

Publications

2019

2019 - Ground based hyperspectral imaging for extensive mango yield estimation, *Computers and Electronics in Agriculture*

Salvador Gutiérrez, Alexander Wendel, James Underwood

Fruit yield estimation in orchard blocks is an important objective in the context of precision agriculture, as it makes it easier for the farmer to plan ahead and efficiently use resources. Nevertheless, its implementation is labour-intensive and involves the manual counting of the fruit present in the trees. While colour (RGB) has been widely shown to be successful and arguably sufficient for yield estimation in orchards, hyperspectral imaging (HSI) shows promise for more nuanced tasks such as disease detection, cultivar classification and fruit maturity estimation. Therefore, it is important to ask how appropriate is HSI for the task of yield estimation, with a view to performing all of these tasks with just one sensor. This paper presents a novel mango yield estimation pipeline using ground based line-scan HSI acquired from an unmanned ground vehicle. Hyperspectral images were collected on a commercial mango orchard block in December 2017 and pre-processed for illumination compensation. After tree delimitation and mango pixel identification, an optimisation process was carried out to obtain the best models for fruit counting, using mango counts obtained by manually counting the fruit on-tree, and using state-of-the-art RGB techniques for yield estimation. Models were validated and tested on hundreds of trees, and subsequently mapped. In testing, determination coefficients reached values of up to 0.75 against field counts (predicting 18 trees) and 0.83 against RGB mango counts (predicting 216 trees). These results suggest that line-scan HSI can be used to accurately estimate yield in orchards, especially in scenarios in which this technology is already chosen for the determination of other traits.

[Paper at COMPAG](#)

2018

2018 - Maturity estimation of mangoes using hyperspectral imaging from a ground based mobile platform, *Computers and Electronics in Agriculture*

Alexander Wendel, James Underwood, Kerry Walsh

Monitoring the maturity of fruit in commercial orchards can help growers optimise the time of harvest. Dry matter content (DM) of fruit is used as an indicator of mango maturity, measured in-field with a hand-held spectrometer. This approach is labour intensive, limiting the extent to which DM variability can be measured across an orchard block, which would enable selective harvesting. This paper proposes an alternative approach that utilises a hyperspectral camera, LIDAR sensor and navigation system mounted to a ground vehicle to predict fruit DM individually for hundreds of trees in a mango orchard block. First, the challenges faced due to tree geometry and shadows in mango orchards are addressed. Then the ability to predict DM at a distance using hyperspectral imaging (411.3–867.0 nm) was demonstrated. Two regression methods, partial least squares (PLS) and a convolutional neural network (CNN) were compared and tested against DM results from a hand-held NIR spectrometer using harvested ($n = 468$, $= 2.32\%$ w/w) and on-tree fruit ($n = 662$, $= 1.79\%$ w/w). The CNN achieved a cross validation $R^2 = 0.642$ and $RMSE = 1.08\%$ w/w CV in fruit on tree, while PLS achieved $R^2 = 0.582$ and $RMSECV = 1.17\%$ w/w. In order to discriminate mango and non-mango pixels, PLS discriminant analysis (PLSDA) and a CNN were also compared, where both methods achieved good classification performance with a mean $F1 > 0.97$. Having established mango classification and DM prediction performance, hyperspectral data were processed for a full orchard block and projected to world coordinates using AGV position and orientation as provided by the navigation system. Trees were segmented using corresponding LIDAR data, which allowed association of projected DM predictions to individual trees. Repeated scans of the orchard block over two days allowed a measure of repeatability, which was achieved with an $RMSE < 0.29\%$ w/w. The results provide strong evidence that predicting maturity at a distance for all trees in an orchard is feasible using a hyperspectral camera, which will be an important management tool for growers to optimise harvest timing and yield.

[Paper at COMPAG](#)

2018 - Estimation of fruit load in mango orchards: tree sampling considerations and use of machine vision and satellite imagery, *Precision Agriculture*

Nicholas T. Anderson, James Underwood, Moshuir M. Rahman, Andrew Robson, Kerry B. Walsh

In current best commercial practice, pre-harvest fruit load predictions of mango orchards are provided based on a manual count of fruit number on up to 5% of trees within each block. However, the variability in fruit number per tree (coefficient of variation, CV, from 27 to 93% across ten orchards) was demonstrated to be such that the best case commercial sampling practice was inadequate for reliable estimation (to an error of 54–82 fruit/tree, and percentage error, PE, of 10% at a probability of 0.95). These results highlight the need for alternative methods for estimation of orchard fruit load. Pre-harvest fruit load was estimated for a case study orchard of 469 trees using (i) count of a sample of trees, (ii) in-field machine vision and (iii) correlation to a tree spectral index estimated using high resolution satellite imagery. A count of 5% of trees (23) in the trial orchard resulted in a PE of 31% (error of 37 fruit/tree), with a count of 157 trees required to achieve a PE of 10% (error of 12 fruit/tree). Sampling effort to achieve a PE of 10% was decreased by only 10% by sampling from aspatial k-means tree classifications based on machine vision derived fruit counts of all trees. Clustering based on tree attributes of canopy volume and trunk circumference was not helpful in decreasing sampling effort as these attributes were poorly correlated to fruit load ($R^2=0.21$ and 0.17 , respectively). In-field multi-view machine vision-based estimation of fruit load per tree achieved a $R^2=0.97$ and a $RMSE=14.8$ fruit/tree against harvest fruit count per tree for a set of 18 trees (average=88; $SD=82$ fruit/tree), using a faster region convolutional neural network trained the previous season. The relationship between WorldView-3 (WV3) satellite spectral reflectance characteristics of sampled trees and fruit number was characterised by a $R^2=0.66$ and a $RMSE=56.1$ fruit /tree. For this orchard, for which the actual fruit harvest was 56,720 fruit, the estimate based on a manual count of 5% of trees was 47,955 fruit, while estimates based on 20 iterations of stratified sampling (of 5% of trees in each cycle) had variation (SD) of 9597. The machine vision method resulted in an estimate of 53,520 ($SD=1960$) fruit and the remote sensing method, 51,944 ($SD=26,300$) fruit for the orchard.

[Paper at Precision Agriculture](#)

2018 - Machine vision assessment of mango orchard flowering, *Computers and Electronics in Agriculture*

Zhenglin Wang, James Underwood, Kerry B. Walsh

Machine vision assessment of mango orchard flowering involves detection of an inflorescence (a panicle) with flowers at various stages of development. Two systems were adopted contrasting in camera, illumination hardware and image processing. The image processing paths were: (i) colour thresholding of pixels followed by SVM classification to estimate inflorescence associated pixel number (panicle area), and panicle area relative to total canopy area ('flowering intensity') using two images per tree ('dual view'), and (ii) a faster R-CNN for panicle detection, using either 'dual-view' or 'multi-view' tracking of panicles between consecutive images to achieve a panicle count per tree. The correlation coefficient of determination between the machine vision flowering intensity and area estimate (path i) and in field human visual counts of panicles (past 'asparagus' stage) per tree was 0.69 and 0.81, while that between the machine vision (path ii) and human panicle count per tree was 0.78 and 0.84 for the dual and multi-view detection approaches, respectively (n=24), while that for repeat human counts was 0.86. The use of such information is illustrated in context of (i) monitoring the time of peak flowering based on repeated measures of flowering intensity, for use as the start date within heat sum models of fruit maturation, (ii) identification and mapping of early flowering trees to enable selective early harvest and (iii) exploring relationships between flowering and fruit yield. For the current orchard and season, the correlation coefficient of determination between machine vision estimates of panicle area and multi-view panicle count and fruit yield per tree was poor (R^2 of 0.19 and 0.28, respectively, n=44), indicative of variable fruit set per panicle and retention between trees.

