

Project title: Application of novel machine learning techniques and high speed 3D vision algorithms for real time detection of fruit

Project number: SF/TF 170

Project leader: Grzegorz Cielniak, University of Lincoln

Report: Annual report, 10/2020

Previous report: 12/2019

Key staff: Grzegorz Cielniak (University of Lincoln)
Pål From (NMBU)
Micheal Mangan (University of Sheffield)
Charles Whitfield (NIAB EMR)

Location of project: University of Lincoln

Industry Representative: Richard Harnden

Date project commenced: 14/09/2017

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 [YEAR]. 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.

Raymond Kirk

PhD Student

University of Lincoln

Signature



Date 30/10/2020

Report authorised by:

Dr Grzegorz Cielniak

Reader

University of Lincoln

Signature



Date ..30/10/2020

CONTENTS

PROJECT SUMMARY	5
Headline.....	5
Background.....	5
Summary	5
RESEARCH SUMMARY	Error! Bookmark not defined.
Projects.....	6
Future work.....	2
Transferable Skills	3

GROWER SUMMARY

Headline

Working towards autonomous fruit harvesting and accurate long term forecasting.

Background

Novel digital technologies including vision systems, robotics and autonomous systems are seen as potential game changers for the horticulture sector. Vision systems can be used to assess and sense the crop to enable better decision support; robotics and autonomous systems offer new means to drive productivity. These issues apply to all soft and top fruits, but also more widely across the whole fresh produce sector. However, all picking and vision systems are dependent on the development of complex algorithms developed to identify, measure and locate fruit in real time. The development of these systems is not trivial, especially in outdoor environments where the background light level and quality can change within an instant.

Summary

This project aims to progressively implement crucial components required for robust autonomous fruit harvesting. The problem is comprised of five major milestones, Fruit: segmentation, detection/classification, maturity evaluation, quality grading and finally 3D localisation and pose estimation. The new challenge will be to minimise computational requirements to identify fruit whilst maximising processing speed and recognition fidelity. This project will initially focus on strawberry and be anticipated eventually include other soft fruit. Recent work focusses on long and short term tracking of individual fruits spatially and temporally to build a map useful for yield forecasting and online harvesting applications.

SCIENCE SECTION

Project

I am currently involved in the RAS-Berry project listed in the Project Summary section. This involves developing a vision system to detect soft fruits, primarily Strawberries. The end goal of the project is to enable smarter harvesting, counting and yield estimation in-field. Vision systems can be very complex and the infrastructure of fruit harvesting environments/ systems are equally as complex, navigating between speed, environmental, robotic and seasonal constraints. Many current approaches in literature fail to work reliably in commercial 'real life' environments, since they have been tested/developed exclusively in lab environments.

The University of Lincoln has access to the Riseholme campus where 'real world' conditions are simulated on site. This enables the development of algorithms and approaches more suited to wide spread commercial application. Such as the approach I am currently working on. As aforementioned the complexity of the proposed system is great, and with the current direction of the research community, a lot of data is required to be able to develop algorithms/approaches that can robustly solve the problem. Especially with the advent, popularity and great advances being made within the Deep Learning domain.

Many processes in agri-food environments impose temporal restrictions, an example of this can be found when harvesting berries; the optimal time to harvest soft-fruit is when they're at their coldest, ensuring minimal damage occurs to them from plant to supermarket shelves, as well as many other time permitting operations such as providing a yield estimation at the start or middle of the season. Traditional processes have been restricted in areas such as this by availability of labourers and unpredictability of soft fruit growing amongst many other factors. My research will focus on very similar challenges in the berry growing cycle, in an attempt to decrease cost and increase efficiency of agri-food environments.

We design our systems with continuous operation in mind and see this as the primary benefit of implementing robotic systems in agri-food environments. By utilising autonomous mobile robotic platforms (primarily Thorvald, Figure below) we can digitise a wide variety of tasks and increase daily efficiency and lower energy use.

My research will focus on the digitisation of soft-fruit harvesting, fruit counting and yield estimation, through machine learning, deep learning and computer vision. Computer vision in combination with deep learning allows us to gain deeper insight into the environments our systems will operate in, for example we can train a detector to separate berries into maturity classes and only harvest ripe Grade A berries.



My research to date has been successful in minimising the impact of luminance (variability of lighting conditions on model performance over multiple days) through the use of single-stage detectors with bio-inspired modifications at the input level (Figure below). This development was to try and leverage the benefits of continuous deployment of the robotic platforms in a commercial setting; more simply the purpose of the detector was for fruit harvesting, the major benefit of fruit harvesting on a robotic platform is that it can harvest berries all day long without worry of damaging them as you would with human pickers, however variable conditions in farming conditions affect our model performance, so using bio-inspired features we managed to show we can be more invariant to these changes. We published this paper at Sensors journal (<https://doi.org/10.3390/s20010275>), the model architecture is contained in the Figure below. We achieved a performance of 0.793 (F1) in this paper and a near realtime frame rate. More in-depth evaluation can be found in the paper.

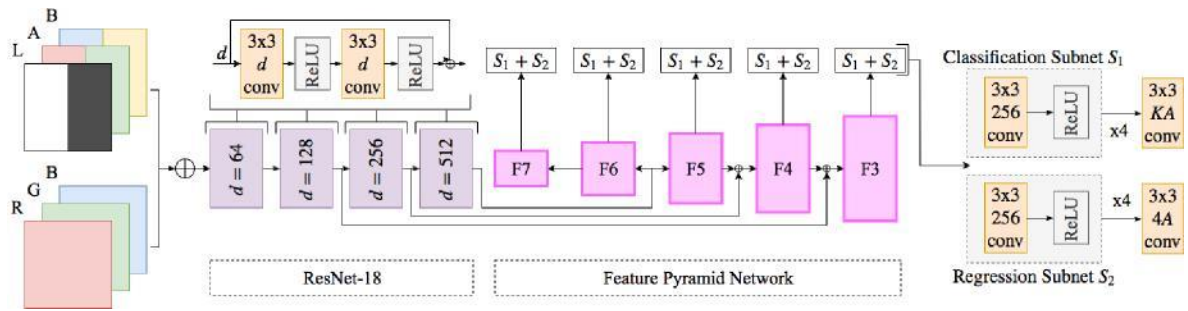


Figure 2: L*a*b*Fruits: A Rapid and Robust Outdoor Fruit Detection System Combining Bio-Inspired Features with One-

Moving on from the single-shot detector we've had success with two-stage architectures such as Hybrid Task Cascade RCNN (Kai Chen 2019) and Mask-RCNN (Kaiming He 2017) due to them fast becoming the standard and also due to the fact they learn pixel level representations of each object, rather than just learning it's bounding box co-ordinates. We've improved our accuracy up to 90% on some datasets utilising more filtering and SOTA approaches. Example images shown in Figure 3. These pixel-level features allow us to more accurately locate each berries position in the real world. We went to translate our Hybrid Task Cascade model results to a cartesian 3 degrees of freedom robotic arm and have managed to show an average picking rate of 3.1s per berry in laboratory conditions.




Figure 3: Object detection, segmentation and localization using deep learning based detectors.

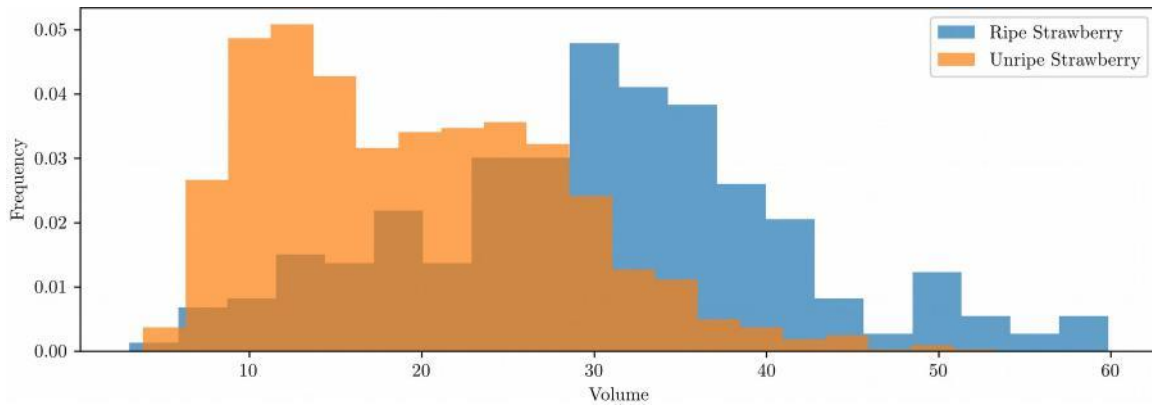
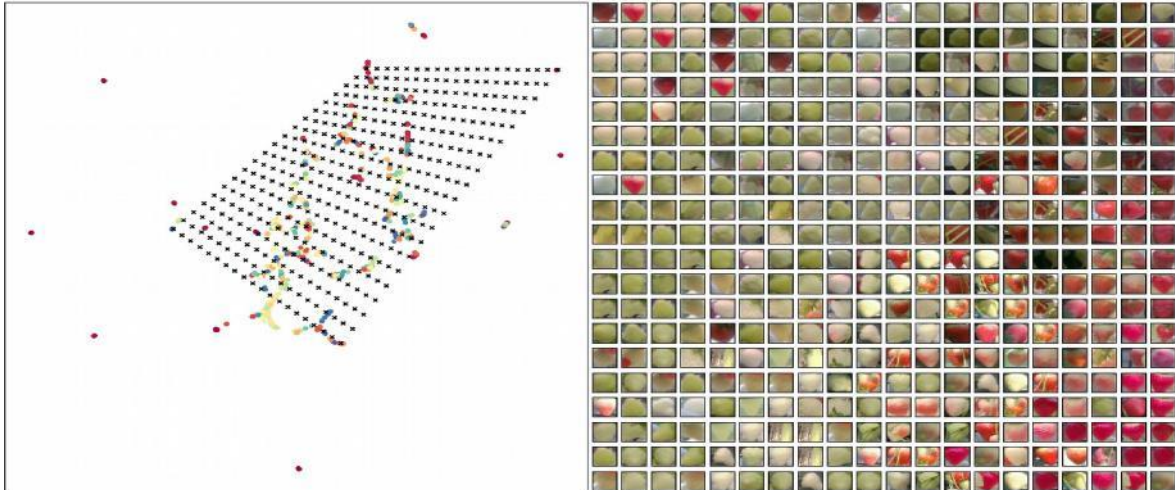


Figure 4: Volume frequency of berries in one row of Riseholme data 2019.

We have since improved our foundational detection systems in terms of performance and accuracy. We can now process 100 frames per second on a Quadro GPU (bbox) and for segmentation detectors we can, in conjunction, calculate phenotypic properties detailed in our UKRAS and upcoming volume estimation paper (<https://doi.org/10.31256/Uk4Td6I>). This process is extremely useful for many horticultural processes primarily yield estimation and current weight on crop.

To further improve our berry system (detection + phenotyping + localisation in GPS coordinates) tracking across images is crucial. We noted in our last AHDB submission the development of a tracking by re-identification model. Since then we have created models based on popular frameworks and evaluated them in the field. The Bayesian tracker has been extended to better work with berries and associative features such as appearance. The figure below details a reprojection from latent 2D space (reduced from 128 dimension space) to strawberry features (figure 5). You can see in the XY coordinates (black dots) the associated learnt identification features. Showing that the berries contain useful information for both re-identification and for tracking purposes. Currently using our Bayesian tracker we can reach performances of up to 0.82 percent, largely dependent on the accuracy of the underlying detection algorithm.

Figure 5: Latent 2D space projection of high dimensional features using UMAP



Future work

For harvesting applications, our naive pick-by-detection strategy was in-effective for dense cluster and highly occluded examples (Harvester pictured in Figure 6), so now we're focusing on dealing with these edge cases. Furthermore, these experiments brought us to the realisation that static detection (detection at a single time point) of produce in-field is not sufficient for intelligent control of harvesting systems, yield estimators, disease control and forecasting systems. Moving forward we're collecting a large scale Strawberry dataset that will provide data suitable for training many different types of artificially intelligent systems, including temporally tracked single berries, dense spatial capture of berries in-row on a mobile platform capturing both environmental data and localisation data (for tracking) and finally alternating camera-views and sensors for generalising to viewpoint and providing quality depth maps used for learning shape and quality attributes of berries. This realisation led to the initial development of a fruit counter using an instance segmentation detections and a Bayesian tracker (statistical inference utilising Bayes' theorem to predict where objects are based on prior observations) to count berries within a fruit row which we successfully demonstrated on Countryfile (<https://www.bbc.co.uk/programmes/m0009v0>).

Figure 6: Strawberry harvester mounted on the Thorvald platform harvesting berries using the object detection approaches mentioned above.



Our work will now focus on finalising the temporal aspects of our underlying system to complete the requirements of a foundational berry detection system that will enable harvesting, fruit monitoring, disease management, labour management, yield prediction and counting applications. The evaluation of the system will occur on data we collect towards the end of 2020, in which we start to process video data over image data. We will also look to finalise our phenotypic paper and publish a Lincoln strawberry detection challenge to bolster future work in this field.

Transferable Skills

Since starting the project I have been involved in numerous smaller projects that have involved high level commercial software implementations. I have also worked closely with industry/academic partners such as Garford Farms Machinery, SAGA Robotics, Thorvald and NMBU.

I have also built up connections with many UK and Norwegian strawberry farmers who assisted me in detailing specifics of the problems at hand. They have also given me access to their farms for future experiments.

While on the project I have spoken numerous times publically for presentations such as LGM and at conferences. I presented my recent work at the Berry Gardens Annual Research Conference in December 2019, showcasing the major advantages of smart computer vision based agri-tech solutions.

I presented in February at the Institution of Agricultural Engineers (IagrE) presentation evening and won the award for the best talk.

At the start of 2020 I presented my work on non-destructive volume and mass estimation of strawberries from images at UKRAS and will present a final system in November to the berry gardens consortium.