bridge for all the valves, in addition to a half H-bridge that is common for all valves. That way two diodes for each valve is avoided and only two more transistors are needed. This makes the circuit less complex and easier to control.

Another advantage for the PWM control is the need of just one voltage source. A solution to remove one voltage source for the spike and hold driver is to use a voltage regulator to produce the hold voltage. The problem here is that when the number of valves increases, several regulators are needed as the current becomes larger. The reduction in components influences the cost of the final PCB as well. Another significant advantage is that the PWM solution is more flexible. If the valve is replaced, the only requirement for the new valve is that the spike voltage needed does not exceed the initial design specifications. Thus only software adjustment is required instead of modifying the circuit.

B. Negative spike time

The negative spike time of the PWM circuit must be carefully chosen. The electrical characteristics of the valve are presented in Table I. Using the first order response of the RL circuit as in Equation 3 the exact time can be derived. The timing and currents are presented in Table II. As calculated before the time for closing the valve under ideal conditions is 0.0463 ms. In practice the resolution in time may have to be limited. The important part is to have the current close to zero so the plunger is not activated again. The residue current will be burnt off over the diodes in the transistors.

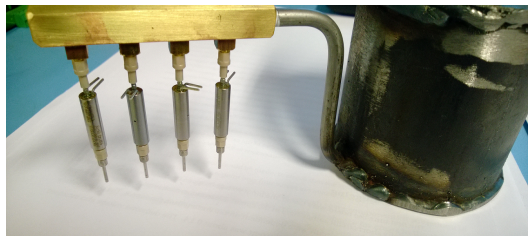

Fig. 7. Experimental setup with valves and pressurized liquid container, for early experiments with droplet formation, as shown in Figure 6 and 8.

To test how the closing time influences the droplet, a test rig was set up. This was done with a black and white high speed camera, PROMON 501 from AOS Technolohies. To provide sufficient light for shooting with 1000 fps a LED panel was placed behind the nozzle pointing directly at the camera. The valve was operated by the PCB controlled from a computer. The rest of the setup consisted of a pressurized liquid container with water and tubing, as shown in Figure 7. The pressure was set to 0.4 bar, which produces droplets with an initial velocity of about 4 m/s. In this experiment regular tap water was used.

The main test was to see how the time resolution affected the droplet, initially two spike times were chosen. The requirement was that both times should be realistic regarding how the system will be programmed for the field. For the first test the spike duration was set to 0.1 ms, while for the second test it was 0.05 ms. The calculations represented in Table II shows the theoretical times for discharging the energy in the coil and the theoretical residue current for the experimental times. A spike duration of 0.1 ms should result in a current of -99.15 mA, while 0.05 ms result in an undershoot of -8.31 mA. A current larger than 87.5 mA is enough to hold the valve open. Thus the larger spike duration may cause the valve to start opening again. In this test it was of interest to see how such an undershoot affects the droplets properties. An undershoot of -8.31 mA should not be enough to actuate the valve at all, thus the difference should be observable.

C. Results

The experimental setup was designed with one valve with a pressure of 0.4 bar. The spike voltage was set to 24 V and the hold voltage to 3.96 V. The only difference in the two tests was the negative spike duration. Figure 6a shows the end of the droplet using a spike duration of 0.1 ms, while Figure 6b shows the end of the droplet when using 0.05 ms for the spike duration. For the first test a thin secondary tail is observable

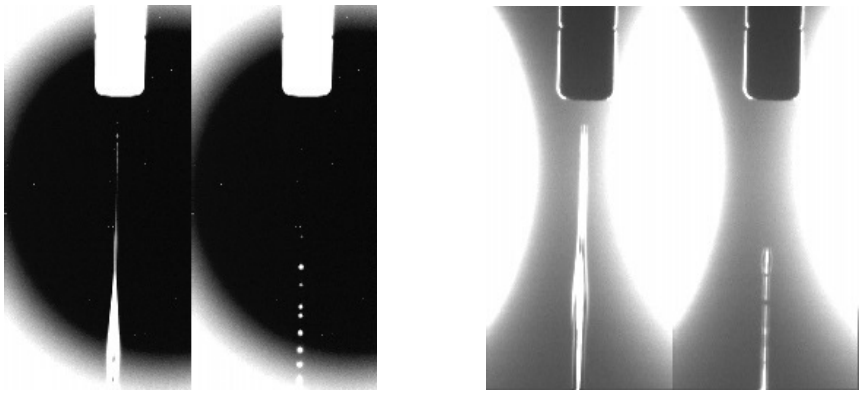

(a) Test 1, tail, $U_{hold} = 3.96\,\text{V}$, $t_{spike} = 0.1\,\text{ms}$ (b) Test 2, tail, $U_{hold} = 3.96\,\text{V}$, $t_{spike} = 0.05\,\text{ms}$

Fig. 6. High speed footage of the droplet tail with a pressure of 0.4 bar, 1000 fps. Figure (a) show the extra tail resulting from the reverse spike overshoot.

before it breaks into many small satellite droplets. This is however avoided in the second test. Common for both tests is that the filament is beginning to break up. The length of the filament makes it unstable as described previously. The breakup of the tail in Figure 6a is similar to filament breakup, but because it is so much thinner than the filament it breaks up faster and to smaller droplets.

Under ideal conditions the satellite droplets will overtake the main droplet, but in practice, the robot will be moving and the presence of wind may affect the satellites differently than the main droplet. Thus the filament breakup and the tail breakup should minimized. This is to reduce the possibility of satellite droplets not merging with the main droplet and missing the target.

III. Discussion

For the autonomous weed control application to work, the DOD system must be very accurate. The accuracy does not solely depend on target precision, but the presence of satellite droplets and their behavior. It is crucial that the satellites coalesce with the main droplet, or that they both hit the same spot. Therefore the droplets tail should be minimized, as the tail will split up in much smaller droplets than the filament.

The usual method for driving solenoid valves in this kind of a application is by a spike and hold driver circuit, with two diodes in reverse series to discharge the energy in the coil. However the described method is based on a full H-bridge, for PWM control. The maximum voltage is chosen as the spike voltage, thus a long spike will open the valve before the PWM control limits the voltage to the hold value. When closing the valve the energy in the coil is discharged with a significant negative spike. The spike time has been calculated using the first order response of a RL-circuit when the inductance and resistance of the valve is known.

Tests of how this control strategy performs confirms the theory, as a long spike time resulted in a thin secondary tail,

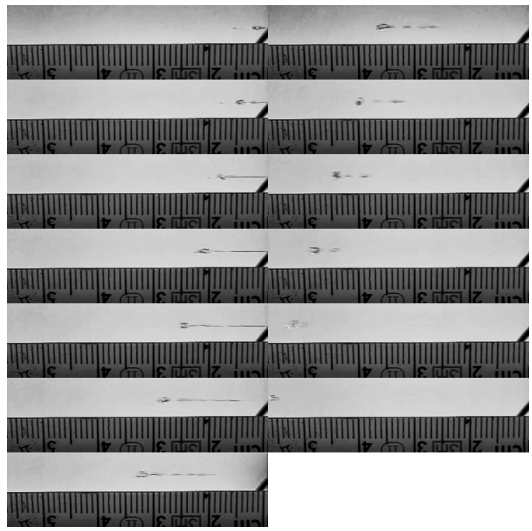

Fig. 8. 1200 fps image sequence of a water droplet with initial velocity of 4 m/s (illustration taken from ref. [17]). The filament first break up to satellites, which then drift in the wake of the main droplet, and join the droplet.

while a spike time of appropriate length avoided this. The long spike time started to actuate the plunger when it should be closed, but the small undershoot for the appropriate time did not actuate the plunger at all.

Thus the spike time of 0.05 ms is close enough to the theoretical time for closing the valve. A possible solution for decreasing the filament break up in this application is to use a larger nozzle and shorter on-time. That way the circumference-to-length ratio is increased for droplets with the same volume. The time before the filament breaks up is thus increased,

and the filament will break up in larger and fewer droplets. Manipulation of the liquid to increase the Ohnesorge number is another solution for decreasing filament breakup. However, this is likely to influence the Weber number and stability of the droplet.

The main disadvantage of PWM control is the increase in time for discharging the energy, as the diodes can be chosen with a higher voltage level. This is not as easy with the PWM solution, as the higher the voltage, the more robust the components must be. This is due to the increase in voltage when using only one source will influence the robustness of the components, especially the transistors, used in the PWM control.

However there have not been observed any significant negative effects of the slower closing of the PWM solution compared to the Schottky closing. Thus an increase in the voltage is not necessary.

Breakup of the filament was observed in the tests, but this will occur regardless of the control strategy, and might be decreased with increased nozzle diameter. The theory and calculations regarding the valve control was confirmed by the experimental setup. The tail was avoided although the closing time is increased compared to the diode solution, thus the advantages of the PWM control strategy may be exploited. This includes a more flexible design with regards to the valves and fewer components are needed for the circuit. The main focus is to make sure that the negative spike time does not undershoot too much as this will create a tail that should be minimized.

Theory regarding the Ohnesorge number and filament breakup finds that a liquid with low surface tension and high viscosity reduces the filament breakup as the Ohnesorge number increases. However when the droplets have to travel a significant distance before hitting their targets it is important that the droplets do not disintegrate. The increased stability of the droplet leads to more satellites due to increased filament breakup. Under ideal conditions the satellites from filament breakup will merge with the main droplet, but this may not be the case in the field. Clearly the best solution would be to avoid filament breakup while maintaining a stable droplet in air. A compromise between the Ohnesorge number and the Weber number is of importance when shooting droplets a significant distance. This is due to the requirement of a stable droplet throughout the whole flight, while trying to minimize the filament breakup.

IV. Conclusion

A valve controller has been developed for drop on demand weed control, using a full H-bridge design and PWM voltage regulation to generate the spike and hold voltages. In contrast with common design practices with solenoid drives, we have not included the discharge diodes. The solenoid discharge is instead done by applying a reverse voltage to the solenoid.

The timing of the reverse voltage has to be calculated using the solenoid inductance given from the datasheet. If the reverse spike is held longer the solenoid may reopen and dispense a secondary tail, which will create additional satellite droplets. If the reverse spike is not long enough the residue current will discharge over the protective diodes in the H-bridge drivers.

The design results in fewer components per solenoid, but demands accurate timing of the reverse voltage spike. The PWM allows for arbitrary spike and hold voltages up to the supply voltage, which for this project has been 24 V. The experiments also illustrate the filament breakup and its connection with the Ohnesorge number, while the Weber number is essential to the stability of droplets in flight.

Acknowledgment

The authors would like to thank Alexander Klungerbo for conducting experiments to conclude what parameters affect the distance traveled for droplets.

References

[1] S. Christensen, H. T. Søgaard, P. Kudsk, M. Nørremark, I. Lund, E. S. Nadimi, and R. Jørgensen, "Site-specific weed control technologies," *Weed Research*, vol. 49, no. 3, 2009.

[2] W. Lee, D. Slaughter, and D. Giles, "Robotic weed control system for tomatoes," *Precision Agriculture*, vol. 1, no. 1, 1999.

[3] A. T. Nieuwenhuizen, "Automated detection and control of volunteer potato plants," Ph.D. dissertation, Wageningen University, 6 2009.

[4] H. Midtiby, S. K. Mathiassen, K. Andersson, and R. Jørgensen, "Performance evaluation of a crop / weed discriminating microsprayer," *Computers and Electronics in Agriculture*, vol. 77, no. 1, pp. 35–40, 2011.

[5] D. Slaughter, D. Giles, and D. Downey, "Autonomous robotic weed control systems: A review," *Computers and Electronics in Agriculture*, vol. 61, no. 1, pp. 63–78, 2008.

[6] T. Utstumo and J. T. Gravdahl, "Implementation and comparison of attitude estimation methods for agricultural robotics," in *proceedings of the 4th IFAC Conference on Modelling and Control in Agriculture, Horticulture and Post Harvest Industry (AgriControl 2013)*, Aalto University, School of Electrical Engineering, Espoo, Finland, August 28-30, 2013.

[7] P. K. Jensen, I. Lund, and D. Nuyttens, "Spray liquid distribution and biological efficacy of commercially available nozzles used for precision weed control," *Biosystems engineering*, vol. 116, no. 4, pp. 316–325, 2013.

[8] J. Watton, *FLUID POWER SYSTEMS: Modeling, simulation, analog and microcomputer controll*. Prentice Hall International (UK) Ltd., 1989.

[9] D. Vadillo, T. Tuladhar, A. Mulji, S. Jung, S. Hoath, and M. Mackley, "Evaluation of the inkjet fluid's performance using the "cambridge trimaster" filament stretch and break-up device," *Journal of Rheology*, vol. 54, no. 2, pp. 261–282, 2010.

[10] S. D. Hoath, S. Jung, and I. M. Hutchings, "A simple criterion for filament break-up in drop-on-demand inkjet printing," *Physics of Fluids*, vol. 25, no. 2, 2013.

[11] H. Dong, W. W. Carr, and J. F. Morris, "An experimental study of drop-on-demand drop formation," *Physics of Fluids*, vol. 18, no. 7, p. 072102, 2006.

[12] J. Castrejón-Pita, G. Martin, S. Hoath, and I. Hutchings, "A simple large-scale droplet generator for studies of inkjet printing," *Review of Scientific Instruments*, vol. 79, no. 7, p. 075108, 2008.

[13] J. Eggers and E. Villermaux, "Physics of liquid jets," *Reports on progress in physics*, vol. 71, no. 3, p. 036601, 2008.

[14] A. Wierzba, "Deformation and breakup of liquid drops in a gas stream at nearly critical weber numbers," *Experiments in Fluids*, vol. 9, no. 1-2, pp. 59–64, 1990.

[15] A. D. Baylis, "Why glyphosate is a global herbicide: strengths, weaknesses and prospects," *Pest Management Science*, vol. 56, no. 4, pp. 299–308, 2000.

[16] H. T. Søgaard, I. Lund, and E. Graglia, "Real-time application of herbicides in seed lines by computer vision and micro-spray system," in *American Society of Agricultural and Biological Engineers, ASABE Annual International Meeting*, 2006.

[17] A. T. Klungerbo, "Drop-on-demand i presisjonsjordbruk," Master's thesis, NTNU, 2013.

2.3 Conference Paper - Non-linear Model Predictive Control

T. Utstumo, T. Berge, and J. Gravdahl. Non-linear model predictive control for constrained robot navigation in row crops. In *Proceedings of the IEEE International Conference on Industrial Technology, Sevilla, Spain, March 17-19, 2015.*

Non-linear Model Predictive Control for constrained robot navigation in row crops

Trygve Utstumo*† , Therese W. Berge‡ , Jan Tommy Gravdahl*
*Department of Engineering Cybernetics, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway.
‡Bioforsk - Norwegian Institute for Agricultural and Environmental Research, Plant Health and Plant Protection Division, NO-1430 Ås, Norway. †Adigo AS, NO-1405 Langhus, Norway. E-mail: trygve.utstumo@adigo.no

***Abstract*—Vehicles which operate in agricultural row crops, need to strictly follow the established wheel tracks. Errors in navigation where the robot sways of its path with one or more wheels may damage the crop plants.**

The specific focus of this paper is on an agricultural robot operation in row cultures. The robot performs machine vision detecting weeds within the crop rows and treats the weeds by high precision drop-on-demand application of herbicide.

The navigation controller of the robot needs to follow the established wheel tracks and minimize the camera system offset from the seed row. The problem has been formulated as a Nonlinear Model Predictive Control (NMPC) problem with the objective of keeping the vision modules centered over the seed rows, and constraining the wheel motion to the defined wheel tracks.

The system and optimization problem has been implemented in Python using the Casadi framework. The implementation has been evaluated through simulations of the system, and compared with a PD controller. The NMPC approach display advantages and better performance when facing the path constraints of operating in row crops.

I. Introduction

An agricultural robot for weed control in row-crops is under development. The weed control is done by drop-on-demand herbicide application, where the weed is first identified by a camera system and then targeted by a drop-on-demand nozzle array. The focus of the project is on computer vision, robot integration and navigation [1].

The project share its ambition of autonomous weed control with many other research projects on robotic weed control. The review by Slaughter 2008 presents an overview of the field [2].

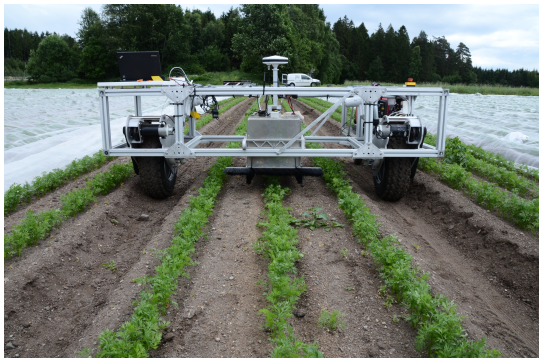

Fig. 1. The wheeled mobile robot developed for weed control in row crops.

The robotic platform has two front wheels with electro motors and two rear castor wheels, Figure 1. The robot has a monocular downward facing RGB camera primarily used for two purposes: Classification of crop and weed plants as part of the spray-on-demand system [3], and for visual odometry measurements as input to the localization filter and crop-row detection.

The visual crop-row estimate will be fused with a forward looking camera for crop-row detection. This information forms the input to the row following controller.

Crop-row following has been well explored within the field of agricultural robotics, and similar applications can be found in [4].

Application of Non-linear Model predictive control (NMPC) in agriculture has been described in [5], and [6] where an actuated trailed implement is controlled to follow field rows.

To the authors knowledge there has not been publications on NMPC applications for robotics in row crops with specific constraints on the wheels to minimize crop damage.

A. Minimizing crop damage

A review of autonomous navigation in agriculture is presented in [7], where the performance of various approaches are compared.

A study of guiding principles in design of robots for agriculture revealed that the most important factor for the end users were minimizing crop damage [8].

Other controllers described in the review article [7], does not incorporate the constraints on navigation directly. This motivates our research on using an NMPC based design where the path constraints can be directly implemented in the controller, adding an additional barrier against damaging the crop.

B. Wheel tracks in row cultures

The production method for most vegetables is row cultures, where the plant rows are set with a fixed intermediate distance between the wheel tracks. The centre to centre distance of the wheel tracks are typically 1.65 m to 1.80 m in European agriculture, 4.

The robots and vehicles operating in the field are restricted to these wheel tracks to avoid damage of the crop, as illustrated in Figure 2.

II. Modelling and simulation

The robot can be modelled as a unicycle-like robot assuming non-slip conditions. Differentially steered robot designs are

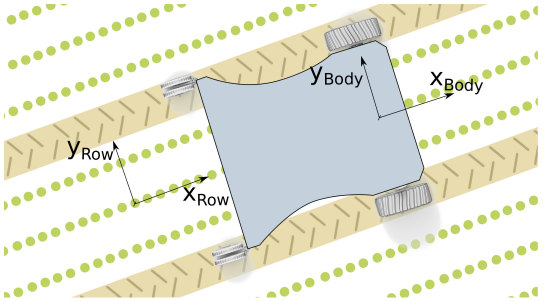

Fig. 2. The robot illustrated has a path-following algorithm seeking to maintain the camera module centered over the seed row. A sideways error will lead the controller to turn to correct the error. A small turn can easily lead the rear castor wheels to enter the sow bed and damage the crop.

very common in many applications, and numerous publications present kinematic and dynamic models for this class of robots.

Industrial motor controllers and commercial robot platforms normally provide the control inputs as linear and angular velocity set-points, not as torque or voltage set-points. A dynamic model with the motor controllers included and velocities as inputs will be advantageous when it comes to implementing the actual robot.

Such a model has been presented [9] and the formulation has been used as the basis for modelling and simulation in this paper. The robot in our project differs from the model schematic presented in [9]: In contrast to the schematic, this robot has the differential drive wheels in front, and trailing castor wheels in rear. Consistency with the schematic is maintained by describing relevant parameters with negative sign, as shown in Figure 3. The camera and spray unit has been mounted centrally at the virtual wheel axis in field experiments, which leaves the tracking point, h, at $a = 0$.

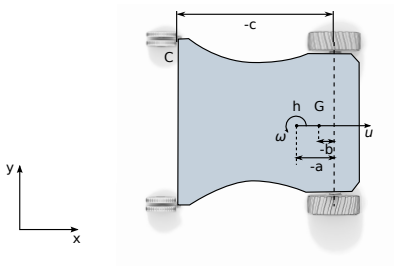

Fig. 3. The unicycle model from [9] can be applied to the robot with differential drive and passive castor wheels. The robot tracking point is located in h, which is the origin of the robot frame. The point h is a distance $-a$ behind the virtual front wheel axis. G is the center of gravity, $-b$ behind the virtual front wheel axis. The track width of the robot is d, and the rear castor wheels, C, is at a distance $-c$ from the virtual wheel axis. The robot has a forward velocity, u, and an angular velocity ω.

The dynamic model is written as [9]:

$$\dot{\mathbf{x}} = f(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\theta}) + \boldsymbol{\delta} \tag{1}$$

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \\ \dot{u} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} u\cos\psi - a\omega\sin\psi \\ u\sin\psi + a\omega\cos\psi \\ \omega \\ \frac{\theta_3}{\theta_1}\omega^2 - \frac{\theta_4}{\theta_1}u \\ -\frac{\theta_5}{\theta_2}u\omega - \frac{\theta_6}{\theta_2}\omega \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{1}{\theta_1} & 0 \\ 0 & \frac{1}{\theta_2} \end{bmatrix} \begin{bmatrix} u_{ref} \\ \omega_{ref} \end{bmatrix} + \begin{bmatrix} \delta_x \\ \delta_y \\ 0 \\ \delta_u \\ \delta_\omega \end{bmatrix} \tag{2}$$

where (x, y) is the position in the base frame, and ψ is the robot orientation or heading, (u, ω) are the linear and angular velocities and u_{ref} and ω_{ref} are the input signals of the system: linear and angular velocity.

The two vectors are identified model parameters, and parametric uncertainties. The uncertainty vector

$$\boldsymbol{\delta} = [\ \delta_x\ \delta_y\ 0\ \delta_u\ \delta_\omega\ ]^T \tag{3}$$

represent slip velocities and effects of uneven ground with its first two elements. The two last elements are functions of physical parameters as mass, inertia, wheel and tire diameters, parameters of the motors, and wheel ground interaction forces.

The parameters in the vector $\boldsymbol{\theta}$ are functions of the robots physical parameters, such as its mass m, inertia I_Z about G, the electrical resistance R_a of the DC motors with motor constant k_a, the friction coefficient B_e, reduction gear inertia I_e, radius of the wheel r, nominal radius of the tire R_t, and the distances b and d. The model assumes a PD motor control loop with gains $k_{PT} > 0$, k_{PR}, k_{DT} and k_{DR}. The equations for $\boldsymbol{\theta}$ were presented in [9], and methods for online parameter identification and an adaptive controller has been presented [10]. The parameter equations are reproduced here for reference:

$$\begin{aligned}
\theta_1 &= \left(\frac{R_a}{k_a}(mR_t r + 2I_e) + 2r_k DT\right)(2rk_{PT}) \\
\theta_2 &= \left(\frac{R_a}{k_a}(I_e d^2 + 2R_t r(I_z + mb^2)) + 2rdk_{DR}\right)/(2rdk_{PR}) \\
\theta_3 &= \frac{R_a}{k_a} mbR_t/(2k_{PT}) \\
\theta_4 &= \frac{R_a}{k_a}\left(\frac{k_a k_b}{R_a} + B_e\right)(rk_{PT}) + 1 \\
\theta_5 &= \frac{R_a}{k_a} mbR_t/(dk_{PR}) \\
\theta_6 &= \frac{R_a}{k_a}\left(\frac{k_a k_b}{R_a} + B_e\right) d/(2rk_{PR}) + 1
\end{aligned} \tag{4}$$

For the simulations in this paper, we have used a set of parameters identified from indoor experiments with a robot comparable to our:

$$\boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \\ \theta_5 \\ \theta_6 \end{bmatrix} = \begin{bmatrix} 0.19 \\ 0.14 \\ 0.02 \\ 1.00 \\ 0.16 \\ 1.00 \end{bmatrix} \tag{5}$$

The disturbances in the $\boldsymbol{\delta}$ have been left at zero for the simulations shown here, while it does provide a possible input for perturbing the system in simulation.

The model has been implemented in Python using the Casadi framework [11] and a Runge Kutta 4 integrator scheme

for time simulation.

For this implementation, the robot is assumed to navigate relative to a local coordinate frame aligned with the crop row, as illustrated in Figure 2.

A. Reference PD controller

A PD controller has been implemented for reference and comparison with the NMPC controller. The controller has been tuned to stay within the constraints when operating with a tracking error less than half the allowed region. That is $y \in [-\frac{\tau_w}{2}, \frac{\tau_w}{2}]$.