[Paper at COMPAG](#)

2018 - Light interception modelling using unstructured LiDAR in avocado orchards, *Computers and Electronics in Agriculture*

Fred Westling, James Underwood, Samuel Örn

In commercial fruit farming, managing the light distribution through canopies is important because the amount and distribution of solar energy that is harvested by each tree impacts the production of fruit quantity and quality. It is therefore an important characteristic to measure and ultimately to control with pruning. We present a solar-geometric model to estimate light interception in individual avocado (*Persea americana*) trees, that is designed to scale to whole-orchard scanning, ultimately to inform pruning decisions. The geometry of individual trees was measured using LiDAR and represented by point clouds. A discrete energy distribution model of the hemispherical sky was synthesised using public weather records. The light from each sky node was then ray traced, applying a radiation absorption model where rays pass the point cloud representation of the tree. The model was validated using ceptometer energy measurements at the canopy floor, and model parameters were optimised by analysing the error between modelled and measured energies. The model was shown to perform well qualitatively well through visual comparison with tree shadows in photographs, and quantitatively well with $R^2 = 0.854$, suggesting it is suitable to use in the context of agricultural decision support systems, in future work.



[Paper at arXiv](#)

[Paper at COMPAG](#)

2018 - Object Detection for Cattle Gait Tracking, *IEEE International Conference on Robotics and Automation (ICRA)*

John Gardenier, James Underwood, Cameron Clark

Lameness in cattle is a health issue where gait is modified to minimise pain. Cattle are currently visually assessed for locomotion score, which provides the degree of lameness for individual animals. This subjective method is costly in terms of labour, and its level of accuracy and ability to detect small changes in locomotion that is critical for early detection of lameness and associated intervention. Current automatic lameness detection systems found in literature have not yet met the ultimate goal of widespread commercial adoption. We present a sensor configuration to record cattle kinematics towards automatic lameness detection. This configuration features four Time of Flight sensors to view cattle from above and from one side as they exit an automatic rotary milking dairy. Two dimensional near infrared images sampled from 223 cows passing through the system were used to train a Faster R-CNN to detect hooves (F1-score = 0.90) and carpal/tarsal joints (F1-score = 0.85). The depth images were used to project these detected key points into Cartesian space where they were tracked to obtain individual trajectories per limb. The results show that kinematic gait features can be successfully obtained as a first and important step towards objective, accurate, automatic lameness detection.



2018-Gardenier-I...ait-Tracking.pdf

[ICRA conference website](#)

2018 - Fruit Load Estimation in Mango Orchards: A Comparison of Methods, IEEE International Conference on Robotics and Automation (ICRA) - Workshop on Robotic Vision and Action in Agriculture

James Underwood, Moshiur Rahman, Andrew Robson, Kerry Walsh, Anand Koirala, Zhenglin Wang

Abstract—The fruit load of entire mango orchards was estimated well before harvest using (i) in-field machine vision on mobile platforms and (ii) WorldView-3 satellite imagery. For in-field machine vision, two imaging platforms were utilized, with (a) day time imaging with LiDAR based tree segmentation and multiple views per tree, and (b) night time imaging system using two images per tree. The machine vision approaches involved training of neural networks with image snips from one orchard only, followed by use for all other orchards (varying in location and cultivar). Estimates of fruit load per tree achieved up to a $R^2 = 0.88$ and a $RMSE = 22.5$ fruit/tree against harvest fruit count per tree ($n = 18$ trees per orchard). With satellite imaging, a regression was established between a number of spectral indices and fruit number for a set ($n=18$) of trees in each orchard (example: $R^2 = 0.57$, $RMSE = 22$ fruit/tree), and this model applied across all tree associated pixels per orchard. The weighted average percentage error on packhouse counts (weighted by packhouse fruit numbers) was 6.0, 8.8 and 9.9% for the day imaging system, night imaging machine vision system and the satellite method, respectively, averaged across all orchards assessed. Additionally, fruit sizing was achieved with a $RMSE = 5$ mm (on fruit length and width). These estimates are useful for harvest resource planning and marketing and set the foundation for automated harvest.



2018-Underwood...o orchards.pdf

[ICRA workshop website](#)



2017

2017 - Extrinsic Parameter Calibration for Line Scanning Cameras on Ground Vehicles with Navigation Systems Using a Calibration Pattern, Sensors

Alex Wendel, James Underwood

Line scanning cameras, which capture only a single line of pixels, have been increasingly used in ground based mobile or robotic platforms. In applications where it is advantageous to directly georeference the camera data to world coordinates, an accurate estimate of the camera's 6D pose is required. This paper focuses on the common case where a mobile platform is equipped with a rigidly mounted line scanning camera, whose pose is unknown, and a navigation system providing vehicle body pose estimates. We propose a novel method that estimates the camera's pose relative to the navigation system. The approach involves imaging and manually labelling a calibration pattern with distinctly identifiable points, triangulating these points from camera and navigation system data and reprojecting them in order to compute a likelihood, which is maximised to estimate the 6D camera pose. Additionally, a Markov Chain Monte Carlo (MCMC) algorithm is used to estimate the uncertainty of the offset. Tested on two different platforms, the method was able to estimate the pose to within 0.06 m/1.05 degrees and 0.18 m/2.39 degrees. We also propose several approaches to displaying and interpreting the 6D results in a human readable way.



[Paper at MDPI Sensors](#)

2017 - Fruit Detection and Tree Segmentation for Yield Mapping in Orchards, The University of Sydney, PhD Thesis

Suchet Bargoti

Accurate information gathering and processing is critical for precision horticulture, as growers aim to optimise their farm management practices. An accurate inventory of the crop that details its spatial distribution along with health and maturity, can help farmers efficiently target processes such as chemical and fertiliser spraying, crop thinning, harvest management, labour planning and marketing. Growers have traditionally obtained this information by using manual sampling techniques, which tend to be labour intensive, spatially sparse, expensive, inaccurate and prone to subjective biases. Recent advances in sensing and automation for field robotics allow for key measurements to be made for individual plants throughout an orchard in a timely and accurate manner. Farmer operated machines or unmanned robotic platforms can be equipped with a range of sensors to capture a detailed representation over large areas. Robust and accurate data processing techniques are therefore required to extract high level information needed by the grower to support precision farming.

This thesis focuses on yield mapping in orchards using image and light detection and ranging (LiDAR) data captured using an unmanned ground vehicle (UGV). The contribution is the framework and algorithmic components for orchard mapping and yield estimation that is applicable to different fruit types and orchard configurations. The framework includes detection of fruits in individual images and tracking them over subsequent frames. The fruit counts are then associated to individual trees, which are segmented from image and LiDAR data, resulting in a structured spatial representation of yield.

The first contribution of this thesis is the development of a generic and robust fruit detection algorithm. Images captured in the outdoor environment are susceptible to highly variable external factors that lead to significant appearance variations. Specifically in orchards, variability is caused by changes in illumination, target pose, tree types, etc. The proposed techniques address these issues by using state-of-the-art feature learning approaches for image classification, while investigating the utility of orchard domain knowledge for fruit detection. Detection is performed using both pixel-wise classification of images followed instance segmentation, and bounding-box regression approaches. The experimental results illustrate the versatility of complex deep learning approaches over a multitude of fruit types.

