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:	Final Report, September 2019
Previous report:	Annual Report, September 2018
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

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.

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 Howard

Senior Lecturer

King's College London

Signature .... Date ....

Report authorised by:

Dr. Jim Dimmock

Resource Management Scientist

AHDB

Signature .... Date ....

Dr. Georgina Key

Resource Management Scientist

AHDB

Signature .... Date ....

Contents

Grower Summary

- This is the final report for AHDB PhD studentship HNS/PO 194 – *“GROWBOT: A Grower-Reprogrammable Robot for Ornamental Plant Production Tasks”*.
- The research outputs from this project have been published in top robotics venues, IEEE International Conference on Robotics and Automation (ICRA), IEEE International Conference on Robotics and Intelligent Systems (IROS), and International Journal of Robotics Research (IJRR). Presenting at these conferences and publishing in a top-tier journal has helped give exposure to the GROWBOT project, and will hopefully spur others, globally, to take up the research agenda of expanding the use of collaborative robots for increasingly dextrous tasks with more general users.
- In addition to research outputs, a number of public engagement activities have been undertaken. Schools workshops, community magazine articles, public speaking events, and engagement with the wider horticultural research community across Europe have been valuable activities for spreading the GROWBOT message of highly adaptive robotic systems offering automation of long-standing problem tasks, not for replacement of human works but for improving throughput capabilities of SME’s and improving working conditions for the employees.

Headline

Horticulture faces significant challenges which conventional automation technology has struggled to address due to capability, economic, and expertise limitations. GROWBOT research has focused on robotic systems that learn directly from growers rather than being explicitly programmed, and has shown that this is a promising approach toward widespread horticultural automation.

Background

The original project brief specified by the AHDB summarised the objectives for this project as (emphasis added):

> *[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.*
>
> *Focus will be given to improving **efficiency** and competitiveness in small/medium scale businesses, typically processing **relatively small batches** of a wide variety of plants, as*

opposed to the traditional large facilities specialised in processing large volumes of single-varieties. For this reason, the project will investigate ways of automating tasks that are usually ***difficult to achieve at small scale****, such as taking and inserting cuttings, grading and collating plant specimens.*

As can be seen from this brief, the core focus was to develop the technology required to better support small/medium scale businesses producing products in relatively small batch sizes. Being small/medium business as well as producing relatively small batched products presents a significant barrier to automation using conventional systems due to lacking an economy of scale to offset the capital costs of setting up the automation.

In addition to batch size limitations, the fragile and variable nature of horticultural products leads to further constraints on automation options, as many systems will lack the capability to manipulate the organic matter directly and reliably without damaging it.

These limitations lead to a strong dependence on manual labour for production processes which, given larger batch sizes and more robust products, might otherwise be automated, e.g. packaging.

Labour security is an issue for the horticulture industry as it is becoming increasingly difficult for to source the casual labour force required for these manual processes. This is a problem that is both affected by short-term ongoing political developments around freedom of movement for people, but also a longer-term downward trend of people entering the industry and people travelling for the casual labour roles offered in horticulture.

Given these labour challenges, and the difficulty of deploying conventional automation into horticultural businesses which depend on flexible, small-batch production processes, it is clear that new approaches to horticultural production are required.

Modern robotic systems offer a potential solution. Combining general-purpose robot arms which are designed using a safety-first approach for use around non-experts, with intelligent learning systems that allow the robot to adaptively automate a wide variety of tasks could enable the automation of the low-batch, repetitive, labour-intensive tasks found in horticulture.

Summary

Robots that learn to perform tasks by observing demonstrations provided by people, i.e. *Learning from Demonstration* (LfD), are capable of performing tasks which would be difficult to explicitly program by hand, even by experts. An issue of using learning systems directly with end-users (i.e. the growers on farms), rather than deploying them with the help of experts, is that the performance of the system strongly depends on the quality of data provided by the

person doing the teaching, and so detailed knowledge on *how* the robot learns is often still critical for effective robot learning to take place.

In order to provide examples to the robot which are informative, it is useful for the person doing the teaching to have some idea of what the robot requires for learning to take place. Does showing it how to perform a task once mean it will learn all possible variations? If not one example, how many examples should we provide? Are some examples more informative than others? The answer to these questions generally depends on the type of learning method deployed with the robot, and the nature of the task being automated. It may not even be possible for the robot to learn the desired task given the selected learning method.

A deeper understanding of the interaction between teacher and learner is thus required to make these systems usable for non-experts. It is preferable to enable non-experts use these systems, as they are far more numerous than robotics systems experts, it is far less costly than employing specialised skilled experts, and non-experts may still offer domain expertise for the process they would like automated (e.g. growers wanting to automate their grading processes could bring detailed knowledge of where best to grab a plant to avoid damaging it). Enabling robot learning systems to work with non-experts would thus provide a large pool of users capable of deploying the systems effectively, in a cost-effective way for businesses, and allow them to use their domain expertise to identify use-cases for robot deployment.

Given the potential benefit of non-expert users of robot learning systems, and the increasing availability of commercial off-the shelf (COTS) robot hardware, GROWBOT has thus sought to improve the efficiency of how people teach robot learners, and make the robot learners more effective at using the data provided by the person. As this is completely new approach to horticultural automation, new research and engineering methods are required, and so the research outputs from GROWBOT are not a one-shot solution to the issues faced by industry, but ground-breaking steps toward satisfying the needs of industry.

This research has thus developed (i) a framework which allows researchers and engineers better design robotic learning systems which interact and are programmed by non-expert users, and (ii) an improved learning algorithm which is designed to be more robust to teaching errors made by people teaching robot systems. Both outputs are supported through extensive experimentation, both under laboratory conditions as well as testing directly with end-users on grower sites.