The controller has been implemented as a P controller driving both the sideways tracking error and the heading to zero. For small heading angles, the time-derivative of y approximates to the heading ψ, and the controller can be though of as a PD controller. The velocity reference is constant.

$$v_{ref} = 0.3\text{m}\,\text{s}^{-1} \tag{6}$$

$$\omega_{refPD} = -k_p y - k_d \psi \tag{7}$$

where the tuning constants has been set to:

$$k_p = 0.70 \quad k_d = 0.49 \tag{8}$$

III. Formulating the optimal control problem

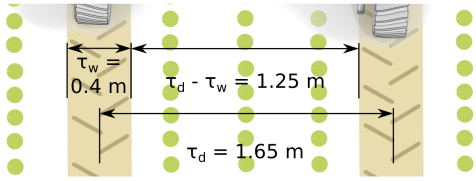

Fig. 4. The feasible wheel track area where a robot can operate without damaging the crop, based upon experience from field test during spring 2014.

The objective for the robot is to maintain a constant velocity while centring the camera systems over each crop row. At the same time, the rear castor wheels should be constrained to the wheel tracks, not to damage the crop.

The position of the rear castor wheels can be expressed in vector notation in the BODY frame, $\mathbf{w}^B$, and rotated into the ROW frame, $\mathbf{w}^R$, to find the wheels' positions. The calculation for the left rear castor wheel becomes:

$$\mathbf{w}_l^B = \begin{bmatrix} c \\ \frac{d}{2} \end{bmatrix} \tag{9}$$

$$\mathbf{w}_l^R = \mathbf{R}(\psi)\mathbf{w}_l^B + \begin{bmatrix} x \\ y \end{bmatrix}^R \tag{10}$$

$$= \begin{bmatrix} \cos(\psi) & \sin(\psi) \\ -\sin(\psi) & \cos(\psi) \end{bmatrix} \begin{bmatrix} c \\ \frac{d}{2} \end{bmatrix} + \begin{bmatrix} x \\ y \end{bmatrix}^R \tag{11}$$

$$= \begin{bmatrix} c\cos(\psi) + \frac{d}{2}\sin(\psi) + x \\ -c\sin(\psi) + \frac{d}{2}\cos(\psi) + y \end{bmatrix} \tag{12}$$

Only the y-component is of interest in the ROW frame, and the constraint for the left castor wheel can be written with respect to the distance between two tracks τ_d and the width of each wheel track τ_w, illustrated in Figure 4, as:

$$\frac{-\tau_d - \tau_w}{2} \leq -c\sin(\psi) + \frac{d}{2}\cos(\psi) + y \leq \frac{\tau_d - \tau_w}{2} \tag{13}$$

If the robot track width and the row track width are equal, $\tau_d = d$, the constraint will be symmetric for the left and right wheel, for small deviations of y, and it will be sufficient to consider one wheel constraint.

The quadratic cost function describes the robots deviation from the row centre line, and deviation from the reference velocity:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \int_{t_0}^{T} (y - y_{row})^2 + (u - u_{setpoint})^2 dt \\ \text{subject to} \quad & \dot{\mathbf{x}} = f(\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\theta}) \\ & \frac{-\tau_d - \tau_w}{2} \leq c\sin(\psi) + \frac{d}{2}\cos(\psi) + y \leq \frac{\tau_d - \tau_w}{2} \end{aligned} \tag{14}$$

A. Inverse proportional limit on the feasible control space

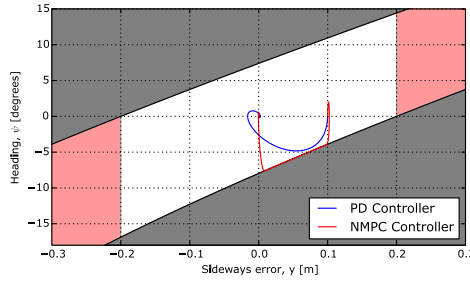

Fig. 5. The allowed range of steering angles ψ as a function of the tracking error. The regions marked in red, are outside the track width. If the robot enters the red region the front wheels are outside the wheel tracks, and the robot will diverge from the row. The NMPC controller with constraints can only be used within the white region. Two trajectories starting with a sideways error of 0.1 meter are shown in the plot, one with the NMPC controller and one with the PD controller. Note that the NMPC controller closely follows the constraint in steering angle.

The kinematics of the system leaves us in a special case when we apply the constraint to the rear castor wheel. An analogy of this scenario is driving a forklift alongside a wall. As the forklift approaches the wall, it will be increasingly impossible to turn away from the wall, as its rear end needs to swing out to turn.

The robot and the wheel tracks scenario leave us with the same situation. The allowed and stable region of steering, or vehicle angles, are shown in Figure 5.

A typical path following algorithm using a PID-type controller, would proportionally increase its control input as the error increases. The feasible control region for this problem leaves us with a set of controls which does not fit well with a proportional control strategy: As the tracking error increase, say to 0.1 m, the maximum vehicle angle to counter the error is reduced by a factor of two. This limitation on the feasible state

space can be taken as an argument for implementing a NMPC controller to utilize the limited control space optimally.

IV. Implementation

The system has been implemented in Python using the Casadi framework [11]. The chosen method for solving the optimization problem (14), is a direct multiple shooting method.

The infinite horizon problem is reformulated to a finite discretized nonlinear problem. The time horizon has been limited to 5 seconds, and the control input has been discretized to N = 20 steps.

The second step is to parameterize the system of differential algebraic equations (DAE) by *multiple shooting*. We also exploit the quadratic form of the cost function, by using the Gauss-Newton method to solve the sequential quadratic program (SQP). The implementation follow the details given in [12]. A more advanced implementation applied to automobile collision avoidance, follows the same implementation approach [13].

The QP problem is solved at each iteration by using the IPOPT library [14]. The QP problem is initialized with the last state at every iteration as an aid to the QP solver.

The system has been simulated without disturbances, $\boldsymbol{\delta} = \mathbf{0}$, from various perturbed initial conditions to investigate the system behaviour. The target velocity has been set to $u_{setpoint} = 0.3\,\mathrm{m\,s^{-1}}$ and the crop row is at the origin of the coordinate system, $y_{row} = 0$

V. Results

The NMPC controller has tested by starting the system in several different initial conditions. Figure 7 show the system recovering from a small tracking error, with the NMPC controller and the reference PD controller. Note the increasing curve of ψ as the robot gains increasingly more headroom to navigate. Figure 8 show the states of the NMPC controller in the time domain. This scenario is also illustrated in the time domain in Figure 8 and with respect to the constraints in Figure 5.

The NMPC controller use two iterations to converge to the optimal trajectory, and steers the robot in the opposite direction with the first control input, as shown in Figure 6. Figure 9 and 10 show the system from an initial condition close to the boundary constraint. The recovery of the robot is significantly slower, and the amplitude of ψ is limited by the path constraint.

In addition to the displayed figures, the NMPC controller has been initialized in several infeasible initial conditions. The trajectory then diverge from the desired trajectory for as long as the path constraint can be met. These scenarios break the assumption of small y deviations, which the symmetry assumption of the path constraint relies on.

VI. Discussion

Looking at Figure 7 it is interesting to see the increasing correction in heading, as the robot approach the desired trajectory. This is the opposite of behaviour of a PID based controller, as illustrated by the reference PD controller. The NMPC controller maintains the rear castor wheels on the path constraint, until the target is reached. The PD controller is not able to converge as quickly without violating the constraints.

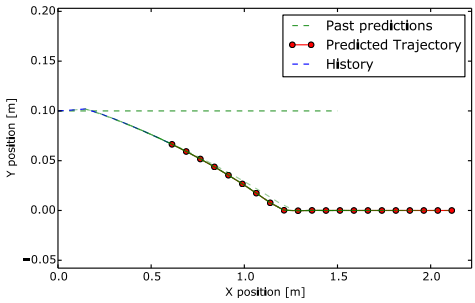

Fig. 6. With the system initialized at $[x, y] = [0\ 0.1]$ the NMPC controller use the first to iterations to converge to the optimal solution. The first prediction can be seen as the green dotted line, and the current prediction is shown as the red dots.

Figures 9 and 10 illustrate the characteristic of the constrained controller: If the tracking error is sufficiently large, the controller is left with close to no headroom for navigation, and the convergence is slow. If the front wheels are at, or outside the boundary the solution will diverge. The reference PD controller is outside its intended region of operation, and it violates the constraint in heading angle.

For operation outside the feasible region, the constraint should be reformulated without the symmetry assumption for the castor wheel constraints. This assumption relies on small deviations in angle and lateral position, and will be increasingly inaccurate outside the wheel tracks.

Another control strategy should be implemented to handle operation outside the feasible region. Some alternative solutions may be:

- Drop the velocity reference from the cost function, and add a quadratic term in the heading, ψ. This will allow the robot to reverse back into the wheel tracks and correct it's heading and position.
- Drop or expand the path constraint to allow the robot to quickly get back to the path, and accept damage to the crop for a short section.
- Switch to a different type of controller with desired dynamics and let that bring the robot back into the feasible region.

In a real field implementation these strategies need to be evaluated with practical considerations in mind.

The implementation of the inequality constraints are relatively straight forward, within the multiple shooting method. One can easily imagine applying such a strategy to vehicles and robots with more complex kinematics or environments, an example pointing in that direction is presented in [5].

The oscillations in ω_{ref} as the system reaches the trajectory is caused by the NMPC controller exploiting its knowledge of the motor controller and system dynamics to maximize the system response. This behaviour may be problematic when faced with inaccuracies in the estimated parameters $\boldsymbol{\theta}$. For example: If the robot is significantly lighter than estimated; this may lead to oscillations in the control. A term to dampen control inputs can be considered in the cost function.

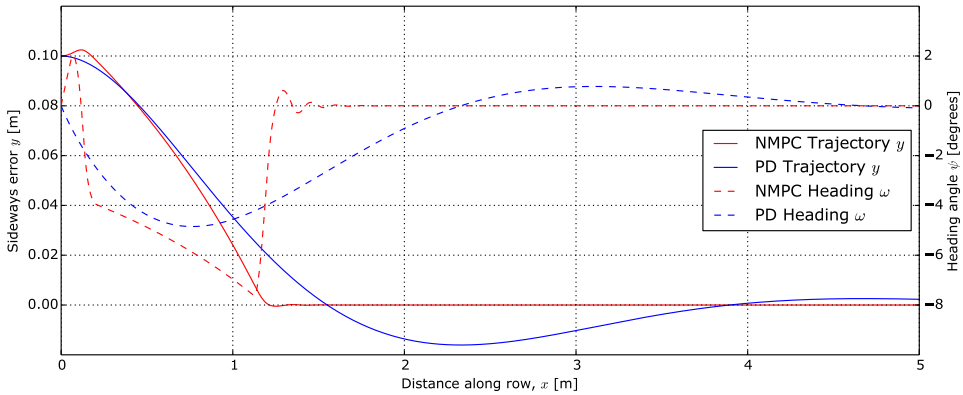

Fig. 7. A comparison of the reference PD controller with the NMPC controller, with sideways error and heading angle as the robot moves along the row. The robot was initialized at $[x, y] = [0\ 0.1]$. The NMPC controller follow the constraint in heading angle, and converge faster than the PD controller. The same trajectories are also illustrated in figure 5.

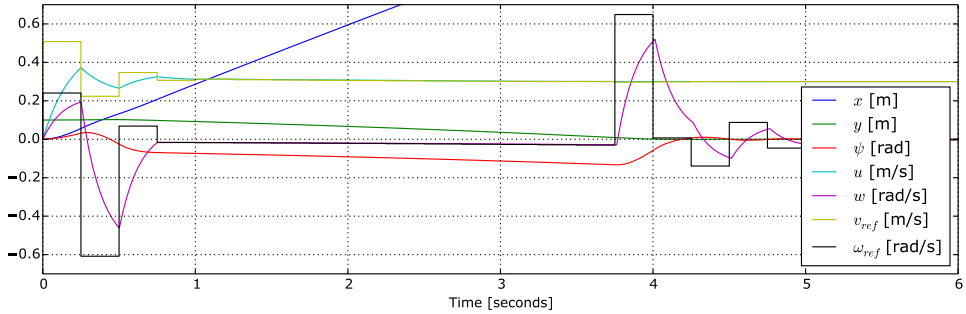

Fig. 8. The robot states and NMPC control inputs shown in the time domain, when initialized at $[x, y] = [0\ 0.1]$. After the robot has accelerated, the rear castor wheel follow the edge of the constraint until the y target is reached just before $t = 4s$. Note that the robot heading, ψ, increase in amplitude, as more headroom for navigation is available. When the target is reached it quickly steers the heading back to zero. The NMPC controller oscillates here to maximize the response from the motor controllers.

The computational performance of the algorithm has not been systematically evaluated, but the run-time is consistently below 50 ms per iteration. Further optimizations may be implemented for real-time applications, and there exists code generation tools within the Casadi framework, which may be useful, [12].

VII. Conclusion

A crop row following controller has been formulated with special focus on constraining the motion of the trailing castor wheels to the wheel tracks. The implementation uses Nonlinear Model Predictive Control (NMPC) with a direct multiple shooting method, and a Gauss-Newton quadratic objective.

The implementation is flexible with regards to expressing the constraints and it can be suitable for real-time implementations. The controller needs to be expanded to operate on a global frame with an arbitrary model of the crop row, and the implementation needs to be verified in experiments.

The kinematic limitations of a trailed castor wheel with path constraints has been investigated. The limited range of feasible control inputs can be an argument for applying constrained model based control, such as this NMPC application, over other control methods. An NMPC approach will better utilize the available control room, and in row crops the NMPC controller can provide safety against damaging the crop.

Acknowledgment

The authors would like to thank the Casadi open source community and Joel Andersson and Greg Horn for their useful input through the Casadi online user forum. The authors would also like to thank Lars Imsland for his valuable input on the NMPC implementation.

References

[1] T. Utstumo and J. T. Gravdahl, "Implementation and comparison of attitude estimation methods for agricultural robotics," in *Agricontrol*, vol. 4, no. 1. IFAC, 2013, pp. 52–57.

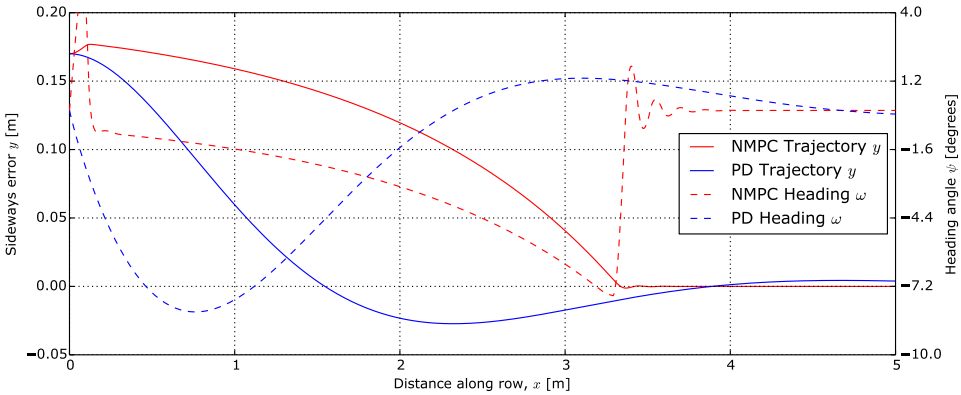

Fig. 9. The position xy plot for the robot, when initialized at $[x, y] = [0\ 0.17]$. This initial condition is close to the boundary, and the space available for navigation is very tight. The controller is slow to reach it's target. The reference PD controller is now far outside its intended region of operation and it goes far outside the constraints. In an actual application, this would result in damage to the neighbouring crop row.

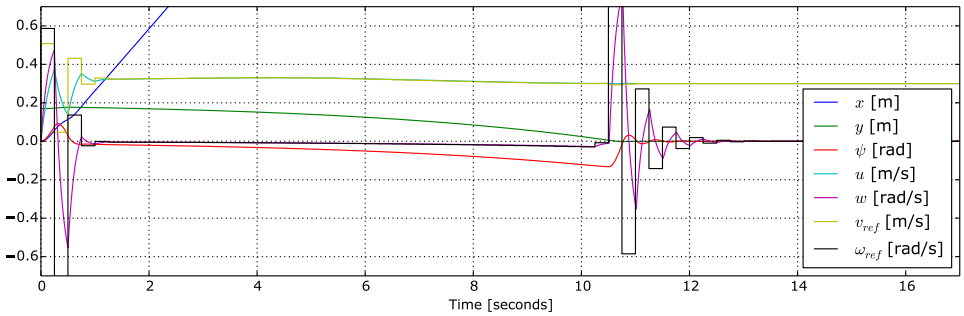

Fig. 10. The robot is initialized at $[x, y] = [0\ 0.17]$, close to the constraint. Note especially the limited amplitude of ψ in the first seconds of this trajectory. As before the NMPC controller use two iterations to converge to the optimal solution, where it steers the robot in the opposite direction.

[2] D. Slaughter, D. Giles, and D. Downey, "Autonomous robotic weed control systems: A review," *Computers and Electronics in Agriculture*, vol. 61, no. 1, pp. 63–78, 2008.

[3] F. Urdal, T. Utstumo, S. A. Ellingsen, J. K. Vatne, and T. Gravdahl, "Design and control of precision drop-on-demand herbicide application in agricultural robotics," in *The 13th international Conference on Control, Automation, Robotics and Vision, Singapore*. IEEE, 2014.

[4] F. Dong, W. Heinemann, and R. Kasper, "Development of a row guidance system for an autonomous robot for white asparagus harvesting," *Computers and Electronics in Agriculture*, vol. 79, no. 2, pp. 216–225, 2011.

[5] T. Kraus, H. J. Ferreau, E. Kayacan, H. Ramon, J. De Baerdemaeker, M. Diehl, and W. Saeys, "Moving horizon estimation and nonlinear model predictive control for autonomous agricultural vehicles," *Computers and Electronics in Agriculture*, vol. 98, pp. 25–33, 2013.

[6] J. Backman, T. Oksanen, and A. Visala, "Navigation system for agricultural machines: Nonlinear model predictive path tracking," *Computers and Electronics in Agriculture*, vol. 82, pp. 32–43, 2012.

[7] H. Mousazadeh, "A technical review on navigation systems of agricultural autonomous off-road vehicles," *Journal of Terramechanics*, vol. 50, no. 3, pp. 211–232, 2013.

[8] C. Sørensen, R. N. Jørgensen, J. Maagaard, K. K. Bertelsen, L. Dalgaard, and M. Nørremark, "Conceptual and user-centric design guidelines for a plant nursing robot," *Biosystems Engineering*, vol. 105, no. 1, pp. 119–129, 2010.

[9] D. La Cruz *et al.*, "Dynamic modeling and centralized formation control of mobile robots," in *IECON 2006-32nd Annual Conference on IEEE Industrial Electronics*, 2006, pp. 3880–3885.

[10] F. N. Martins, W. C. Celeste, R. Carelli, M. Sarcinelli-Filho, and T. F. Bastos-Filho, "An adaptive dynamic controller for autonomous mobile robot trajectory tracking," *Control Engineering Practice*, vol. 16, no. 11, pp. 1354–1363, 2008.

[11] J. Andersson, "A General-Purpose Software Framework for Dynamic Optimization," PhD thesis, Arenberg Doctoral School, KU Leuven, Department of Electrical Engineering (ESAT/SCD) and Optimization in Engineering Center, Kasteelpark Arenberg 10, 3001-Heverlee, Belgium, October 2013.

[12] M. Diehl, H. G. Bock, J. P. Schlöder, R. Findeisen, Z. Nagy, and F. Allgöwer, "Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations," *Journal of Process Control*, vol. 12, no. 4, pp. 577–585, 2002.

[13] J. V. Frasch, A. Gray, M. Zanon, H. J. Ferreau, S. Sager, F. Borrelli, and M. Diehl, "An auto-generated nonlinear mpc algorithm for real-time obstacle avoidance of ground vehicles," in *Control Conference (ECC), 2013 European*. IEEE, 2013, pp. 4136–4141.

[14] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.

2.4 Conference Paper - Adaptive Control

J. Dørum, T. Utstumo, and J. Gravdahl. Experimental comparison of adaptive controllers for trajectory tracking in agricultural robotics. In *Proceedings of the 19th International Conference on System Theory, Control and Computing, ICSTCC, Cheile Gradistei, Romania, October 14-16, 2015.*

Experimental Comparison of Adaptive Controllers for Trajectory Tracking in Agricultural Robotics

Jarle Dørum, Tryve Utstumo, Jan Tommy Gravdahl
Department of Engineering Cybernetics
Norwegian University of Science and Technology
NO-7491 Trondheim, Norway

***Abstract*—This paper describes the development of several controllers to handle a trajectory tracking problem for a differentially wheeled robot. Both simulations and tests on a real robot were performed. A simple kinematic controller has been implemented to calculate desired velocities based on current position and trajectory. In order to also consider the current velocities, i.e. the dynamics of the system, the output of this controller was used as input to a dynamic controller derived from a nonlinear model. The dynamic controller was made adaptive by using an on-line parameter estimation scheme to estimate the unknown parameters of the nonlinear model. Lastly, a direct model reference adaptive controller (MRAC) based on a linear model was derived and implemented as an alternative to the adaptive dynamic controller.**

***Index Terms*—Mobile robots, differentially wheeled robots, trajectory tracking, adaptive control, system identification, agricultural robotics,**

I. Introduction

This paper presents part of the ongoing research for developing an agricultural robot that autonomously navigates in row crops while identifying and precision spraying individual weed leaves with herbicide. The robot is a differentially steered robot with two rear mounted caster wheels, and may be modeled as a unicycle-like robot. A picture of the prototype during testing in row crops is shown in Fig. 1.

Previous research on the project includes development of a precision drop-on-demand nozzle for herbicide application [1], a model predictive row controller [2] to minimize potential crop damage during operation and attitude estimation in agricultural robotics [3].

The nozzle array presented in [1] is intended to only be slightly wider than the row crops, meaning that the robot has to follow the row crops precisely. A small offset could mean that the weed is out of reach for the nozzles, leaving the weed untreated. This motivates the research in this paper to find a trajectory tracking controller that minimizes the tracking error.

Another aspect to consider is changing physical properties of the robot. For example, the weight of the robot will change as herbicide and fuel is consumed. To ensure satisfactory performance at all times, several adaptive approaches that update the controller gains continuously have been tested.

Unicycle-like robots are used extensively in all kinds of fields and numerous models and controllers have been described in publications. In this paper a nonlinear model proposed in [4] has been used for simulations. The same

Fig. 1. A picture of the prototype robot on a field test.

model was also used in [5] to develop an adaptive dynamic controller, which has been implemented and tested here but with a different adaptation law. In [6] an adaptive controller using adaptive backstepping is presented. [7] developed a model reference adaptive controller (MRAC) for the tracking problem, but only simulations were performed. A similar direct MRAC has been derived here and implemented on the robot for testing.

The most important contribution of this paper is the comparison of two different adaptive controllers implemented on the same robot. The author is not aware of any previous implementations of the MRAC controller presented here on a real robot.

Different approaches to row crop guidance systems has been thoroughly explored and reviewed in [8]. However, this paper focuses merely on tracking a smooth and well defined trajectory without considering how to obtain the trajectory. The results obtained should be applicable to most unicycle-like robots.

