Project title:	GROWBOT: A Grower-Reprogrammable Robot for Ornamental Plant Production Tasks.
Project number:	HNS/PO 194
Project leader:	Aran Sena, King's College London
Report:	Annual report, September 2018
Previous report:	Annual report, September 2017
Key staff:	Academic Supervisor (1 st): Dr. Matthew J. Howard Academic Supervisor (2 nd): Dr. Elizabeth Sklar AHDB Research Manager: Dr. Jim Dimmock
Location of project:	King's College London
Industry Representative:	Bruce Harnett (Kernock Park Plants), Jamie Dewhurst (J&A Growers)
Date project commenced:	October 2016

© Agriculture and Horticulture Development Board 2019. All rights reserved

DISCLAIMER

While the Agriculture and Horticulture Development Board seeks to ensure that the information contained within this document is accurate at the time of printing, no warranty is given in respect thereof and, to the maximum extent permitted by law the Agriculture and Horticulture Development Board accepts no liability for loss, damage or injury howsoever caused (including that caused by negligence) or suffered directly or indirectly in relation to information and opinions contained in or omitted from this document.

© Agriculture and Horticulture Development Board 2018. No part of this publication may be reproduced in any material form (including by photocopy or storage in any medium by electronic mean) or any copy or adaptation stored, published or distributed (by physical, electronic or other means) without prior permission in writing of the Agriculture and Horticulture Development Board, other than by reproduction in an unmodified form for the sole purpose of use as an information resource when the Agriculture and Horticulture Development Board or AHDB Horticulture is clearly acknowledged as the source, or in accordance with the provisions of the Copyright, Designs and Patents Act 1988. All rights reserved.

All other trademarks, logos and brand names contained in this publication are the trademarks of their respective holders. No rights are granted without the prior written permission of the relevant owners.

AUTHENTICATION

We declare that this work was done under our supervision according to the procedures described herein and that the report represents a true and accurate record of the results obtained.

Aran Sena

PhD Student

King's College London

Signature .... Date ....

Dr. Matthew J. Howard

Senior Lecturer

King's College London

Signature .... Date ....

Report authorised by:

Dr. Jim Dimmock

Resource Management Scientist

Agriculture and Horticulture Development Board

Signature .... Date ....

© Agriculture and Horticulture Development Board 2019. All rights reserved

CONTENTS

© Agriculture and Horticulture Development Board 2019. All rights reserved

GROWER SUMMARY

Headline

Robotics offers many opportunities for alleviating labour challenges in the horticulture industry. Current generation robots have made robotics much more accessible, with simplified programming interfaces and improved saftey features; however these developments fall short for the horticulture industry in many cases due to the complexity and variety of product being handled. Using learning systems in combination with robotics is a possible solution, which allows robots to perform tasks that would be difficult to program directly. Using the learning approach of “Learning from Demonstration” a robot can learn to perform repetitive tasks from a human teacher, without the use of any programming languages. A caveat to Learning from Demonstration is that the performance of the robot system is dependent on the quality of the teaching provided by the person. This can be challenging for robotics and machine learning experts, and even more so for novice users such as growers. The research presented in this project has made advancements in understanding how to assess, guide, and improve novice teaching performance, with the aim of placing advanced robotic capabilities directly in the hands of growers.

Background

Horticulture as an industry faces significant labour challenges. These challenges are present both in the near term, with the current uncertain political climate, and in the long term, with a general downward trend of people entering the industry[1,2].

GROWBOT is attempting to address these labour challenges by reducing the dependency on human labour for horticultural businesses through robotic automation, specifically using “*learning from demonstration”*. In addition to the business-focused reasons for pursuing robotic automation, many tasks in horticultural production are physically demanding and often in difficult environmental conditions. By developing automation systems that can be used

[1]:https://www.theguardian.com/society/2015/jul/01/sue-biggs-rhs-horticultural-timebomb

[2]:https://www.rhs.org.uk/education-learning/careers-horticulture/horticulture-matters/

© Agriculture and Horticulture Development Board 2019. All rights reserved

directly by the grower, when and where it is required, this technology offers benefits to both workers and businesses.

As best stated in the top-level goal of the GROWBOT project brief[3]:

[GROWBOT] will investigate ways in which non-expert users (i.e., those without technical expertise in robot programming and control), but that are nevertheless skilled in plant processing, can use robots in their work, to relieve them of the more repetitive, labour-intensive tasks encountered.