The second contribution of this thesis is a tree segmentation approach to detect the individual trees that serve as a standard unit for structured orchard information systems. The work focuses on trellised trees, which present unique challenges for segmentation algorithms due to their intertwined nature. LiDAR data are used to segment the trellis face, and to generate proposals for individual tree trunks. Additional trunk proposals are provided using pixel-wise classification of the image data. The multi-modal observations are fine-tuned by modelling trunk locations using a hidden semi-Markov model (HSMM), within which prior knowledge of tree spacing is incorporated.

The final component of this thesis addresses the visual occlusion of fruit within geometrically complex canopies by using a multi-view detection and tracking approach. Single image fruit detections are tracked over a sequence of images, and associated to individual trees or farm rows, with the spatial distribution of the fruit counting forming a yield map over the farm. The results show the advantage of using multi-view imagery (instead of single view analysis) for fruit counting and yield mapping.

This thesis includes extensive experimentation in almond, apple and mango orchards, with data captured by a UGV spanning a total of 5 hectares of farm area, over 30 km of vehicle traversal and more than 7,000 trees. The validation of the different processes is performed using manual annotations, which includes fruit and tree locations in image and LiDAR data respectively. Additional evaluation of yield mapping is performed by comparison against fruit counts on trees at the farm and counts made by the growers post-harvest. The framework developed in this thesis is demonstrated to be accurate compared to ground truth at all scales of the pipeline, including fruit detection and tree mapping, leading to accurate yield estimation, per tree and per row, for the different crops. Through the multitude of field experiments conducted over multiple seasons and years, the thesis presents key practical insights necessary for commercial development of an information gathering system in orchards.



[Thesis at USYD Library](#)

2017 - Multi-Modal Obstacle Detection in Unstructured Environments with Conditional Random Fields, *arXiv*

Mikkel Kragh, James Underwood

Reliable obstacle detection and classification in rough and unstructured terrain such as agricultural fields or orchards remains a challenging problem. These environments involve large variations in both geometry and appearance, challenging perception systems that rely on only a single sensor modality. Geometrically, tall grass, fallen leaves, or terrain roughness can mistakenly be perceived as non-traversable or might even obscure actual obstacles. Likewise, traversable grass or dirt roads and obstacles such as trees and bushes might be visually ambiguous. In this paper, we combine appearance- and geometry-based detection methods by probabilistically fusing lidar and camera sensing using a conditional random field. We apply a state-of-the-art multi-modal fusion algorithm from the scene analysis domain and adjust it for obstacle detection in agriculture with moving ground vehicles. This involves explicitly handling sparse point cloud data and exploiting both spatial, temporal, and multi-modal links between corresponding 2D and 3D regions. The proposed method is evaluated on a diverse dataset, comprising a dairy paddock and a number of different orchards gathered with a perception research robot in Australia. Results show that for a two-class classification problem (ground and non-ground), only the camera leverages from information provided by the other modality. However, as more classes are introduced (ground, sky, vegetation, and object), both modalities complement each other and improve the mean classification score. Further improvement is achieved by introducing recursive inference with temporal links between successive frames.



[Paper at arXiv](#)

2017 - Efficient In-Field Plant Phenomics for Row-Crops with an Autonomous Ground Vehicle, *Journal of Field Robotics*

James Underwood, Alex Wendel, Brooke Schofield, Larn McMurray, Rohan Kimber

The scientific areas of plant genomics and phenomics are capable of improving plant productivity, yet they are limited by the manual labor that is currently required to perform in-field measurement, and a lack of technology for measuring the physical performance of crops growing in the field. A variety of sensor technology has the potential to efficiently measure plant characteristics that are related to production. Recent advances have also shown that autonomous airborne and manually driven ground-based sensor platforms provide practical mechanisms for deploying the sensors in the field. This paper advances the state-of-the-art by developing and rigorously testing an efficient system for high throughput in-field agricultural row-crop phenotyping. The system comprises an autonomous unmanned ground-vehicle robot for data acquisition and an efficient data post-processing framework to provide phenotype information over large-scale real-world plant-science trials. Experiments were performed at three trial locations at two different times of year, resulting in a total traversal of 43.8 km to scan 7.24 hectares and 2423 plots (including repeated scans). The height and canopy closure data were found to be highly repeatable ($r^2 = 1.00$ $N = 280$, $r^2 = 0.99$ $N = 280$, respectively) and accurate with respect to manually gathered field data ($r^2 = 0.95$ $N = 470$, $r^2 = 0.91$ $N = 361$, respectively), yet more objective and less-reliant on human skill and experience. The system was found to be a more labor-efficient mechanism for gathering data, which compares favorably to current standard manual practices.



[Full paper at JFR](#)

2017 - Illumination Compensation in Ground Based Hyperspectral Imaging, *ISPRS Journal of Photogrammetry and Remote Sensing*

Alex Wendel, James Underwood,

Hyperspectral imaging has emerged as an important tool for analysing vegetation data in agricultural applications. Recently, low altitude and ground based hyperspectral imaging solutions have come to the fore, providing very high resolution data for mapping and studying large areas of crops in detail. However, these platforms introduce a unique set of challenges that need to be overcome to ensure consistent, accurate and timely acquisition of data. One particular problem is dealing with changes in environmental illumination while operating with natural light under cloud cover, which can have considerable effects on spectral shape. In the past this has been commonly achieved by imaging known reference targets at the time of data acquisition, direct measurement of irradiance, or atmospheric modelling. While capturing a reference panel continuously or very frequently allows accurate compensation for illumination changes, this is often not practical with ground based platforms, and impossible in aerial applications. This paper examines the use of an autonomous unmanned ground vehicle (UGV) to gather high resolution hyperspectral imaging data of crops under natural illumination. A process of illumination compensation is performed to extract the inherent reflectance properties of the crops, despite variable illumination. This work adapts a previously developed subspace model approach to reflectance and illumination recovery. Though tested on a ground vehicle in this paper, it is applicable to low altitude unmanned aerial hyperspectral imagery also. The method uses occasional observations of reference panel training data from within the same or other datasets, which enables a practical field protocol that minimises in-field manual labour. This paper tests the new approach, comparing it against traditional methods. Several illumination compensation protocols for high volume ground based data collection are presented based on the results. The findings in this paper are applicable not only to robotics or agricultural applications, but most very low altitude or ground based hyperspectral sensors operating with natural light.



[Full paper at ISPRS](#)

2017 - Deep Fruit Detection in Orchards, *IEEE International Conference on Robotics and Automation (ICRA) 2017*

Suchet Bargoti, James Underwood

An accurate and reliable image based fruit detection system is critical for supporting higher level agriculture tasks such as yield mapping and robotic harvesting. This paper presents the use of a state-of-the-art object detection framework, Faster R-CNN, in the context of fruit detection in orchards, including mangoes, almonds and apples. Ablation studies are presented to better understand the practical deployment of the detection network, including how much training data is required to capture variability in the dataset. Data augmentation techniques are shown to yield significant performance gains, resulting in a greater than two-fold reduction in the number of training images required. In contrast, transferring knowledge between orchards contributed to negligible performance gain over initialising the Deep Convolutional Neural Network directly from ImageNet features. Finally, to operate over orchard data containing between 100-1000 fruit per image, a tiling approach is introduced for the Faster R-CNN framework. The study has resulted in the best yet detection performance for these orchards relative to previous works, with an F1-score of >0.9 achieved for apples and mangoes.