II. Mathematical Model

The model used for simulations and some of the controller designs in this paper was presented in [4]. It is given as

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \\ \dot{u} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} u\cos\psi - a\omega\sin\psi \\ u\sin\psi + a\omega\cos\psi \\ \omega \\ \frac{\theta_3}{\theta_1}\omega^2 - \frac{\theta_4}{\theta_1}u \\ -\frac{\theta_5}{\theta_2}u\omega - \frac{\theta_6}{\theta_2}\omega \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{1}{\theta_1} & 0 \\ 0 & \frac{1}{\theta_2} \end{bmatrix} \begin{bmatrix} u_{ref} \\ \omega_{ref} \end{bmatrix} + \begin{bmatrix} \delta_x \\ \delta_y \\ 0 \\ \bar{\delta}_u \\ \bar{\delta}_\omega \end{bmatrix} \tag{1}$$

where $\boldsymbol{h} = \begin{bmatrix} x & y \end{bmatrix}^T$ is the position, ψ is the heading angle, u is forward velocity, ω is angular velocity and a is the distance

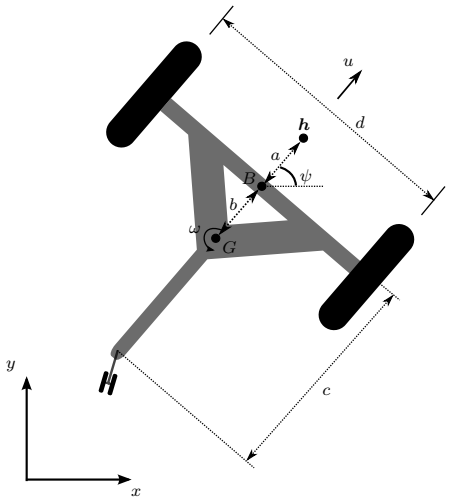

Fig. 2. Drawing of unicycle-like robot with centered rear mounted caster wheel, similar to the robot all tests were performed on.

from center of wheel axis to $\boldsymbol{h}$ as shown in Fig. 2. Motor inputs are given as velocities instead of torque values, which means that the motor controller is assumed to have a PID controller or similar. $\boldsymbol{\theta}$ is a collection of physical parameters derived in [4] and included here for reference:

$$\begin{aligned}
\theta_1 &= \left(\frac{R_a}{k_a}(mR_t r + 2I_e) + 2rk_{DT}\right)/(2rk_{PT}) \\
\theta_2 &= \left(\frac{R_a}{k_a}(I_e d^2 + 2R_t r(I_z + mb^2)) + 2rdk_{DR}\right)/(2rdk_{PR}) \\
\theta_3 &= \frac{R_a}{k_a} mbR_t/(2k_{PT}) \\
\theta_4 &= \frac{R_a}{k_a}\left(\frac{k_a k_b}{R_a} + B_e\right)/(rk_{PT}) + 1 \\
\theta_5 &= \frac{R_a}{k_a} mbR_t/(dk_{PR}) \\
\theta_6 &= \frac{R_a}{k_a}\left(\frac{k_a k_b}{R_a} + B_e\right) d/(2rk_{PR}) + 1
\end{aligned} \tag{2}$$

where R_a is motor resistance, k_a motor torque multiplied by gear ratio, k_b motor voltage multiplied by gear ratio, r wheel radius, I_e motor moment of inertia, B_e motor viscous friction coefficient, $k_{PT}, k_{DT}, k_{PR}, k_{DR}$ are PID motor controller gains, I_z moment of inertia about vertical axis at center of mass, m mass.

$$\boldsymbol{\delta} = \begin{bmatrix}\delta_x & \delta_y & 0 & \bar{\delta}_u & \bar{\delta}_\omega\end{bmatrix}^T \tag{3}$$

represent the uncertainties of the system caused by wheel slips and forces exerted by the caster wheel. For the purpose of this paper it has been assumed that $\boldsymbol{\delta} = 0$.

III. Controller Design

In many cases unicycle-like robots operate at low speeds and often inhibit low moment of inertia. In other words, the dynamics of u and ω are so fast that in many cases one may simplify $u \approx u_{ref},\ \omega \approx \omega_{ref}$ and only study the kinematic model given by

$$\begin{bmatrix}\dot{x} \\ \dot{y} \\ \dot{\psi}\end{bmatrix} = \begin{bmatrix}u_{ref}\cos\psi - a\omega_{ref}\sin\psi \\ u_{ref}\sin\psi + a\omega_{ref}\cos\psi \\ \omega_{ref}\end{bmatrix} \tag{4}$$

For larger robots operating at higher speeds the dynamics cannot simply be ignored. In the next sections, various controller designs are considered.

A. Trajectory Tracking Controller

Let $\boldsymbol{h_d}(t) = \begin{bmatrix}x_d(t) & y_d(t)\end{bmatrix}^T$ denote the time varying reference trajectory for the robot. Only the reference position is considered, he the The tracking error is defined as $\tilde{\boldsymbol{h}} = \begin{bmatrix}x_d - x & y_d - y\end{bmatrix}^T$. The kinematics from (1) may be written as

$$\begin{bmatrix}\dot{x} \\ \dot{y} \\ \dot{\psi}\end{bmatrix} = \begin{bmatrix}\cos\psi & -a\sin\psi \\ \sin\psi & a\cos\psi \\ 0 & 1\end{bmatrix}\begin{bmatrix}u \\ \omega\end{bmatrix} \tag{5}$$

Recalling that $\boldsymbol{h} = [x\ \ y]^T$

$$\dot{h} = \begin{bmatrix}\dot{x} \\ \dot{y}\end{bmatrix} = \begin{bmatrix}\cos\psi & -a\sin\psi \\ \sin\psi & a\cos\psi\end{bmatrix}\begin{bmatrix}u \\ \omega\end{bmatrix} = \boldsymbol{A}\begin{bmatrix}u \\ \omega\end{bmatrix} \tag{6}$$

Multiplying (6) with the inverse of $\boldsymbol{A}$ gives

$$\begin{bmatrix}u \\ \omega\end{bmatrix} = \boldsymbol{A}^{-1}\begin{bmatrix}\dot{x} \\ \dot{y}\end{bmatrix} = \begin{bmatrix}\cos\psi & \sin\psi \\ -\frac{1}{a}\sin\psi & \frac{1}{a}\cos\psi\end{bmatrix}\begin{bmatrix}\dot{x} \\ \dot{y}\end{bmatrix} \tag{7}$$

A controller based on inverse kinematics is proposed in [5]:

$$\begin{bmatrix}u^c_{ref} \\ \omega^c_{ref}\end{bmatrix} = \begin{bmatrix}\cos\psi & \sin\psi \\ -\frac{1}{a}\sin\psi & \frac{1}{a}\cos\psi\end{bmatrix}\begin{bmatrix}\dot{x}_d + l_x\tanh\left(\frac{k_x}{l_x}\tilde{x}\right) \\ \dot{y}_d + l_y\tanh\left(\frac{k_y}{l_y}\tilde{y}\right)\end{bmatrix} \tag{8}$$

where $k_x, k_y > 0$ are controller gains and $l_x, l_y > 0$ are saturation constants. u^c_{ref} and ω^c_{ref} are desired forward and angular velocities, respectively. The controller is shown in [5] to have an asymptotically stable equilibrium at the origin $\tilde{\boldsymbol{h}} = \begin{bmatrix}0 & 0\end{bmatrix}^T$ under the assumption of $u = u^c_{ref}$ and $\omega = \omega^c_{ref}$. The name trajectory tracking controller and kinematic controller both refer to the same controller for the rest of this paper.

Note that it is not necessary to explicitly control the desired heading, as discussed in [5]. Due to the non-holonomic nature of a differentially wheeled mobile robot, the heading will be tangent to the trajectory path given small position errors. Any heading deviation from the tangent will cause a change in position errors, so proving stability for the position will be sufficient.

B. Dynamic Controller

The kinematic controller will work adequately as long as the dynamics of the system are fast enough, i.e. the assumption of $u \approx u^c_{ref}$ and $\omega \approx \omega^c_{ref}$ is reasonable. In cases where the dynamics are too slow to be ignored or high precision tracking

is required, the kinematic controller alone may no longer be sufficient.

Consider the dynamic part of (1)

$$\begin{bmatrix}\dot{u}\\ \dot{\omega}\end{bmatrix} = \begin{bmatrix}\frac{\theta_3}{\theta_1}\omega^2 - \frac{\theta_4}{\theta_1}u\\ -\frac{\theta_5}{\theta_2}u\omega - \frac{\theta_6}{\theta_2}\omega\end{bmatrix} + \begin{bmatrix}\frac{1}{\theta_1} & 0\\ 0 & \frac{1}{\theta_2}\end{bmatrix}\begin{bmatrix}u_{ref}\\ \omega_{ref}\end{bmatrix} \tag{9}$$

Rearranging gives

$$\begin{bmatrix}u_{ref}\\ \omega_{ref}\end{bmatrix} = \begin{bmatrix}\theta_1\dot{u} - \theta_3\omega^2 + \theta_4 u\\ \theta_2\dot{\omega} + \theta_5 u\omega + \theta_6\omega\end{bmatrix} \tag{10}$$

Which may be written as

$$\begin{bmatrix}u_{ref}\\ \omega_{ref}\end{bmatrix} = \begin{bmatrix}\theta_1 & 0\\ 0 & \theta_2\end{bmatrix}\begin{bmatrix}\dot{u}\\ \dot{\omega}\end{bmatrix} + \begin{bmatrix}0 & 0 & -\omega^2 & u & 0 & 0\\ 0 & 0 & 0 & 0 & u\omega & \omega\end{bmatrix}\boldsymbol{\theta} \tag{11}$$

Motivated by the inverse dynamics in (11), [5] proposes the controller given as

$$\begin{bmatrix}u_{ref}\\ \omega_{ref}\end{bmatrix} = \begin{bmatrix}\theta_1 & 0\\ 0 & \theta_2\end{bmatrix}\begin{bmatrix}\sigma_1\\ \sigma_2\end{bmatrix} + \begin{bmatrix}0 & 0 & -\omega^2 & u & 0 & 0\\ 0 & 0 & 0 & 0 & u\omega & \omega\end{bmatrix}\boldsymbol{\theta} \tag{12}$$

where

$$\boldsymbol{\sigma} = \begin{bmatrix}\sigma_1\\ \sigma_2\end{bmatrix} = \begin{bmatrix}\dot{u}^c_{ref} + k_1\tilde{u}\\ \dot{\omega}^c_{ref} + k_2\tilde{\omega}\end{bmatrix}, \qquad \begin{aligned}\tilde{u} &= u^c_{ref} - u\\ \tilde{\omega} &= \omega^c_{ref} - \omega\end{aligned} \tag{13}$$

and $k_1, k_2 > 0$ are constant gains. In order to implement (12), the values of $\boldsymbol{\theta}$ must be known. Measuring or otherwise obtaining the parameters needed to calculate $\boldsymbol{\theta}$ may prove hard, thus the need to estimate $\boldsymbol{\theta}$ becomes a necessity. Replacing $\boldsymbol{\theta}$ with the estimate $\hat{\boldsymbol{\theta}}$ in (12) gives

$$\boldsymbol{\nu_{ref}} = \hat{\boldsymbol{D}}\boldsymbol{\sigma} + \boldsymbol{E}\hat{\boldsymbol{\theta}} \tag{14}$$

where

$$\begin{aligned}&\boldsymbol{\nu_{ref}} = \begin{bmatrix}u_{ref}\\ \omega_{ref}\end{bmatrix},\ \boldsymbol{E} = \begin{bmatrix}0 & 0 & -\omega^2 & u & 0 & 0\\ 0 & 0 & 0 & 0 & u\omega & \omega\end{bmatrix},\\ &\hat{\boldsymbol{D}} = \begin{bmatrix}\hat{\theta}_1 & 0\\ 0 & \hat{\theta}_2\end{bmatrix},\ \hat{\boldsymbol{\theta}} = \begin{bmatrix}\hat{\theta}_1 & \hat{\theta}_2 & \hat{\theta}_3 & \hat{\theta}_4 & \hat{\theta}_5 & \hat{\theta}_6\end{bmatrix}^T\end{aligned} \tag{15}$$

Following is a stability analysis similar to what was done in [5]. (11) may be written as

$$\boldsymbol{\nu_{ref}} = \boldsymbol{D}\dot{\boldsymbol{\nu}} + \boldsymbol{E}\boldsymbol{\theta} \tag{16}$$

Similarly, (14) is written

$$\boldsymbol{\nu_{ref}} = \hat{\boldsymbol{D}}\boldsymbol{\sigma} + \boldsymbol{E}\hat{\boldsymbol{\theta}} = \boldsymbol{G}\boldsymbol{\theta} - \boldsymbol{G}\tilde{\boldsymbol{\theta}} = \boldsymbol{D}\boldsymbol{\sigma} + \boldsymbol{E}\boldsymbol{\theta} - \boldsymbol{G}\tilde{\boldsymbol{\theta}} \tag{17}$$

where $\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}$ and

$$\boldsymbol{G} = \begin{bmatrix}\sigma_1 & 0 & -\omega^2 & u & 0 & 0\\ 0 & \sigma_2 & 0 & 0 & u\omega & \omega\end{bmatrix},\quad \boldsymbol{D} = \begin{bmatrix}\theta_1 & 0\\ 0 & \theta_2\end{bmatrix} \tag{18}$$

(13) may be written as

$$\boldsymbol{\sigma} = \begin{bmatrix}\sigma_1\\ \sigma_2\end{bmatrix} = \dot{\boldsymbol{\nu}}^{\boldsymbol{c}}_{\boldsymbol{ref}} + \boldsymbol{K}\tilde{\boldsymbol{\nu}} \tag{19}$$

where $\boldsymbol{K} = \mathrm{diag}(k_1, k_2)$. Combining (16), (17) and (19)

$$\boldsymbol{D}\dot{\boldsymbol{\nu}} + \boldsymbol{E}\boldsymbol{\theta} = \boldsymbol{D}\dot{\boldsymbol{\nu}}^{\boldsymbol{c}}_{\boldsymbol{ref}} + \boldsymbol{D}\boldsymbol{K}\tilde{\boldsymbol{\nu}} + \boldsymbol{E}\boldsymbol{\theta} - \boldsymbol{G}\tilde{\boldsymbol{\theta}} \tag{20}$$

$$\dot{\tilde{\boldsymbol{\nu}}} = -\boldsymbol{K}\tilde{\boldsymbol{\nu}} + \boldsymbol{D}^{-1}\boldsymbol{G}\tilde{\boldsymbol{\theta}} \tag{21}$$

where $\dot{\tilde{\boldsymbol{\nu}}} = \dot{\boldsymbol{\nu}}^{\boldsymbol{c}}_{\boldsymbol{ref}} - \dot{\boldsymbol{\nu}}$ describes the error dynamics of the system. For this analysis, $\boldsymbol{\theta}$ is considered known, i.e. $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$, reducing (21) to

$$\dot{\tilde{\boldsymbol{\nu}}} = -\boldsymbol{K}\tilde{\boldsymbol{\nu}} \tag{22}$$

Consider the following Lyapunov-like function

$$V = \frac{1}{2}\tilde{\boldsymbol{\nu}}^T\boldsymbol{P}\tilde{\boldsymbol{\nu}} \tag{23}$$

where $\boldsymbol{P} = \boldsymbol{P}^T > \boldsymbol{0}$. Differentiating (23) along the solution of (22) gives

$$\dot{V} = -\tilde{\boldsymbol{\nu}}^T\boldsymbol{P}\boldsymbol{K}\tilde{\boldsymbol{\nu}} < 0\ \ \forall\ \ \tilde{\boldsymbol{\nu}} \neq 0 \tag{24}$$

Which means that $\dot{V}$ is negative definite and global asymptotic stability can be concluded.

C. On-line Parameter Estimation

The dynamic controller given by (14) needs a good estimate $\hat{\boldsymbol{\theta}}$ in order to perform well. One approach to estimate $\hat{\boldsymbol{\theta}}$ is to log a test run with sufficiently excited input signal and use an off-line system identification technique, e.g. least-squares method. Another approach is to estimate $\hat{\boldsymbol{\theta}}$ on-line using an adaptation law $\dot{\hat{\boldsymbol{\theta}}}$. This section shows the derivation of $\dot{\hat{\boldsymbol{\theta}}}$ using the gradient method, which is motivated by the minimization of a cost function.

Consider (10) written on the form

$$\boldsymbol{\nu_{ref}} = \boldsymbol{\varphi}^T\boldsymbol{\theta} = \begin{bmatrix}\dot{u} & 0 & -\omega^2 & u & 0 & 0\\ 0 & \dot{\omega} & 0 & 0 & u\omega & \omega\end{bmatrix}\boldsymbol{\theta} \tag{25}$$

where $\boldsymbol{\nu_{ref}} = [u_{ref}\ \ \omega_{ref}]^T$. Filtering both sides gives

$$\frac{\boldsymbol{\nu_{ref}}}{\Lambda(s)} = \frac{\boldsymbol{\varphi}^T\boldsymbol{\theta}}{\Lambda(s)} = \frac{1}{\Lambda(s)}\begin{bmatrix}su & 0 & -\omega^2 & u & 0 & 0\\ 0 & s\omega & 0 & 0 & u\omega & \omega\end{bmatrix}\boldsymbol{\theta} \tag{26}$$

Which may be written as the parametric model

$$\boldsymbol{z} = \boldsymbol{\Phi}^T\boldsymbol{\theta} \tag{27}$$

where $z = \frac{\boldsymbol{\nu_{ref}}}{\Lambda(s)}$, $\boldsymbol{\Phi}^T = \frac{\boldsymbol{\varphi}^T}{\Lambda(s)}$ and $\Lambda(s)$ is chosen to be a Hurwitz polynomial of degree one, e.g. $\Lambda(s) = s + 1$. Note that $\boldsymbol{z}$ and $\boldsymbol{\Phi}$ are available measurements, while $\boldsymbol{\theta}$ is unknown. An estimate of $\boldsymbol{z}$ denoted $\hat{\boldsymbol{z}}$ is generated as

$$\hat{\boldsymbol{z}} = \boldsymbol{\Phi}^T\hat{\boldsymbol{\theta}} \tag{28}$$

where $\hat{\boldsymbol{\theta}}$ is the currently best estimate of $\boldsymbol{\theta}$. A normalized estimation error is defined as

$$\boldsymbol{\epsilon} = (\boldsymbol{M}^T\boldsymbol{M})^{-1}(\boldsymbol{z} - \hat{\boldsymbol{z}}) = (\boldsymbol{M}^T\boldsymbol{M})^{-1}(\boldsymbol{z} - \boldsymbol{\Phi}^T\hat{\boldsymbol{\theta}}) \tag{29}$$

where $\boldsymbol{M}^T\boldsymbol{M} = \boldsymbol{I} + \boldsymbol{N_s}^T\boldsymbol{N_s}$ is a diagonal matrix that normalizes the estimation error, and $\boldsymbol{N_s}^T\boldsymbol{N_s}$ is another diagonal matrix for design of the normalized signal. The reason for this normalization is to ensure boundedness, i.e.

$$\boldsymbol{\Phi}\boldsymbol{M}^{-1} \in \mathcal{L}_\infty \tag{30}$$

If $\boldsymbol{\Phi} \in \mathcal{L}_\infty$, then $\boldsymbol{M} = \boldsymbol{I}$ is sufficient. If it is not, choosing

$$\boldsymbol{M}^T\boldsymbol{M} = \boldsymbol{I} + \boldsymbol{\Phi}^T\boldsymbol{\Phi} \tag{31}$$

will ensure (30) is satisfied [9, p. 172]. An instantaneous cost function $\boldsymbol{J}(\hat{\boldsymbol{\theta}})$ is defined as

$$\boldsymbol{J}(\hat{\boldsymbol{\theta}}) = \frac{1}{2}\boldsymbol{\epsilon}^T\boldsymbol{M}^T\boldsymbol{M}\boldsymbol{\epsilon} = \frac{1}{2}(\boldsymbol{z} - \boldsymbol{\Phi}^T\hat{\boldsymbol{\theta}})^T(\boldsymbol{M}^T\boldsymbol{M})^{-1}(\boldsymbol{z} - \boldsymbol{\Phi}^T\hat{\boldsymbol{\theta}}) \tag{32}$$

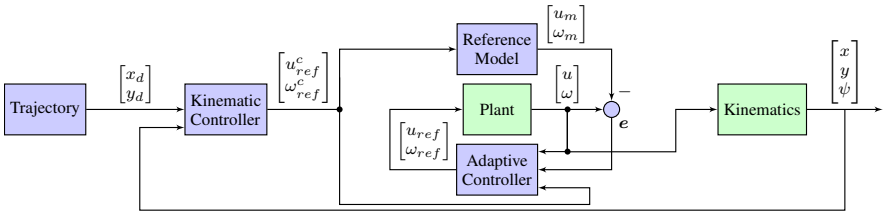

Fig. 3. Block diagram of the model reference adaptive controller.

The gradient of (32) is

$$\Delta \boldsymbol{J}(\hat{\boldsymbol{\theta}}) = -\boldsymbol{\Phi}(\boldsymbol{M}^T\boldsymbol{M})^{-1}(\boldsymbol{z} - \boldsymbol{\Phi}^T\hat{\boldsymbol{\theta}}) = -\boldsymbol{\Phi}\boldsymbol{\epsilon} \tag{33}$$

Motivated by this, the following adaptation law for generating $\hat{\boldsymbol{\theta}}(t)$ is proposed

$$\dot{\hat{\boldsymbol{\theta}}} = -\boldsymbol{\Gamma}\Delta \boldsymbol{J}(\hat{\boldsymbol{\theta}}) = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\epsilon} \tag{34}$$

where $\boldsymbol{\Gamma} = \boldsymbol{\Gamma}^T >$ is a diagonal gain matrix. According to [9, p. 175], (34) ensures that

1) $\hat{\boldsymbol{\theta}}, \boldsymbol{\epsilon} \in \mathcal{L}_\infty$
2) $\boldsymbol{\epsilon}, \boldsymbol{N_s}^T\boldsymbol{\epsilon}, \dot{\hat{\boldsymbol{\theta}}} \in \mathcal{L}_\infty$

independent of the boundedness properties of $\boldsymbol{\Phi}$. In other words, both parameters and estimation errors should remain bounded. It does not, however, ensure that $\tilde{\boldsymbol{\theta}}(t) = \boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t) \to 0$ as $t \to \infty$. To ensure that the parameters $\hat{\boldsymbol{\theta}}$ do in fact converge to their actual value $\boldsymbol{\theta}$, $\boldsymbol{\Phi}$ must be persistently excited (PE), i.e., it satisfies [9, p. 254]

$$\alpha_1 \boldsymbol{I} \geq \frac{1}{T_0}\int_t^{t+T_0} \boldsymbol{\Phi}(\tau)\boldsymbol{\Phi}^T(\tau)\,\mathrm{d}\tau \geq \alpha_0 \boldsymbol{I}, \quad \forall t \geq 0 \tag{35}$$

for some $T_0, \alpha_0, \alpha_1 \geq 0$. It is in general difficult to show that $\boldsymbol{\Phi}$ is PE for an input signal $\boldsymbol{\nu_{ref}}$, and especially in a case like this where $\boldsymbol{\Phi}$ has some nonlinear elements.

D. Adaptive Dynamic Controller

The results from section III-B and section III-C may be combined to form an adaptive dynamic controller. The on-line parameter estimation operates independently from the dynamic controller and vice versa, making it a modular design. This may prove beneficial in cases where parameter estimation is only needed parts of the time, or if it is desirable to run parameter estimation without running the dynamic controller. The control laws are given by

$$\boldsymbol{\nu_{ref}} = \hat{\boldsymbol{D}}\boldsymbol{\sigma} + \boldsymbol{E}\hat{\boldsymbol{\theta}}, \quad \dot{\hat{\boldsymbol{\theta}}} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\epsilon} \tag{36}$$

where the notation is the same as in section III-B and section III-C.

E. Direct Model Reference Adaptive Controller

In this section, a simple direct Model Reference Adaptive Controller (MRAC) scheme as shown in Fig. 3 is derived. The concept is to design a model of similar structure to the plant (robot), let the tracking reference be an input to the model, and make the output of the plant track the output of the model. For the direct MRAC approach, this is made possible by developing adaptation laws for the controller gains directly without having to identify actual system parameters.