In order to deploy robots on grower's sites effectively in the near future, it is proposed that this approach of designing robot learner systems for manual task automation with a focus on increasing non-expert understanding of the robot learning process and deploying learning algorithms which minimize the impact of poor teaching behaviour is a critical step toward success.

Financial Benefits

Direct analysis of financial benefits from robotic automation on grower sites was not an objective for the GROWBOT project; however, a rough estimation has been provided in previous reports. This estimation has been based on, and refined through, discussions with growers and information gathered from government reports.

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. National Insurance contributions, pension entitlements, holiday hours, sick leave, etc. lead to an annual cost closer to £16,000 per year, shortening the breakeven/payback period. Payback periods can be even lower again, as there are many factors beyond the cost price which affect payback, e.g. a robot can work consistently for long hours, and can work with higher precision. Of course, these benefits can only be realised if the robot system is *capable* of automating the target process. A challenge faced in the horticulture industry is that many of the involved processes which are not already handled by conventional automation equipment such as conveyors and palletization robots are not feasibly automated by conventional systems without *significant* engineering effort and cost, thus eliminating any economic benefit over maintaining a casual workforce.

The financial benefits offered by robotic automation are not exclusively gained through labour *replacement* as might conventionally be believed. Gains may also be made by making more *effective* use of people. People will continue to be far more capable and cost-effective at dynamic dexterous manipulation across a wide array of tasks than robots for the foreseeable future, so if robot systems can be improved to handle the simpler and repetitive tasks that currently drain workers' time, the workers can then focus on activities which provide a greater value-add to the product or production process. Indeed, based on multiple grower site visits, it was often observed that there were many processes being manually done that *existing* collaborative robot systems could help with, but either a lack of awareness or risk-aversion had prevented the companies exploring a collaborative robot option.

Finally, 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, with similar difficulties experienced in the Netherlands, Germany, Australia, and the US.

Action Points

Horticulture is a diverse industry, with each business facing unique challenges. To help accelerate development of robotic automation, work with researchers in whatever capacity you can to help them gather real-world, industry-relevant information.

While automation of directly handling plant material is currently not achievable with end-user programmable systems (at least not easily, or reliably), there are many basic processes that have been observed during grower visits which would benefit from automation with collaborative robots (e.g. during packaging processes, palletisation, etc.). Engage with researchers and collaborative robot distributors to better understand how these robots are deployed in practice.

Science Section

Introduction

This section provides summaries of the published research outputs. Research objectives, experimental methods, results, and discussions are briefly presented along with relevant figures. These findings can be broadly grouped under the topics of "improving teachers" and "improving learners".

"Teaching Human Teachers to Teach Robot Learners"

Initial research into improving the teaching quality of non-experts was published in "*Teaching Human Teachers to Teach Robot Learners"* in ICRA 2018 [1]. When using LfD based approaches for teaching robots new skills, a challenged face by the teacher is understanding whether the robot can generalise the given examples to unseen conditions. If the teacher has shown a grasp action for location A, can they tell if the robot would be able to repeat the action but for location B, i.e. can the robot *generalise* the skill from the examples? This research sought to gain an understanding on how well a teacher can make this assessment, and how the learning process can be made clearer to the teacher to improve their assessment of the robot learner.

Materials and methods

A planar maze-style human study experiment was designed for this research. While this is a somewhat abstract task, this is a commonly found *type* of problem found in production processes, whereby we are faced with a grid of repetitive reaching or grasping actions. It will be seen in following research that this experiment correlates reasonably well with a grading-style task found in horticulture.

Participants were tasked with showing a robot how to move it's end-effector from a start-zone to a goal point, without hitting any obstacles, by physically guiding it through the motions required, see Figure 1.

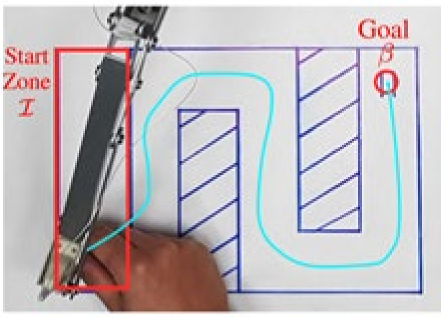

Figure 1: Planar maze-style experiment, showing a typical path from the start zone to the goal point which the robot could take.

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).

In each case, the participant had to provide as many demonstrations to the robot *as they thought necessary* for the robot to be able to fully perform the task, i.e. reach the goal point from anywhere inside the start zone. In the NG case, participants needed to rely on their own instinct for the learner's current performance level. In the VFNG case, the participant was provided with a visualisation of the learner's current performance level, which was generated by sampling the learnt model at a grid of starting positions and testing their feasibility. In the VFRG case, participants were first provided a rule set to give them an idea of how expert teaching would be conducted and then also provided visual feedback.

In these experiments, the learner collects the participant's demonstrations and learns a model of the demonstrated task which it can use to generate new trajectories. *Performance* is then measured by a teaching efficiency measure defined to be the level of performance achieved by the learner, normalised by the number of demonstrations given by the teacher. The performance level of the learner is taken as the ratio of *acceptable* trajectories and *unacceptable* trajectories out of 140 sampled start points in the start zone, see the circles in Figure 2. An unacceptable trajectory is one which leaves the maze area or collides with an obstacle, and an acceptable trajectory is one which links the specified start point to the goal point, which are coloured red and green respectively in Figure 2.

An example of a good teaching sequence is shown in Figure 2, where it can be seen that the learner has completely learned the task with 5 demonstrations.

