

Project title:	Safe Human Robot Interaction
Project number:	unknown
Project leader:	Dr. Paul Baxter and Alexander Gabriel
Report:	Annual report, 12 2020
Previous report:	Annual report, 04/2020
Key staff:	
Location of project:	University of Lincoln, United Kingdom
Industry Representative:	Richard Harnden, Berry Garden Growers
Date project commenced:	8 1 2018

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 2020. 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.

Dr. Paul Baxter Senior Lecturer in Computer Science University of Lincoln Signature ... Date21/01/2021.... Dr Nicola Bellotto Associate Professor in Computer Science University of Lincoln SignatureNicola Bellotto..... Report authorised by: [Name] [Position] [Organisation] Signature Date [Name]

[Position] [Organisation]

Signature Date

CONTENTS

Headline1
Background1
Summary1
Financial Benefits1
Action Points1
Introduction2
Methodology2
Data Processing2
Belief-Desire-Intention Agent3
Evaluation Strategy4
Results5
Evaluation of the Perception System4
Initial BDI System Evaluation5
Discussion8
Conclusions8
Glossary8
References9
Appendices10

GROWER SUMMARY

Headline

Save 20% labour costs and reduce the time between picking and processing by letting your fruit pickers concentrate on picking fruit while your robots transport the produce.

Background

Fruit production is labour intensive and relies heavily on migrant workers. Socio-economic changes (e.g., Brexit) pose challenges for this field and make the case to reduce this reliance on manual work. Automation can help, but automation solutions aren't yet commercially available. The agricultural environment poses several challenges to both Robotics as well as Human-Robot-Interaction that must be overcome before this technology can be considered mature enough to be applied in a productively in agricultural settings. This work contributes to this effort by developing solutions that enable comfortable, safe and efficient Human-Robot-Interaction.

Summary

This thesis is part of the RASBerry research project. The project aims to develop an autonomous fleet of robots for in-field transportation. Specifically, the robots are expected to aid human fruit pickers by transporting crates from the picker's point of work to locations outside the field or poly-tunnel. Introduction of robots into this workspace will significantly reduce the costs of producing berries and is the first step towards fully autonomous agricultural systems.

Within the RASBerry research project, this thesis is concerned with the safe interaction of humans and robots, specifically the recognition and estimation of human behaviour and its interpretation as commands given to the robot. The results of this thesis will let the robot better prioritize its navigation goals and allow for a comfortable interaction between human and robot co-workers.

Financial Benefits

A robotic fruit transport system could save 20% labour costs and 10% land use [From2018].

Action Points

- Support research and development by allowing researchers access to small parts of fruit fields and personnel.
- Rent/Buy a robotic fruit transport system once commercially available

1

SCIENCE SECTION

Introduction

Introducing robots into a human working space can increase efficiency but should not come at the cost of comfort or safety. To achieve this balance in a challenging setting like agriculture, a robot needs to understand the intentions behind their co-workers' behaviour and basic communication. Gestures form an ideal medium to maintain reliability in adverse circumstances but are limited to situations where the human has their hands free. Additional clues from the environment as well as behaviour analysis can be used to estimate their state. Our interpretation of intentions [Gabriel2019b] sees them as the meaning [Fleischman2005] of, explanation [Tahboub2006] for, or idea [Youn2007] behind an action, plan or utterance. In our agricultural setting, workers pick berries into crates in a poly-tunnel environment. The robot is acting in a supporting role, supplying the human with empty crates, taking away full crates and staying out of the way the rest of the time. To facilitate the robot's autonomy, we created in integrated sensor data processing pipeline and Belief-Desire-Intention (BDI) [Bratman1987] agent system. The general motivation for this system is that in order to 'understand' the intentions of their human interaction partner (from observable behaviour) and to generate appropriate responses, the robot should consider both the environmental context but also its own goals (or 'desires'): this supports our use of a BDI architecture.

Methodology

Data Processing

The robot perceives its environment through a stereo RGB-D camera, a thermal camera, 2D and 3D LIDAR (Light Detection and Ranging), as well as differential GPS and odometry. It can also receive its co-workers' location either provided by GPS or ultrasonic localization and is supplied with predefined topological and laser scan maps.