Consider a simplified, linear model of the dynamics of (1) given by

$$\begin{bmatrix} \dot{u} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} au + bu_{ref} \\ c\omega + d\omega_{ref} \end{bmatrix} \tag{37}$$

where a, b, c, d are unknown system parameters (not the same as those introduced in Fig. 2). In this case, u and ω are considered decoupled and will be analyzed separately. A reference model u_m for u is chosen to be

$$\dot{u}_m = -a_m u + b_m u^c_{ref} \tag{38}$$

Laplace transforming (37) and (38) gives

$$u = \frac{b}{s-a}u_{ref}, \quad u_m = \frac{b_m}{s+a_m}u^c_{ref} \tag{39}$$

The following control law is proposed

$$u_{ref} = -k^*_u u + l^*_u u^c_{ref} \tag{40}$$

Inserting (40) into (39) gives

$$u = \frac{bl^*_u}{s - a + bk^*_u}u^c_{ref}, \quad u_m = \frac{b_m}{s+a_m}u^c_{ref} \tag{41}$$

It is desirable to make the transfer functions of (41) equal. Choosing

$$l^*_u = \frac{b_m}{b}, \quad k^*_u = \frac{a + a_m}{b} \tag{42}$$

ensures equal transfer functions. However, it is not possible to implement since the values of a and b are unknown. Instead of using the control law (40), a control law using estimates of k^*_u and l^*_u is proposed

$$u_{ref} = -k_u(t)u + l_u(t)u^c_{ref} \tag{43}$$

where $k_u(t)$ and $l_u(t)$ are the currently best estimates of k^*_u and l^*_u, respectively. Adding and subtracting $b(-k^*_u u + l^*_u u^c_{ref})$ to $\dot{u}$ yields

$$\dot{u} = au + bu_{ref} + b(-k^*_u u + l^*_u r) - b(-k^*_u u + l^*_u r) \tag{44}$$

which, after combining with (42), may be written as

$$\dot{u} = -a_m u + b_m u^c_{ref} u + b(k^*_u u - l^*_u u^c_{ref} + u_{ref}) \tag{45}$$

Laplace transforming (45) gives

$$u = \underbrace{\frac{b_m}{s+a_m}u^c_{ref}}_{=u_m} + \frac{b}{s+a_m}(k^*_u u - l^*_u u^c_{ref} + u_{ref}) \tag{46}$$

Define the tracking error $e_u = u - u_m$ to obtain

$$e_u = \frac{b}{s+a_m}(k_u^* u - l_u^* u_{ref}^c + u_{ref}) \tag{47}$$

Since k_u^* and l_u^* are unknown, our best estimate of the tracking error $\hat{e}_u$ is

$$\hat{e}_u = \frac{b}{s+a_m}(k_u(t)u - l_u(t)u_{ref}^c + u_{ref}) \tag{48}$$

Inserting u_{ref} from (43) into (48) simply gives $\hat{e}_u = 0$, i.e. the estimated tracking error is zero. Note that the estimation error $\epsilon_u = e_u - \hat{e}_u = e_u = u - u_m$ is equal to the tracking error e_u. Combining (43) and (47) while defining the gain parameter estimation errors $\tilde{k}_u(t) = k_u(t) - k_u^*$, $\tilde{l}_u(t) = l_u(t) - l_u^*$ gives

$$e_u = \frac{b}{s+a_m}(-\tilde{k}_u u + \tilde{l}_u u_{ref}^c) \tag{49}$$

$$\dot{e}_u = -a_m e_u + b(-\tilde{k}_u u + \tilde{l}_u u_{ref}^c) \tag{50}$$

Consider the Lyapunov-like function

$$V = \frac{1}{2}e_u^2 + \frac{|b|}{2\gamma_1}\tilde{k}_u^2 + \frac{|b|}{2\gamma_2}\tilde{l}_u^2 \tag{51}$$

with $\gamma_1, \gamma_2 > 0$. Differentiating (51) along the solution of (49) gives

$$\begin{aligned}\dot{V} = &- a_m e_u^2 + |b|\tilde{k}_u\left(-e_u u\,\mathrm{sgn}(b) + \frac{1}{\gamma_1}\dot{k}\right) \\ &+ |b|\tilde{l}_u\left(e_u u_{ref}^c\,\mathrm{sgn}(b) + \frac{1}{\gamma_2}\dot{l}\right)\end{aligned} \tag{52}$$

Choosing

$$\dot{k}_u = \gamma_1 e_u u\,\mathrm{sgn}(b), \quad \dot{l}_u = -\gamma_2 e_u u_{ref}^c\,\mathrm{sgn}(b) \tag{53}$$

ensures

$$\dot{V} = -a_m e_u^2 \leq 0 \tag{54}$$

Thus it is shown that $\dot{V}$ is negative semi definite, V has an upper bound $V(0)$ and bounded below by zero, i.e. $0 \leq V(t) \leq V(0)$. From the boundedness of $V(t)$ and (51), it is clear that $e_u, \tilde{k}_u, \tilde{l}_u \in \mathcal{L}_\infty$. u_{ref}^c, the output of the kinematic controller (8), is bounded, so $u_{ref}^c \in \mathcal{L}_\infty$. The transfer functions of (41) are in $\mathcal{L}_1$, and it follows from [9, p. 80] that $u, u_m \in \mathcal{L}_\infty$. This means that all signals of (49) are bounded and $\dot{e}_u \in \mathcal{L}_\infty$. From [9, p. 74] it follows that since $V(t)$ is bounded from below and non-increasing, it has a finite limit as $t \to \infty$, denoted V_∞. It can also be seen that

$$\begin{aligned}||e_u||_2 &= \left(\int_0^\infty e_u^2(\tau)\mathrm{d}\tau\right)^{1/2} = \left(\int_0^\infty -\frac{1}{a_m}\dot{V}(\tau)\mathrm{d}\tau\right)^{1/2} \\ &= \left(\frac{1}{a_m}(V(0) - V_\infty)\right)^{1/2}\end{aligned} \tag{55}$$

which is clearly bounded, so that $e_u \in \mathcal{L}_2$. Finally, [9, p. 80] shows that since $\dot{e}_u, e_u \in \mathcal{L}_\infty$ and $e_u \in \mathcal{L}_2$, then $e_u(t) \to 0$ as $t \to \infty$.

The results obtained show that the tracking objective of making the output of the plant $u(t)$ track the output of the reference model $u_m(t)$ is achieved. It does not, however, guarantee that $k_u(t), l_u(t) \to k_u^*, l_u^*$ as $t \to \infty$, i.e. the poles of the plant may differ from those of the reference model. This should be of less concern, since $k_u(t), l_u(t)$ are bounded and the true values of k_u^*, l_u^* are not of any real importance.

In a very similar manner the same results are found for ω. A summary of the control laws are given in Table I.

A modification that was done to provide for a more robust implementation was to add a small feedback loop to (53) to get

$$\dot{k}_u = \gamma_1 e_u u\,\mathrm{sgn}(b) - \alpha k_u, \quad \dot{l}_u = -\gamma_2 e_u u_{ref}^c\,\mathrm{sgn}(b) - \beta l_u \tag{56}$$

where $0 < \alpha << 1$ and $0 < \beta << 1$.

TABLE I. Control Laws for The MRAC

Plant	Reference Model	Control Law
$u = \frac{b}{s-a}u_{ref}$	$u_m = \frac{b_m}{s+a_m}u_{ref}^c$	$u_{ref} = -k_u(t)u + l_u(t)u_{ref}^c$, $\dot{k}_u = \gamma_1 e_u u\,\mathrm{sgn}(b)$ $\dot{l}_u = -\gamma_2 e_u u_{ref}^c\,\mathrm{sgn}(b)$ $e_u = u - u_m$
$\omega = \frac{d}{s-c}\omega_{ref}$	$\omega_m = \frac{d_m}{s+c_m}\omega_{ref}^c$	$\omega_{ref} = -k_\omega(t)\omega + l_\omega(t)\omega_{ref}^c$, $\dot{k}_\omega = \gamma_3 e_\omega \omega\,\mathrm{sgn}(d)$ $\dot{l}_\omega = -\gamma_4 e_\omega \omega_{ref}^c\,\mathrm{sgn}(d)$ $e_\omega = \omega - \omega_m$

IV. Simulations And Real Runs

Simulations were performed using Matlab/Simulink, while real runs were performed on the robot shown in Fig. 4. The dimensions of the robot are approximately $c = 1.1\,\mathrm{m}$ and $d = 1.0\,\mathrm{m}$ using the notation from Fig. 2. The robot is running Robot Operating System (ROS) and all controllers were implemented in C++. The motor controller has a low level PID controller that uses individual motor velocities as setpoints, and motor acceleration can be saturated to ensure slower dynamics. Without limits on acceleration the dynamics were so fast that all controllers had equal performance. For both simulations and tests on the real robot the following figure eight trajectory was used:

$$\begin{aligned}x_d(t) &= r_e\sin(2\omega_e t) \\ y_d(t) &= r_e(\cos(\omega_e t) - 1)\end{aligned} \tag{57}$$

For both simulations and real runs $r_e = 0.6\,\mathrm{m}$ and $\omega_e = 0.3\,\mathrm{rad/s}$. The distance from wheel axle to $\boldsymbol{h}$ was chosen to be $a = 0.10\,\mathrm{m}$. The simulated system uses $\boldsymbol{\theta} = \begin{bmatrix}1.0 & 0.4 & 0.2 & 1.1 & 0.2 & 0.9\end{bmatrix}^T$ and initial estimates $\hat{\boldsymbol{\theta}}_\mathbf{0} = \begin{bmatrix}1.0 & 1.0 & 0 & 1.0 & 0 & 1.0\end{bmatrix}^T$ are used for the dynamic and adaptive dynamic controllers. The on-line parameter estimation method was tested in simulations where $\boldsymbol{\theta}$ is known to ensure that the estimated $\hat{\boldsymbol{\theta}}$ converges to its actual value $\boldsymbol{\theta}$. It was found that $\hat{\boldsymbol{\theta}}$ does indeed converge correctly while attempting to track the figure eight, which means that the input signal is sufficiently excited to ensure convergence.

A comparison of all controllers is shown in Fig. 8. The estimated parameters during a real run are shown in Fig. 5 and Fig. 6. Controller inputs during a real run are shown in Fig. 7. It is clear from Fig. 8 (a) and (e) that the kinematic controller alone does not provide sufficient performance in this case. The dynamic controller shows good performance given

Fig. 4. A picture of the robot used for the tests. The robot uses a rear caster wheel.

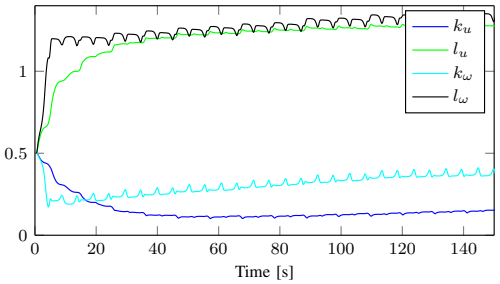

Fig. 5. Controller gains for the MRAC during a real run. All controller gains have initial values of 0.5.

fairly accurate estimates $\hat{\boldsymbol{\theta}}$. The adaptive dynamic controller appears to be able to improve upon the performance of the dynamic controller as $\hat{\boldsymbol{\theta}}$ adapts (shown in Fig. 6).

In Fig. 8 (i) and (j) the motor acceleration saturation limit was increased to make the system a bit faster. It is interesting that the MRAC improves greatly when the dynamics are faster, while the adaptive dynamic controller has almost identical performance to the case with slower dynamics.

V. Conclusion

Two different adaptive dynamic controllers for tracking a trajectory were implemented on a differentially wheeled robot and compared with non-adaptive kinematic and dynamic controllers. The MRAC configuration, which the author has been unable to find previous papers presenting real implementations of, delivered the best performance of all controllers on a system with fairly slow dynamics, while the other adaptive dynamic controller had equal or better performance on a very slow system.

References

[1] F. Urdal, T. Utstumo, S. A. Ellingsen, J. K. Vatne, and T. Gravdahl, "Design and control of precision drop-on-demand herbicide application in agricultural robotics," in *The 13th international Conference on Control, Automation, Robotics and Vision, Singapore*. IEEE, 2014.

[2] T. Utstumo, T. Berge, and J. T. Gravdahl, "Non-linear model predictive control for constrained robot navigation in row crops," *in the Proceedings of the 2015 IEEE International Conference on Industrial Technology (ICIT 2015), Seville, Spain*, March 17-19 2015.

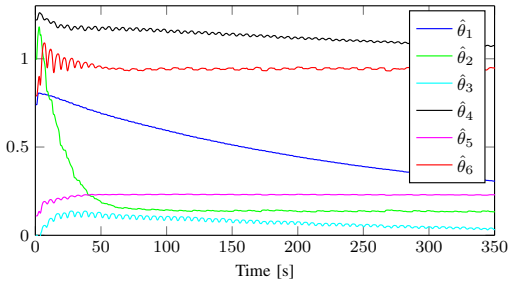

Fig. 6. Estimated system parameters during a real run with the adaptive dynamic controller with $\hat{\boldsymbol{\theta}}_\mathbf{0} = [0.74\ 1.00\ 0.00\ 1.22\ 0.11\ 0.79]^T$ and $\boldsymbol{\Gamma} = \text{diag}(1.0, 1.0, 1.0, 1.0, 1.0, 1.0)$.

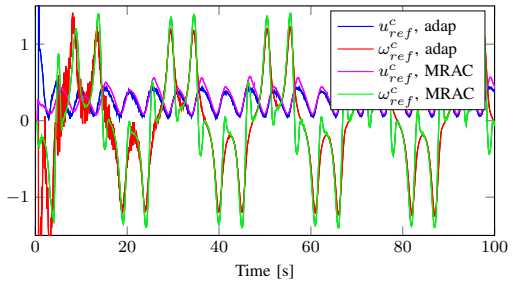

Fig. 7. Motor controller inputs from the MRAC and the adaptive dynamic controller during a test run with the slow dynamics. The inputs become smoother as the parameters adapt.

[3] T. Utstumo and J. T. Gravdahl, "Implementation and comparison of attitude estimation methods for agricultural robotics," *Agricontrol. Vol. 4. No. 1. 2013*, pp. 52–57, 2013.

[4] C. De La Cruz and R. Carelli, "Dynamic Modeling and Centralized Formation Control of Mobile Robots," *IECON 2006 - 32nd Annual Conference on IEEE Industrial Electronics*, pp. 3880–3885, Nov. 2006. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4153091

[5] F. N. Martins, W. C. Celeste, R. Carelli, M. Sarcinelli-Filho, and T. F. Bastos-Filho, "An adaptive dynamic controller for autonomous mobile robot trajectory tracking," *Control Engineering Practice*, vol. 16, no. 11, pp. 1354–1363, Nov. 2008. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0967066108000373

[6] T. Fukao, H. Nakagawa, and N. Adachi, "Adaptive tracking control of a nonholonomic mobile robot," *Robotics and Automation, IEEE Transactions on*, vol. 16, no. 5, pp. 609–615, Oct 2000.

[7] E. Canigur and M. Ozkan, "Model reference adaptive control of a nonholonomic wheeled mobile robot for trajectory tracking," in *Innovations in Intelligent Systems and Applications (INISTA), 2012 International Symposium on*, July 2012, pp. 1–5.

[8] F. Dong, W. Heinemann, and R. Kasper, "Development of a row guidance system for an autonomous robot for white asparagus harvesting," *Computers and Electronics in Agriculture*, vol. 79, no. 2, pp. 216 – 225, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0168169911002304

[9] P. Ioannou and J. Sun, *Robust Adaptive Control*. Dover Publications, 2013. [Online]. Available: https://books.google.no/books?id=ffavAAAAQBAJ

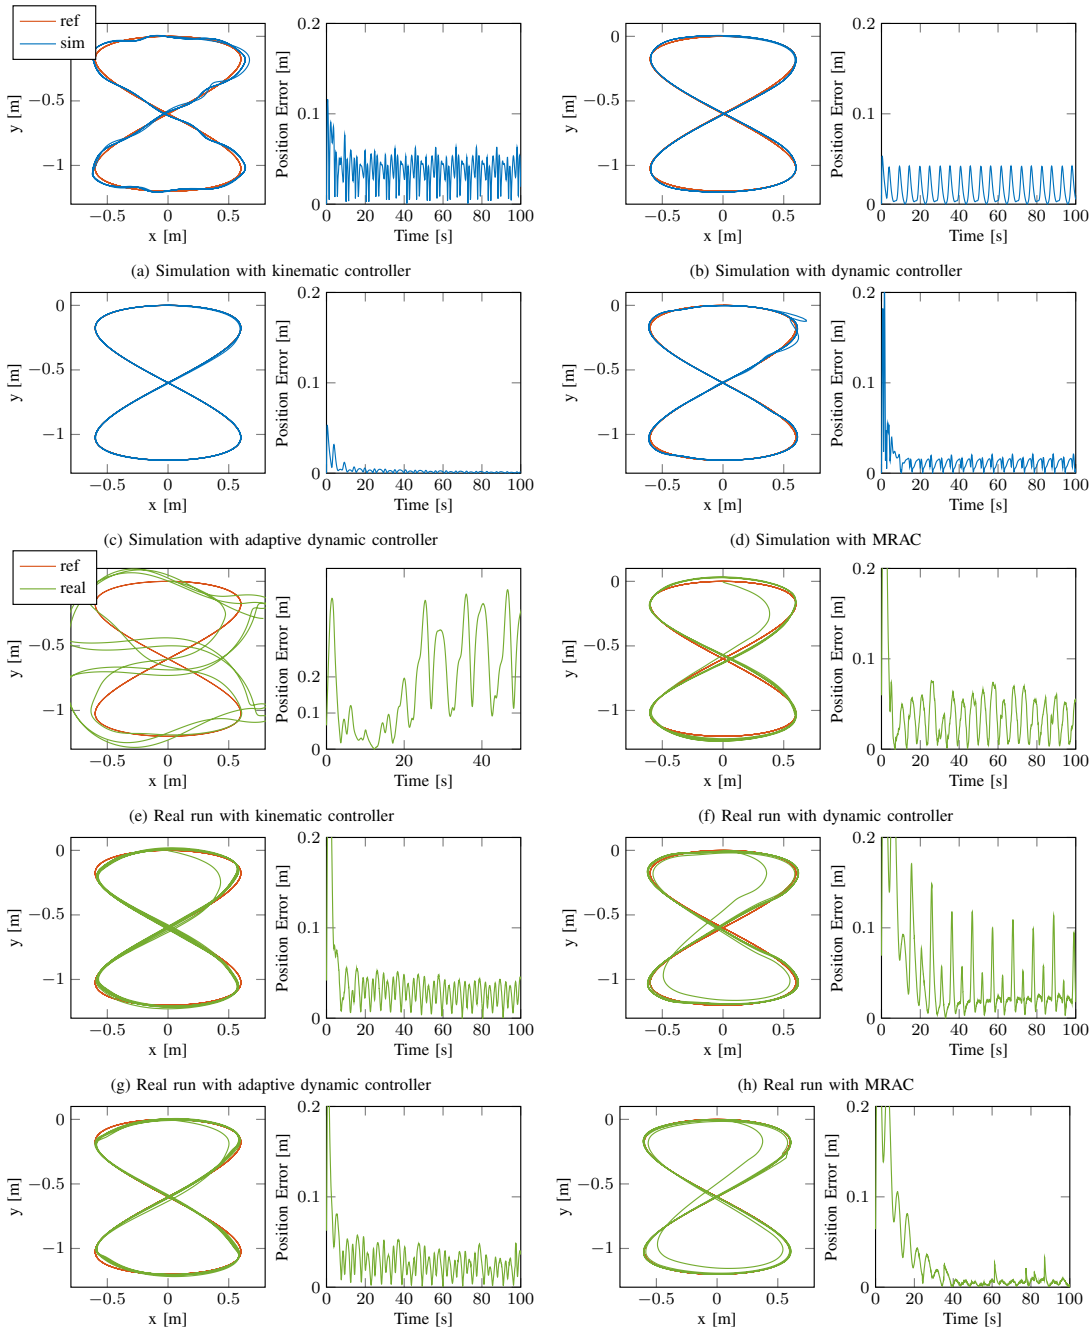

(a) Simulation with kinematic controller

(b) Simulation with dynamic controller

(c) Simulation with adaptive dynamic controller

(d) Simulation with MRAC

(e) Real run with kinematic controller

(f) Real run with dynamic controller

(g) Real run with adaptive dynamic controller

(h) Real run with MRAC

(i) Real run with adaptive dynamic controller on a faster system

(j) Real run with MRAC on a faster system

Fig. 8. Comparison of all the controllers with a relatively slow system and two selected runs on a faster system.

2.5 Journal Paper - Delayed Measurements

M. H. Arbo, T. Utstumo, E. Brekke, and J. T. Gravdahl. Unscented multi-point smoother for fusion of delayed displacement measurements: Application to agricultural robots. *Modeling, Identication and Control*, 38(1):1, 2017.

Unscented Multi-point Smoother for Fusion of Delayed Displacement Measurements: Application to Agricultural Robots

Mathias Hauan Arbo [1] Trygve Utstumo [12] Edmund Brekke [1] Jan Tommy Gravdahl [1]

[1] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, N-7491 Trondheim, Norway. E-mail:* {Mathias.Arbo, Edmund.Brekke, Jan.Tommy.Gravdahl}@itk.ntnu.no

[2] *Adigo AS, N-1405 Langhus, Norway. E-mail:* Trygve@adigo.no

Abstract

Visual Odometry (VO) is increasingly a useful tool for robotic navigation in a variety of applications, including weed removal for agricultural robotics. The methods of evaluating VO are often computationally expensive and can cause the VO measurements to be significantly delayed with respect to a compass, wheel odometry, and GPS measurements. In this paper we present a Bayesian formulation of fusing delayed displacement measurements. We implement solutions to this problem based on the unscented Kalman filter (UKF), leading to what we term an unscented multi-point smoother. The proposed methods are tested in simulations of an agricultural robot. The simulations show improvements in the localization RMS error when including the VO measurements with a variety of latencies.

Keywords: Robot Navigation, Sensor fusion, Agricultural Robotics

1 INTRODUCTION

Figure 1: A prototype of the Asterix robot.

In recent years Agricultural Robotics has been an increasingly important research topic, and there are numerous publications presenting unmanned ground vehicles and robot platforms, (Bogue, 2016; Biber et al., 2012; Jørgensen et al., 2007; Molstad et al., 2014; Grimstad et al., 2015, 2016). Most agricultural platforms rely heavily on GPS for navigation. RTK GPS systems may provide localization accuracy of ±0.02 m under ideal conditions. The cost of an RTK GPS system may hinder the adaption of mobile robots in agriculture, and the signal conditions often reduce the position accuracy.

Visual-aided navigation may provide several benefits for agricultural robots, and lessen the dependence on expensive RTK GPS systems. The work presented here has sprung out from a research project by Adigo AS, in cooperation with the Norwegian University of Science and Technology and the Norwegian Institute of Bioeconomy Research (NIBIO).

An agricultural robot for high precision drop-on-demand herbicide application for row crops is under development. The robot uses a downward facing camera to identify weeds and a nozzle array applies the

herbicide as the robot passes over them, (Urdal et al., 2014; Utstumo and Gravdahl, 2013; Utstumo et al., 2015).

Weed control is a vital part of agriculture, and autonomous robotic weed control has become an important research area. The review by Slaughter et al. (2008) illustrates the potential of robotic weed control, and presents several similar systems to the one which is the focus of this paper.

The same images used for plant classification may also be used to compute visual odometry (VO) measurements. In this paper, VO measurements are assumed to be frame-to-frame rotation and displacements that the robot has undergone between two overlapping images.

The VO technique considered finds a set of matching features in two subsequent images. These features are used to reconstruct the movement of the camera by minimization of the transformation that matches one set to the other. This is similar to wheel odometry (WO) which uses the wheel encoders to reconstruct the movement of the robot. Unlike the WO measurements, VO will be unaffected by any skidding and sliding of the wheels. Both VO and WO are relative displacements with respect to a previous state of the system. Absolute measurements, such as GPS, provide measurements with respect to a known coordinate frame.

The feature identification and matching algorithms introduces a latency between when a picture is taken, and when the measurement becomes available. We refer to this time as processing delay. This processing delay may be more than one second in non-optimized implementations, to an order of milliseconds in implementations as described in Forster et al. (2014), or with dedicated hardware and tight coupling of inertial and VO measurements as in Goldberg and Matthies (2011).

The processing delay for the VO measurement is characteristically different from the other available sensors such as compass, GPS, wheel encoders, which are considered instantaneous in this paper. The process of detecting and matching feature points, and solving for the displacement measurement, is not necessarily fixed. It may vary with the number of feature points processed.

The camera is triggered by a hardware GPIO line. This allows us to know when a picture is taken with the same time reference as the GPS, compass and other measurements.

For the fusion of relative displacement measurements, a method called stochastic cloning (Mourikis et al., 2007) was considered. This method "clones" the state estimates when a measurement should have arrived, augmenting the state vector used by the filter, thus maintaining the cross-covariances between the current state of the system and the time when a measurement should have arrived. Practical examples of stochastic cloning can be found in Romanovas et al. (2013) for visual-inertial/magnetic data fusion, in Van der Merwe and Wan (2004) under the name latency compensation, and in Mourikis et al. (2009) for spacecraft entry, descent and landing.

Similar to Van der Merwe and Wan (2004), we use the UKF as opposed to the EKF. The EKF often suffers from providing covariance estimates that are lower than the actual covariance, something which can be very detrimental in precision agriculture where the decision to spray or not may be directly based on the covariance estimate. Instead we use the Unscented Kalman filter (UKF), which through the unscented transform (UT) deals more directly with non-linearities by choosing a strategic set of sample points (Julier and Uhlmann, 1997).

The main contributions of this paper are

- A presentation of a Bayesian framework for delayed displacement fusion, which build upon the method of stochastic cloning.
- An implementation of a novel fixed-point smoother termed unscented multi-point smoother.
- A description of the relation between stochastic cloning and the proposed framework.