While using learning from demonstration with low-cost, general purpose robotic systems offers a promising automation option for growers, there are a number of challenges that must be addressed before these systems can be widely used in the horticulture industry to alleviate labour concerns. GROWBOT seeks to make robots more effective at performing horticultural tasks, while at the same time making growers more effective at using robots.

Summary

Robotics presents horticulture with many possible solutions to labour issues encountered by lots of growers; however horticulture also presents robotics many significant challenges. Robots are used extensively in many industries due to their speed, accuracy and strength capabilities; however taking advantage of these traits typically requires the workspace of the robot being highly controlled (i.e. the positions and locations of all objects around the robot are known to high degree of accuracy). This level of workspace control, or structure, is often not possible in horticulture, as the robot must interact with organic material which is inherently uncertain in shape and appearance, and consumer demands result in growers having to regularly adapt their production processes. Examples of where high levels of automation *are* found in horticulture can be seen in precision agriculture systems for high-volume repetitive tasks, such as sticking machines (which plant cuttings) used in high-volume production lines (see Figure 1 in the appendix). While highly effective, machines such as these represent specific use-cases that do not address much of the manual work which is being performed in the horticulture industry, particularly in sites with varying products and batch sizes. Commonly found manual tasks which are often difficult to automate include: plant grading (lots of plant

[3]:https://horticulture.ahdb.org.uk/project/growbot-grower-reprogrammable-robot-ornamental-plant-production-tasks-phd-studentship

© Agriculture and Horticulture Development Board 2019. All rights reserved

varieties being handled); packaging (can be challenging even in high-volume growing when there are numerous customers with different packaging requirements, Figure 3, Appendix); and plant management for bottlenecks in production which cannot be easily addressed (e.g. loading/unloading machines, transferring between stations etc.).

Learning from Demonstration (LfD, also known as Programming by Demonstration), is a method for robot programming, where a robot is programmed to do a task by observing *demonstrations* provided by the person using it. As an example, if you wanted to show a robot how to pick up an object, you would demonstrate how the robot should do this by physically moving the robot limbs through the motions required, while the robot recorded all the information it has access to (joint angles/cameras/etc.). Using the recorded data, the robot then tries to extract a control policy, using various machine learning methods, that it believe will let it perform the task the person wants done. When this procedure is successful, LfD allows people who are not robotics experts to “teach” robots how to perform complex tasks, all without any actual programming required. LfD has been used to teach robots everything from conventional industrial pick-and-place tasks, to dynamic tasks such as flipping pancakes and catching thrown objects.

LfD can be used with a variety of robot systems; however the primary robot platforms considered in GROWBOT are *collaborative robot* systems. Here, collaborative robot systems specifically refers to robots which have been designed to be safe in close proximity to people, both through imposed speed and load limitations, as well as advanced collision-detecting sensing. These robots are typically designed to be *back-driveable*, allowing users to move their joints with minimum effort, as required by LfD. This is a favourable form of interaction, as it has been shown to be intuitive for non-experts, and reduces some of the technical challenges which can arise when the robot is attempting to learn a task from observations. The robot used in GROWBOT is a model called Sawyer from Rethink Robotics (Figure 2 Appendix).

It is hoped that the combination of relatively low-cost, safe robotic systems with intelligent learning systems, will enable growers to deploy *flexible automation* on their site, to address challenges as they arise. This goal is based on observations that tasks encountered by growers often change crop-to-crop, season-to-season, and customer-to-customer.

© Agriculture and Horticulture Development Board 2019. All rights reserved

While LfD has much to offer horticulture in terms of providing an automation option that can be readily deployed when it is needed, there are still many limitations that have prevented its wider-spread use in industry, with many robots instead still predominantly relying on simplifying programming interfaces for users. Many of these challenges surround the interaction between the robot and the person teaching it. As the robot is learning a task from the teacher, it's performance is dependent on the quality of the data provided by the teacher, sometimes referred to as the *pedagogical quality* of the data. As the person teaching it is assumed to be a *domain* expert, but not a robotics expert (i.e. they have knowledge about horticulture, but not robotics), exactly what constitutes *good* demonstrations and how this influences how the robot learns the tasks can be difficult for the user to understand. It has been recognised in machine learning literature that an effective way to improve a person's teaching abilities is to improve their knowledge of what the *learner understands so far*. Providing the teacher insight to what the learner currently understands is often referred to as providing *transparency* into the learning system.

Therefore, to improve the performance of LfD-based systems, there are two-sides to consider: the learning abilities of the robot, where focus must be placed on the robot systems ability to extract task information from sensor signals; but then also the teaching abilities of the person using the system, where focus must be placed on how to improve the *transparency* of the robot learner.