[ICRA version](#)

[Original arxiv version](#)

2017 - Image Segmentation for Fruit Detection and Yield Estimation in Apple Orchards, *Journal of Field Robotics*, Special Issue in Agricultural Robotics

Suchet Bargoti, James Underwood

Ground vehicles equipped with monocular vision systems are a valuable source of high resolution image data for precision agriculture applications in orchards. This paper presents an image processing framework for fruit detection and counting using orchard image data. A general purpose image segmentation approach is used, including two feature learning algorithms; multi-scale Multi-Layered Perceptrons (MLP) and Convolutional Neural Networks (CNN). These networks were extended by including contextual information about how the image data was captured (metadata), which correlates with some of the appearance variations and/or class distributions observed in the data. The pixel-wise fruit segmentation output is processed using the Watershed Segmentation (WS) and Circular Hough Transform (CHT) algorithms to detect and count individual fruits. Experiments were conducted in a commercial apple orchard near Melbourne, Australia. The results show an improvement in fruit segmentation performance with the inclusion of metadata on the previously benchmarked MLP network. We extend this work with CNNs, bringing agrovision closer to the state-of-the-art in computer vision, where although metadata had negligible influence, the best pixel-wise F1-score of 0.791 was achieved. The WS algorithm produced the best apple detection and counting results, with a detection F1-score of 0.858. As a final step, image fruit counts were accumulated over multiple rows at the orchard and compared against the post-harvest fruit counts that were obtained from a grading and counting machine. The count estimates using CNN and WS resulted in the best performance for this dataset, with a squared correlation coefficient of $r^2 = 0.826$.



[Full paper at Journal of Field Robotics](#)

[Original arxiv version](#)

2016

2016 - Image Based Mango Fruit Detection, Localisation and Yield Estimation Using Multiple View Geometry, *Sensors*

Madeleine Stein, Suchet Bargoti, James Underwood

This paper presents a novel multi-sensor framework to efficiently identify, track, localise and map every piece of fruit in a commercial mango orchard. A multiple viewpoint approach is used to solve the problem of occlusion, thus avoiding the need for labour-intensive field calibration to estimate actual yield. Fruit are detected in images using a state-of-the-art faster R-CNN detector, and pair-wise correspondences are established between images using trajectory data provided by a navigation system. A novel LiDAR component automatically generates image masks for each canopy, allowing each fruit to be associated with the corresponding tree. The tracked fruit are triangulated to locate them in 3D, enabling a number of spatial statistics per tree, row or orchard block. A total of 522 trees and 71,609 mangoes were scanned on a Calypso mango orchard near Bundaberg, Queensland, Australia, with 16 trees counted by hand for validation, both on the tree and after harvest. The results show that single, dual and multi-view methods can all provide precise yield estimates, but only the proposed multi-view approach can do so without calibration, with an error rate of only 1.36% for individual trees.



[HTML Version from Sensors](#)

[PDF Version from Sensors](#)

2016 - Mapping almond orchard canopy volume, flowers, fruit and yield using lidar and vision sensors, *Computers and Electronics in Agriculture*

James Underwood, Calvin Hung, Brett Whelan, Salah Sukkarieh

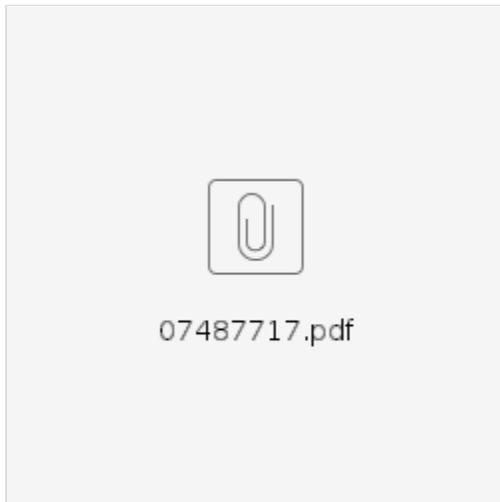
This paper present a mobile terrestrial scanning system for almond orchards, that is able to efficiently map flower and fruit distributions and to estimate and predict yield for individual trees. A mobile robotic ground vehicle scans the orchard while logging data from on-board lidar and camera sensors. 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.



2016 - Self-supervised weed detection in vegetable crops using ground based hyperspectral imaging, *IEEE International Conference on Robotics and Automation (ICRA)*

Alex Wendel, James Underwood

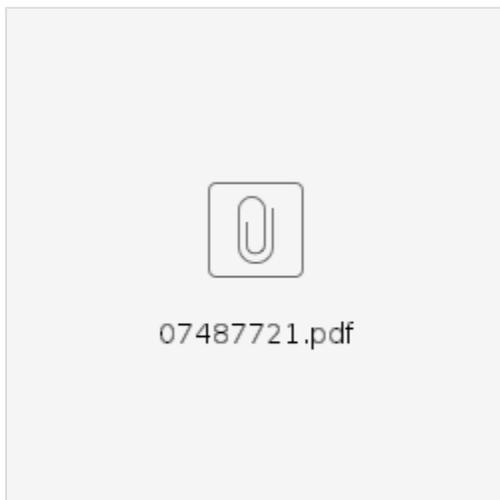
A critical step in treating or eradicating weed infestations amongst vegetable crops is the ability to accurately and reliably discriminate weeds from crops. In recent times, high spatial resolution hyperspectral imaging data from ground based platforms have shown particular promise in this application. Using spectral vegetation signatures to discriminate between crop and weed species has been demonstrated on several occasions in the literature over the past 15 years. A number of authors demonstrated successful per-pixel classification with accuracies of over 80%. However, the vast majority of the related literature uses supervised methods, where training datasets have been manually compiled. In practice, static training data can be particularly susceptible to temporal variability due to physiological or environmental change. A self-supervised training method that leverages prior knowledge about seeding patterns in vegetable fields has recently been introduced in the context of RGB imaging, allowing the classifier to continually update weed appearance models as conditions change. This paper combines and extends these methods to provide a self-supervised framework for hyperspectral crop/weed discrimination with prior knowledge of seeding patterns using an autonomous mobile ground vehicle. Experimental results in corn crop rows demonstrate the system's performance and limitations.



2016 - Image Classification with orchard metadata, *IEEE International Conference on Robotics and Automation (ICRA)*

Suchet Bargoti, James Underwood

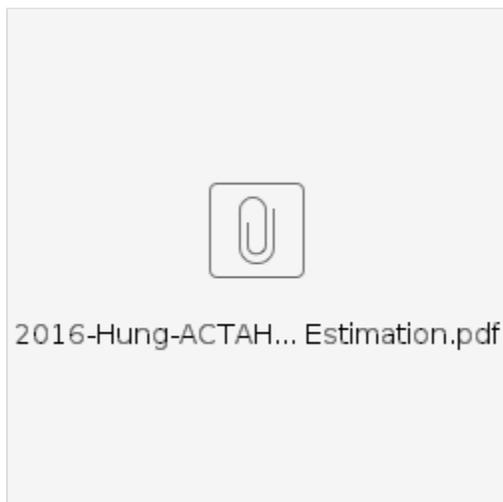
Low cost and easy to use monocular vision systems are able to capture large scale, dense data in orchards, to facilitate precision agriculture applications. Accurate image parsing is required for this purpose, however, operating in natural outdoor conditions makes this a complex task due to the undesirable intra-class variations caused by changes in illumination, pose and tree types, etc. Typically these variations are difficult to explicitly model and discriminative classifiers strive to be invariant to them. However, given the presence of structure, in both the orchard and how the data was obtained, a subset of these factors of variations can correlate with readily available metadata, including extrinsic experimental information such as the sun incidence angle, position within farm, etc. This paper presents a method to incorporate such metadata to aid scene parsing based on a multi-scale Multi-Layered Perceptron (MLP) architecture. Experimental results are shown for pixel segmentation over data collected at an apple orchard, leading to fruit detection and yield estimation. The results show a consistent improvement in segmentation accuracy with the inclusion of metadata under different network complexities, training configurations and evaluation metrics.



