Assuring autonomy of robots in soft fruit production

Muhammad Khalid, Leonardo Guevara, Marc Hanheide and Simon Parsons
Lincoln Centre for Autonomous Systems (L-CAS), School of Computer Science
University of Lincoln, Lincoln, LN6 7TS, UK.
{mkhalid, lguevara, mhanheide, sparsons}@lincoln.ac.uk

***Abstract*—This paper describes our work to assure safe autonomy in soft fruit production. The first step was hazard analysis, where all the possible hazards in representative scenarios were identified. Following this analysis, a three-layer safety architecture was identified that will minimise the occurrence of the identified hazards. Most of the hazards are minimised by upper layers, while unavoidable hazards are handled using emergency stops. In parallel, we are using probabilistic model checking to check the probability of a hazard's occurrence. The results from the model checking will be used to improve safety system architecture.**

***Index Terms*—agricultural robotics, human-robot interaction, hazard analysis.**

Fig. 1: A Thorvald robot as used in our work

I. Introduction

The UK food supply chain network, from farm to fork, has an average total worth of £108 billion every year and employs around 4 million people (close to 12% of the workforce). The current relatively low level of productivity can be enhanced using Robotics and Autonomous System (RAS) [3]. RAS, in combination with other digital technologies, can have a very positive impact on overall food production by enabling higher production [4]. This is because RAS can work for longer than human workers, and can deal with weather conditions that humans find unpleasant [5]. This increased productivity means that the use of RAS could potentially add £58 billion to the food sector of the UK economy [5]. In the current, post Brexit, scenario in the UK, the food production industry anticipates, and indeed has already experienced, a shortage of labour. This has led to increased demand for RAS equipment [2] while also meaning that, unlike in other sectors, there is no significant danger of increased automation displacing human workers.

Soft fruit makes up 21.3% of the value of all the fruit and vegetables grown in the UK, with strawberries contributing almost 12.5% (£274 million). Soft fruit is thus an important part of the horticulture sector in the UK. Soft fruit production is also very labour-intensive — see for example [1] — and these higher labour costs, compared to other areas of horticulture, mean that RAS can be particularly beneficial. It is for these reasons that we are focussed on the use of robots in soft fruit, particularly strawberry, production.

The use of RAS can increase production, but for the near future, RAS in soft fruit production will have to work alongside humans, and in the agricultural environment [5] this means that there is considerable risk. We believe that the risks involved in using RAS for soft fruit production can be minimised by through a process of hazard identification and mitigation, and that is the work that we are engaged in.

This project is supported by the Assuring Autonomy International Programme, a partnership between Lloyd's Register Foundation and the University of York (2020-2022).

II. Overview

Our work is designing techniques that can contribute to the safe autonomy robots that assist in strawberry production, particularly focusing on safe human-robot interaction. The robots used in this work are Thorvalds, robots that are medium-sized, see Figure 1, but large enough to potentially cause damage to a human co-worker.

We are focusing on four scenarios in a farm setting that is sketched in Figure 2:

- **UV treatment**: Robots deploy UV light to kill powdery mildew. The UV treatment is performed at night time when there are no farm workers in action. As UV light is dangerous for humans no humans should have access to the polytunnels where the plants are during the UV treatment. There is no close interaction with the robot during UV treatment.
- **Logistics**: Robots bring empty trays to fruit pickers, collect full trays and take them to the collection point. The pickers have close interaction with the robot during logistic operations, putting full trays on robot etc.
- **Scouting**: In scouting, the robot traverses the polytunnels to collect data (photos of plants) using RGB cameras. This data helps in predicting yield, making treatment decisions and helping to plan harvesting.
- **Automated picking**: The robot, equipped with a picking arm will be used either for fully automated picking or work alongside other pickers. The robot may come in close interaction with a human during fruit picking.

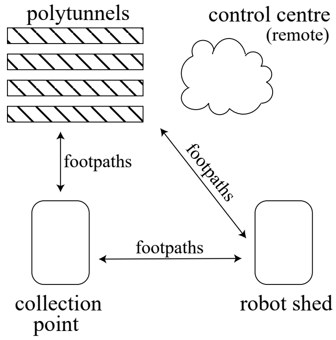

Fig. 2: The major components of a typical fruit farm.