The remainder of this paper is structured as follows: in Section 2 the Bayesian formulation of delayed displacement measurement fusion is presented. The dependencies in the basic model and the delayed displacement scenario are presented with Hidden Markov models. These are used to formulate a Bayesian filter capable of fusing the delayed displacement measurement. In Section 3 the unscented multi-point smoother is implemented. First Gaussian assumptions are applied to the filter algorithm, such that the algorithm is described using an augmented state vector for a multivariate Gaussian distribution. Then the unscented transform is briefly described for evaluating the necessary expectations and covariances. In Section 4 the robot simulation setup is described, followed by simulation results. In Section 5 potential future work is outlined.

2 BAYESIAN FORMULATION

In this section we present the Bayesian formulation of fusion with delayed displacement measurements, using graphical models to indicate stochastic dependencies.

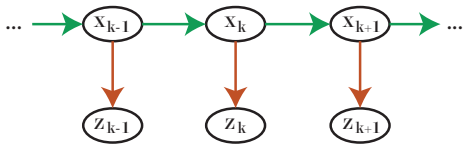

Figure 2: Hidden Markov Model of the basic system (1)-(2) with state vector $\boldsymbol{x}$, and measurement vector $\boldsymbol{z}$. Green arrows represent (1) and red represent (2).

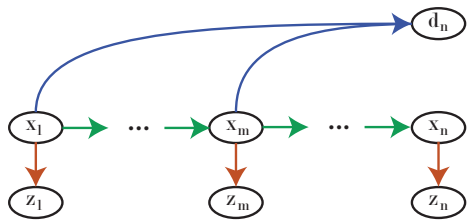

Figure 3: Hidden Markov model showing a delayed displacement measurement $\boldsymbol{d}_n$ being dependent on the states $\boldsymbol{x}_l$ and $\boldsymbol{x}_m$, but arriving at t_n. Green arrows represent (1), red arrows represent (2), and blue arrows represent (5).

2.1 Basic System

The basic underlying system considered in this paper is

$$\boldsymbol{x}_{k+1} = f(\boldsymbol{x}_k) + \boldsymbol{w}_k \tag{1}$$
$$\boldsymbol{z}_k = h(\boldsymbol{x}_k) + \boldsymbol{v}_k \tag{2}$$
$$\boldsymbol{w}_k = \mathcal{N}(\boldsymbol{w}_k; 0, Q_k) \tag{3}$$
$$\boldsymbol{v}_k = \mathcal{N}(\boldsymbol{v}_k; 0, R_k) \tag{4}$$

where $\boldsymbol{x} \in \mathbb{R}^L$ is the state vector of pose, forward velocity, and rotational velocity, $\boldsymbol{z}$ are the measurements from the fast sensors, compass, GPS, and WO. The process noise $\boldsymbol{w}_k$ and measurement $\boldsymbol{v}_k$ is considered additive Gaussian. In Fig. 2 a Hidden Markov model is presented where there is no delay on any measurements, and there are no displacement measurements. The green arrows indicate the transition model, (1), showing how the next state is only conditioned on the previous state. The red arrow indicates the measurement model, (2), and how the likelihood of the measurement is only conditioned on the current state. We desire to find the probability density function (pdf) of the current state given the measurements, $p(\boldsymbol{x}_k|\boldsymbol{z}_{0:k})$. For the basic system this can be done with a recursive Bayesian filter (Thrun et al., 2005).

2.2 Delayed Displacement Measurements

The delayed displacement measurement model is

$$\boldsymbol{d}_n = \tilde{h}(\boldsymbol{x}_l, \boldsymbol{x}_m) + \tilde{\boldsymbol{v}}_n \tag{5}$$
$$\tilde{\boldsymbol{v}}_n = \mathcal{N}(\tilde{\boldsymbol{v}}_n; 0, \tilde{R}_n) \tag{6}$$

where $\boldsymbol{d}_n$ is the delayed displacement measurement dependent on two previous states $\boldsymbol{x}_l$ and $\boldsymbol{x}_m$, and $\tilde{\boldsymbol{v}}_n$ is Gaussian noise. The Hidden Markov model of the system is given in Fig. 3. At time t_l the first picture is taken, and at time t_m the second picture is taken. The displacement measurement becomes available at time t_n. We desire to find the distribution $p(\boldsymbol{x}_n|\boldsymbol{z}_{0:n}, \boldsymbol{d}_n)$.

As the delayed displacement measurement is dependent on two previous states, optimal fusion of $\boldsymbol{d}_n$ requires a joint distribution of $\boldsymbol{x}_l$, $\boldsymbol{x}_m$ and $\boldsymbol{x}_n$.

In reality, the displacement measurements are dependent on the underlying features' poses relative to the camera's pose at time t_l and t_m. The simplification of the displacement measurement as a function of the states is based on Mourikis et al. (2007).

With an initial distribution $p(\boldsymbol{x}_0)$, from t_0 until t_l, the measurements $\boldsymbol{z}_{0:l}$ are fused iteratively using the recursive Bayes filter up to $p(\boldsymbol{x}_l|\boldsymbol{z}_l)$. From there on the filter must maintain $\boldsymbol{x}_l$. This can be done by augmenting the state vector by a "clone" $\boldsymbol{x}_l^c$ of $\boldsymbol{x}_l$ at time t_l. This leads to a joint density as

$$p(\boldsymbol{x}_l, \boldsymbol{x}_l^c|\boldsymbol{z}_l) = p(\boldsymbol{x}_l^c|\boldsymbol{x}_l)p(\boldsymbol{x}_l|\boldsymbol{z}_l) \tag{7}$$

where $p(\boldsymbol{x}_l^c|\boldsymbol{x}_l)$ is a Dirac delta distribution. Thus, even if $p(\boldsymbol{x}_l|\boldsymbol{z}_l)$ is Gaussian, the joint density will not be Gaussian unless a regularization constant ϵ is used. Under Gaussian assumptions,

$$p(\boldsymbol{x}_l, \boldsymbol{x}_l^c|\boldsymbol{z}_l) = \mathcal{N}(\boldsymbol{x}_l^c; \boldsymbol{x}_l, \epsilon I)\mathcal{N}(\boldsymbol{x}_l; \hat{\boldsymbol{x}}_{l|l}, P_{l|l}). \tag{8}$$

A just as straightforward, and in principle more general, solution to the delayed fusion problem follows from "extending" the state vector when the first prediction after time t_l is done:

$$p(\boldsymbol{x}_l, \boldsymbol{x}_{l+1}|\boldsymbol{z}_l) = p(\boldsymbol{x}_{l+1}|\boldsymbol{x}_l)p(\boldsymbol{x}_l|\boldsymbol{z}_l) \tag{9}$$

which under Gaussian assumptions is

$$p(\boldsymbol{x}_l, \boldsymbol{x}_{l+1}|\boldsymbol{z}_l) = \mathcal{N}(\boldsymbol{x}_{l+1}; f(\boldsymbol{x}_l), Q_l)\mathcal{N}(\boldsymbol{x}_l; \hat{\boldsymbol{x}}_{l|l}, P_{l|l}). \tag{10}$$

This is referred to as extending, and occurs as an alternative to prediction at timestep t_{l+1}. In this spirit, we can make further extensions as required. In particular, we also extend at time t_{m+1} to arrive at the joint distribution $p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_{m+1}|\boldsymbol{z}_m)$, which, when iteratively fusing $\boldsymbol{z}$ with a recursive Bayes filter leads to $p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_n|\boldsymbol{z}_n)$.

When a picture has been taken at a timestep, the next predict step in the filter is replaced with an extend step. A system where the camera is not synchronized with the filter will face the additional challenge of estimating the clock synchronization and trigger time.

With the joint distribution $p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_n|\boldsymbol{z}_n)$, the delayed displacement measurement $\boldsymbol{d}_n$ is fused using Bayes theorem according to

$$p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_n|\boldsymbol{z}_n, \boldsymbol{d}_n) = \mu p(\boldsymbol{d}_n|\boldsymbol{x}_l, \boldsymbol{x}_m) p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_n|\boldsymbol{z}_n) \tag{11}$$

where μ is the normalization constant and $p(\boldsymbol{d}_n|\boldsymbol{x}_l, \boldsymbol{x}_m)$ is the likelihood of the displacement measurement. Note that the likelihood is specified conditional on both $\boldsymbol{x}_l$ and $\boldsymbol{x}_m$, but not on $\boldsymbol{x}_n$. If no other displacement measurements depend on the states $\boldsymbol{x}_l$ and $\boldsymbol{x}_m$, the states are marginalized from the joint distribution by

$$p(\boldsymbol{x}_n|\boldsymbol{z}_n, \boldsymbol{d}_n) = \int_{\boldsymbol{x}_l} \int_{\boldsymbol{x}_m} p(\boldsymbol{x}_l, \boldsymbol{x}_m, \boldsymbol{x}_n|\boldsymbol{z}_n, \boldsymbol{d}_n) \mathrm{d}\boldsymbol{x}_m \mathrm{d}\boldsymbol{x}_l \tag{12}$$

which under Gaussian assumptions is simply done by omitting the parts of the expectation vector and covariance matrix associated with $\boldsymbol{x}_l$ and $\boldsymbol{x}_m$.

In Alg.1 the delayed displacement fusing filter algorithm is outlined. This algorithm is a Bayesian on-demand smoother. When a picture has been taken, the state vector is extended with the predicted state. And after updating with a displacement measurement, any unnecessary old states are marginalized. The list of state indices when the pictures have been taken is denoted K. The set N contains the current index i only if a displacement measurement is available. The augmented state vector $\boldsymbol{x}_S$ contains all the states maintained in the joint distribution.

The purpose of the list S is to keep track of the corresponding state indices. In the algorithm $\cup$ is used to append an index, and $\setminus$ is used to remove indices. The function $ind(\boldsymbol{d})$ returns the indices associated with a displacement measurement $\boldsymbol{d}$ (e.g. l and m) that can be marginalized. The function ind and the list S are used to handle the case when a picture is used for more than one displacement measurement. This is generally the case, as an image generates displacement measurements both together with the preceeding and the succeeding image.

Algorithm 1 Bayesian delayed displacement fusion

Require: $p(\boldsymbol{x}_0)$, K, N
1: $S \leftarrow \{0\}$
2: **for** $i = 1$ **to** $i = \infty$ **do**
3: *Predict:*
4: **if** $i \in K$ **then** {*picture taken*}
5: $p(\boldsymbol{x}_S, \boldsymbol{x}_i|\boldsymbol{z}_{0:i-1}) = p(\boldsymbol{x}_i|\boldsymbol{x}_{i-1})p(\boldsymbol{x}_S)$
6: $S \leftarrow S \cup \{i\}$
7: **else**
8: $p(\boldsymbol{x}_S, \boldsymbol{x}_i|\boldsymbol{z}_{0:i-1}) = \int_{\boldsymbol{x}_{i-1}} p(\boldsymbol{x}_i|\boldsymbol{x}_{i-1})p(\boldsymbol{x}_S)\mathrm{d}\boldsymbol{x}_{i-1}$
9: $S \leftarrow S \setminus \{i-1\}$
10: $S \leftarrow S \cup \{i\}$
11: **end if**
12: *Update:*
13: **if** $i \in N$ **then** {*displacement available*}
14: $p(\boldsymbol{x}_S|\boldsymbol{z}_{0:i}, \boldsymbol{d}_i) = \mu p(\boldsymbol{z}_i, \boldsymbol{d}_i|\boldsymbol{x}_S)p(\boldsymbol{x}_S|\boldsymbol{z}_{0:i-1})$
15: $A \leftarrow ind(\boldsymbol{d}_i)$
16: $p(\boldsymbol{x}_i|\boldsymbol{z}_{0:i}, \boldsymbol{d}_i) = \int_{\boldsymbol{x}_A} p(\boldsymbol{x}_S|\boldsymbol{z}_{0:i}, \boldsymbol{d}_i)\mathrm{d}\boldsymbol{x}_A$
17: $S \leftarrow S \setminus A$
18: **else**
19: $p(\boldsymbol{x}_S|\boldsymbol{z}_{0:i}, \boldsymbol{d}_i) = \mu p(\boldsymbol{z}_i|\boldsymbol{x}_i)p(\boldsymbol{x}_S|\boldsymbol{z}_{0:i-1})$
20: **end if**
21: **end for**

K is updated by the camera triggering before each iteration and N is updated before the update step by the VO module indicating a displacement measurement being ready. The filter acts similar to a fixed-point smoother with the capability of formulating the points to be smoothed on-demand, a "multi-point" smoother.

3 UNSCENTED MULTI-POINT SMOOTHER

In this section we look at an unscented multi-point smoother based on Alg.1. First we describe more thoroughly what the Gaussian assumptions means for our augmented state vector $\boldsymbol{x}_S$, then we present important properties of the unscented transformation, and show how the method in this paper relates to stochastic cloning.

3.1 Gaussian Assumption

To implement Alg.1 with a UKF as the underlying filter, one must define how the extension method, (10), under Gaussian assumptions, behaves for the augmented state vector $\boldsymbol{x}_S$. Consider the case where we

are extending at time t_k

$$p(\boldsymbol{x}_S, \boldsymbol{x}_k|\boldsymbol{z}_{k-1}) = p(\boldsymbol{x}_k|\boldsymbol{x}_S)p(\boldsymbol{x}_S|\boldsymbol{z}_{k-1}) = \mathcal{N}\left(\begin{bmatrix}\boldsymbol{x}_S \\ \boldsymbol{x}_k\end{bmatrix}; \begin{bmatrix}\mathrm{E}(\boldsymbol{x}_S) \\ \mathrm{E}(\boldsymbol{x}_k)\end{bmatrix}, \begin{bmatrix}\mathrm{Var}(\boldsymbol{x}_S) & \mathrm{Cov}(\boldsymbol{x}_S, \boldsymbol{x}_k) \\ \mathrm{Cov}(\boldsymbol{x}_k, \boldsymbol{x}_S) & \mathrm{Var}(\boldsymbol{x}_k)\end{bmatrix}\right). \tag{13}$$

The notation $\mathrm{Cov}(\boldsymbol{x}_S, \boldsymbol{x}_k)$ is shorthand for

$$\mathrm{Cov}(\boldsymbol{x}_S, \boldsymbol{x}_k) = \begin{bmatrix}\mathrm{Cov}(\boldsymbol{x}_{S_1}, \boldsymbol{x}_k) \\ \vdots \\ \mathrm{Cov}(\boldsymbol{x}_{k-1}, \boldsymbol{x}_k)\end{bmatrix} \tag{14}$$

where S_1 indicates the oldest index in the list of state indices S. The newest index is always $k-1$. The mean and covariance of $\boldsymbol{x}_S$ is maintained by the filter. We must evaluate the parts associated with $\boldsymbol{x}_k$. In the following section we describe $\mathrm{E}(\boldsymbol{x}_k)$ and $\mathrm{Var}(\boldsymbol{x}_k)$ with the unscented transform. For the cross-covariances between $\boldsymbol{x}_S$ and $\boldsymbol{x}_k$, (14), consider S^+ to be a list of state indices containing the same indices as S, with the last index, $k-1$ replaced with k. Then the cross-covariance between $\boldsymbol{x}_S$ and $\boldsymbol{x}_{S^+}$ is by definition

$$\mathrm{Cov}(\boldsymbol{x}_S, \boldsymbol{x}_{S^+}) = \begin{bmatrix}\mathrm{Var}(\boldsymbol{x}_{S_1}) & \dots & \mathrm{Cov}(\boldsymbol{x}_{S_1}, \boldsymbol{x}_k) \\ \vdots & \ddots & \vdots \\ \mathrm{Cov}(\boldsymbol{x}_{k-1}, \boldsymbol{x}_{S_1}) & \dots & \mathrm{Cov}(\boldsymbol{x}_{k-1}, \boldsymbol{x}_k)\end{bmatrix}. \tag{15}$$

Hence, evaluating the cross-covariance between $\boldsymbol{x}_S$ and $\boldsymbol{x}_{S^+}$, and extracting the rightmost columns of size L, gives us (14). This augmented system keeps all but the last state the same between each iteration.

$$\boldsymbol{x}_{S^+} = f^a(\boldsymbol{x}_S) = \begin{bmatrix}\boldsymbol{x}_{S_1} \\ \boldsymbol{x}_{S_2} \\ \vdots \\ f(\boldsymbol{x}_{k-1})\end{bmatrix} + \begin{bmatrix}0 \\ 0 \\ \vdots \\ \boldsymbol{w}_{k-1}\end{bmatrix} \tag{16}$$

$$\boldsymbol{z}_k = h^a(\boldsymbol{x}_S^+) + \boldsymbol{v}_k = h(\boldsymbol{x}_k) + \boldsymbol{v}_k. \tag{17}$$

The augmented version of the displacement measurement model $\tilde{h}$ must know which states in the augmented state vector $\boldsymbol{x}_S$ correspond to the imminent displacement measurement

$$\boldsymbol{d}_n = \tilde{h}^a(\boldsymbol{x}_S) + \tilde{\boldsymbol{v}}_n. \tag{18}$$

3.2 Unscented Transform

In this section we briefly introduce the unscented transform and the properties that need to be evaluated for Alg.1.

The basic premise of the unscented transform is that it is easier to approximate a probability distribution than a nonlinear function (Julier and Uhlmann, 2002). This is done by propagating a set of deterministically chosen sigma-points through the nonlinear function. For a nonlinear transformation

$$\boldsymbol{y} = f(\boldsymbol{x}) \tag{19}$$

where $\boldsymbol{x} \in \mathbb{R}^L$ is a Gaussian random variable of dimension L. Sigma-points can be defined by the selection scheme

$$\mathcal{Y}_0 = \mathrm{E}(\boldsymbol{x}) \tag{20}$$

$$\mathcal{Y}_h = \mathrm{E}(\boldsymbol{x}) + \sqrt{(L+\lambda)\mathrm{Var}(\boldsymbol{x})_i} \tag{21}$$

$$\mathcal{Y}_j = \mathrm{E}(\boldsymbol{x}) - \sqrt{(L+\lambda)\mathrm{Var}(\boldsymbol{x})_i} \tag{22}$$

$$w_0 = 1/(L+\lambda) \tag{23}$$

$$w_{h,j} = 1/2(L+\lambda) \tag{24}$$

for $h = 1, ..., L$, $i = 1, ..., L$, and $j = L+1, ..., 2L+1$, where λ is a tuning parameter defining the spread of the sigma-points around the expected value (Van der Merwe, 2004).

The subscript i of the matrix $\mathrm{Var}(\boldsymbol{x})$ denotes its ith column vector. From Van der Merwe (2004), under the Gaussian assumption, the mean, covariance and cross-covariance are constructed from the sigma-points by

$$\mathrm{E}(\boldsymbol{y}) = \sum_{i=0}^{2L+1} w_i \boldsymbol{\mathcal{Y}}_i, \tag{25}$$

$$\mathrm{Var}(\boldsymbol{y}) = \sum_{i=0}^{2L+1} w_i \left(\boldsymbol{\mathcal{Y}}_i - \mathrm{E}(\boldsymbol{y})\right)\left(\boldsymbol{\mathcal{Y}}_i - \mathrm{E}(\boldsymbol{y})\right)^T, \tag{26}$$

and

$$\mathrm{Cov}(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=0}^{2L+1} w_i \left(\boldsymbol{\mathcal{X}}_i - \mathrm{E}(\boldsymbol{x})\right)\left(\boldsymbol{\mathcal{Y}}_i - \mathrm{E}(\boldsymbol{y})\right)^T. \tag{27}$$

These methods allow us to evaluate the expectation, covariance, and cross-covariance of the augmented system described in the previous section, and the unscented multi-point smoother can be constructed.

3.3 Remark on Stochastic Cloning

If the underlying filter is an EKF, then $f(\boldsymbol{x}_k) \approx F_k \boldsymbol{x}_k$. The cloning procedure of (7) without a regularization constant constructs the augmented expectation state vector $\hat{\boldsymbol{x}}_{S|k}$, and covariance matrix $P^S_{k|k}$ of the filter according to

$$\hat{\boldsymbol{x}}_{S|k} = \begin{bmatrix} \hat{\boldsymbol{x}}_{k|k} \\ \hat{\boldsymbol{x}}_{k|k} \end{bmatrix} \tag{28}$$

$$P_{S|k} = \begin{bmatrix} P_{k|k} & P_{k|k} \\ P_{k|k} & P_{k|k} \end{bmatrix}. \tag{29}$$

This does not describe a Gaussian distribution as the covariance matrix is singular. When performing the prediction step, an UKF will fail as it relies on the Cholesky factorization of the state covariance matrix, which is only unique on nonsingular matrices. On the other hand, an EKF can readily perform the prediction step with an augmented transition model giving the expectation state vector and covariance matrix given by

$$\hat{\boldsymbol{x}}_{S^+|k} = \begin{bmatrix} \hat{\boldsymbol{x}}_{k|k} \\ F_k \hat{\boldsymbol{x}}_{k|k} \end{bmatrix} \tag{30}$$

and

$$P_{S^+|k} = \begin{bmatrix} P_{k|k} & P_{k|k} F_k^T \\ F_k P_{k|k} & F_k P_{k|k} F_k^T + Q_k \end{bmatrix}, \tag{31}$$

which is identical to the extension method described in this paper. It is interesting to note that by Schur's complement, we require non-zero process noise to ensure that the distribution remains Gaussian through an extend step.

3.4 Remark on Smoothing

The filter algorithm, Alg.1, performs optimal fixed-point smoothing of the augmented states under the assumption that the model can be described by (1), (2) and (5). If the transformations are highly nonlinear, the unscented transform will also encounter problems. As such, it is debatable whether the smoothing procedure is beneficial. When absolute measurements are available from e.g. GPS or a compass, it was observed that the covariance associated with an old augmented state diminished to the point where the delayed displacement measurement was assumed to be more accurate than it actually was. To remedy this, for the simulation with high process delay and low camera framerate, the covariance associated with the oldest augmented state had a lower threshold it could not decrease below. This was used on the pose of the robot as these were the states affected by the absolute measurements.

4 ROBOT SIMULATION

4.1 System

The robot is modeled as a unicycle-like robot with no-slip conditions. The kinematic model for the system is based on the model by Cruz and Carelli (2006), with a change where the castor wheel is at the rear end of the robot,

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \\ \dot{u} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} u\cos\psi - a\omega\sin\psi \\ u\sin\psi + a\omega\cos\psi \\ \omega \\ 0 \\ 0 \end{bmatrix} \tag{32}$$

where x and y is the robot's east-north position, ψ is yaw, u is the forward speed, and ω is the yaw rate. The kinematic model (32) was implemented by using the forward Euler method and Gaussian process noise was added:

$$\dot{\boldsymbol{x}} = f_c(\boldsymbol{x}) \tag{33}$$

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + f_c(\boldsymbol{x}_k)(t_{k+1} - t_k) + \boldsymbol{w}_k. \tag{34}$$

We assume that the onboard sensors are not delayed, except for the visual odometry measurements, when setting up the observation mapping $h(\boldsymbol{x}_k)$, (2). We model the onboard RTK-GPS, compass, and wheel odometry, as direct measurements of the position, (x, y), the heading (ψ), and linear and angular velocities (u, ω) respectively. The measurements are assumed to be affected by additive Gaussian measurement noise.

The measurement $\boldsymbol{d}_n$ is rotated by the oldest picture frame,

$$\boldsymbol{d}_n = \begin{bmatrix} R(-\psi_k) & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_m - x_k \\ y_m - y_k \\ \psi_m - \psi_k \end{bmatrix} + \tilde{\boldsymbol{v}}_n \tag{35}$$

where $R(\cdot)$ is the 2D rotation matrix. Using the simulink models and controller of Martins (2013) the robot was simulated with a timestep of 0.05 s, following a figure eight trajectory. As the robot is heavy and slow moving, the process noise is negligible. The acronym JUKF is used for the delayed displacement fusing UKF in reference to the joint distribution. The kinematic controller has an upper bound of 1.5 m/s in these simulations.

For the simulations, the errors of the filter estimates are

$$\begin{bmatrix} \tilde{x}_k \\ \tilde{y}_k \\ \tilde{\psi}_k \\ \tilde{u}_k \\ \tilde{\omega}_k \end{bmatrix} = \boldsymbol{x}_k - \hat{\boldsymbol{x}}_{k|k}. \tag{36}$$

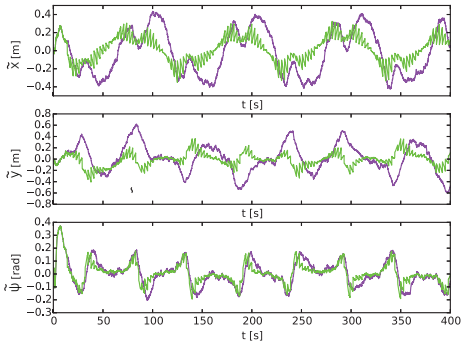

Figure 4: Simulation of the robot showing pose error with respect to time. Simulated with $\tau_C = 3s$ and $\tau_D = 2s$. Purple indicates UKF, green indicates JUKF.