2016 - Autonomous Intelligent System for Fruit Yield Estimation, *Acta Horticulturae*

Calvin Hung, James Underwood, Juan Nieto, Salah Sukkarieh

The growing global population is generating increasing demands on food production. Challenges posted by climate change and reduction in agricultural workforce make it more difficult to meet the demands. Robotics and automation are expected to promote sustainability in horticulture by increasing efficiency and reducing labour costs. This paper summarises the progress towards in-field fruit counting using the autonomous intelligent systems from The University of Sydney. The system consists of ground based robots and processing software. The focus of this study is on the application of fruit yield estimation. The robots collected image data, which were processed automatically using algorithms in a software pipeline. The first stage of the pipeline uses a generic fruit segmentation algorithm, demonstrated on apple, mango, lychee and almond orchards to classify fruit pixels in the image. The second stage performs fruit detection by finding circular clusters of fruit pixels. Finally, a fruit count estimate is produced by tallying the number of circles in images spaced at 0.5 m intervals along rows of the orchard. The estimates were compared to ground-truth yield provided by the grower after harvest, using a weighing and counting machine, and a positive correlation with $R=0.81$ was found.



http://www.actahort.org/books/1130/1130_82.htm

2016 - Tree Centric Localisation in Almond Orchards, *Acta Horticulturae*

James Underwood, Gustav Jagbrant, Juan Nieto, Salah Sukkarieh

Robotics and intelligent sensing systems can provide useful information to improve yield and quality in specialty crop production. A key requirement for treecrop applications is the ability to associate sensed data to the individual trees in an orchard. A mobile ground robot with a scanning lidar (laser range sensor) is used to build a three dimensional (3D) model of an orchard and algorithms are derived to automatically detect and segment each tree. The height profile of each canopy is used to match the tree to a previously obtained database, to determine the location of the robot in the orchard, and to associate newly obtained agronomic data to the existing database. Experiments were conducted over 16 months in a 2.3 ha section of an almond orchard in Mildura, Victoria. An average tree segmentation accuracy of 99.1% was obtained, and the localisation accuracy was 98.2% for data obtained one full year apart. The method is sufficiently accurate to provide a feasible mechanism for localisation and data management in orchard environments.



http://www.actahort.org/books/1130/1130_92.htm

2016 - Trunk Localisation in Trellis Structured Orchards, *Acta Horticulturae*

Suchet Bargoti, James Underwood, Juan Nieto, Salah Sukkarieh

Information gathering and processing in horticulture helps optimise control processes and can enable more precise farm management. Robotics and automation are helping make high resolution, timely, farm wide measurements for tasks such as yield estimation, crop health and soil analysis. An efficient means of storing and processing such information is to discretise it to individual trees. To automate this process, an unmanned ground vehicle was deployed at a commercial apple orchard near Melbourne, Australia. The robot captured three dimensional (3D) laser range data and image data over orchard rows spanning an area of 1.6 ha. The area contained different apple cultivars on two types of trellis systems, a vertical I-trellis structure and a modern Güttingen V-trellis structure. Initially, tree trunk candidates (representative of the individual trees) were detected within the 3D laser range data. These candidates were then projected onto images taken at the corresponding locations to confirm their presence. By repeating this over individual orchard rows, a tree inventory was built over the farm. The experimentation was done at different times of the year and for different apple cultivars and trellis structures. A trunk localisation accuracy ranging from 89-96% was obtained during the pre-harvest season and there was near perfect performance (99% accuracy) during the flowering season, which is sufficient for building a tree inventory over a trellis structured orchard.



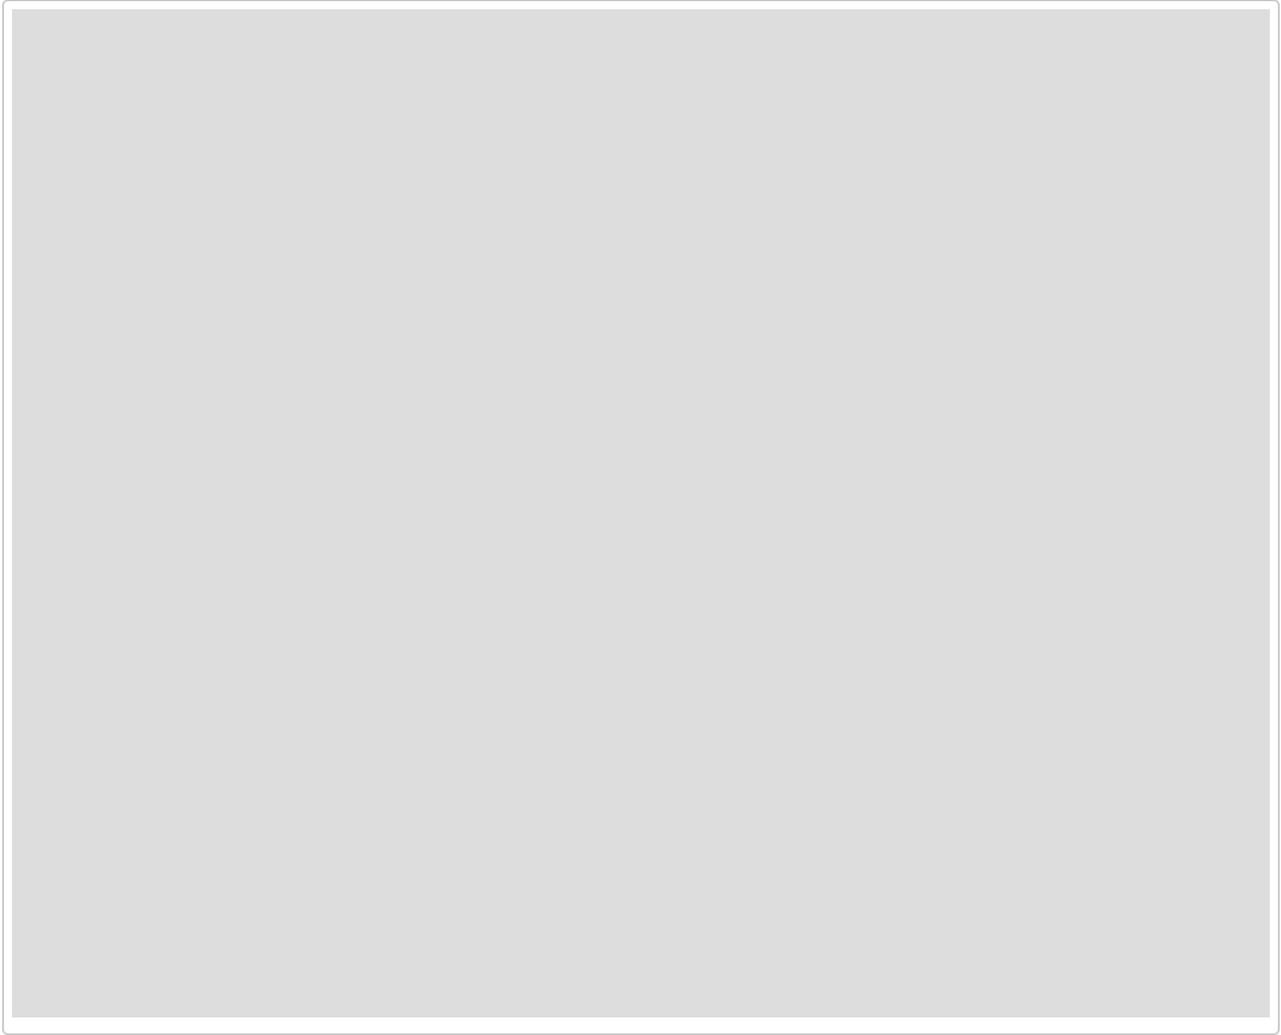
http://www.actahort.org/books/1130/1130_93.htm

