Predicting fruit harvest with drones and artificial intelligence

October 17 2019

Credit: University of Cambridge

Outfield Technologies is a Cambridge-based agri-tech start-up company which uses drones and artificial intelligence, to help fruit growers maximise their harvest from orchard crops.

Outfield Technologies' founders Jim McDougall and Oli Hilbourne have been working with Ph.D. student Tom Roddick from the Department's Machine Intelligence Laboratory to develop their technology capabilities

to be able to count the blossoms and apples on a tree via drones surveying enormous [apple](#) orchards.

"An accurate assessment of the blossom or estimation of the harvest allows growers to be more productive, sustainable and environmentally friendly", explains Outfield's commercial director Jim McDougall.

"Our aerial imagery analysis focuses on yield estimation and is really sought after internationally. One of the biggest problems we're facing in the fruit sector is accurate yield forecasting. This system has been developed with growers to plan labour, logistics and storage. It's needed throughout the industry, to plan marketing and distribution, and to ensure that there are always apples on the shelves. Estimates are currently made by growers, and they do an amazing job, but orchards are incredibly variable and estimates are often wrong by up to 20%. This results in lost income, inefficient operations and can result in substantial amount of wastage in unsold crop."
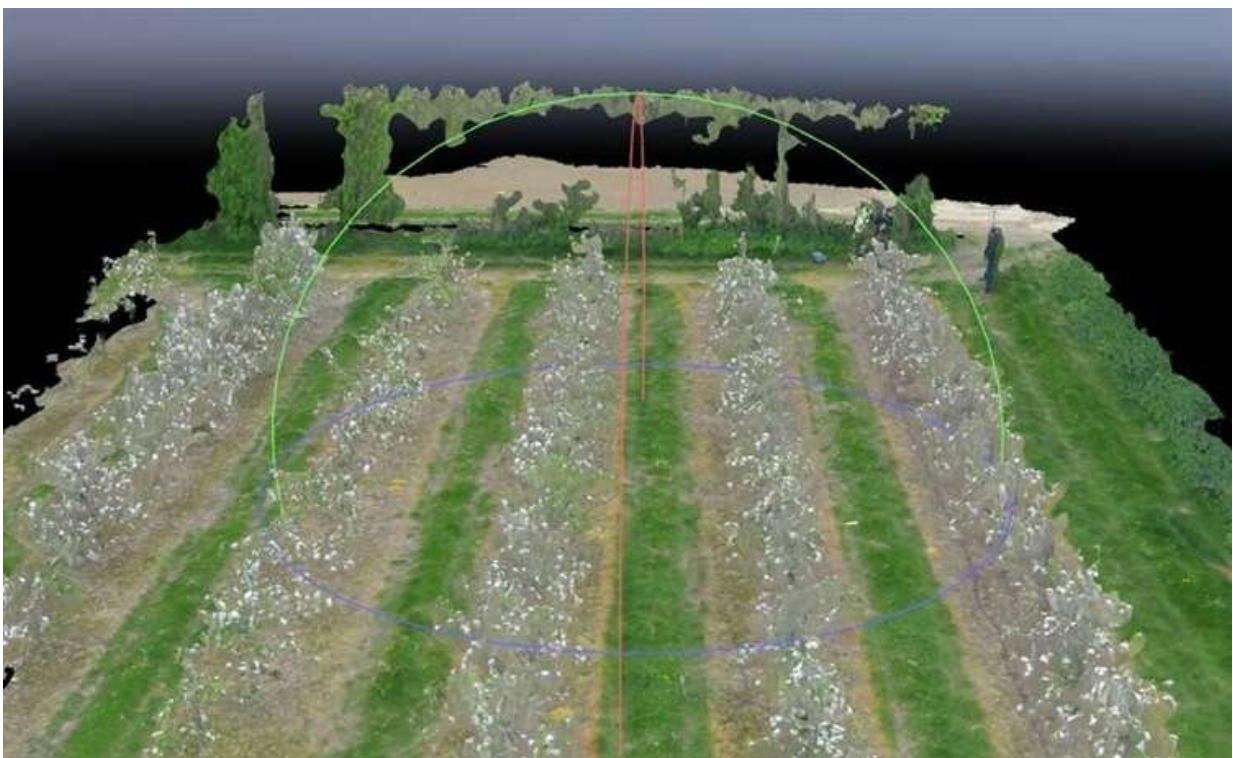
3-D computer reconstruction of a UK orchard flowering in April 2019. Credit: University of Cambridge