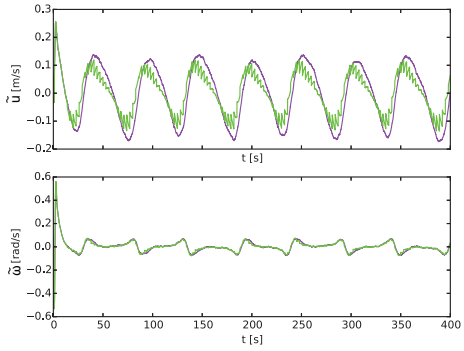

Figure 5: Simulation of the robot showing velocity error with respect to time. Simulated with $\tau_C = 3s$ and $\tau_D = 2s$. Purple indicates UKF, green indicates JUKF.

4.2 Simulation Results

In Test 1 the time between images, τ_C, is 0.25 s, and the processing delay, τ_D is 0.2 s. The RMS errors are given in Tab. 1. The JUKF fusing the delayed displacement measurements has a lower RMS error than the UKF not fusing the delayed displacement measurement.

In Test 2 the time between images is 2 s, and the processing delay is 1 s. The RMS errors are given in Tab. 2. The JUKF fusing the delayed displacement measurement has a lower RMS error than the UKF not fusing the delayed displacement measurement.

In Test 3 the time between images is 3 s, and the processing delay is 2 s. The RMS errors are given in Tab. 3. See Fig. 4 for the error in pose with respect to time, and Fig. 5 for error in velocities with respect to time. The purple line is the UKF, the green is the JUKF. Note that whenever a VO measurement is fused, the error in JUKF decreases, causing a sawtooth effect.

In Test 4 the time between images is 5 s, and the processing delay is 3 s. The RMS errors are given in Tab. 4. In this test a covariance threshold was applied to the covariances associated with the oldest x, y and ψ. Without the covariance threshold, the JUKF did not have lower RMS errors than the UKF. At this frame rate, the pictures are not expected to overlap for the robot under typical field conditions. This test was included to show that the algorithm show improvement in the localization RMS error for cases beyond the worst expected latencies.

Table 1: RMS errors of the simulation Test 1, $\tau_C = 0.25s$, $\tau_D = 0.2s$.

	x	y	ψ	u	ω
UKF	0.272287	0.239938	0.089938	0.100554	0.045322
JUKF	0.040214	0.043980	0.034718	0.017865	0.032077

Table 2: RMS errors of the simulation Test 2, $\tau_C = 2s$, $\tau_D = 1s$

	x	y	ψ	u	ω
UKF	0.262771	0.233008	0.098997	0.101097	0.058464
JUKF	0.093007	0.115435	0.065567	0.051329	0.053759

5 Future Work and Discussion

The algorithm proposed in this paper handles a mild case of out-of-sequence measurements, in the sense that displacement measurements are only processed after several GPS and compass measurements have been processed, i.e. out-of-sequence. The key to handling this is the *ind* function, declaring which augmented states are to be marginalized.

In the simulations presented here, we have only considered GPS and compass measurements in addition to VO. In practice, most platforms will also have inertial

Table 3: RMS errors of the simulation Test 3, $\tau_C = 3s$, $\tau_D = 2s$

	x	y	ψ	u	ω
UKF	0.251910	0.249378	0.099966	0.101343	0.058235
JUKF	0.138441	0.117433	0.081633	0.072851	0.056340

Table 4: RMS errors of the simulation Test 4, $\tau_C = 5s$, $\tau_D = 4s$.

	x	y	ψ	u	ω
UKF	0.278994	0.252959	0.098941	0.102308	0.057945
JUKF	0.265563	0.244570	0.09856	0.097869	0.057328

sensors available, which may be included in the measurement model. When using MEMS sensors, it would also be natural to include the gyro and accelerometer biases in the filter.

Finally, this work employed a full-state representation as opposed to an error-state representation, which has been used for visual odometry as described in Mourikis and Roumeliotis (2007). In continuation of this research, it may be interesting to investigate how error-state smoothing relates to the Bayesian paradigm.

6 Conclusion

This paper has proposed a Bayesian framework for the fusion of delayed displacement measurements. This led to a multi-point smoother capable of defining fixed-points to be smoothed on demand according to the GPS, WO, and compass measurements. This filtering technique may prove useful in other scenarios where fixed-points to be smoothed are formulated while the filter is running.

By using the unscented transform, the filter was able to use the smoothing effects of the fusion method for its benefits. The filter maintains estimates of the state the robot was in when a picture was taken, smoothed by the absolute measurements. The VO fusing methods improved the localization RMS error for a variety of latencies.

References

Biber, P., Weiss, U., Dorna, M., and Albert, A. Navigation System of the Autonomous Agricultural Robot "BoniRob". *Workshop on Agricultural Robotics: Enabling Safe, Efficient, and Affordable Robots for Food Production*, 2012. pages 1–7.

Bogue, R. Robots poised to revolutionise agriculture. *Industrial Robot: An International Journal*, 2016. 43(5):450–456. doi:10.1108/IR-05-2016-0142.

Cruz, C. D. L. and Carelli, R. Dynamic Modeling and Centralized Formation Control of Mobile Robots. In *IECON 2006 - 32nd Annual Conference on IEEE Industrial Electronics*. pages 3880–3885, 2006. doi:10.1109/IECON.2006.347299.

Forster, C., Pizzoli, M., and Scaramuzza, D. SVO: Fast semi-direct monocular visual odometry. *Proceedings - IEEE International Conference on Robotics and Automation*, 2014. pages 15–22. doi:10.1109/ICRA.2014.6906584.

Goldberg, S. B. and Matthies, L. Stereo and IMU assisted visual odometry on an OMAP3530 for small robots. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2011. doi:10.1109/CVPRW.2011.5981842.

Grimstad, L., Pham, C. D., Phan, H. T., and From, P. J. On the design of a low-cost, light-weight, and highly versatile agricultural robot. *Proceedings of IEEE Workshop on Advanced Robotics and its Social Impacts, ARSO*, 2016. 2016-March(October). doi:10.1109/ARSO.2015.7428210.

Grimstad, L., Phan, H. N. T., Pham, C. D., and Bjugstad, N. Initial field-testing of Thorvald , a versatile robotic platform for agricultural applications. *IEEE International Workshop on Advanced Robotics and its Social Impacts*, 2015.

Jørgensen, R. N., Sørensen, C. G., Pedersen, J. M., Havn, I., Jensen, K., Søgaard, H. T., and Sørensen, L. Hortibot: A system design of a robotic tool carrier for high-tech plant nursing. *CIGR Ejournal*, 2007. IX:Manuscript ATOE 07 006.

Julier, S. J. and Uhlmann, J. K. New extension of the kalman filter to nonlinear systems. 1997. doi:10.1117/12.280797.

Julier, S. J. and Uhlmann, J. K. Reduced sigma point filters for the propagation of means and covariances through nonlinear transformations. *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, 2002. 2:887–892 vol.2. doi:10.1109/ACC.2002.1023128.

Martins, F. Velocity-based dynamic model and adaptive controller for differential steered mobile robot. Mathworks Fileexchange 44850, 2013.

Van der Merwe, R. *Sigma-point Kalman filters for probabilistic inference in dynamic state-space models*. Ph.D. thesis, Oregon Health & Science University, 2004.

Van der Merwe, R. and Wan, E. A. Sigma-Point Kalman Filters for Integrated Navigation. In *Proceedings of the 60th Annual Meeting of the Institute of Navigation (ION)*. pages 641–654, 2004.

Molstad, L., Reent Köster, J., Bakken, L., Dörsch, P., Lien, T., Overskeid, Ø., Utstumo, T., Løvås, D., and Brevik, A. A field robot for autonomous laser-based n2o flux measurements. In *EGU General Assembly Conference Abstracts*, volume 16. page 16109, 2014.

Mourikis, A. I. and Roumeliotis, S. I. A multi-state constraint Kalman filter for vision-aided inertial navigation. *Proceedings - IEEE International Conference on Robotics and Automation*, 2007. pages 3565–3572. doi:10.1109/ROBOT.2007.364024.

Mourikis, A. I., Roumeliotis, S. I., and Burdick, J. W. SC-KF mobile robot localization: A stochastic cloning Kalman filter for processing relative-state measurements. *IEEE Transactions on Robotics*, 2007. 23(4):717–730. doi:10.1109/TRO.2007.900610.

Mourikis, A. I., Trawny, N., Roumeliotis, S. I., Johnson, A. E., Ansar, A., and Matthies, L. Vision-aided inertial navigation for spacecraft entry, descent, and landing. *IEEE Transactions on Robotics*, 2009. doi:10.1109/TRO.2009.2012342.

Romanovas, M., Schwarze, T., Schwaab, M., and Traechtler, M. Stochastic Cloning Kalman Filter for Visual Odometry and Inertial / Magnetic Data Fusion. *Information Fusion (FUSION), 2013*, 2013. pages 1434–1441.

Slaughter, D. C., Giles, D. K., and Downey, D. Autonomous robotic weed control systems: A review. *Computers and Electronics in Agriculture*, 2008. 61(1):63–78. doi:10.1016/j.compag.2007.05.008.

Thrun, S., Burgard, W., and Fox, D. *Probabilistic robotics*. MIT press, 2005.

Urdal, F., Utstumo, T., Vatne, J. K., Ellingsen, S. A. A., and Gravdahl, J. T. Design and control of precision drop-on-demand herbicide application in agricultural robotics. In *2014 13th International Conference on Control Automation Robotics and Vision, ICARCV 2014*. IEEE, Singapore, pages 1689–1694, 2014. doi:10.1109/ICARCV.2014.7064570.

Utstumo, T., Berge, T. W., and Gravdahl, J. T. Nonlinear model predictive control for constrained robot navigation in row crops. *2015 IEEE International Conference on Industrial Technology (ICIT)*, 2015. (July):357–362. doi:10.1109/ICIT.2015.7125124.

Utstumo, T. and Gravdahl, J. T. Implementation and comparison of attitude estimation methods for agricultural robotics. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 2013. 4(PART 1):52–57. doi:10.3182/20130828-2-SF-3019.00051.

2.6 Journal Paper - Robotic In-Row Weed Control

T. Utstumo, F. Urdal, A. Brevik, J. Dørum, J. Netland, Ø. Overskeid, T. W. Berge, and J. T. Gravdahl. Robotic in-row weed control for vegetables. *Computers and Electronics in Agriculture, Submitted for review, February 28, 2018.*

Robotic in-row weed control in vegetables

Trygve Utstumo[a,b,*], Frode Urdal[a], Anders Brevik[a], Jarle Dørum[a], Jan Netland[c], Øyvind Overskeid[a], Therese W. Berge[c], Jan Tommy Gravdahl[b]

[a] *Adigo AS, Berghagan 3, NO-1405 Langhus, Norway*
[b] *Dept. of Engineering Cybernetics, Norwegian University of Science and Technology (NTNU), NO-7034 Trondheim, Norway*
[c] *Norwegian Institute of Bioeconomy Research (NIBIO), P.O. Box. 115, NO-1431 Ås, Norway*

Abstract

Vegetables and other row-crops represent a large share of the agricultural production. There is a large variation in crop species, and a limited availability in specialized herbicides. The robot presented here utilizes systematic growing techniques to navigate and operate in the field. By the use of machine vision it seperates seeded vegetable crops from weed, and treat each weed within the row with individual herbicide droplets, without affecting the crop. This allows for the use of herbicides that would otherwise harm the crop and results in a significant reduction in herbicide use.

The robot is tailored to this purpose with cost, maintainability, efficient operation and robustness in mind. The three-wheeled design is unconventional, and the design maintains maneuverability and stability with the benefit of reduced weight, complexity and cost.

Indoor pot trials with four weed species demonstrated that the Drop-on-Demand system (DoD) could control the weeds with as little as 7.6 µg glyphosate or 0.15 µg iodosulfuron per plant. The results also highlight the importance of liquid characteristics for droplet stability and leaf retention properties. The common herbicide glyphosate had no effect unless mixed with suitable additives. A field trial with the robot was performed in a carrot field, and all the weeds were effectively controlled with the DoD system applying 5.3 µg of glyphosate per droplet. The robot and DoD system represent a paradigm shift to the environmental impact and health risks of weed control, while providing a valuable tool to the producers.

Keywords: precision agriculture, drop-on-demand, weed control, agricultural robotics,

1. Introduction

The production of row crops represent a significant portion of the overall food production in the world. This production is composed of large variety of crops of which each individual crop has a smaller volume. In contrast to major crops such as corn, soy and cereal, the vegetable crops have a smaller selection of available herbicides. In the past 20 years we have seen a significant increase in herbicide resistant weeds (Heap, 2014), while the availability of herbicides has been reduced by regulations due to health and environmental concern. The end result is an increasingly challenging situation for farmers who are left with fewer efficient herbicides.

Weed control is one of the most important factors in all agricultural production. Weeds compete with crop plants for moisture, nutrients and sunlight and will have a significant negative impact on yield without sufficient weed control. Typical weed control methods for row crops include a combination of pre-emergence herbicide application, pre-emergence tillage, mechanical row harrowing and post-emergence herbicide application - if a selective herbicide or crop resistance is available, (Slaughter et al., 2008; Fennimore et al., 2016).

In 2008, the European Commission withdrew the approval for several herbicides, among them herbicides with Propachlor as the active ingredient (European Commission, 2008). The herbicide was a health risk and had been documented contaminating ground water and harmful to aquatic life. The consequence to farmers of some cabbages and rutabaga was that they lost access to their most effective herbicide. In Norway this spurred a joint project with farmers and the Norwegian Extension Service in the search for alternative weed control methods, which one could say marked the start of the work presented here.

The weed that occur in between rows, inter-row weeds, can be controlled by row-harrowing, flaming or shielded spraying. Whereas the in-row weeds pose a greater challenge for the farmers. In lack of selective post-emergence herbicides they are left with few other options than manual in-row hoeing by hand, which is much more expensive than conventional spraying.

*Corresponding author
Email address: trygve@adigo.no (Trygve Utstumo)

In the past 10 - 20 years we have seen a significant push to bring new methods to the farmers to control in-row weeds. And for transplanted crops, there are methods available with vision-controlled in-row harrowing such as the *Garford Robocrop In-row weeder, Steketee IC Weeder* and *F. Poulsen Engineering Robovator*. The transplanted crops are relatively sparse and allow for these methods, as well as selective spraying where two notable examples are the companies *BlueRiver Technologies* and *Ecorobotix*.

Seeded crops present a greater challenge as there isn't enough room in between crop plants to allow for a mechanical hoe to pass in and out of the crop row. Herbicide application either requires a selective herbicide which does not harm the crop, or a better resolution application to not affect the crop. DoD herbicide application, Figure 1, is one of the most promising technologies for controlling weeds in the plant row (Fennimore et al., 2016; Slaughter et al., 2008). The resolution in this paper is taken to the extreme by controlling individual droplets of herbicide, Figure 1.

The essence of DoD spraying is to detect the weeds within the plant row, and selectively shoot droplets of herbicide on those weed leaves. By targeting only the weed leaves, the crop and soil are left unaffected, which allows for the use of broad spectre herbicides that would normally harm the crop.

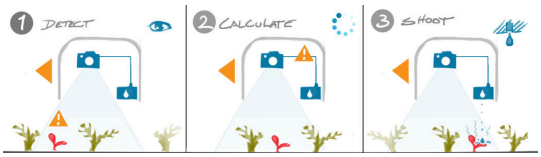

Figure 1: Visualization on Drop-on-Demand herbicide application.

Figure 2: The 2017 Asterix robot prototype in field trials in Central Norway.

We have focused much of our attention to carrots, as we consider it a good example of the more challenging crops. It is a seeded culture which account for 6.25 % of Europe's harvested area for vegetables, with 2.6 million Ha. It is a high value crop with a gross production value for Europe above 3 billion USD in 2014 (FAO, 2014).

Carrot competes poorly with weeds especially in the early stages, as documented by Swanton et al. (2010) in a field trial in Ontario, Canada. The critical weed-free period for carrots was found to be 450 growing-degree-days (3 to 6 weeks at 10 to 20°C), or until the carrot plants have reached the six-leaf stage.

While there are commercially available products for in-row mechanical hoeing, we are not aware of other commercially viable projects providing a DoD weed control system. This paper will present the newly developed autonomous robot platform shown in Figure 2, and a novel system for drop-on-demand (DoD) application of herbicide. Finally, successful results from laboratory and field tests are reported.

We also present a system for flushing the valves, and handling excess spray liquid.

2. State of the art

The available products for guided hoeing and selective thinning are paving the way for further advances in automatic weed control in speciality crops. Our attention will be focused on precision-spray application targeting individual weeds - a domain which is yet to see its first commercially available solution.

One of the first demonstrations of a Precision-Spray robot was by Lee, Slaughter, and Giles as early as 1999. They developed a robot for controlling weeds in tomato crops. The robot was equipped with an Cohu RGB camera which information was digitized to 256x240 pixels at 8 bit per channel. The processing was done by a 200 MHz Pentium Pro CPU running MSDOS. The system recognized 73 % of the tomato plants and 69 % of the weeds, and was able to treat 48 % of the weeds at a speed of 0.8 km/h.

Nearly 20 years has passed since then, and while the robots has become incrementally better, we are yet to see weeding robots make an impact on the use of herbicides in agriculture. A thorough overview of this field can be found in Fennimore et al. (2016) or Slaughter et al. (2008), while we here will focus on a few relevant technical aspects.

2.1. Drop-On-Demand herbicide application

A challenge presented by Lee et al. (1999) is to increase the accuracy, precision and efficacy of the herbicide application. This effort involves everything from the design of the droplet forming mechanism, the fluid dynamics of the droplets, the droplets retention on the weed leaves, the choice of active ingredient, to the motion estimation and targeting algorithm.

Most of the previously presented systems for DoD herbicide application has either used adapted industrial print-heads (Lund and Mathiassen, 2010; Midtiby et al., 2011) or an array of solenoid valves and needles (Søgaard and Lund, 2005; Lee et al., 1999; Nieuwenhuizen, 2009) to form droplets. There is also a presented paper by Basi et al. (2012) where a pneumatic valve is presented for better dosing and formation of individual droplets. The fluid dynamics of the in-flight droplets has been investigated by Lund and Mathiassen (2010) and Lund and Olsen (2010). They describe the disintegration of droplets and the effects of altering the viscosity and surface tension of the fluid. We expanded on this and also explored the effect of the electrical control signal to the solenoid valve on the droplet formation in our experiments presented in Urdal et al. (2014).

Lund and Mathiassen (2010) and Lund et al. (2006) demonstrated that herbicide droplets formulated with glyphosate (27µg per plant) can effectively control *Solanum nigrum* L., (Black Nightshade) a weed which is resistant to most selective herbicides. Midtiby et al. (2011) presented a simulated row crop trial where plants passed under the system on a conveyor belt at 0.5 m/s. The system was able to effectively control weeds larger than 11x11 mm, which gave good results on *Brassica napus* L. (oilseed rape) and to some extent *Tripleurospermum inodorum*(L) Sch. Bip. (Scentless Mayweed). Koukiasas et al. (2016) demonstrated that *Galium aparine* L. (Cleavers) is effectively controlled with 19.3 µg of glyphosate per plant.

2.2. Leaf Classification

Weed and crop classification has largely followed the classical approach of segmenting plant material from the background soil, for subsequent classification based on shape, color and texture features. Several systems have incorporated a Near-Infrared (NIR) channel to enhance the soil segmentation, e.g.: Nieuwenhuizen (2009). These classifiers has been demonstrated with high accuracy. They are however highly reliant on shape features and do not demonstrate satisfactory robustness when challenged by overlapping leaves and irregularities such as specular reflection from water droplets. There has been much effort invested in improving these algorithms (Fennimore et al., 2016). One example is Haug et al. (2014), who was able to circumvent the reliance on segmenting individual plants by implementing a form of sliding-window classifier. As a result, the classifier was robust to overlapping leaves.

In recent years there has been an important shift in Computer Vision towards deep learning. In nearly all domains we see classification tasks being taken over by deep convolutional neural networks (Deep CNN). These methods are also making their way into weed detection. One out of several examples is Milioto et al. (2018), who demonstrate pixel-wise semantic segmentation into weed and crop.

2.3. State-of-the art in Agricultural Robotic Platforms

There is a significant body of research and industrial push towards robotization in agriculture. There are philosophies towards automating tractors, building specialized robots for each task and towards making highly versatile and modular robots. A selection of comparable robots that have been presented for weed control is shown in Figure 3.

(a) DeepField Robotics BoniRob, photo courtesy of Bosch AG

(b) Ecorobotix, photo courtesy of Ecorobotix Ltd

(c) Thorvald II platform, photo courtesy of Saga Robotics AS.

(d) AgBot II, photo courtesy of Queensland University of Technology

Figure 3: A selection of other robot platforms presented in litterature.

Modularity has been uphold as an important design criteria for the Armadillo (Nielsen et al., 2012), Naïo Dino from *Naïo Technologies* and the Thorvald II platform (Grimstad and From, 2017) which can be customized to different configurations. Thorvald II, BoniRob, (Figure 3a and 3c), and Naïo Dino have drive and steering on all four wheels. This enables holonomic control of the robot: The robot can navigate freely in all directions, handle tight environments such as greenhouses and the front and rear wheels can follow the same tracks through a turn. This comes at the cost of having 8 motors for steering and drive.

The AgBot II shown in Figure 3d and presented by Bawden et al. (2017) is a robot platform for weed control, set up with differential drive front wheels and two rear castor wheels. The design emphasizes modularity and ease of on-site assembly of the system. A docking container covered with solar panels provide the power needs for charging, and the system has been tested with a range of chemical and mechanical weed control implements (McCool et al., 2018).

A more minimalistic approach has been taken by the Swiss company Ecorobotix who are working on a fully solar powered robot, Figure 3b, which applies a micro-dose of herbicide by two robotic parallel arms.

The systems described above are intended to be a representative selection, and not an exhaustive review of the field, as there are several other systems that could have been mentioned.

3. System requirements and specification

In this section we will present the requirements for the robot. We have performed experiments and data collection in cooperation with vegetable producers for ten years. Through this work we have built up an understanding of the challenges at hand and how a robotic system for weed control can generate value for the farmer.

We have worked with producers of carrots, leeks, salad, cabbages, bush beans and spinach in Norway and Germany, while we have the most experience with carrots. The farmers we have cooperated with run a combination of conventional and organic production.

We envision the robot to be a tool for the farmer that integrate well with their existing growing practices, that does not require alteration on the cultivation practices. With regards to in-row weeding, the challenge and need is stronger with seeded cultures, than with transplanted cultures. A requirement for our system is to work with seeded row cultures. A set of design requirements are listed in Table 1, and detailed in the following paragraphs.

Table 1: Main technical specifications of the robot in-row weeding system

Description	Value	Unit
Vehicle mass	300	kg
Nominal speed	0.8	m/s
Transport speed	1.4	m/s
Nominal incline	5	degree
Max incline[1]	40	degree
Track width	1.6 - 2.1	m
DoD operation width	168	mm
DoD resolution	6	mm

[1] Max incline for short time loads, e.g. trailer ramps, thresholds etc.

3.1. Cultivation methods

We are using carrots as a proxy for a larger group of row cultures with comparable cultivation methods. Carrots are seeded in three rows on a flat lifted bed, or on two ridges in between the wheel tracks, as illustrated in Figure 4. Each row is typically double- or triple-seeded in carrot cultivation, a triple row has typically 5 cm spacing between the seed lines. The track width vary between producers, i.e. the width from center of the left wheel to the center of the right. One producer will typically run all their equipment at the same track width and it is usually in the range of 1.6 to 2.1 m. Some larger productions use wider machinery for bed forming and seeding. Distance between tracks are then triple or wider, apart with continuous beds or ridges between. This is more common in transplanted

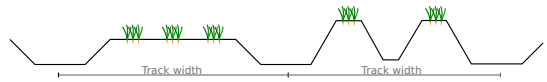

Figure 4: Carrots are typically triple-seeded in three rows on a lifted bed (left), or on two ridges between the tracks in the field (right). The track width is the distance across the row, measured from the center of the wheels of the machinery.

crops such as salads and cabbages, while it is also used by some producers of carrots, turnips, spinach etc.

The system must have one DoD unit for each crop row, and be height adjustable to adapt for the different cultivation methods. The width of the crop rows and seed-lines define the operational area for the DoD array. We need to control weeds in the crop row with a sufficient margin to overlap with the conventional tools for inter row weed control, such as guided harrows and weed brushes.