Improving transparency with the aim of achieving higher task performance is not a trivial task. Primarily, on the robot side, it fundamentally does not know what the task *actually is*; it must *infer* this from the user's demonstrations. On the user side, communicating the robot's understanding of the task can be problematic due to the high-dimensionality of the robot sensor data.

Early work addressing these issues has been presented at a top robotics conference, the IEEE International Conference on Robotics and Automation (ICRA) in May of this year. My paper, "Teaching Human Teachers to Teach Robot Learners", studied the effects of transparency on non-expert teachers, and analysed the behavioural teaching patterns observed as the teachers tried to understand how the robot was learning a point-to-point reaching task. Here we showed a great improvement in task performance when the user was provided a transparent representation of the robot's learning process (~180% teaching

© Agriculture and Horticulture Development Board 2019. All rights reserved

efficiency improvement versus a control case), and observed that natural teaching behaviour for non-experts was often far from optimal.

While these are promising early results, the experiment used was highly constrained, by design, and so the follow-on work from this study is now in progress to see if these results can be extended and expanded upon in a more complex real-world scenario. This follow-on study has so far involved on-site data collection at a grower facility with participants programming an industrial robot system to perform a task they often perform temselves, with many participants from the seasonal labour workforce. Submission of these results for journal publication is targeted for the end of this year (2018).

As mentioned earlier, the improvement of robot performance and improvement of teaching data is a dual problem. Moving forward into my third year, this will be expanded upon further by using the results gathered so far on understanding and improving user teaching behaviours to guide a robot task self-improvement procedure. This will allow the robot system to not only learn the task as demonstrated by the user, but to optimize it further based on the robot's on abilities, and the understanding of task constraints that the robot has extracted from the demonstration data.

Financial Benefits

A very rough cost/benefit analysis is provided in my annual report for 2017, where I described that by taking a naïve view we can compare the £40,000 fixed cost of the robot against the recurring £14,217 cost of the human labourer, this represents a ~2.8 year break-even/payback period. Some feedback from growers on this report indicates I likely underestimated the cost of agricultural labour given National Insurance contributions, pension entitlements, holiday hours, sick leave, absenteenism, etc., leading to an annual cost closer to £16,000 per year, further enhancing support for flexible automation. In reality there are many factors beyond the cost price which affect payback, e.g. a robot can work consistently for long hours, can work with higher precision. In a particularly impressive display of this, the collaborative robotics company Universal Robotics released a case study recently detailing a SME manufacturing company deploying 4 of their robots for machine tending tasks, and achieved a payback period of 34 *days,* with the company crediting the robots uptime per day

© Agriculture and Horticulture Development Board 2019. All rights reserved

(21 hours of operation, daily) and the ability to re-task the robot without the need to deploy extensive safety guarding.

In addition to the potential financial benefits, there is the remaining issue that labour is becoming more difficult to secure for growers. This is not an exclusive issue to the UK market, with similar difficulties experienced in the Netherlands, Germany, Australia, and the US.

Action Points

During my research the past two years I have visited many growers in the UK, and recently in the Netherlands, and have met and discussed automation issues with some US growers.

Similar to my previous action point recommendation, a common observation has been that there are many tasks being done in growing sites that could be labelled as “low hanging fruit” for automation by currently available collaborative robot systems. From these discussions and visits, it appears these opportunities for automation are often missed due to a lack of awareness of the automation options available. This is reasonable, as the world of collaborative robotics has evolved a lot in the past 5 years, and quite removed from traditional automation systems considered in horticulture, but as a basic call to action I would advise growers to review the collaborative robot market to see if there are robots which could meet their needs.

While the collaborative robots currently available are somewhat limited in the types of task they can perform (primarily pick-and-place type tasks), this basic set of skills still covers a wide variety of commonly encountered tasks. Furthermore, by having robots *in-situ,* this opens doors to more readily collaborate with researchers.

© Agriculture and Horticulture Development Board 2019. All rights reserved

SCIENCE SECTION

This section details results from a paper I presented earlier this year at the IEEE International Conference on Robotics and Automation, titled *"Teaching Human Teachers to Teach Robot Learners".*

Introduction

Using Learning from Demonstration to teach robot learners *generalisable skills* relies on having effective human teachers. This paper aims to address two problems commonly observed in demonstration data sets that arise due to poor teaching strategies; undemonstrated states and ambiguous demonstrations.