During simulation, the robot determines its position using simulated laser scans and odometry. The location of the human is supplied by a picker-simulation engine and abstracted using Qualitative Trajectory Calculus (QTC) [Van de Weghe2006].

Figure 1. Overview of Data Processing System highlighting different processing stages.



As shown in **Figure 1**, the video is first pre-processed using OpenPose [Cao2017] to extract joint positions. Those are further processed to extract joint angles before both are fed into a naïve classifier that produces pose labels for each frame of the video individually, based on predefined prototype poses. Series of frames are classified using voting rounds where each frame contributes a single vote towards a pose. Whichever pose first wins 10 votes, labels the round. This action label is then memorized in a Knowledge Base (KB) to be used in Intention Recognition in the next stage. As Knowledge Base we chose OpenCog's AtomSpace [Goertzel2014].

The movement samples used in our first evaluations are part of a new dataset for Action Recognition in agri-robotics, that we collected in the summer of 2019. It contains samples of behaviours, such as picking of berries and carrying of crates, and samples of gestures to communicate with the robot. The samples were recorded from 10 different subjects between morning and early afternoon in a poly-tunnel environment featuring strawberry plants carrying ripe fruit.

The Intention Recognition (IR) stage uses the so generated action labels together with contextual information like the positions of human, robot and ripe berries, as well as past experiences, to infer the co-worker's intention. Inferred intentions are then stored in the KB to be used by the BDI system.

Belief-Desire-Intention Agent

Figure 2. Overview of the BDI agent system.



The BDI system (**Figure 2**) chooses intentions (plans to reach a goal) from its desires (abstract goals) based on its beliefs as captured in the Knowledge Base (KB).

The use of a BDI system separates reasoning about which goals to achieve from managing the execution of said goals. This allows the robot to consider more contextual information when deciding which goal to follow and leads to an aesthetic analogue to our idea of human motivation, intention and action on the robot. In our system, plans are represented as ordered tree structures with executable actions as leaves.

In this model, plans consist of a series of actions. Actions each have a set of preconditions and expected consequences, which combine to form the preconditions and expected consequences of a plan. When the agent decides which of its desires are applicable in each situation, it searches the KB for patterns of beliefs that match its desires' preconditions. If successful, it generates a corresponding intention. When idle, the robot chooses the intention with the highest utility per time interval and performs the next action in its plan.

Evaluation Strategy

The visual conditions in poly-tunnels vary a lot depending on the ripeness of the berries. Bigger plants obstruct the field of view and obscure the human fruit picker. Our time for experiments including humans is thusly constrained to the warmer part of the year. To make matters worse, experiments including robots and humans take a lot of time and are costly to set up.

Experiments in simulation on the other hand are relatively cheaper, easier to set up, and can often be automated to run unattended at night. Unfortunately, they can't reproduce the natural environment and human reactions nearly close enough to produce dependable results but in times of SARS-Cov-2 human interaction is drastically limited to slow the spread of this disease and so we cannot set up experiments involving human participants.

Our evaluation strategy is thus twofold. First, we evaluated the performance of the perception part of the system. After successfully concluding this phase, our evaluation of the whole system takes place in simulation, making use of the data gathered in phase one to keep conditions as consistent as possible with the real world.

Results

Evaluation of the BDI System

Our experiments evaluate how the agent system handles varying human behaviour in common service situations.

The two scenarios in question involve initially delivering a crate to a newly arrived co-worker, as well as exchanging a full crate for a new empty crate to allow an already working human to continue their job (**Figure 3**).

Figure 3. Left: Delivery Scenario: Meet the human, wait for 2.5 seconds, leave. Right: Exchange Scenario: Meet the human, wait for 5 seconds, leave.



In these scenarios, the robot is faced with different levels of cooperation on the side of its human co-workers. We explore three different interaction patterns with decreasing levels of human cooperation (**Figure 4**).