3.2. Robot and operational requirements

Setting requirements for the robotic platform is more about interpreting the producers needs, than it is an exact science. Together with producers we have envisioned several use cases and scenarios for the robot. The robot is designed to be a highly specialized tool for in-row weed control, focused on that task alone. The focus allows for a tailored and lightweight robotic platform.

The design requirement for the robot is a gross weight under 300 kg, both with regards to minimizing soil compaction and to limit the risk in human robot interaction. A target nominal operation speed of 0.8 m/s was selected on the basis of safety, area coverage and timing requirements imposed on the DoD system.

The fields and operation area is normally relatively flat, and we have set a nominal 5 degree incline specification. This will allow sufficient headroom for the variety of fields, and to some extent account for wet or high friction soil conditions. The extreme climbing requirement at 40 degree has been chosen to allow loading on and off trailers, and climbing over thresholds to access the field.

The system should come at a low cost of adaptation. The system does not require significant new infrastructure, and the robot is able to operate continuously throughout a working day. The robot is able to transport itself between fields, and for longer distances it can easily be loaded on a trailer with ramps.

3.3. Operator health and environment

The handling, loading and mixing of herbicides integrate with existing work flows, and does not present additional exposure of herbicides to the operator or environment. Variable rate application presents a challenge to predicting

the amount of herbicide required for a field. Excess herbicide is to be diluted and dispensed according to label, or transferred to a process for hazardous waste.

4. Robot design

The overall design goal is to make a specialized robot, best adapted to the task at hand: Efficient weed control in row crops. This implies that we have not attempted to design a modular and versatile robot - rather a simple, robust, maintainable and cost efficient system for DoD weed control in row crops.

4.1. Three wheeled design

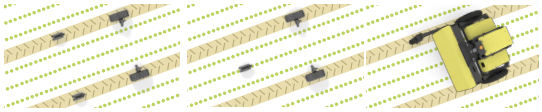

Figure 5: The off center 3-wheel configuration of the robot, allows for a lighter design with fixed wheel suspension, and only two motorized axes.

A common cost-effective robot design is using differential drive front wheels and rear castor wheels as the robot presented by Bawden et al. (2017). If you were to reduce to one castor wheel, Figure 5, the conventional design is to center the rear wheel. This is obviously not a good design for row crops, as the wheel would damage the crop - one could say it would *trample the salad*. Therefore we propose to use one off-center rear castor wheel. This necessitates a design with special care to weight distribution and stability.

A four-wheel design would require wheel suspension for the wheels to maintain ground contact when moving over uneven ground. The operational speed of the vehicle is sufficiently low such that we do not require a wheel suspension system from a vibration perspective. Three wheel design are not common, and we believe they have been disregarded in design of agricultural robot, as a symmetric design would not be suitable for operations in row crops.

By designing the system with a asymmetrical three-wheel configuration, we obtain a minimal wheel configuration while maintaining the systems suitability for operations in row crops. By designing the robot ground up, we obtain a highly cost-effective robot with a minimum of movable parts and good handling capabilities and stability.

4.2. Hybrid drive

The robot should be able to operate nearly continuously. The power requirements are outside of what can be delivered by solar panels on the available surface area, and the robot should require a minimum of additional infrastructure for its operation. Therefore the Asterix robot has been designed as a hybrid vehicle with a 48 V DC backbone and a four-stroke generator. The generator satisfies the average energy demands, while a battery bank ensures sufficient power during peak demand.

4.3. Concept for herbicide application

The resolution of the DoD array is determined by the crop culture and types of weeds we seek to control. The system is required to operate in the dense seed line of carrots, and effectively control weeds at early stages, including grass weeds with thin leaves. We have balanced these requirements, the technical feasibility and cost of a nozzle array, and arrived at a 6 mm lateral spacing of the DoD nozzles.

The longitudinal resolution is governed by the velocity of the robot and the DoDs maximum dispensing frequency. We have designed the system to maintain a 6 mm resolution at 0.8 m/s. The DoD modules are setup with 28 nozzles, giving an operational width of 168 mm, which leaves a margin for the row harrows towards the seed lines.

The spray controller and nozzle array is described further in (Urdal et al., 2014).

4.4. Camera system and vision processing

The vision unit employs an Nvidia Jetson TK1, with an embedded camera unit using the Omnivision 4682 4MP sensor. The computer, camera and LED flash is embedded in the DoD unit, making the unit fully modular and compact design, Figure 6. The DoD modules are mounted on a height adjustable beam, to account for variations in height between the wheel tracks and crop rows.

The vision pipeline is illustrated in Figure 7. The raw images are debayered to the RGB and HSV color space. The Hue and Saturation channels are used for segmentation of plant material from soil, which forms the mask. The mask is processed to separate individual leaves, and reduce noise. For each connected component in the resulting image, we compute a feature vector based on shape, texture and color. A support vector machine classifies each feature vector as either weed or crop which is used to generate a spray map. The corresponding crop map is used to mask out a safety margin in the spray map.

The spray map is transferred to the spray controller, which continuously estimates the relative motion by integrating the wheel encoder signals received over CAN-bus. The time-stamp of the spray-map, and the motion estimation is used to localize the current position of the valve array relative to the spray map. When a valve enters an active cell of the spray map, the valve is triggered and a droplet is dispensed on the weed, accounting for the droplet flight path and vehicle velocity.

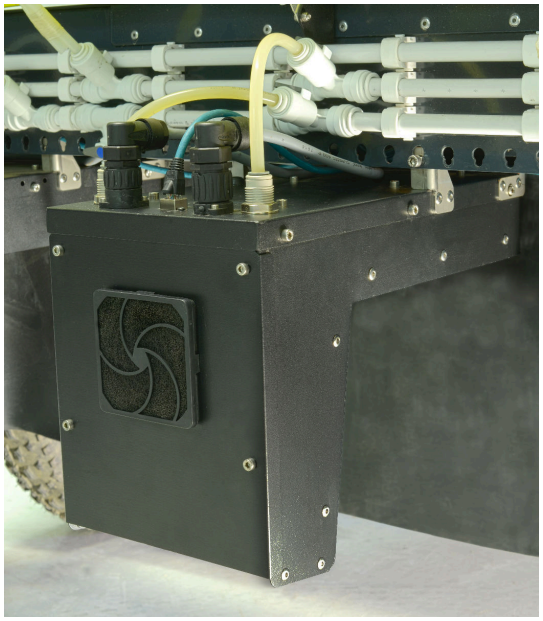

Figure 6: The Blythii module is a self-contained module for the machine vision and droplet application. The interface to the robot is the supply and return line for spray liquid, 48 VDC, CAN-bus and ethernet.

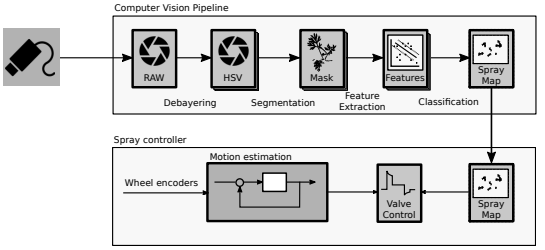

Figure 7: Every 200 ms the Blythii module captures an image, segments and classifies the image to generate a spray map. This is transferred to the Spray Controller which estimates its motion and excites the solenoid valves according to the spray map.

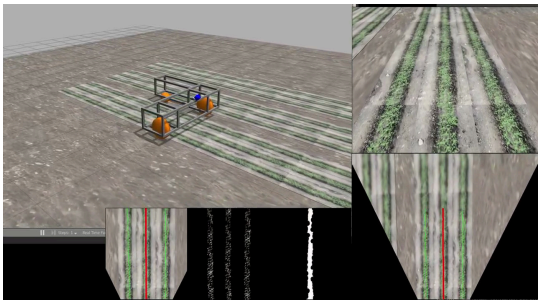

Figure 8: The simulator of the system is used to develop the row following algorithms, and headland turning. The upper right shows the view of the forward facing camera. Using a homography transform with the assumption of a flat surface in front of the robot, we convert the image to an ortonormal view as shown in the lower images.

4.5. Navigation unit

The navigation unit is based on the same computer hardware as the DoD modules, with the addition of peripherals for connectivity through mobile LTE/3G, WiFi when available, a GPS module and a forward facing camera for row detection. It also has an embedded CAN-bus module to connect to the backbone and command the Brushless DC (BLDC) motor controller.

The computer runs the open source Robotic Operating System (ROS), and for localization we utilize the Extended Kalman filter in the ROS package robot_pose_ekf. We use the forward facing camera to detect the seed lines: We assume a flat surface in front of the robot, and perform a homography transform of the image to obtain an orthonormal perspective. The image is segmented using a threshold on Green over Red*Blue channels, which become the input for a Hough Transform detecting straight lines in the image. We group the resulting line candidates to left, center and right, filter away outliers. The remaining line candidates are forwarded as measurements to a dedicated extended Kalman Filter estimating the current crop row location and heading in the global reference frame. The process is illustrated in the screenshot from the ROS / Gazeebo simulator in Figure 8. In our experience the flat surface assumption holds well in most field conditions, for the limited field of view that we operate with.

We have previously presented a non-linear model predictive control algorithm for row following Utstumo et al. (2015). The purpose of the controller is to prevent the rear castor wheel from damaging the crop, by limiting the steering control input. The complex implementation and our experience with path following controllers presented in Dørum et al. (2015), has led us to utilizing a simple line following algorithm, which better handle the context switching between row following and navigation in the headlands to enter the next row.

4.6. Valve flushing and management of excess spray liquid

While we have performed experiments to test and improve our precision and accuracy, we have experienced issues with clogged nozzles and residues of the spray liquid inside the valves. In the development, we have emphasized robustness and repeatability of our valve system, while this has not been given much attention in the literature.

The herbicides are to a varying degree corrosive liquids, and will leave residues if they are left in the valve system. The residue may prevent the valves from sealing properly, and we are left with a leaking valve, potentially damaging

crop plants. To obtain a robust and reliable DoD system we have implemented measures to counter these effects. Our most important tool in this context is to regularly flush the valves, and ensure that the valves are clean when left unused for an extended period of time.

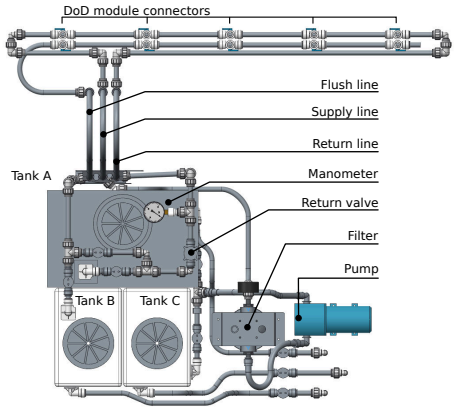

Figure 9: The spray liquid system supplies pressurized spray liquid to the DoD modules. The liquid continuously circulates through the DoD modules to ensure that the liquid is properly mixed.

In Figure 9, the system and its functionality is illustrated. The spray liquid in the main tank is continuously circulated by the pump. The liquid goes through a filter to remove any particles that could clog the valves and nozzles. The liquid circulates out using the supply line to all the DoD units, where some of the liquid end up as droplets deposited on weed leaves. The bulk of the liquid circulate back through the return valve, which regulates the pressure in the system.

To flush and clean the valves and supply lines, we have fitted the robot with two extra tanks. Tank A holds the spray liquid, while Tank C holds clean water for flushing the system. The return liquid from a flushing operation goes to tank B.

With a variable rate application of herbicide, it is required to do an estimate of the herbicide use through a field, and minimize the number of refills and the remaining amount left at the end of treating a field. Any remains in the herbicide tank will either need to enter a waste management system, or for some herbicides they may be properly deposited on organically active soil, as specified by their label.

5. Efficacy of single herbicide droplets

The objective of these pot trials was to find a liquid suitable for DoD application and with good weed control properties at the relevant growth stages, i.e. two - five true leaves.

5.1. Materials and methods

Table 2: In each of the two pot trials performed on February 4 (A) and 12 (B) 2016, there was one control (water) and four different herbicide treatments. The Active Ingredient (A.I.) of both Roundup Flex Plus and Glyfonova Plus is glyphosate, while Hussar OD is based on iodosulfuron. Each liquid was diluted in two steps, and the plants were treated with 3 droplets of 1.1 µg volume each.

	Trade name	**A.I. conc** g/L	**1. dil.** %	**2. dil.** %	**A.I. conc** g/L	**Dose plant** µg
A1[1]	-					
A2	Glyfonova plus	360	2.8	2.0	0.20	0.61
A3	Glyfonova plus	360	2.8	25.0	2.52	7.56
A4	Hussar OD	100	10.0	0.5	0.05	0.15
A5	Hussar OD	100	10.0	10.0	1.00	3.00
B1[1]	-					
B2	Glyfonova plus	360	2.8	2.0	0.20	0.61
B3	Glyfonova plus	360	2.8	25.0	2.52	7.56
B4	Roundup Flex Plus	480	2.1	2.0	0.20	0.61
B5	Roundup Flex Plus	480	2.1	25.0	2.52	7.56

[1] Same base solution as A- and B-, no A.I.

In these trials, we have focused on finding an appropriate application and dose for four common weeds in carrot crops. *Tripleurospermum inodorum* (L.) Sch. Bip, (Scentless Mayweed), is a challenging weed as it carries visual similarities to the carrot leaves in early stages. In addition it is resistant towards aclonifen, which is the most commonly used herbicide in carrots, so it requires an additional herbicide, like metribuzin, to be controlled by conventional spraying. The first leaves of *Chenopodium album* L., (Fat-hen), have a hairy and waxy-coated surface. Water droplets typically bounce off its leaves, and it presents an important adhesion test for our DoD system. *Poa annua* L., (Annual meadowgrass), is an annual grass weed and *Stellaria media* (L.) Vill., (Common Chickweed), is an annual broadleaf weed.

The technical setup is analogous to what was presented by Urdal et al. (2014). We are using the same control circuit and the same solenoid valve and nozzle (INKX0514300A and INZA4710975H) from *The Lee Company*.

This experiment was performed as two separate pot trials, which were sprayed on February 4 (Trial 1) and February 12 (Trial 2) in 2016. As an extension to previous studies on DoD herbicide application, which use glyphosate as the active ingredient, we have included iodosulfuron in our experiments. The herbicide solutions are described in Table 2. In each of the pot trials, there was a control which received three droplets of our base solution containing only the blue dye and additives for liquid properties. Each herbicide was first diluted with water, and then diluted with the base solution, as described in Table 2. In total there was 200 pots with one weed plant per pot. There was 5 pots for each combination of species and liquid formulation, (4 x 10). There was one weed plant per pot, and there were 5 replicate pots for each combination of weed species and

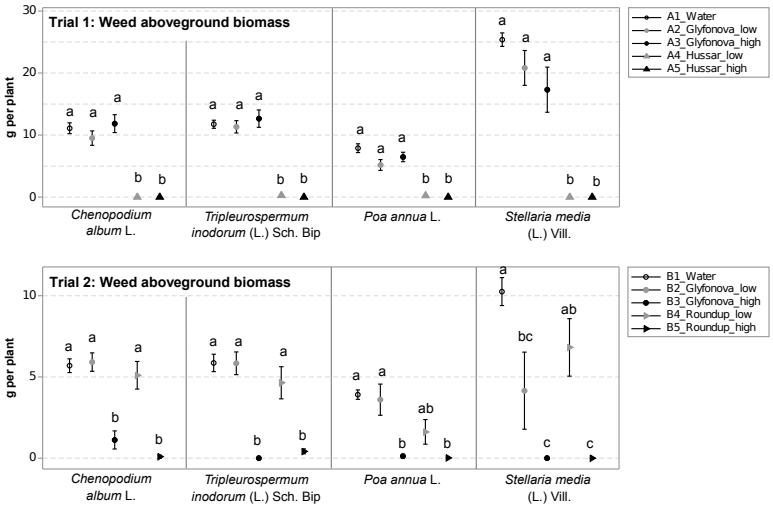

Figure 10: The fresh weight of the four weed species in the two pot trials, 25 (trial 1) and 17 (trial 2) days after herbicide treatment (cf. Table 2). A1 and B1 serve as controls, represented by the leftmost bar in each plot. The error bars are one standard deviation from the mean, and the letters above indicate the grouping by the Tukey test. In trial 1 only the two treatments with iodosulfuron, A4 and A5, gave significant reduction in weed biomass compared to the control, A1 (water only). In trial 2, the two treatments with high dose of glyphosate, B3 and B5, gave significant reduction in weed biomass.

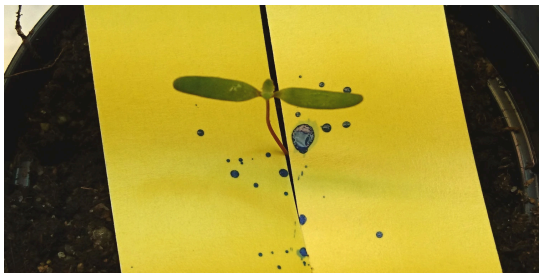

Figure 11: *Chenopodium album* L. after DoD application of liquid A3 with glyphosate. The liquid had very poor leaf retention properties, and satellite droplets were formed mid-air. No liquid is visible on the leaves, while the yellow liquid sensitive test strips highlight the droplets by turning blue. The droplets have either disintegrated in-air or bounced off the leaves.

Figure 12: *Stellaria media* (L.) Vill. after DoD application of liquid A5 with iodosulfuron. The three applied droplets have spread and achieved good contact with the leaf.

liquid formulation. In total 200 pots (5 replicates x 4 weed species x 5 liquid formulations x 2 trials) were included.

The droplet volume was measured to 1.16 µL by dispensing 1000 droplets into a container on a digital scale. This measure was verified by dispensing 1000 droplets into a 1.5 ml graduated test-tube.

The above ground biomass for each pot was cut and weighed on March 1. The fresh-weight datasets of the aboveground biomass per species and trial were analysed separately using ANOVA GLM considering replicate pots and liquid formulations as random and fixed factors, respectively. The resulting means were compared using Tukey test at significance level = 0.05. Model assumptions like normal distribution of residuals and equal variances were tested by Anderson-Darling test and visual inspection of residual plots.

5.2. Results and observations

Treatments A2 and A3 with glyphosate in the first trial had poor droplet quality. We observed satellite droplets being formed mid-air and poor leaf retention. Nearly no herbicide was retained on the leaves as can be seen in the example

in Figure 11. The A4 and A5 solutions with iodosulfuron produced well formed droplets and had high leaf retention, as shown in Figure 12. The liquid viscosity and surface tension was adjusted by using high speed photography and experience from earlier experiments (Urdal et al., 2014) for the following pot trials. Four liquids with glyphosate was tested in Trial 2, where the leaf retention and droplet performance was good.

In trial 1, the glyhosate treatments, i.e. A2 and A3, gave no weed control effect, whereas the iodosulfuron treatments, i.e. A4 and A5, gave very good control Figure 10. The lack of effect of the glyphosate droplets was unexpected.

The groups A2, B2, B4 and A3, B3, B5 have the same active ingredient and dose, but different liquid properties. The liquid formulation in A2 and A3 treatments had poor leaf retention properties. After having revisited the liquid viscosity and surface tension, the high doses of glyhosate formulations in trial 2, i.e. B3 and B5, demonstrated a good ability to control weeds. With 7.56 µg glyphosate per plant we could effectively control the four weed species.

6. Field trial on efficacy of Drop-on-Demand

Figure 13: The robot at the start of the field trial. Only the center of the three carrot rows was used in the experiment.

Throughout the past years the Asterix robots has been in the field with a team of vegetable farmers in Norway. To perform an end-to-end test with a new build of the robot and to document the efficacy of the system, we set up a trial in 2017.

6.1. Materials and methods

A part of the field was seeded with carrot in late August after the regular harvest specifically for this trial. The trial had two different treatments, herbicide (glyhosate) and unsprayed control, and each treatment was replicated ten times in a randomized block design. The plots treated were 2 m long, and were laid out along the crop rows. The areas assessed were 12 cm wide and 1 m long. The field is shown in Figure 13 and two plots before and after treatment are shown in Figure 14.

The trial was designed to evaluate the effect of the DoD system in the field. To eliminate errors from misclassification, all weed and carrot plants were treated. The trial represents an end-to-end test of the camera system, plant detection, generation of the spray-map, motion estimation, droplet target and shooting and the overall robot system.

The robot treatment was done September 28, and the plots were surveyed for number of carrot and weed plants (by species), by a skilled and experienced person in weed assessment October 2, and October 19. The observed weeds and their average occurrence in the plots on October 2 is presented in Table 3, on average there was 548 carrot plants per m². Images were recorded of all plots with a handheld camera the two latter days (Figure 16). Since the air temperature was relatively low, the four days between robot treatment and first weed assessment was considered unproblematic.

Table 3: The weeds observed in the field trial October 2 with their occurence per m² as an average over the 10 plots surveyed.

English name	**Latin name**	$\frac{\textbf{plants}}{\textbf{m}^2}$
Scentless Mayweed	*Tripleurospermum inodorum* (L.) Sch. Bip.	186.4
Annual meadowgrass	*Poa annua* L.	126.5
Fanweed	*Thlaspi arvense* L.	73.5
Field pansy	*Viola arvensis* Mur.	9.8
Prickly sowthistle	*Sonchus asper* (L.) Hill	4.5
Purple Deadnettle	*Lamium purpureum* L.	1.5
Fumitory	*Fumaria officinalis* L.	1.5
Storksbill	*Erodium cicutarium* (L.) L'Hér.	0.8

We have utilized the images to estimate the relative green index (RGI) of each plot, to evaluate the efficacy of the treatment, in conjunction with the field observations. The RGI is computed by first computing the Triangular Greenness Index (TGI) (Raymond Hunt et al., 2011) of each pixel: $TGI = 1.0Green - 0.39Red - 0.61Blue$. We then segment the images using the average Otsu threshold value for all the images (Otsu, 1979). The RGI is the number of pixels above threshold divided by total pixels.

6.2. Droplet volume and herbicide liquid

The spray mixture was the same as in treatment B3 of the pot trial described in Table 2. The active ingredient was glyphosate in a concentration of 2.52 g/L. The droplet sizes was estimated by shooting 1000 droplets through 7 individual nozzles into an empty container on a digital scale. The scale has a precision of 0.1 g, and the resulting volume per droplet is estimated to 2.1 µL per droplet.

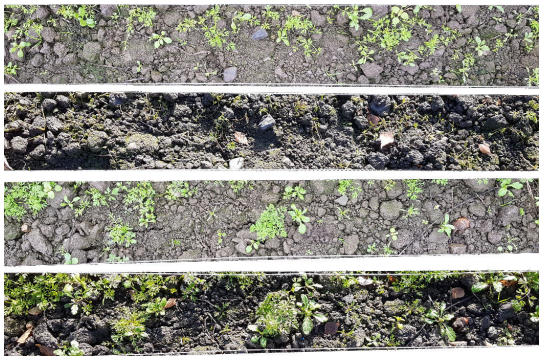

Figure 14: Two uppermost images: same plot (plot No. 1001) before and 17 days after glyphosate application with the robot. Two bottom images: untreated plot (plot No. 1002) October 2 and 19.

6.3. Experimental results

On Oct 19, the control plots were considered to have a substantial weed infestation, and the surveyor opted to document the plots by images, rather than counting individual weeds. The treated plots were surveyed and the only observed weeds were 12 seedlings of *Poa annua* L., (Annual meadowgrass), an average 10 plants per m^2. It is possible that the *P. annua* seedlings were too small to have been detected by the vision system, or accurately targeted by the DoD system at the time of treatment. Their size however indicate that they have emerged after treatment.

We have analyzed the difference in RGI between the two treatments (glyphosate and unsprayed control) by a pairwise Tukey method and 95 % confidence interval. Model assumptions like normal distribution of residuals and equal variances were tested by Anderson-Darling test and visual inspection of residual plots. The two groups were significantly different, and on average the treated plots had a reduction in RGI by 6.3 % and the untreated an increase by 5.6 % of the image area, as shown in Figure 15.

The RGI measure will include leaves that have died, but still green enough to pass the threshold. To evaluate the systems ability to control weeds in the field, we rely on both the RGI assessment and the field observations by the surveyor. All the treated plots were surveyed after treatment, and *P. annua* (Annual meadowgrass) was the only weed present. The *P. annua* seedlings had likely emerged after treatment. The RGI measures show an increase in green index for all untreated plots, thus we consider the trial a demonstration of successful weed control.