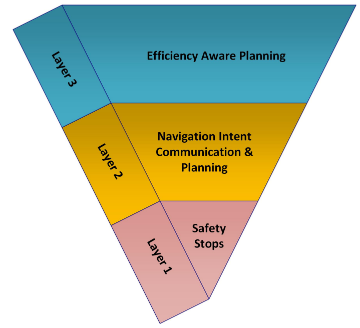

Fig. 3: The safety system architecture.

III. Progress

A. Hazard analysis and mitigation

The first step in our work was to identify all hazards in the four scenarios and categorise them [7]. Having identified the hazards, we could start to design mitigation strategies to avoid the hazards or, at least, minimise their occurrence. It quickly became clear that mitigation would depend upon three key functionalities — the ability to detect people, identify their intentions and predict their motion [6].

B. Functionalities

Detecting the presence of humans is the keystone of safety in a farm context. While the robots are equipped with a mechanical stop, meaning that they will not hurt a human co-worker through collision, activating this will shut the robot down, hurting efficiency. Detecting people at range using 2D and 3D LiDAR will allow the robot to perform a more graceful avoidance, meaning the mechanical stop does not need to be invoked. In addition, in UV treatment, any human within 7m of the robot may suffer UV burns, so detection at range is essential. Having detected people, reliably predicting their motion allows the robot to more efficiently navigate around them rather than stopping and waiting for them to move away, and being able to determine human intentions — signaled using gestures — will further improve robot efficiency.

C. Safety architecture

Assuming the ability to detect people and predict motion, the safety architecture in Figure 3 can help to ensure safety. The architecture is made up of three connected layers where, each layer is designed to address safety interaction at a different level, and higher layers aim to reduce the activation of lower layers. The aim of layer three is to plan routes which minimize the probabilities of interaction with human workers which share the work space. If an interaction is detected, this layer will re-plan in order to avoid the robot having to pause for a long time. In case an interaction occurs, layer two introduces human-to-robot and robot-to-human communication to both make the robot behaviour more comprehensible to the humans and to allow the robot to infer more precisely human intentions in order to increase the fluency of planned interactions and prevent the human and robot getting too close to one another. When a close interaction is about to happen, layer two is also responsible for ensuring a safety by reducing robot speed, performing evasive maneuvers or pausing operations. Finally, if layer two fails to ensure a safe close interaction, the layer one will activate emergency stops in case of imminent physical contact. These stops can be activated by LiDAR readings or by anomalies detected though soft sensors mounted on the robot structure.

D. Probabilistic model checking

In order to validate and enhance the safety features of the robot we are using the probabilistic model checker PRISM to model the human-robot interactions as Markov Decision Processes. The resulting probability models of each agricultural scenario predict the probabilities of the failures identified during the hazard analysis help us to assess the effectiveness of the robot safety system architecture. This analysis will be complemented with experiments in a soft-fruit farm setting.

IV. Conclusions

We have performed a hazard analysis on four scenarios that span the soft fruit production process, and based on this analysis are designing a safety system architecture that provides a layered approach to dealing with these hazards. Probabilistic model checking allows us to quantify the risks faced in deployment.

References

[1] L. Chen. Study on picking system for strawberry harvesting robots. *Beijing: Chinese Agriculture University*, 2005.

[2] F. Chiacchio, G. Petropoulos, and D. Pichler. The impact of industrial robots on eu employment and wages: A local labour market approach. Technical report, Bruegel working paper, 2018.

[3] T. Duckett, S. Pearson, S. Blackmore, B. Grieve, W.-H. Chen, G. Cielniak, J. Cleaversmith, J. Dai, S. Davis, C. Fox, et al. Agricultural robotics: the future of robotic agriculture. *arXiv preprint arXiv:1806.06762*, 2018.

[4] J. Maier. Made smarter—review 2017. *Department for Business EIS. The Stationery Office. London*, 2017.

[5] S. B. R. Maul and W. Maul. Taking control: Robots and risk, emerging risk report. *Lloyd's of London.*, 2019.

[6] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.

[7] R. Woodman, A. F. Winfield, C. Harper, and M. Fraser. Building safer robots: Safety driven control. *The International Journal of Robotics Research*, 31(13):1603–1626, 2012.