2015

2015 - A Pipeline for Trunk Detection in Trellis Structured Apple Orchards, *Journal of Field Robotics*

Suchet Bargoti, James Underwood, Juan Nieto, Salah Sukkarieh

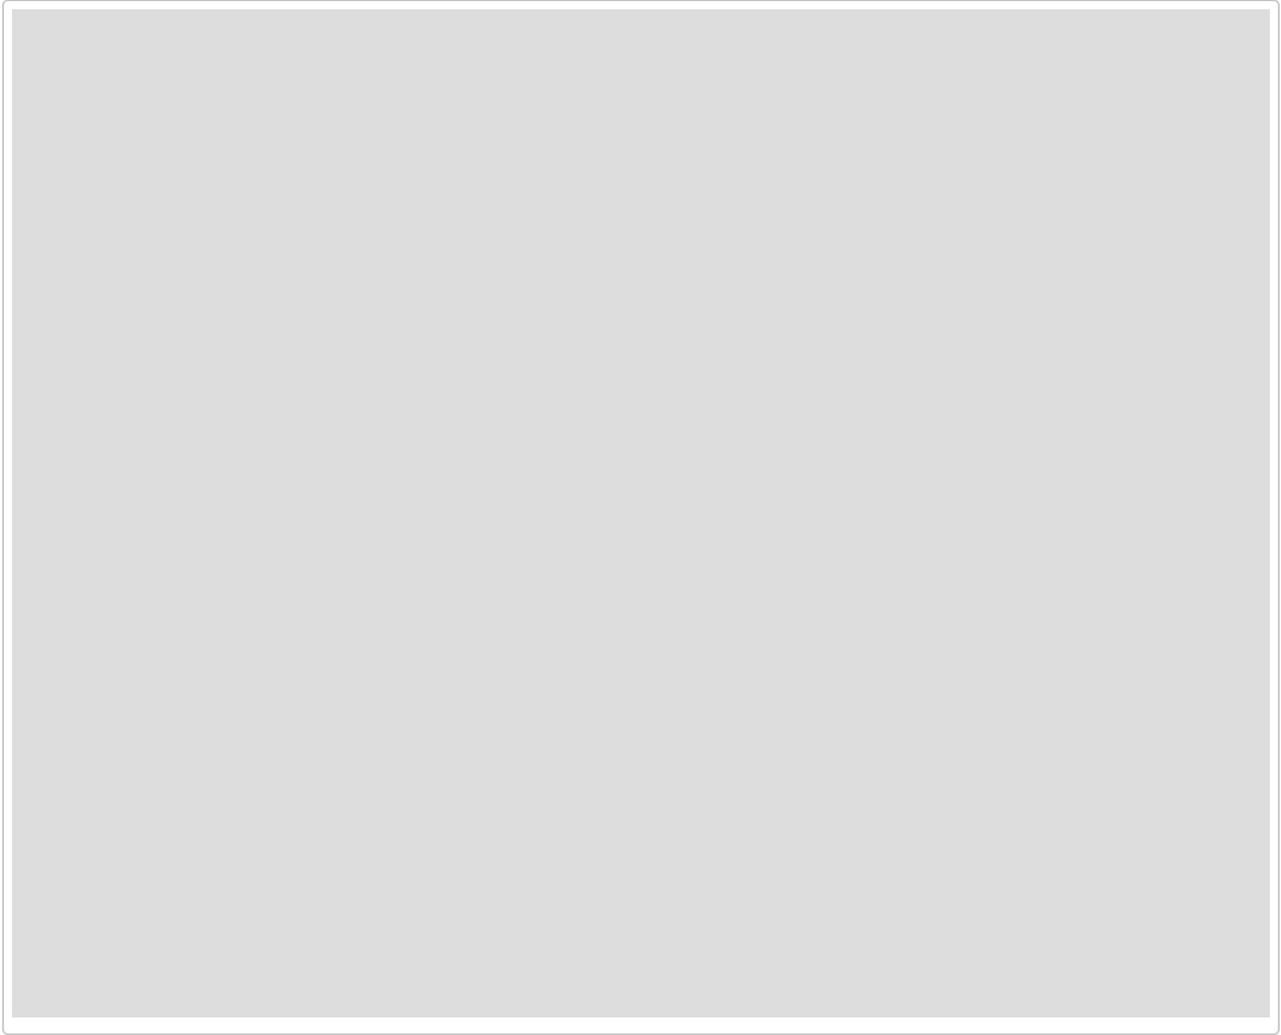
The ability of robots to meticulously cover large areas while gathering sensor data has widespread applications in precision agriculture. For autonomous operations in orchards, a suitable information management system is required, within which we can gather and process data relating to the state and performance of the crop over time, such as distinct yield count, canopy volume, and crop health. An efficient way to structure an information system is to discretize it to the individual tree, for which tree segmentation/detection is a key component. This paper presents a tree trunk detection pipeline for identifying individual trees in a trellis structured apple orchard, using ground-based lidar and image data. A coarse observation of trunk candidates is initially made using a Hough transformation on point cloud lidar data. These candidates are projected into the camera images, where pixelwise classification is used to update their likelihood of being a tree trunk. Detection is achieved by using a hidden semi-Markov model to leverage from contextual information provided by the repetitive structure of an orchard. By repeating this over individual orchard rows, we are able to build a treemap over the farm, which can be either GPS localized or represented topologically by the row and tree number. The pipeline was evaluated at a commercial apple orchard near Melbourne, Australia. Data were collected at different times of year, covering an area of 1.6 ha containing different apple varieties planted on two types of trellis systems: a vertical I-trellis structure and a Güttingen V-trellis structure. The results show good trunk detection performance for both apple varieties and trellis structures during the preharvest season (87–96% accuracy) and near perfect trunk detection performance (99% accuracy) during the flowering season.



2015 - A Pipeline for Trunk Localisation using LiDAR in Trellis Structured Orchards, *Field and Service Robotics*