In this paper, we presented a method for overcoming these issues, through the use of visual feedback and simple heuristic rules. The objective of these methods is to guide novice users to more effectively teach robot learners to generalise a task. The proposed methods intend to offer the user a more *transparent understanding* of the robot learner's model state during the teaching phase, to create a more interactive and robust teaching process.

Materials and methods

The experiment in this paper was designed around a simple two-dimensional point-to-point reaching task. The participant was tasked with providing demonstrations to the robot on how to reach the end point, from anywhere inside a valid start region. The robot would then learn a model for the task from the user demonstrations, and this model was tested for how well it covered all possible starting positions for the task. Due to the type of machine learning method used, if the participant provided insufficient demonstrations, or poorly distributed demonstrations, the model would not successfully generate a suitable trajectory for every point tested, i.e. the model would fail to *generalise.*

The hypotheses tested in this experiment were that teaching performance could be improved through increasing the transparency of the learning system, and that a non-expert's teaching ability could be improved through a simple set of heuristic rules, designed to guide the participant toward a good demonstration set.

Teaching performance was defined by a cost function that considered the complexity of the task, the actual performance of the learning system, and the number of demonstrations

© Agriculture and Horticulture Development Board 2019. All rights reserved

required to achieve the performance observed. Transparency on the learning system was provided through a visualisation of the current model, possible in this case due to the constrained (two dimensional) nature of the task.

A single-factor, three-phase repeated measures study was conducted with *n=30* participants. This experiment compared three test conditions: No Guidance (NG), Visual Feedback No Guidance (VFNG), Visual Feedback Rule Guidance (VFRG).

Results

Both VFNG and VFRG modes showed a statistically significant improvement of user teaching efficiency versus the NG condition, with an improvement of approximately 180% over the NG case. While VFNG and VFRG showed similar performance versus each other; the results indicate different teaching behaviours, with users providing demonstrations in an alternating pattern in VFNG, and in a progressive movement pattern in VFRG.

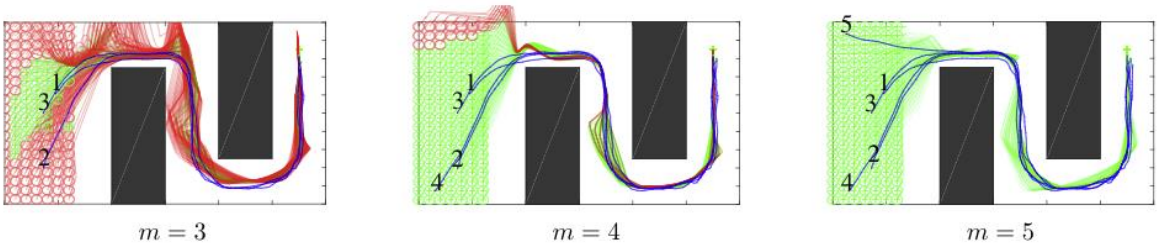

Fig. 1: Example of a participant provided demonstration sequence with the visualisation shown to participants. M is number of demonstrations given at the presented snapshot. Red and green regions indicate test results from sampling the initiation set after m demonstrations, with green lines representing successfully generated trajectories, and red representing failed trajectories. Blue lines indicate demonstrations provided by the participant, with the numbers indicating the order they are provided.

Discussion

The results gathered in this study support the concept of transparency as a tool for improving LfD system performance and indicates that the type of feedback and guidance provided to the user can noticeable influence their teaching behaviour $F(2,58)=7.952$, $p=0.001$. The 180% improvement in teaching performance can be attributed to the participants being able to see how the robot is learning with each additional demonstration they provide. The effect of this can be seen in Figure 1, where it can be seen for m=4 the participant has chosen to provide a demonstration in the region with greatest number of failures (around demonstration number

© Agriculture and Horticulture Development Board 2019. All rights reserved

4), and for m=5 they have provided a demonstration for the remaining failure region. Without this visualisation, it would be unclear to the participant where exactly the robot is struggling. With a knowledge of how the underlying machine learning is working, an expert could assume that the 5th demonstration needs to be placed at the top of the test grid, as all other demonstrations have been provided for the lower portion; however it is unlikely a novice user would identify this need. Looking at the behaviour of the participants while teaching the robot, it can be seen from the pattern of where they place demonstrations that their natural interaction with the robot is suboptimal, with participants often not providing enough demonstrations to successfully teach the task, and providing them in poor locations; however when provided the feedback, their teaching behaviour improves to be more similar to an expert teaching behaviour.