In the *first pattern*, human a co-worker first waves to the robot to signal for service, then waits for the robot to start approaching before approaching the robot themself. When robot and

human meet, the picker deposits their full crate on the robot (only in the exchange scenario) and takes a new empty crate, before walking back to where they started. After arriving they pick berries.

In the second pattern, the human does not wave to or wait for the robot. Instead, they simply approach the robot before continuing as in the first pattern.

In the *third pattern*, the human waves, but doesn't approach the robot. Instead, they expect the robot to drive all the way to them before they start depositing/collecting a crate.

Figure 4. Left: Delivery behaviour patterns. Right: Exchange behaviour patterns.



Both the delivery and exchange scenarios are initiated by the human, either by signalling to the robot or by approaching it, but they require a different response. The robot can make the distinction based on prior observed human behaviour (whether the person has been seen picking berries).

The recorded human behaviours in poly-tunnels, are passed as input to the proposed system (see **Figure 2**), acting on a simulated environment. Given 10 recorded subjects, 3 simulations per subject are performed (given a stochastic simulation), resulting in 30 simulations per scenario.

Scenario	Success Rate	Stop Order Distance Meeting Distance [m]	Waiting Time Service Time [s]
Delivery	1.00	μ: 0.77 σ: 0.001	μ: 0.10 σ: 0.006
standard		μ: 0.16 σ: 0.002	μ: 8.15 σ: 0.282

Table 1. Experimental Results Evalution of Behavioural Tolerance

Delivery	1.00	μ: 0.76 σ: 0.001	μ: 0.07 σ: 0.003
no call		μ: 0.15 σ: 0.002	μ: 8.05 σ: 0.221
Delivery	1.00	μ: 0.76 σ: 0.000	μ: 0.07 σ: 0.006
no approach		μ: 0.37 σ: 0.004	μ: 11.69 σ: 0.098
Exchange	1.00	μ: 0.76 σ: 0.001	μ: 0.11 σ: 0.017
standard		μ: 0.15 σ: 0.001	μ: 9.37 σ: 0.687
Exchange	1.00	μ: 0.76 σ: 0.001	μ: 0.15 σ: 0.080
no call		μ: 0.14 σ: 0.001	μ: 9.29 σ: 0.702
Exchange	1.00	μ: 0.76 σ: 0.000	μ: 0.08 σ: 0.004
no approach		μ: 0.38 σ: 0.003	μ: 13.03 σ: 0.043

Table 1 shows for each scenario the success rate as well as the mean (μ) and standard deviation (σ) of the metrics of evaluation: distance at which the stop command is issued, actual meeting distance, waiting time and time to service. The robot's speed averages around 0.74 m/s which corresponds to a slow walking pace.

Success rate is defined as the share of experiment runs that ended with the robot successfully interpreting the situation and performing the expected actions.

Stop Order Distance is the distance at which the agent system tells the robot navigation system to stop moving. The variance for this metric can be interpreted as an indicator of how much reasoning affects the agent's reaction time.

The Meeting Distance is the distance between picker and robot after both came to a halt. As can be seen from the difference to the stop order distance, the navigation system takes a while to stop the robot (there is an additional safety mechanism that cuts power to the motors before the robot can hit something). This distance is larger in the "no approach" cases due to the picker staying still.

Waiting time is the time between the human signalling to the robot and the robot starting to approach the human. Time to Service is the time it takes the robot to perform the correct action. It is measured from the moment the robot begins to follow a new intention to the completion of that intention. This time consists mainly of the time it takes to meet the human (4-8s), and the time for the delivery (2.5s) or exchange (5s) of crates.

Discussion

The agent system is successful in responding correctly to different patterns of human behaviour. The absence of human signalling for example does not lead to the robot erroneously deciding the human wants to pass, thanks to its capability of taking context (the presence of berries, the fill state of the co-worker's crate) into account. The robot's movement speed at a slow walking pace, and its attention to human orientation (it only approaches a person if they face the robot) contribute to a safer working environment as well as to the comfort of human co-workers, which the short waiting times also contributes to. Short waiting and service times additionally ensure efficiency.

This evaluation of the overall system is limited in its significance due to only taking place in simulation, but the high percentage of correctly performed behaviours is promising.