Suchet Bargoti, James Underwood, Juan Nieto, Salah Sukkarieh

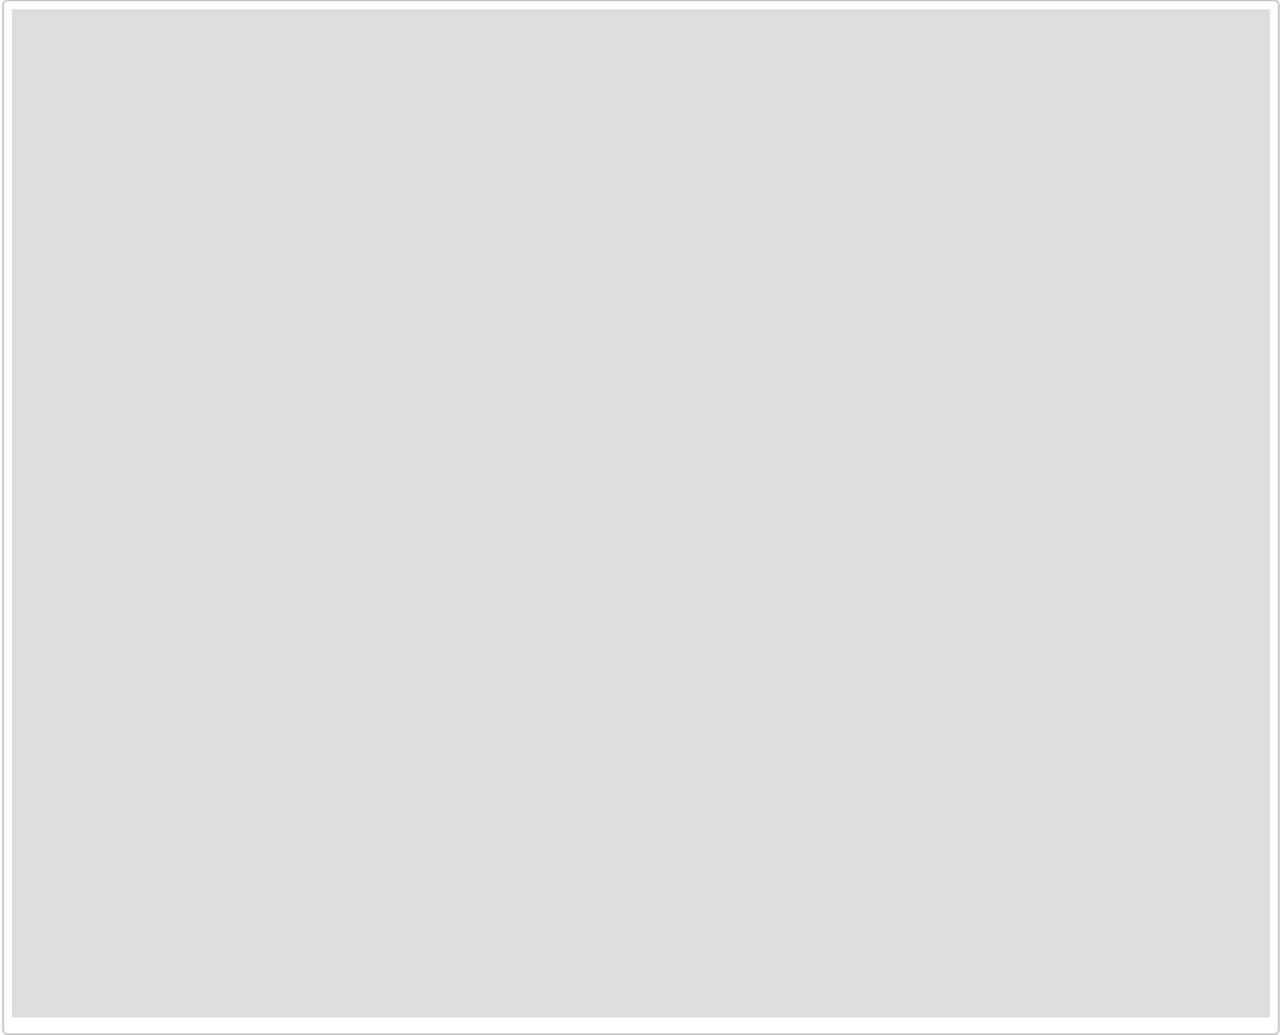
Autonomous operation and information processing in an orchard environment requires an accurate inventory of the trees. Individual trees must be identified and catalogued in order to represent their distinct measures such as yield count, crop health and canopy volume. Hand labelling individual trees is a labour-intensive and time-consuming process. This paper presents a trunk localisation pipeline for identification of individual trees in an apple orchard using ground based LiDAR data. The trunk candidates are detected using a Hough Transform, and the orchard inventory is refined using a Hidden Semi-Markov Model. Such a model leverages from contextual information provided by the structured/repetitive nature of an orchard. Operating at an apple orchard near Melbourne, Australia, which hosts a modern Guttingen V trellis structure, we were able to perform tree segmentation with 89% accuracy.



2015 - Lidar-based Tree Recognition and Platform Localization in Orchards, *Journal of Field Robotics*

James Underwood, Gustav Jagbrant, Juan Nieto, Salah Sukkarieh

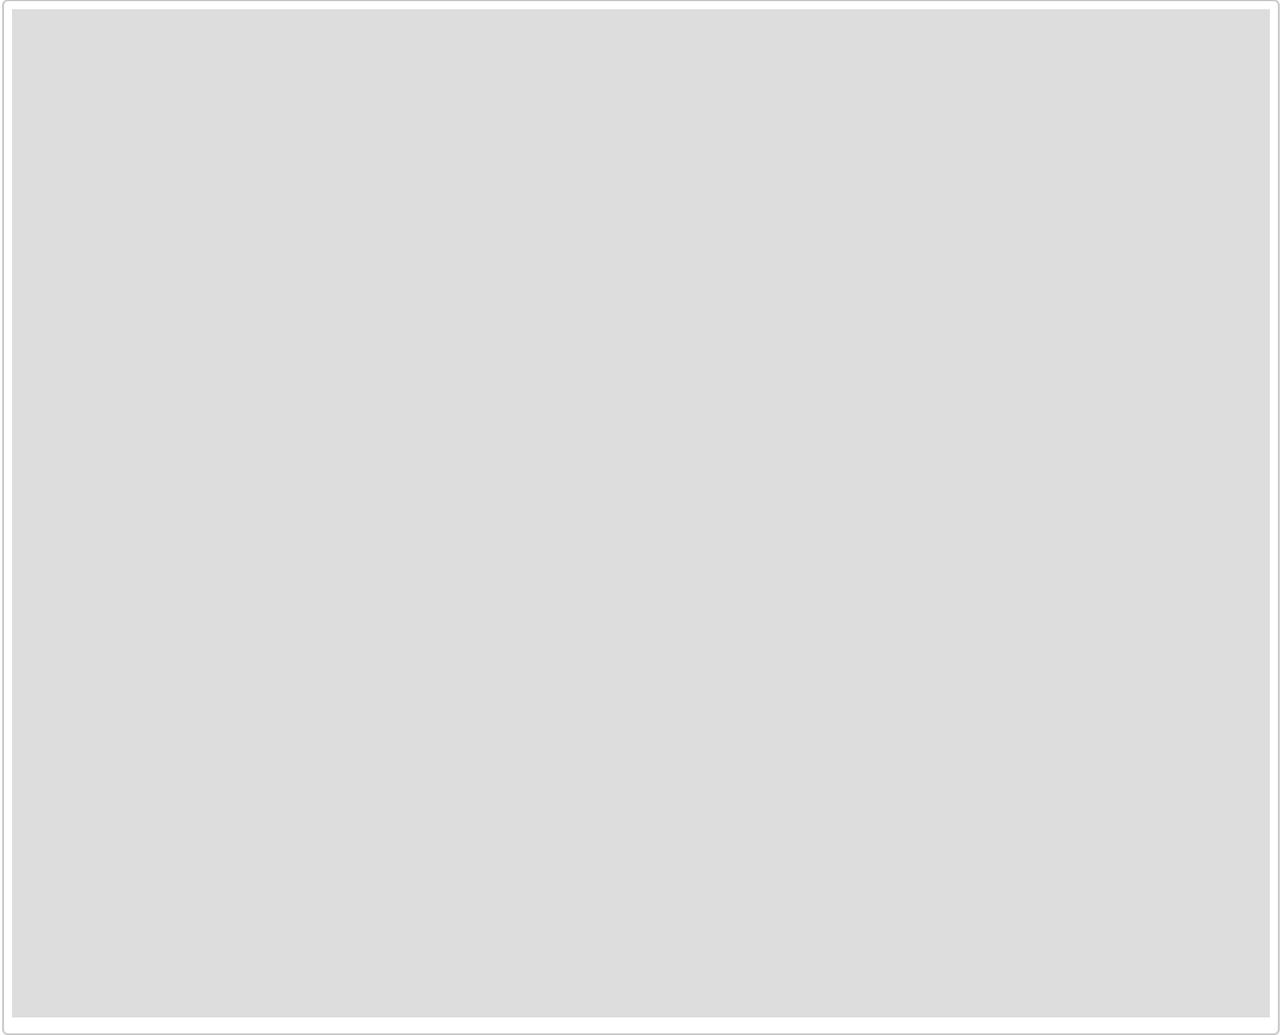
We present an approach to tree recognition and localization in orchard environments for tree-crop applications. The primary objective is to develop a pipeline for building detailed orchard maps and an algorithm to match subsequent lidar tree scans to the prior database, enabling correct data association for precision agricultural applications. Although global positioning systems (GPS) offer a simple solution, they are often unreliable in canopied environments due to satellite occlusion. The proposed method builds on the natural structure of the orchard. Lidar data are first segmented into individual trees using a hidden semi-Markov model. Then a descriptor for representing the characteristics or appearance of each tree is introduced, allowing a hidden Markov model based matching method to associate new observations with an existing map of the orchard. The tree recognition method is evaluated on a 2.3 hectare section of an almond orchard in Victoria, Australia, over a period spanning 16 months, with a combined total of 17.5 scanned hectares and 26 kilometers of robot traversal. The results show an average matching performance of 86.8% and robustness both to segmentation errors and measurement noise. Near perfect recognition and localization (98.2%) was obtained for data sets taken one full year apart, where the seasonal variation of appearance is minimal.