Conclusions

The results gathered so far are encouraging, but, as previously mentioned, these findings must now be extended to a more complex task which is closer to that encountered in real-world conditions. This extension has begun, with a data-collection experiment conducted in early May at a grower site, on a task similar to one performed on-site with an actual collaborative robot system.

Knowledge and Technology Transfer

I have demonstrated the collaborative robot used in my research to growers during my data collection activities this year.

Appendices

© Agriculture and Horticulture Development Board 2019. All rights reserved

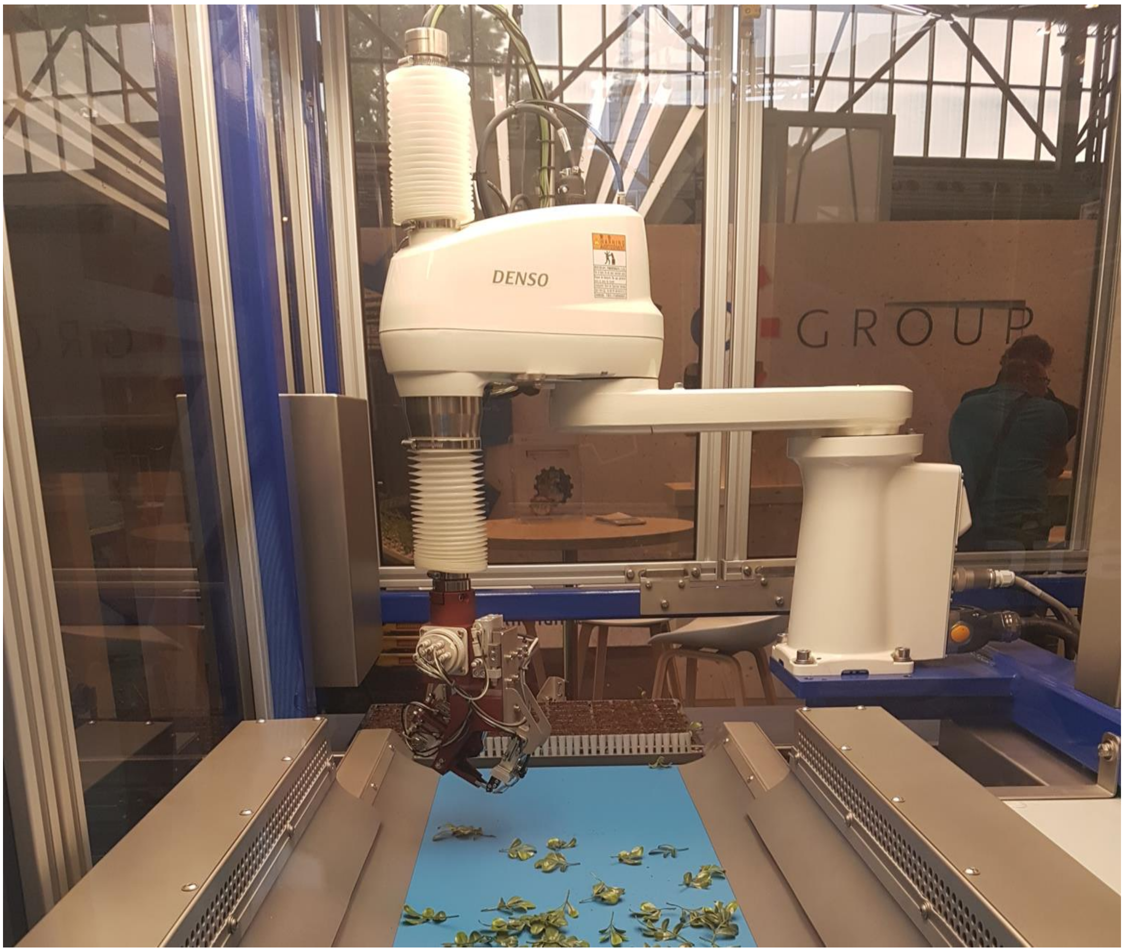

Figure 1: Precision agriculture system used for planting cuttings (foreground) into standard plant trays (background). Uses a SCARA robot for high-speed pick and place with a custom gripper and vision system.

© Agriculture and Horticulture Development Board 2019. All rights reserved

Figure 2: Rethink Robotics collaborative Sawyer robot at grower site.

© Agriculture and Horticulture Development Board 2019. All rights reserved

Figure 3: Various customer packaging options at grower site.

© Agriculture and Horticulture Development Board 2019. All rights reserved