Outfield's identification methods are an excellent application of the research that Ph.D. student Tom Roddick, supervised by Professor Roberto Cipolla, is working on. Tom is part of the Computer Vision and Robotics Group which concentrates on artificial intelligence and Machine Learning, using Deep Learning methods, via Artificial Neural Networks (ANNs).

ANNs are computing systems modelled loosely after the human brain, that are designed to recognise patterns. They interpret sensory data by labelling or clustering raw input. The patterns they recognise are numerical, into which all real-world data, be it images, sound, text or time series, is translated.

Such systems "learn" to perform tasks by analysing examples, generally without being programmed with task-specific rules. For example, in image recognition, the ANN might learn to identify images that contain apples by analysing example images that have been manually labelled as "apple" or "no apple" and using the results to identify apples in other images. They do this without any prior knowledge of apples, for example, apple colours or shapes. Instead, they automatically generate identifying characteristics from the examples that they process.

ANNs cluster and classify by first detecting simple patterns in the data, such as edges in images or sounds in speech, and then gradually build a hierarchy of concepts until complex features like faces or sentences emerge in the data. The original goal of the ANN approach was to solve

problems in the same way that a human brain would. However, over time, attention has moved to performing specific tasks. ANNs have been used on a variety of tasks, including computer vision, speech recognition, machine translation, social network filtering, playing board and video games and medical diagnosis.

During his Ph.D., Tom has been working on autonomous driving, looking at street scenes captured on camera, annotating and labelling each element. He pinpoints where the cars are, the pedestrians, the kerb etc. To do this he uses a tool called semantic segmentation to label each individual pixel to give a high level understanding of what is going on. Outfield need to identify apples and blossom in their orchard photographs and one way to do that is to use this semantic segmentation method.

Another aspect of Outfield's data collection method is to pinpoint where their drones are at all times and there is another strand of computer vision that concentrates on localisation, working out where you are in the world and what you are looking at. Alumnus Kesar Breen, independent [machine learning](#) and computer vision consultant, has taken time out of his busy schedule to advise Jim and Oli. Kesar has helped them with an overview of the technologies that they could be using for the orchard modelling and analysis to find out where the crops are, and drafted a potential algorithm to do this, with time frames and requirements. Kesar says, "Outfield is doing some very cool stuff working with some interesting but proven technologies, on an important business problem. I think it is very likely to be commercially viable."

Talking about his work with Outfield, Tom says, "Outfield's semantic segmentation needs have some very specific subtleties that are very interesting from a research point of view for example. I am used to looking at images to identify large objects such as cars, which are easy to spot, but what Outfield has are these huge aerial view images of orchards

that are millions and millions of pixels and it wants to detect each blossom flower or each piece of fruit to calculate how many of them there are. I have been looking at how to do that efficiently and robustly to be able to distinguish between things like; is this an apple on a tree? Or is this an apple on the ground?"

Jim says, "The UK has some of the best technology and the best technology scientists in the world. We are currently beta testing which includes using the model with crops other than apples.

"We have a robust plan for the next two to three years, and we are opening an investment round in October 2019 to close in Q1 2020. This will allow us to bring onboard more of the team full time and test the products out at scale in New Zealand and the UK over the next year."

Provided by University of Cambridge

Citation: Predicting fruit harvest with drones and artificial intelligence (2019, October 17) retrieved 10 April 2024 from
https://techxplore.com/news/2019-10-fruit-harvest-drones-artificial-intelligence.html