7. Discussion

The lab trials demonstrate that the four selected weeds can be effectively controlled by DoD application of herbicide

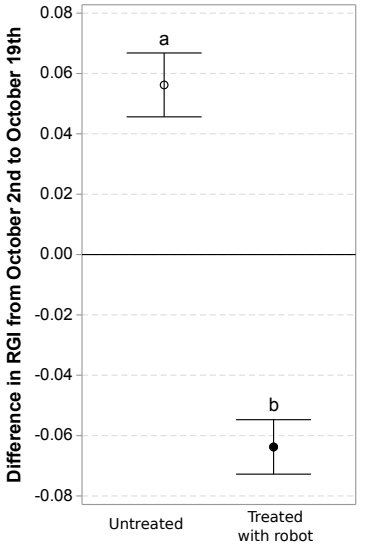

Figure 15: The change in relative green index in the treated and untreated plots, with error bars of one standard error. The letter indicate the grouping by Tukey HSD test. There was growth in the untreated plots, and a reduction in green index in the treated plots.

with doses as low as 7.56 µg glyphosate or 0.15 µg iodosulfuron, per plant. In field conditions we have demonstrated total weed control with the system using droplets with 5.3 µg glyphosate content.

7.1. Reduced herbicide application

We sprayed approximately 10 % of the area with our droplet spacing of 6 mm in the field trial, which is analogue to an application of 191 gram glyphosate per hectare. The label application for Glyfonova Plus ranges from 540 g/ha to 2880 g/ha depending on the types of weeds and weed pressure (Cheminova AS, 2015). This yields a herbicide saving in the range of 73 to 95 % comparing with label glyphosate application. Based on our findings from the lab trials, we expect to reduce the droplet size to 1 µL and that the weeds will cover less than 5 % of the area we treat, reducing our glyphosate application rate to below 50 g/ha.

A more relevant comparison is towards the commonly used selective herbicides in carrots today. The combination of aclonifen and metribuzin is the most common application in conventional production of carrots, with a maximum of 3 treatments with a total application of 1050 g/ha aclonifen and 106 g/ha metribuzin.

A treatment scheme with the robot and the DoD system, would consist of 2-3 treatments in combination with mechanical weed control in between the rows. Building on the experience from the lab and field trials, we would esti-

mate a total application of 50 - 150 g/ha glyphosate. This represent a ten-fold reduction in applied herbicide.

7.2. Reducing environmental and health risks

The amount of herbicide used is the main factor regarding the benefits towards health and environmental impact. All the herbicides we are dealing with are toxic to aquatic organisms, and to a varying extent they pose a health risk. Metribuzin is toxic if swallowed, and aclonifen is known to cause allergic skin reactions and is suspected of causing cancer.

The herbicides we replace them with have less severe health risks associated with them, even with the health impact of glyphosate being under heavy debate it is clear that the selective herbicides such as aclonifen and metribuzin pose a greater health risk. Iodosulfuron and glyphosate does not have health risk classifications beyond its potential to causing eye damage.

A DoD system produces larger droplets than a regular sprayer, and most importantly it does not produce aerosols. This reduces the exposure to operators and people near the field. The design of the robot and its implementation in the producers workflow will have to minimize the operators exposure to the herbicides, and enable a safe waste management system for excess herbicide and empty containers.

7.3. Impact on need for manual weeding in vegetable crops

While the environmental and health benefits of the system are significant, the DoD method will not see adaptation with vegetable producers unless it provides value to the producer. From experience we know that producers are frequently having to resort to manual weeding. They may be dealing with herbicide tolerant weeds or the weather conditions have not allowed for efficient herbicide application. Manual weeding is very much a candidate for automation, it strikes two out of three on the phrase "Dull, Dirty and Dangerous". The labour is inherently seasonal, and finding skilled labour willing and able to take on the work is challenging - and vulnerable to changes in immigration legislation as many are migrant workers.

A DoD robot can increase the quality of weeding, reduce the reliance on seasonal workers and improve food quality as the product is not affected by the herbicides.

8. Conclusion

The robot presented here has been designed with the specific task of Drop on Demand herbicide application in mind. The robot is tailored to this purpose with cost, maintainability, efficient operation and robustness in mind. The three-wheeled design is unconventional, and the design maintains maneuverability and stability with the benefit of reduced weight, complexity and cost. The robot and DoD-system is adjustable to account for differences in cultivation methods, number of crop rows, track width and height of the crop row. The forward facing camera, and navigation unit enables row following through the field. A combination of vision and GPS localization detects the end of a row, and aids the navigation in the headlands.

The current DoD modules treat a width of 168 mm with individual droplets of herbicide, spaced 6 mm apart. The efficacy of the DoD method is investigated in lab trials with four weed species, including one grass species and three dicot species. The weeds are effectively controlled by 7.6 µg glyphosate, and 0.15 µg iodosulfuron per plant.

The robot effectively control all weeds in the field trial with a ten-fold reduction of herbicide use. The field trial serves to demonstrate that our DoD system is a capable alternative to conventional spraying. The DoD system can reduce the amount of herbicides used by more than 90 %, utilize herbicides with lower environmental and health risks, reduce or even eliminate the need of manual in-row weeding.

Acknowledgements

The authors would like to thank Jon Frogner who hosted the field trial, Kjell Wærnhus at NIBIO and Per Hammerstad at The Norwegian Extension Service for their help and effort in completing the lab and field trials. We would also like to thank and acknowledge the work of Eirik Guttulsrud, Daniel Løvås, Jan Kåre Vatne, Torgim Lien, Peter Magnus Skarbø, Marius Bjornes and the rest of Adigo AS in developing and testing the robot platform. The work was funded by the Norwegian Research Council trough the IPN project "Innovasjon for bedre ugresskontroll i grønnsaker" (No. 234226/E50), the Industry PhD project (No. 218701), and by Adigo AS.

References

Basi, S., Hunsche, M., Damerow, L., Lammers, P. S., Noga, G., 2012. Evaluation of a pneumatic drop-on-demand generator for application of agrochemical solutions. Crop Protection 40, 121 – 125.

Bawden, O., Kulk, J., Russell, R., McCool, C., English, A., Dayoub, F., Lehnert, C., Perez, T., 2017. Robot for weed species plant-specific management. Journal of Field Robotics 34 (6), 1179–1199. URL http://dx.doi.org/10.1002/rob.21727

Cheminova AS, 2015. Glyfonova Plus. P.O. Box 9, DK-7620 Lemvig, Denmark, herbicide label, https://www.mattilsynet.no/plantevernmidler/etiketter/2004_48_14.pdf.

Dørum, J., Utstumo, T., Gravdahl, J., 2015. Experimental comparison of adaptive controllers for trajectory tracking in agricultural robotics. In: 2015 19th International Conference on System Theory, Control and Computing, ICSTCC 2015 - Joint Conference SINTES 19, SACCS 15, SIMSIS 19, Cheile Gradistei, Romania.

European Commission, 2008. Commission decision of 18 september 2008 concerning the non-inclusion of propachlor in annex i to council directive 91/414/eec. Official Journal of the European Union.
URL http://data.europa.eu/eli/dec/2008/742(2)/oj

FAO, 2014. Faostat database collections. Online.
URL http://faostat.fao.org

Fennimore, S. A., Slaughter, D. C., Siemens, M. C., Leon, R. G., Saber, M. N., 2016. Technology for Automation of Weed Control in Specialty Crops. Weed Technology 30 (04), 823–837.

Grimstad, L., From, P., 2017. The Thorvald II Agricultural Robotic System. Robotics 6 (4), 24.

Haug, S., Michaels, A., Biber, P., Ostermann, J., 2014. Plant classification system for crop /weed discrimination without segmentation. 2014 IEEE Winter Conference on Applications of Computer Vision, WACV, Steamboat Springs, CO, 1142–1149.

Heap, I., 2014. Integrated Pest Management: Pesticide Problems, Vol.3. Springer Netherlands, Dordrecht, Ch. Herbicide Resistant Weeds, pp. 281–301.

Koukiasas, N., Yu, T., Murdoch, A., 2016. Dose response relationship of droplet applications for the leaf-specific weed control in vegetable crops. Aspects of Applied Biology 132, 343–348.

Lee, W. S., Slaughter, D. C., Giles, D. K., 1999. Robotic Weed Control System for Tomatoes. Precision Agriculture 1 (1), 95–113.

Lund, I., Mathiassen, S. K., 2010. Transport and biological efficacy of single herbicide droplets. Aspects of Applied Biology (99), 435–438.

Lund, I., Olsen, H. J., 2010. Disintegration pattern of a liquid jet generated by a Micro-Dosing system. Aspects of Applied Biology (99), 249–254.

Lund, I., Søgaard, H. T., Graglia, E., 2006. Micro-spraying with one drop per weed plant. Third Danish Plant Production Congress, Denmark, 451–452.

McCool, C., Beattie, J., Firn, J., Lehnert, C., Kulk, J., Bawden, O., Russell, R., Perez, T., April 2018. Efficacy of mechanical weeding tools: A study into alternative weed management strategies enabled by robotics. IEEE Robotics and Automation Letters 3 (2), 1184–1190.

Midtiby, H. S., Mathiassen, S. K., Andersson, K. J., Jørgensen, R. N., 2011. Performance evaluation of a crop/weed discriminating microsprayer. Computers and Electronics in Agriculture 77 (1), 35–40.

Milioto, A., Lottes, P., Stachniss, C., 2018. Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in cnns. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia.

Nielsen, S. H., Jensen, K., Jørgensen, R., Bøgild, A., Jacobsen, N., Jørgensen, O., Jæger-Hansen, C., 2012. A low cost, modular robotics tool carrier for precision agriculture research. In: Proceedings 11th International Conference on Precision Agriculture Indianapolis, USA.

Nieuwenhuizen, A. T., 2009. Automated detection and control of volunteer potato plants. Ph.D. thesis, Wageningen University, Wageningen, The Netherlands.

Otsu, N., jan 1979. A Threshold Selection Method from Gray-Level Histograms. IEEE Transactions on Systems, Man, and Cybernetics 9 (1), 62–66.

Raymond Hunt, E., Daughtry, C. S. T., Eitel, J. U. H., Long, D. S., 2011. Remote sensing leaf chlorophyll content using a visible band index. Agronomy Journal 103 (4), 1090–1099.

Slaughter, D. C., Giles, D. K., Downey, D., 2008. Autonomous robotic weed control systems: A review. Computers and Electronics in Agriculture 61 (1), 63–78.

Søgaard, H. T., Lund, I., 2005. Precision agriculture '05. Wageningen Academic Publishers, Ch. Investigation of the accuracy of a machine vision based robotic micro-spray system, pp. 613–619.

Swanton, C. J., O'Sullivan, J., Robinson, D. E., 2010. The Critical Weed-Free Period in Carrot. Weed Science 58 (3), 229–233.

Urdal, F., Utstumo, T., Vatne, J. K., Ellingsen, S. A. A., Gravdahl, J. T., 2014. Design and control of precision drop-on-demand herbicide application in agricultural robotics. In: 2014 13th International Conference on Control Automation Robotics and Vision, ICARCV 2014. IEEE, Singapore, pp. 1689–1694.

Utstumo, T., Berge, T., Gravdahl, J., 2015. Non-linear model predictive control for constrained robot navigation in row crops. In: Proceedings of the IEEE International Conference on Industrial Technology, Sevilla, Spain.

Chapter 3

Concluding Remarks

The work presented here is part of the larger effort by Adigo AS and cooperating partners towards robotic in-row weed control. The robot and Drop-on-Demand (DoD) system are to provide reliable in-row weed control and reduce the overall environmental impact of weed control in row cultures. Robotic weeding has built momentum over the past decade where a wide range of actors in academia and in industry have been working towards more efficient and sustainable weed control in row cultures, (Fennimore et al., 2016). Producers of transplanted crops are successfully utilizing guided hoes and rotating disc implements like the Garford In-Row weeder (Tillett et al., 2008), and are able to demonstrate an economic benefit of robotic weeding in transplanted crops (Fennimore et al., 2014). From experience working with producers, we know that there is a strong pull for robotic methods also for seeded cultures and the producers are eager to see the methods mature and become available.

3.1 Localization and Navigation

The Non-linear Model Predictive Controller (NMPC) presented in Section 2.3 demonstrates the use of NMPC for navigation constrained to the crop rows, while preventing the controller to sway the wheels into the crop row and subsequently damage the crop. The framework is also suitable for implementing other constrains for operation in other environments, such as greenhouses or confined spaces.

The adaptive controller and a Model Reference Adaptive Controller (MRAC) presented in Section 2.4, can provide accurate tracking also with varying field conditions and payload which will influence the controller response.

3.2 Drop-on-Demand Weed Control

The DoD modules treat a width of 168 mm with individual droplets of herbicide, spaced 6 mm apart. The efficacy of the DoD method is investigated in lab trials with four weed species and 8 formulations of herbicide.

Iodsulfuron has not previously been presented as an active ingredient for DoD application. It provides a valuable alternative to the much debated herbicide glyphosate.

Out of the liquid formulations presented in the pot trial in Section 2.6, the iodosulfuron formulations demonstrated the best weed control capabilities. The weeds were effectively controlled by 7.6 µg glyphosate, and 0.15 µg iodosulfuron per plant.

The robot effectively controlled all weeds in the field trial with an estimated application of 191 g/ha glyphosate. This represent more than a ten-fold reduction of herbicide use, and use of a herbicide with lower environmental and health impact than the selective herbicides used today. The field trial serves to demonstrate that a DoD system is a capable alternative to conventional spraying, and may replace manual weeding weeding where weeds have become resistant to the selective herbicides.

3.3 Future Work

The work presented in this thesis represent an important contribution in the transition of DoD weed control from "a promising research concept" towards an applicable tool for producers. The project successfully demonstrate the core technology of DoD weed control by controlling weeds in pot and field trials. The presented robot is designed for the task at hand, and a broader approach to the work process has been considered, with the system for spray liquid handling as one example.

Section 2.1 and 2.5 are small steps towards utilizing visual odometry and inertial measurements for localization. Autonomous operation in agriculture, most commonly seen in tractor auto-steer systems and somewhat in autonomous mobile field robots is heavily reliant on RTK GPS. The GPS systems provide sufficient accuracy for row following, but an autonomous robot will have to rely on additional sensors for situational awareness and collision avoidance.

The advances in Visual Inertial Navigation Systems (VINS) in the past 10 years has been driven by navigation for robotics, and in later years in autonomous cars and in Virtual and Augmented Reality systems utilizing inside-out tracking. These systems demonstrate high accuracy and high framerate tracking on low-power devices. The author expects that we will see utilization of these methods also in agricultural robotics in the coming years. This will aid agricultural robots to achieve situational awareness and higher levels of autonomy.

3.4 The Asterix project

This thesis concludes the scope of the Industrial PhD project, while the larger effort moves ahead led by Adigo AS. There is much work ahead on the path towards full scale in-row weed control, and the project will continue to span the full width from research and engineering to business development.

As part of the Horizon 2020 Work Program in the European Union, the SME instrument funds high-potential innovation developed by small to medium enterprises (SMEs). The European Commision states that the program will support groundbreaking innovative ideas for products, services or processes that are ready to conquer global markets.

Figure 3.1: One of the prototype robots during a parsley field trial in 2018.

The SME instrument has received more then 46 700 applications and is highly competitive. In the period 2014 - 2017 only 8 % of Phase 1 and 4.8 % of Phase 2 applications has received funding.

Adigo completed the SME Instrument Phase 1 in 2017 and in July 2018 the Asterix project was selected for Phase II with a grant of 1.7 million € and a total budget of 2.4 million € for the coming two years[1]. This will be invested towards maturing the technology, piloting and preparations to commercialise the product. The Horizon 2020 SME award is a strong indicator to the projects innovation potential which also serves to attract additional investment.

The project continues with an increased velocity, and the robots will increasingly be roaming vegetable fields Figure 3.1. The robot receives attention from producers all over Europe, who reach out to inquire about the robots availability and applicability to other crops.

[1] EIC SME Instrument data hub, https://sme.easme-web.eu/?p=829983

References

J. Andersson. *A General-Purpose Software Framework for Dynamic Optimization*. PhD thesis, Arenberg Doctoral School, KU Leuven, Department of Electrical Engineering (ESAT/SCD) and Optimization in Engineering Center, Kasteelpark Arenberg 10, 3001-Heverlee, Belgium, October 2013.

M. H. Arbo, T. Utstumo, E. Brekke, and J. T. Gravdahl. Unscented multi-point smoother for fusion of delayed displacement measurements: Application to agricultural robots. *Modeling, Identification and Control*, 38(1):1, 2017.

S. Basi, M. Hunsche, L. Damerow, P. S. Lammers, and G. Noga. Evaluation of a pneumatic drop-on-demand generator for application of agrochemical solutions. *Crop Protection*, 40:121 – 125, 2012. ISSN 0261-2194. doi: https://doi.org/10.1016/j.cropro.2012.04.011.

O. Bawden, J. Kulk, R. Russell, C. McCool, A. English, F. Dayoub, C. Lehnert, and T. Perez. Robot for weed species plant-specific management. *Journal of Field Robotics*, 34(6):1179–1199, 2017. ISSN 1556-4967. doi: 10.1002/rob.21727. URL http://dx.doi.org/10.1002/rob.21727.

Blue River Technologies. Blue River See and Spray. Online, February 2018. URL http://smartmachines.bluerivertechnology.com/.

J. Dørum, T. Utstumo, and J. Gravdahl. Experimental comparison of adaptive controllers for trajectory tracking in agricultural robotics. In *2015 19th International Conference on System Theory, Control and Computing, ICSTCC 2015 - Joint Conference SINTES 19, SACCS 15, SIMSIS 19, Cheile Gradistei, Romania*, 2015. ISBN 9781479984817. doi: 10.1109/ICSTCC.2015.7321294.

Ecorobotix. Precise herbicide application with autonomous robot. Online, January 2018. URL http://www.ecorobotix.com/en/autonomous-robot-weeder/.

European Commission. Commission decision of 18 september 2008 concerning the non-inclusion of propachlor in annex i to council directive 91/414/eec. *Official Journal of the European Union*, 2008. URL http://data.europa.eu/eli/dec/2008/742(2)/oj.

FAO. Faostat database collections. Online, 2014. URL http://faostat.fao.org.

S. A. Fennimore, B. D. Hanson, L. M. Sosnoskie, J. B. Samtani, A. Datta, S. Z. Knezevic, and M. C. Siemens. Field applications of automated weed control: Western hemisphere. In *Automation: The Future of Weed Control in Cropping Systems*, pages 151–169. Springer, 2014.

S. A. Fennimore, D. C. Slaughter, M. C. Siemens, R. G. Leon, and M. N. Saber. Technology for Automation of Weed Control in Specialty Crops. *Weed Technology*, 30(04):823–837, 2016. ISSN 0890-037X. doi: 10.1614/WT-D-16-00070.1.

R. E. Gast. Industry Views of Minor Crop Weed Control. *Weed Technology*, 22(2): 385–388, apr 2008. ISSN 0890-037X. doi: 10.1614/WT-07-103.1.

L. Grimstad and P. From. The Thorvald II Agricultural Robotic System. *Robotics*, 6(4):24, 2017. ISSN 2218-6581. doi: 10.3390/robotics6040024.

G. Guren. Markdag i kålrotprosjektet. *Norwegian Agricultural Extension Service*, June 2010. URL https://www.nlr.no/fagartikler/6899/.

S. Haug, A. Michaels, P. Biber, and J. Ostermann. Plant classification system for crop /weed discrimination without segmentation. *2014 IEEE Winter Conference on Applications of Computer Vision, WACV, Steamboat Springs, CO*, pages 1142–1149, 2014. doi: 10.1109/WACV.2014.6835733.

I. Heap. *Integrated Pest Management: Pesticide Problems, Vol.3*, chapter Herbicide Resistant Weeds, pages 281–301. Springer Netherlands, Dordrecht, 2014. ISBN 978-94-007-7796-5. doi: 10.1007/978-94-007-7796-5_12.

N. Koukiasas, T. Yu, and A. Murdoch. Dose response relationship of droplet applications for the leaf-specific weed control in vegetable crops. *Aspects of Applied Biology*, 132:343–348, 2016.

W. S. Lee, D. C. Slaughter, and D. K. Giles. Robotic Weed Control System for Tomatoes. *Precision Agriculture*, 1(1):95–113, 1999. ISSN 13852256. doi: 10.1023/a:1009977903204.

I. Lund and S. K. Mathiassen. Transport and biological efficacy of single herbicide droplets. *Aspects of Applied Biology*, (99):435–438, 2010.

I. Lund and H. J. Olsen. Disintegration pattern of a liquid jet generated by a Micro-Dosing system. *Aspects of Applied Biology*, (99):249–254, 2010. ISSN 0265-1491.

I. Lund, H. T. Søgaard, and E. Graglia. Micro-spraying with one drop per weed plant. *Third Danish Plant Production Congress, Denmark*, pages 451–452, 2006.

C. McCool, J. Beattie, J. Firn, C. Lehnert, J. Kulk, O. Bawden, R. Russell, and T. Perez. Efficacy of mechanical weeding tools: A study into alternative weed management strategies enabled by robotics. *IEEE Robotics and Automation Letters*, 3(2):1184–1190, April 2018. doi: 10.1109/LRA.2018.2794619.

H. S. Midtiby, S. K. Mathiassen, K. J. Andersson, and R. N. Jørgensen. Performance evaluation of a crop/weed discriminating microsprayer. *Computers and Electronics in Agriculture*, 77(1):35–40, 2011. ISSN 01681699. doi: 10.1016/j.compag.2011.03.006.

A. Milioto, P. Lottes, and C. Stachniss. Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in cnns. In *2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia*, 2018. doi: 10.1007/978-3-319-48036-7_9.

Naïo Technologies. Large-scale vegetable weeding robot. Online, January 2018. URL https://www.naio-technologies.com/en/agricultural-equipment/large-scale-vegetable-weeding-robot/.

S. H. Nielsen, K. Jensen, R. Jørgensen, A. Bøgild, N. Jacobsen, O. Jørgensen, and C. Jæger-Hansen. A low cost, modular robotics tool carrier for precision agriculture research. In *Proceedings 11th International Conference on Precision Agriculture Indianapolis, USA*, 2012.

A. T. Nieuwenhuizen. *Automated detection and control of volunteer potato plants*. PhD thesis, Wageningen University, Wageningen, The Netherlands, 2009.

F. E. Poulsen. Robovator mechanical. Online, January 2018. URL http://www.visionweeding.com/robovator.

D. C. Slaughter, D. K. Giles, and D. Downey. Autonomous robotic weed control systems: A review. *Computers and Electronics in Agriculture*, 61(1):63–78, 2008. ISSN 01681699. doi: 10.1016/j.compag.2007.05.008.

H. T. Søgaard and I. Lund. *Precision agriculture '05*, chapter Investigation of the accuracy of a machine vision based robotic micro-spray system, pages 613–619. Wageningen Academic Publishers, 2005. ISBN 9789076998695.

Steketee. Steketee ic. Online, January 2018. URL http://www.steketee.com/product/Steketee-IC.

C. J. Swanton, J. O'Sullivan, and D. E. Robinson. The Critical Weed-Free Period in Carrot. *Weed Science*, 58(3):229–233, 2010. ISSN 0043-1745. doi: 10.1614/WS-09-098.1.

N. D. Tillett, T. Hague, A. C. Grundy, and A. P. Dedousis. Mechanical within-row weed control for transplanted crops using computer vision. *Biosystems Engineering*, 99(2):171–178, 2008. ISSN 15375110. doi: 10.1016/j.biosystemseng.2007.09.026.

F. Urdal, T. Utstumo, J. K. Vatne, S. A. A. Ellingsen, and J. T. Gravdahl. Design and control of precision drop-on-demand herbicide application in agricultural robotics. In *2014 13th International Conference on Control Automation Robotics and Vision, ICARCV 2014*, pages 1689–1694, Singapore, 2014. IEEE. ISBN 9781479951994. doi: 10.1109/ICARCV.2014.7064570.

T. Utstumo and J. T. Gravdahl. Implementation and comparison of attitude estimation methods for agricultural robotics. *IFAC Proceedings Volumes*, 46(18):52 – 57, 2013. ISSN 1474-6670. doi: https://doi.org/10.3182/20130828-2-SF-3019.00051. 4th IFAC Conference on Modelling and Control in Agriculture, Horticulture and Post Harvest Industry.

T. Utstumo, T. Berge, and J. Gravdahl. Non-linear model predictive control for constrained robot navigation in row crops. In *Proceedings of the IEEE International Conference on Industrial Technology, Sevilla, Spain*, 2015. doi: 10.1109/ICIT.2015.7125124.

T. Utstumo, F. Urdal, A. Brevik, J. Dørum, J. Netland, Ø. Overskeid, T. W. Berge, and J. T. Gravdahl. Robotic in-row weed control for vegetables. *Computers and Electronics in Agriculture*, Submitted for publication 2018.