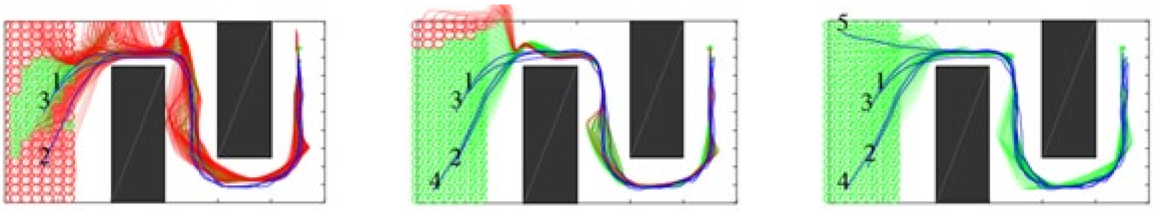

Figure 2: Example of a good teaching sequence. Red trajectories indicate model-generated trajectories which either collide with obstacles, or leave the maze area. Green trajectories indicate model-generated trajectories which satisfactorily link the start point to the goal point.

Results

Both VFNG and VFRG modes showed a statistically significant improvement of user teaching efficiency versus the NG condition, $F(2,58)=7.952$, $p=0.001$, with an improvement of approximately 180% over the NG case.

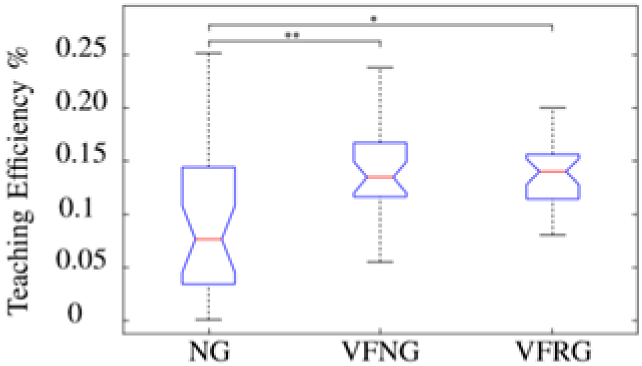

Figure 3: Box plot of results for No Guidance (NG), Visual Feedback/No Guidance (VFNG), Visual Feedback/Rule Guidance (VFRG).

Discussion

The results show a clear benefit in providing a visualised feedback to the teacher during LfD. Figure 3 shows that performance is improved under VFNG and VFRG compared to the NG case. Analysing the data shows a reason for the drop in efficiency in the NG case is due to participants tend to *underestimate* the number of demonstrations required for effective learning to take place. In the NG case, even for participants who overestimate the number of demonstrations required, simply providing *more* data does not guarantee the learner performance will improve, as they may provide multiple *poor quality* demonstrations, e.g. providing multiple demonstrations starting from a similar position will not help the robot learner generalise across the whole start zone.

VFNG and VFRG showed similar performance, see Figure 3. While similar performance levels are achieved, looking at demonstration start locations for each case indicates different teaching behaviours, with users providing demonstrations in an alternating pattern in VFNG, and in a progressive movement pattern in VFRG. As the performance results for VFNG and

VFRG are similar, despite the differences in strategy, this suggests the possibility of offering procedural instructions as teaching strategies to non-expert users in certain contexts. It is not however clear if this similarity in score is due to the rule set, or if a more complex task/robot would result in a spread between VFNG and VFRG performance.

“Quantifying Teaching Behaviour in Robot Learning from Demonstration”

Given the leap in robot performance gained by improving the quality of teaching, the ideas presented in the previous paper were expanded to a journal article, “Quantifying Teaching Behaviour in Robot Learning from Demonstration”, accepted to the IJRR in 2019 [2].

This work presents a theoretical framework which seeks to formalise the empirical observations found in the ICRA publication, and assess the developed ideas on a real-world task with an increase learning challenge.

The presented theoretical framework offers a way to model the whole LfD process in such a way that also accounts for the uncertain, variable nature of incorporating data from a person. The data provided by the person doing the teaching is a product of their own belief of what the robot is capable of, and what it has learned so far. In some cases, like the previous experiment, it is straight-forward for the teacher to develop an understanding of what the robot has learned so far as we can visualise the learning progress as the problem involves two dimensional trajectories and a single pass/fail criterion. In many real-world cases, there will be too many factors to enable a straight-forward visualisation, and evaluating a large number of generalisation test points will not be feasible.

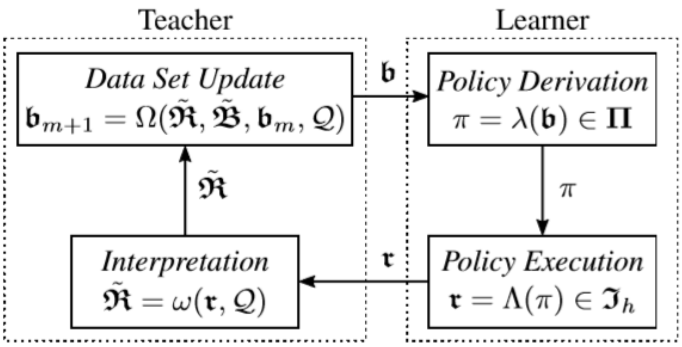

Figure 4: Proposed framework for modelling Learning from Demonstration showing a cycle between the teacher and the learner. For detailed information on the mathematical notation used here, please see [2].

The teacher is thus faced with a greater challenge in trying to determine robot learner capabilities, as they must not only decide how best to teach the robot, but also how best to *query* the robot’s knowledge, and form a belief of what example needs to be given to the robot learner to help it learn the task. This process forms a loop between teacher and learner, with the teacher attempting to best interpret the needs of the learner, and the learner deriving policies from the provided data, see Figure 4.