Conclusions

In the last year we have improved upon the agent system and developed a set of advanced metrics to evaluate the efficiency, safety and comfort of interacting with our robot. Our initial evaluation shows a system that can keep working in the face of degrading cooperation on the human side but requires further experiments on the effects of decreasing sensor performance or decreased availability of contextual information. It also still needs to be tested in the real-world to make sure the promising results of the simulation experiments hold up.

Glossary

RGB-D: Red, Green, Blue, Depth. Image format storing colour and distance values.

LIDAR: Light Detection and Ranging, like RADAR, but using laser light instead of radio waves.

Odometry: Estimation of change in position over time based on motion data

Topological map: Map consisting of interconnected waypoints.

Laser scan map: Map consisting of 3 dimensional points.

Qualitative Trajectory Calculus (QTC): Calculus that enables the abstraction of time-series of positional data into a symbolic representation (e.g., Entity X is moving away, coming closer)

Knowledge Base: A repository of information facilitating a selection of interaction scenarios with the stored information.

Deep Learning: A class of machine learning algorithms utilizing a layered approach to extract progressively complex features from raw data.

Bounding Box: A rectangle (2D) or cuboid (3D) bounding the borders of an area of interest.

References

Bratman, M. et al (1987). Intention, plans, and practical reason, volume 10. Harvard University Press Cambridge, MA.

Cao, Z. et al (2017). Realtime multi-person 2D pose estimation using part affinity fields. Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, 2017-Januar:1302–1310.

Everingham, M., Eslami, S.M.A., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A. (2015) The pascal visual object classes challenge: A retrospective. International Journal of Computer Vision 111(1), 98–136.

Fleischman, M. and Roy, D. (2005). Why Verbs are Harder to Learn than Nouns: Initial Insights from a Computational Model of Intention Recognition in Situated Word Learning. In Proc. of the Annual Meeting of the Cognitive Science Society.

From, P., Grimstad, L., Hanheide, M., Pearson, S. and Cielniak, G. (2018). RASberry -Robotic and Autonomous Systems for Berry Production. Mechanical Engineering Magazine Select Articles, 140 (6). ISSN 0025-6501.

Gabriel, A., Bellotto, N., Baxter, P. (2019a). Towards a Dataset of Activities for Action Recognition on Open Fields. In: to be published in the Proceedings of the UKRAS 2019 Conference on Embedded Intelligence.

Gabriel, A., Coşar, S., Bellotto, N., and Baxter, P. (2019b). A dataset for action recognition in the wild. In Annual Conference Towards Autonomous Robotic Systems, pages 362–374. Springer.

Goertzel, B., Pennachin, C. and Geisweiller, N. (2014). The opencog framework. In Engineering General Intelligence, Part 2, pages 3–29. Springer.

Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L. (2014). Microsoft coco: Common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision – ECCV 2014. pp. 740–755. Springer International Publishing, Cham.

Redmon, J., Farhadi, A. (2018). Yolov3: An incremental improvement. ArXiv.

Tahboub, K.A. (2006). Intelligent human-machine interaction based on Dynamic Bayesian Networks probabilistic intention recognition. Journal of Intelli-gent and Robotic Systems: Theory and Applications, 45(1):31–52.

Tome D., Russell C., and Agapito L. (2017). Lifting from the deep: Convolutional 3D pose estimation from a single image, CVPR, pp. 5689–5698.

Van de Weghe, N. et al (2006). A qualitative trajectory calculus as a basis for representing moving objects in geographical information systems. Control and cybernetics, 35(1):97–119.

Youn, S. and Oh, K. (2007). Intention recognition using a graph representation. International Journal of Applied Science, Engineering and Technology, 4(1):13–18.

Appendices

Example photos of the robot and sensor setup for the initial data collection:



Example photo of robot showing the sensor setup for the poly-tunnel data collection. The human is directing the robot to move to the side:



Example picture of poly-tunnel data collection. Subject picking berries approximately 7 meters from the robot:



Examples of varying light and weather conditions during the recording of the first dataset:



Skeleton extraction confidence values from thermal false colour video. Larger/brighter values indicate better performance:

