Computers and Electronics in Agriculture 130 (2016) 83–96

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Original papers

Mapping almond orchard canopy volume, flowers, fruit and yield using lidar and vision sensors

James P. Underwood *, Calvin Hung, Brett Whelan, Salah Sukkarieh

Australian Centre for Field Robotics, The University of Sydney, NSW 2006, Australia

ARTICLE INFO

Article history:
Received 15 April 2016
Received in revised form 20 September 2016
Accepted 24 September 2016

Keywords:
Robotics
Sensing
Machine vision
Lidar
Multi-sensor fusion
Orchard yield mapping

ABSTRACT

This paper present a mobile terrestrial scanning system for almond orchards, that is able to efficiently map flower and fruit distributions and to estimate and predict yield for individual trees. A mobile robotic ground vehicle scans the orchard while logging data from on-board lidar and camera sensors. An automated software pipeline processes the data offline, to produce a 3D map of the orchard and to automatically detect each tree within that map, including correct associations for the same trees seen on prior occasions. Colour images are also associated to each tree, leading to a database of images and canopy models, at different times throughout the season and spanning multiple years. A canopy volume measure is derived from the 3D models, and classification is performed on the images to estimate flower and fruit density. These measures were compared to individual tree harvest weights to assess the relationship to yield. A block of approximately 580 trees was scanned at peak bloom, fruit-set and just before harvest for two subsequent years, with up to 50 trees individually harvested for comparison. Lidar canopy volume had the strongest linear relationship to yield with $R^2 = 0.77$ for 39 tree samples spanning two years. An additional experiment was performed using hand-held photography and image processing to measure fruit density, which exhibited similar performance ($R^2 = 0.71$). Flower density measurements were not strongly related to yield, however, the maps show clear differentiation between almond varieties and may be useful for other studies.

1. Introduction

Technological improvements in sensing, computing, algorithms and robotics have the potential to increase productivity for commercial farming and efficiency for plant science. For farmers, mobile data gathering systems can provide detailed information to assist their decision making and management processes and the information can plug into decision support software that is capable of recommending particular actions. Eventually it will be possible for these actions to be directly applied using mobile field robotics technology. For plant scientists, mobile data systems can provide high throughput, in-field plant phenomics. This will allow greater capacity for in-field experimentation, where manual labour for in-field data acquisition is currently a limiting factor, leading to yield improvements from genomics and improvements to best-practice for growers.

This paper presents a robotic ground-vehicle information system for almond orchard mapping and per-tree yield estimation. The system continuously records data from a camera and lidar sensor, while the vehicle drives through the orchard. This is combined with processing software that automatically extracts geometric and visual information of each tree and matches the data from scans taken at different times of the year and over multiple seasons. This allows the assessment of flowering and fruit production per individual tree, in a manner that is efficient for scanning whole orchard blocks. This represents a significant increase in resolution compared to the typical practice of weighing the total produce harvested from whole orchard blocks, allowing the variability between individual trees and smaller regions of the orchard to be estimated and mapped.

Three dimensional and image-based sensing have been applied to many aspects of tree-crop precision agriculture. There are many examples of the use of lidar to measure tree canopy geometry (Tumbo et al., 2002; Walklate et al., 2002; Rosell et al., 2009; Rosell and Sanzs, 2012; Wellington and Campoy, 2012; Escolà et al., 2015) and as a proxy for related measurements such as leaf-area-index (Sanz et al., 2013). Alternative range sensors have also been used including ultrasound (Tumbo et al., 2002; Hosainpour et al., 2013), structured light (Rosell-Polo et al., 2015) and stereo vision Rosell and Sanzs (2012), but lidar is popular given

* Corresponding author.
E-mail addresses: james.underwood@sydney.edu.au (J.P. Underwood), brett.whelan@sydney.edu.au (B. Whelan), salah.sukkarieh@sydney.edu.au (S. Sukkarieh).

http://dx.doi.org/10.1016/j.compag.2016.09.014
0168-1699/

the relatively high accuracy and invariance under natural illumination conditions.

Vision sensors have been coupled with machine vision algorithms to estimate fruit and flower densities for individual trees for a number of different crop types (Gongal et al., 2015). Hand-held digital cameras and relatively simple image classification algorithms have been used to estimate flower densities (Adamsen et al., 2000; Aggelopoulou et al., 2011; Thorp and Dierig, 2011); relatively simple algorithms are possible due to the typically high colour contrast exhibited by flowers. Machine vision cameras have also been mounted on tractors to improve the image acquisition process (Hočevar et al., 2014), which allowed flower estimation on larger numbers of trees ($N = 136$). Nevertheless, the process requires manual demarcation of the individual trees within the frames, which limits the efficiency when scaling up to scan entire commercial orchard blocks. Unlike flowers, fruit and nuts are often visually similar to the leaves and surrounding foliage of trees, meaning that more sophisticated machine vision methods are required to automatically detect them. These include artificial neural networks (Gongal et al., 2015; Hung et al., 2013, 2015; Bargoti and Underwood, 2016a,b), which have the advantage of automatically learning appropriate feature descriptions for classification from the data; multi-sensor fusion, which simplifies the classification problem by fusing data from complementary sensors such as vision and thermal cameras (Bulanon et al., 2009) or vision and near-infra-red cameras (Hung et al., 2013); and hybrid approaches that combine colour and shape (Singh et al., 2010; Nuske et al., 2011; Wang et al., 2013). Cameras have also been combined with lidar for tree-crop applications, such as Shalal et al. (2015) and Bargoti et al. (2015), which both address tree trunk detection by combining the geometry sensed by the laser with the visual appearance sensed by the camera.

The size of almond trees is known to be an important factor in estimating the yield (Hill et al., 1987), which motivates canopy geometry sensing. Flower densities are also considered to be relevant to yield, although the relationship is complicated by variability in pollination (e.g. availability of pollinators) and other limitations in how much fruit the tree can bare Dicenta et al. (2005). The potential utility of flower density mapping, as well potentially being able to directly measure fruit density motivates the use of vision.

In order to allow all of these methods to scale up to entire orchards, automated, streamlined data management is also required, which includes software for tree segmentation and detection (Wellington and Campoy, 2012; Shalal et al., 2015; Bargoti et al., 2015; Underwood et al., 2015b) and tree matching (correct data to tree assignment) for repeated scans at different times (Underwood et al., 2015b). There are few whole-system examples that combine geometric and visual sensing, together with efficient mobile data acquisition and automated data processing and management steps that facilitate entire blocks of commercial farms to be efficiently scanned, including comparisons to ground truth yield such as fruit counts or harvest weights. Prominent examples include holistic systems for ground crops (Busemeyer et al., 2013), vineyards (Nuske et al., 2014) and apple orchards (Hung et al., 2015). Each of these systems relies on constraints relating to the specific nature of the target crop and no one system and approach is likely to be adaptable to vastly different crop types. The geometry of ground based crops is well suited to systems that straddle above the crop (Busemeyer et al., 2013; Deery et al., 2014; Underwood et al., 2015a), which enables controlled illumination and provides for an ideal sensor vantage point, but those systems are not applicable to tree crops, which are taller and difficult to straddle. Amongst tree crop applications, algorithms are typically tailored to the appearance and geometry of the specific fruit such as circle detection for apples in Hung et al. (2015) and a customised grape berry detection method in Nuske et al. (2014). Algorithms are also tailored to the orchard configuration, such as the two dimensional trellis fruit wall in Hung et al. (2015), which avoids the need for individual tree segmentation for image masking, which is a key component of our system.

Although several of the components that are necessary for an almond orchard scanning system have been explored in the literature (including the authors' prior work and others), the contribution of this paper is to develop an integrated methodology to create a single, efficient orchard scanning system for almonds. The contribution is the complete system, including the necessary developments to combine the components of tree detection and segmentation with flower and fruit mapping and yield estimation, together with experimentation covering approximately 580 trees at three times of season for two years.

2. Materials and methods

A 2.3-hectare block of a commercial almond orchard was scanned using the "Shrimp" ground vehicle robot, three times per year (flowering, fruit-set, pre-harvest) for two subsequent years. At harvest, individual trees were also photographed with a hand-held digital camera and then selectively harvested and weighed. The colour images and lidar from the robot and the manually taken photographs were post-processed, using custom algorithms and software, to derive measures relating to the canopy volume and the density of flowers and fruit on each tree. The measures were compared with the selective harvest weights, to quantify performance. This section describes the system and sensors, the protocol for data collection, and how the data were processed.

2.1. The "Shrimp" mobile robot

The "Shrimp" mobile ground vehicle robot was designed and built at the Australian Centre for Field Robotics at the University of Sydney in 2009 (see Fig. 1(a)) as a general purpose research platform to study robotic sensing and perception. For this application, we use a subset of sensors: a 2D line scanning lidar (SICK LMS-291) and a machine vision camera (Point Grey Ladybug2, single two mega-pixel camera, with natural illumination only), both face to the right to scan the trees as the vehicle travels continuously forwards at up to 2 m/s (see Fig. 1(b)). A real time kinematic global positioning inertial navigation system (Novatel SPAN RTK GPS/INS) is used for positioning, a gamma radiometer (RS700) was mounted on-board to record passive soil gamma emissions and an electromagnetic induction instrument (EM38) is towed behind the vehicle to measure apparent soil electrical conductivity (ECa). The system includes a computer for data logging.

2.2. Data collection

All data were obtained from a 2.3-hectare section of Lake Cullulleraine Almonds, in Victoria, Australia, shown in Fig. 2. The area includes 10 rows spaced 7.35 m apart and roughly 313.5 m long, with 58 trees per row spaced at 5.5 m. The primary variety in every second row is Nonpareil, with alternating Carmel and Monterey pollinator rows between. Datasets were collected at peak bloom (as estimated by the farm manager), at fruit-set and just prior to harvest, for both the 2014 and 2015 harvest seasons. Several datasets were taken on subsequent days to assess repeatability. All datasets are summarised in Table 1.

To obtain each dataset, the raw data from all sensors were logged continuously, while the vehicle is driven by a remote operator up and back down the rows at speeds of 1–2 m/s depending on