Using the proposed model, ideas that were previously general terms used in LfD can be more formally described, such as common teaching failures encountered during LfD *undemonstrated states, ambiguous demonstrations,* and *failed demonstrations,* see Figure 5 for examples of these teaching failures*.* By formalising features such as these, the framework enables system designers to better design features which focus on overcoming or avoiding these issues. This is shown through an experiment using a real-world system and a challenging manipulation task.

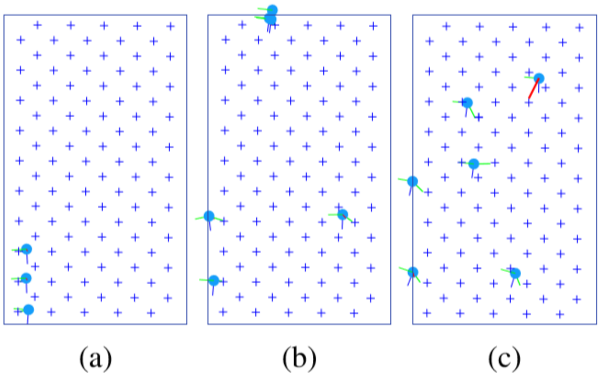

Figure 5: Teaching failures encountered during LfD. (a) undemonstrated states, (b) ambiguous demonstrations, (c) failed demonstrations.

Materials and methods

Using an industrial collaborative robot, 36 participants were tasked with teaching the robot how to pick plants from a 100 plant tray and deposit them into a disposal bowl. The experimental setup is shown on the left of Figure 6.

The group of 36 participants were split into 4 conditions, *No Feedback (NF)* where participants were given no feedback or guidance on how the robot's learning is progressing to simulate a naïve teaching approach, *Replay Feedback (RF)* where the robot would repeat the task for the last demonstration given, *Batch Feedback (BF)* where the robot would perform the task for five designated "generalisation testing" points (shown on the right in Figure 6), and *Selected Feedback (SF)* where the participant was free to ask the robot to perform the task for any plant position at any time to simulate a more natural uninformed teaching interaction. For each condition, participants were given 15 minutes to complete the teaching task, and were asked to do the experiment three times.

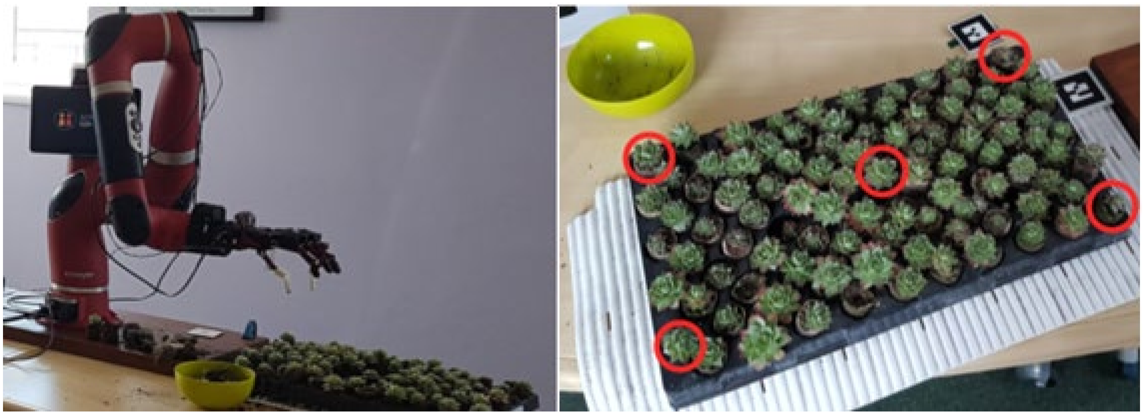

Figure 6: Experimental setup and target plants used in generalisation testing.

A typical demonstration set and the resulting generated trajectories from the learned model are shown in Figure 7. The underlying model used by the robot was a *Task-Parameterised Gaussian Mixture Model*, as used in the previous experiment.

Evaluation of the teaching performance was again done by considering the ratio of successful to unsuccessful trajectories generated by the learned model, normalised by the number of demonstrations given. This provides the teaching efficiency metric required to compare teaching conditions.

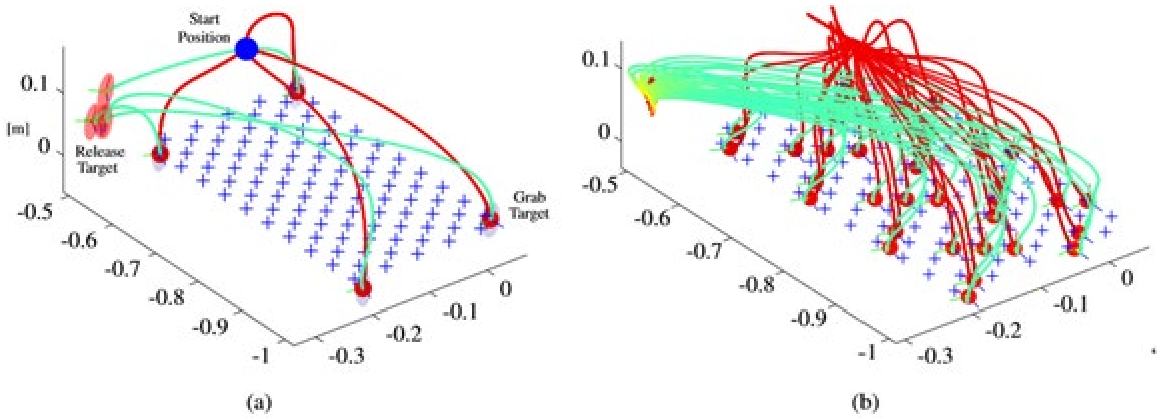

Figure 7: (a) Example demonstration set and (b) resulting generated trajectories for other plant positions.

The impact of providing good or poor quality demonstrations is then shown in Figure 8, which shows the grasp point error (the difference between where the robot hand is at the time it makes a grasp action to where it *should* be) from a top-down view. Here, it can be seen in part (a) that a poor quality demonstration set has been provided as all demonstrations are near the top of the tray, resulting in poor trajectory generation performance near the bottom half of the tray in (b), and in part (c) a better quality demonstration set is provided resulting in more consistent trajectory generation across the tray in (d).

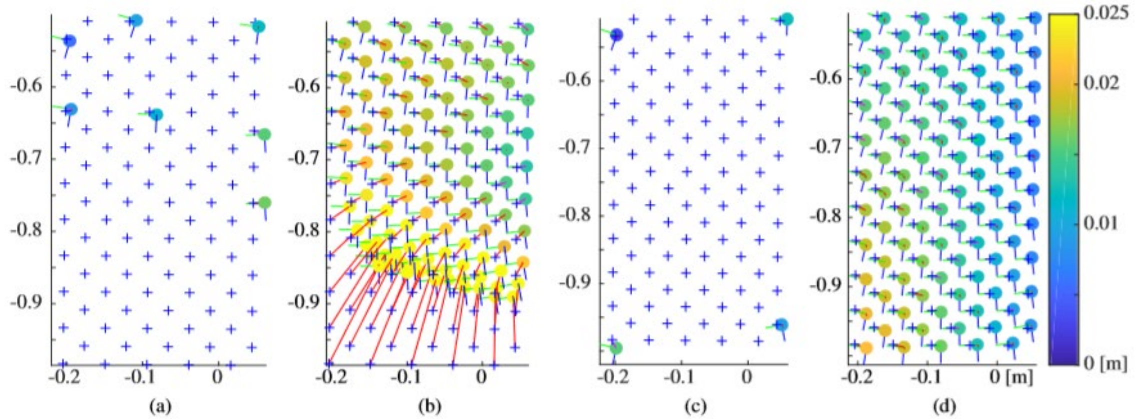

Figure 8: (a) and (b) example of poor quality demonstration set and resulting generated trajectory errors. (b) example of better demonstration set and resulting generated trajectories.

Results

Analysing the teaching efficiency between the four conditions reveals a significant difference in performance, $F(3,32) = 13.864, p = 5.797\times10^{-6}$.

A multiple comparison of means shows that BF provides a 10.76% improvement in efficiency compared to NF ($p = 1.017\times 10^{-5}$), 9.64% compared to RF ($p = 6.2947\times10^{-5}$), and 8.20% compared to SF ($p = 5.8708\times10^{-4}$), see Figure 9.

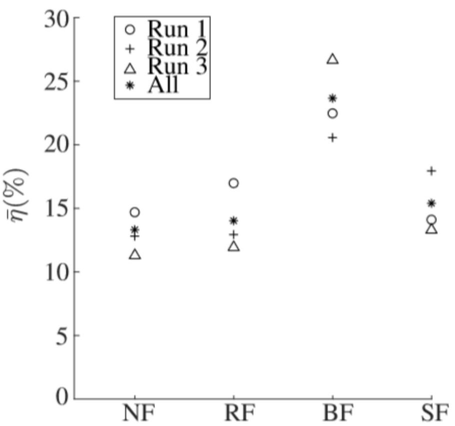

Figure 9: Teaching efficiency results for plant tray experiment.

To gain further insight into the participants' teaching behaviour, the occurrence of (i) demonstration errors, (ii) ambiguous demonstrations, and (iii) undemonstrated states are recorded in Table 1.

***Table 1:** Error counts found in demonstration sets, across all participant demonstrations.*

	NF	RF	BF	SF
Demonstration Errors	4	0	3	3
Undemonstrated States	499	619	827	485
Ambiguous Demonstrations	54	51	6	39

Discussion

The results from this experiment again show strong support for feedback being provided to the teaching being an important element to effective teaching and learning in LfD systems. This is significant, as it reflects the findings from the original maze experiment, but for a real-world manipulation task. Furthermore, it presents the insight that *not just any feedback will do,* but choice of feedback is critical toward actually helping the teacher understand the learner.

Looking at Figure 9, the RF and SF conditions are not significantly better than the NF case. This suggests that repeating what has just been shown offers little insight to the robot learning status, and that the participants were not able to *naturally* query the robot about it's performance to gain a better understanding. A note on this for the RF condition is that a potential benefit of repeating what has been demonstrated can be seen in Table 1, where RF incurred no demonstration errors across all trials, potentially indicating that this feedback condition helps teachers catch errors.

These results point to future work around autonomously querying the robot learner efficiently and effectively in order to provide feedback to the teacher. In addition to improving the performance of learning from demonstration based robotics in real world conditions, the research presented in this work provides foundations for a measured approach toward evaluating the teaching performance of the people who will be tasked with deploying robots in the workplace.

“Improving Task-Parameterised Movement Learning Generalisation with Frame-Weighted Trajectory Generation”

The previous two experiments have focused on gaining a deeper understanding on the interaction between the human teacher and the robot learner. Ultimately the developed understanding has been used to define a framework and feedback approaches that can be used to shape the behaviour of the teacher toward providing better quality demonstrations for the learner. The complementary direction of improving the learner, based on the developed understanding, is taken in “Improving Task-Parameterised Movement Learning Generalisation with Frame-Weighted Trajectory Generation”, presented at IROS 2019 [3].

The research presented in this paper draws from two key ideas. The first is that for repetitive tasks, generalisation of the demonstrated skill from a limited set of demonstrations, without loss of the fine motions required to successfully execute the skill, is highly desirable. By having effective generalisation that maintains local motion features (e.g. precisely how the robot should approach a plant), the robot can apply the learned skill widely given little teaching data. The second is that people often provide suboptimal demonstration sets that either underestimate the amount of demonstrations required for the robot to learn an effective model of the task that can generalise and/or provide ambiguous demonstrations where they repeat what they have already shown to the robot and neglect other regions of the task.

Developing the previously used learning trajectory learning method, *Task-Parameterised Gaussian Mixture Regression (TP-GMR),* an improved method for modelling the task is presented which is shown to be highly effective at generalising to unseen conditions, and helps the robot learner overcome suboptimal demonstration sets.

Figure 10 provides an overview of the proposed learning pipeline. Beginning with collection of demonstration data from the user, the robot the learns a series of local models by projecting the data into a set of local coordinate systems and then separately learning a Gaussian Mixture Model (GMM) in each coordinate system. By analysing the demonstration data correlations in each of the local coordinate systems, a notion of importance can be used to determine which coordinate system is most relevant at each timestep in the trajectory. These importance values can then be used to calculate time-indexed weights that are used when we would like to generate a new trajectory and must recombine the local models. By generating a trajectory with the importance weightings, local structure is prioritised when it is relevant to do so. In the examples in Figure 10, it can be seen that for each demonstration the trajectory is aligned with the red and blue pegs at the beginning and end of the trajectory respectively, and in the generated trajectories toward the right of the figure this local behaviour is maintained, even for the radically different trajectory shown at the very right.

These weightings give the method it's name of Relevance-Weighted Task-Parameterised Gaussian Mixture Regression (αTP-GMR). Please see [3] for it's full details of the method.

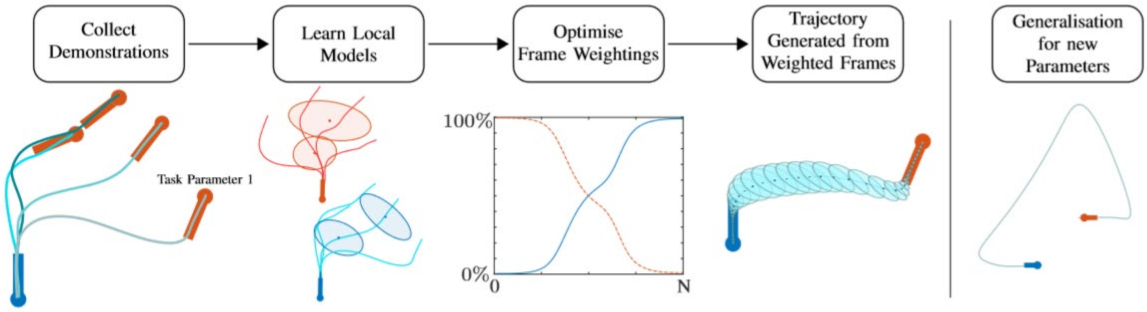

Figure 10: Learning pipeline for αTP-GMR.

Materials and methods

Using two-dimensional simulated reaching data from [5] and data from the real-world experiments in [2], a series of experiments were conducted using the improved model to evaluate generalisation properties for the method.

In the first round of experiments, the reaching data from [5] is used to evaluate the reconstruction ability of the proposed method. This data is shown in left of Figure 10, where each demonstration is a trajectory starting from one of the red pegs and ending at the blue peg. Here, a single demonstration is removed from the set of four demonstrations, and the learner must then attempt to reconstruct all four demonstrations using a model learned with the remaining demonstrations. This is repeated so that each demonstration is left out once (leave-1-out cross-validation). This experiment is repeated using TP-GMR (the original learning method), αTP-GMR, and then a modified Parameteric Gaussian Mixture Model (mPGMM) which is a competing learning method that has been proposed to specifically improve generalisation of task learning [6].

In the second round of experiments, again using the reaching data from [5], the same methods are tested for generalisation to unseen conditions using a grid-search method. Here, a 10m x 10m grid of test points is setup around the original demonstrations, and for each test point a trajectory is generated which must try to link the start point to the original demonstration goal point. The generated trajectories are evaluated based on three criteria, (i) the length of the generated trajectory, (ii) the error at the trajectory end-points (testing whether the trajectories begin and end where they are meant to, and (iii) constraint satisfaction errors (testing whether the trajectories leave their starting positions and enter the goal position aligned with the peg orientations, as seen in the original demonstration data.

In the final round of experiments, the real-world grasping data collected from human participants in [2] is used to see what performance improvement can be achieved when using αTP-GMR versus TP-GMR. All 108 trajectories collected in [2] are used, as this experiment considers the before and after case of using αTP-GMR.

Results

The first experiment results are shown in Table 2, where it can be seen that αTP-GMR has the lowest root mean squared (RMSE) error and lowest standard deviation.

Table 2: *Performance measures for leave-1-out cross validation reconstruction experiment.*

	TP-GMR	mPGMM	αTP-GMR
RMSE (m)	0.279	0.270	**0.200**
Std.	±0.236	±0.229	**±0.210**

Table 3 then shows the results from the second experiment, and Figure 11 visualises the resulting data. It can be seen that αTP-GMR has the lowest trajectory end-points error and lowest constraint satisfaction error count.

Table 3: *Performance measures for different learning methods.*

	Trajectory Length		Trajectory End-Points Errors		Constraint Satisfaction Errors	
	Mean (m)	**Std.**	**Mean (m)**	**Std.**	**Mean**	**Std.**
TP-GMR	**9.53**	±3.68	1.10	±0.74	19.34	±2.59
mPGMM	10.66	±4.84	1.03	±1.30	18.00	±2.76
αTP-GMR	9.73	**±3.40**	**0.04**	**±0.00**	**1.00**	**±0.00**

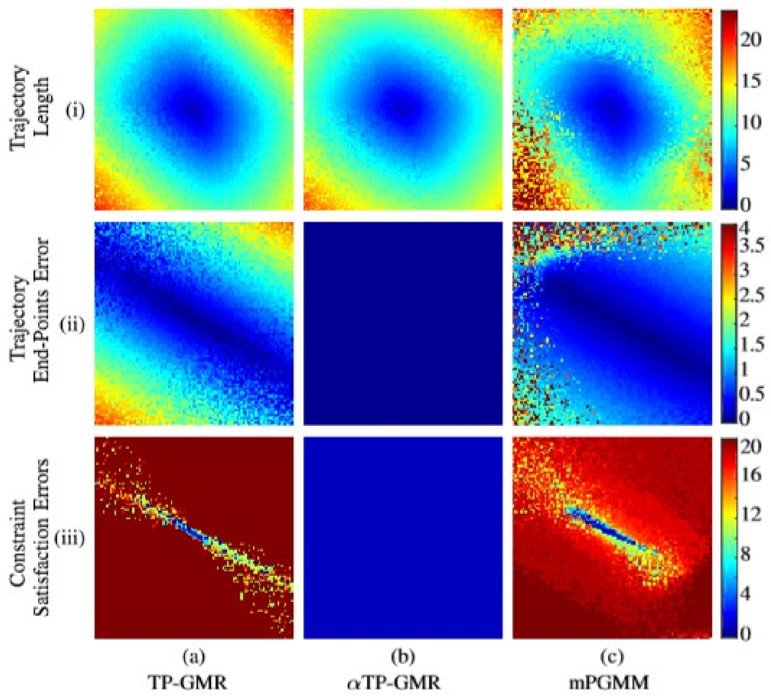

Figure 11: Error analysis plots for each learning method across three metrics.

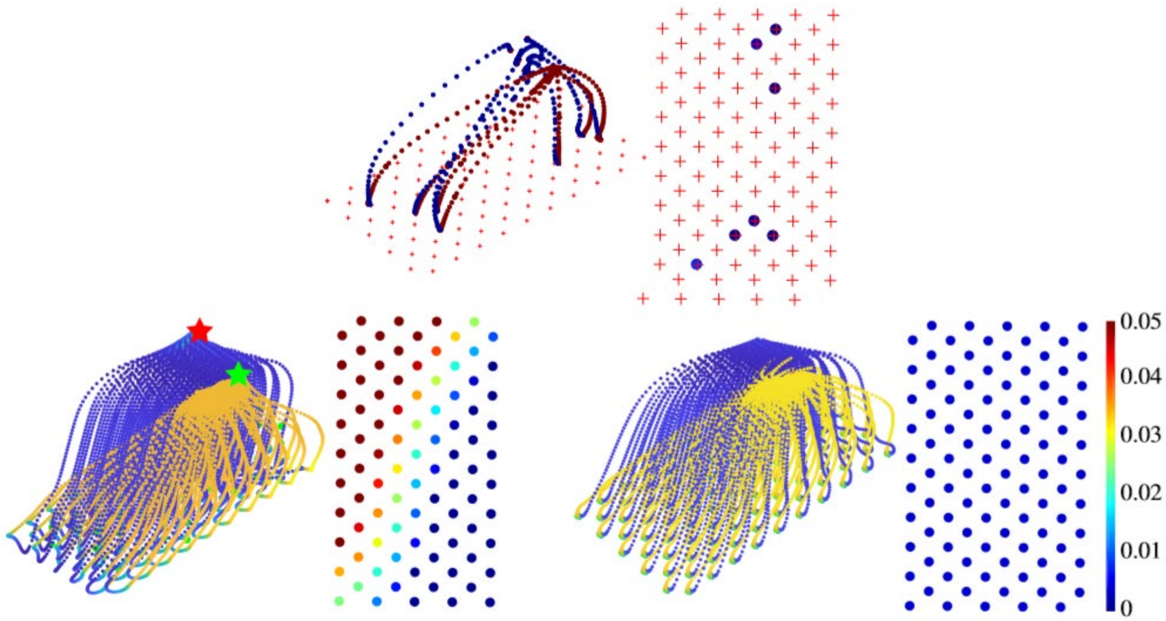

Figure 12: Original demonstrations shown at top. Trajectories generated with TP-GMR bottom left. Trajectories generated with αTP-GMR bottom right.

For the third experiment, an Anderson-Darling test indicated the data was non-normal, and so a non-parametric Wilcoxon Signed-Rank test was used to analyse the data. This indicated that the median of the total error incurred across all test-grabs with αTP-GMR was statistically significantly lower than under conventional TP-GMR, $Z = -8.8584, p < 10^{-18}$, representing a *∼40% reduction in error*. An example of the improvement achieved using αTP-GMR over TP-GMR can be seen in Figure 12.

Discussion

Across all three experiments, the performance of αTP-GMR is clearly superior to the alternative methods.

Observing some output test trajectories using the data from [5] in Figure 13, αTP-GMR can be seen to improve performance in the three test cases over TP-GMR and mPGMM. In the left most test case, the mPGMM path leaves the red peg region from the left side edge rather than the end of the peg. In the middle test case, both mPGMM and TP-GMR exit the red peg from the right side edge, and fail to approach the blue peg from the correct orientation. In the right test case, mPGMM and TP-GMR fail to generate satisfactory trajectories, with both beginning and ending in incorrect locations and both having incorrect approach orientations. In contrast, αTP-GMR produces smooth and correct trajectories for each case.

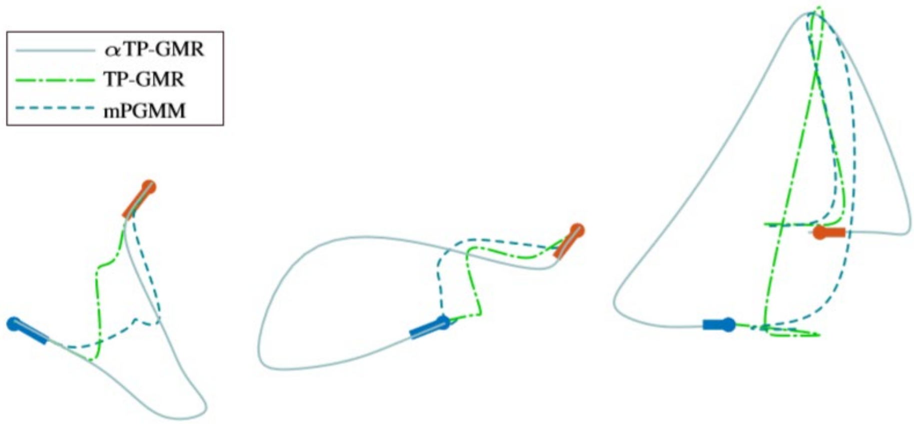

Figure 13: Generalisation abilities of the three methods considered.

Conclusions

As collaborative robot systems move beyond basic record-and-playback operation, or executing simple procedures defined in a list, instead *learning* task representations for more complex skills, the input provided by the human teacher becomes increasingly important.

The research presented here shares the overarching narrative of focusing on the impact the human teacher has on robot system performance. This involves both rethinking the interaction between man and machine, to help people gain a better insight to the inner-workings of the robot learners, as well as seeking ways to improve the resilience of the learner to suboptimal teaching as is often encountered with non-experts.

There are three key contributions from the research undertaken:

1. a new model for learning from demonstration, which better incorporates how people interact with the robots they are teaching,
2. a new method for robot skill learning, which provides significantly improved performance in learning skills which must be generalised to unseen conditions, and
3. several insights into the teaching behaviours of non-experts interacting with robot learning systems, gained through on-site testing directly with growers.

This research presents opportunities for wider deployment of robots in horticultural settings by showing how complex manipulation skills can be learned by a robot when paired with human teachers that are provided with adequate support by the learning system.

By approaching learning from demonstration with this focus on the teacher, new ways of evaluating teaching performance can also be considered. Rather than only considering basic descriptive statistics for the robots such as picks per minute, this approach allows us to query system performance in greater detail, such as questions around the people using robots, more general questions around whether the learning system in use by the robot is performing as expected, and whether the person and the robot are operating together as expected. This level of inspection of the human-robot team opens new directions for training, examination, and certification of a robot-supported workforce.

Knowledge and Technology Transfer

- Grower visits, beyond the project partners.
- GROWBOTICS schools workshops.
    - Raised +£2,500 from awards (King's College London), grants (Science Gallery London), and private sponsorship (Vitacress, New Forest Fruit).
    - 2019: https://www.aransena.com/2019-growbotics
    - 2017: https://www.aransena.com/2017-growbotics
- Magazine/Press articles:
    - AHDB Grower November 2016, AHDB Grower June 2017, AHDB Grower July 2017 (Cover), AHDB Grower October 2017,
    - AHDB Ornamentals Review 2018,
    - CONFOR February 2018,
    - Workshops article on hortidaly https://www.hortidaily.com/article/34481/UK-Vitacress-sponsors-educational-workshops-on-robotics-in-agriculture/
    - Workshops article on fruitnet http://www.fruitnet.com/fpj/article/172118/vitacress-sponsors-robotics-workshops
- Participation in wider horticulture and agriculture research workshops/ conferences/ events:
    - AHDB SmartHort researcher, speaker & panellist 2019
    - RHS PhD symposium 2016, 2018.
    - International Plant Propogators' Society (IPPS) Invited speaker 2018.
    - AHDB Industry Visits 2016, 2017, 2018.
    - AHDB Christmas Conference Invited speaker 2017
    - Chelsea Flower Show "Growing your career breakfast" 2017.
    - British Tomato Growers' Annual Conference 2017.
- Public speaking events:
    - Science museum, Robots event collaborator.
    - KCL 3rd Year Poster Competition – Winner in "Quality" category.

References

[1] Teaching Human Teachers to Teach Robot Learners, A. Sena, Y. Zhao, M. Howard, IEEE International Conference on Robotics and Automation (ICRA), 2018

[2] Quantifying Teaching Behaviour in Robot Learning from Demonstration, A. Sena, M. Howard, International Journal of Robotics Research (IJRR), 2019

[3] Improving Task-Parameterised Movement Learning Generalisation with Frame-Weighted Trajectory Generation, A. Sena, B. Michael, M. Howard, IEEE International Conference on Intelligent Robots and Systems (IROS), 2019

[4] A survey of robot learning from demonstration, BD Argall, S Chernova, M Veloso and B Browning, Robotics and Autonomous Systems, 2009

[5] A tutorial on task-parameterized movement learning and retrieval, S. Calinon, Intelligent Service Robotics, 2015

[6] On improving the extrapolation capability of task-parameterized movement models, S. Calinon, T. Alizadeh, and D. G. Caldwell, IEEE International Conference on Intelligent Robots and Systems (IROS), 2013

Appendices

Experimental data for planar maze experiment:

http://doi.org/10.18742/RDM01-242

Experimental data for plant tray picking experiment:

https://doi.org/10.6084/m9.figshare.8953124