2015 - LiDAR Based Tree and Platform Localisation in Almond Orchards, *Field and Service Robotics*

Gustav Jagbrant, James Underwood, Juan Nieto, Salah Sukkarieh

In this paper we present an approach to tree recognition and localisation in orchard environments for tree-crop applications. The method builds on the natural structure of the orchard by first segmenting the data into individual trees using a Hidden Semi-Markov Model. Second, a descriptor for representing the characteristics of the trees is introduced, allowing a Hidden Markov Model based matching method to associate new observations with an existing map of the orchard. The localisation method is evaluated on a dataset collected in an almond orchard, showing good performance and robustness both to segmentation errors and measurement noise.



2015 - Real-time Target Detection and Steerable Spray for Vegetable Crops, *IEEE International Conference on Robotics and Automation (ICRA), Workshop on Robotics in Agriculture*

James P. Underwood, Mark Calleja, Zachary Taylor, Calvin Hung, Juan Nieto, Robert Fitch, Salah Sukkarieh

This paper presents a system for autonomously delivering a quantum of fluid to individual plants in a vegetable crop field, using a fine spray nozzle attached to a manipulator arm on a mobile robot. This can reduce input cost and environmental impact while increasing productivity for applications such as micro-nutrient provision, thinning and weeding. The Ladybird platform is introduced and a pipeline for targeted spray is presented, including image-based seedling detection, geometry and transformation, camera-arm calibration, inverse kinematics and target sequence optimisation. The system differs from existing approaches due to the continuously steerable nozzle, which can precisely reach targets in a larger workspace.

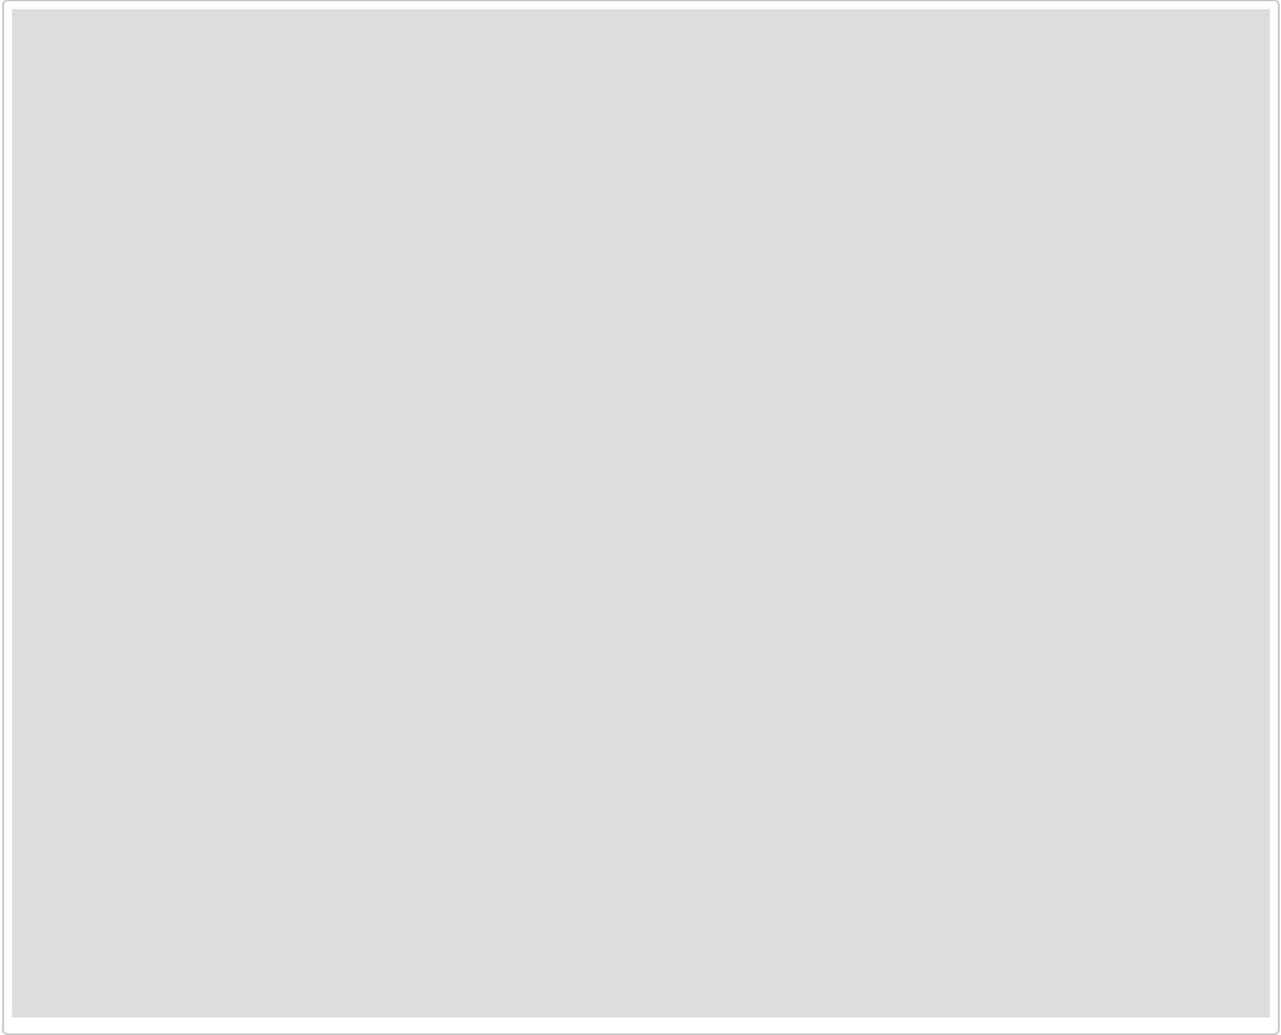


2015-Underwood...AgWs-Spray.pdf

2015 - A Feature Learning Based Approach for Automated Fruit Yield Estimation, *Field and Service Robotics*

Calvin Hung, James Underwood, Juan Nieto, Salah Sukkarieh

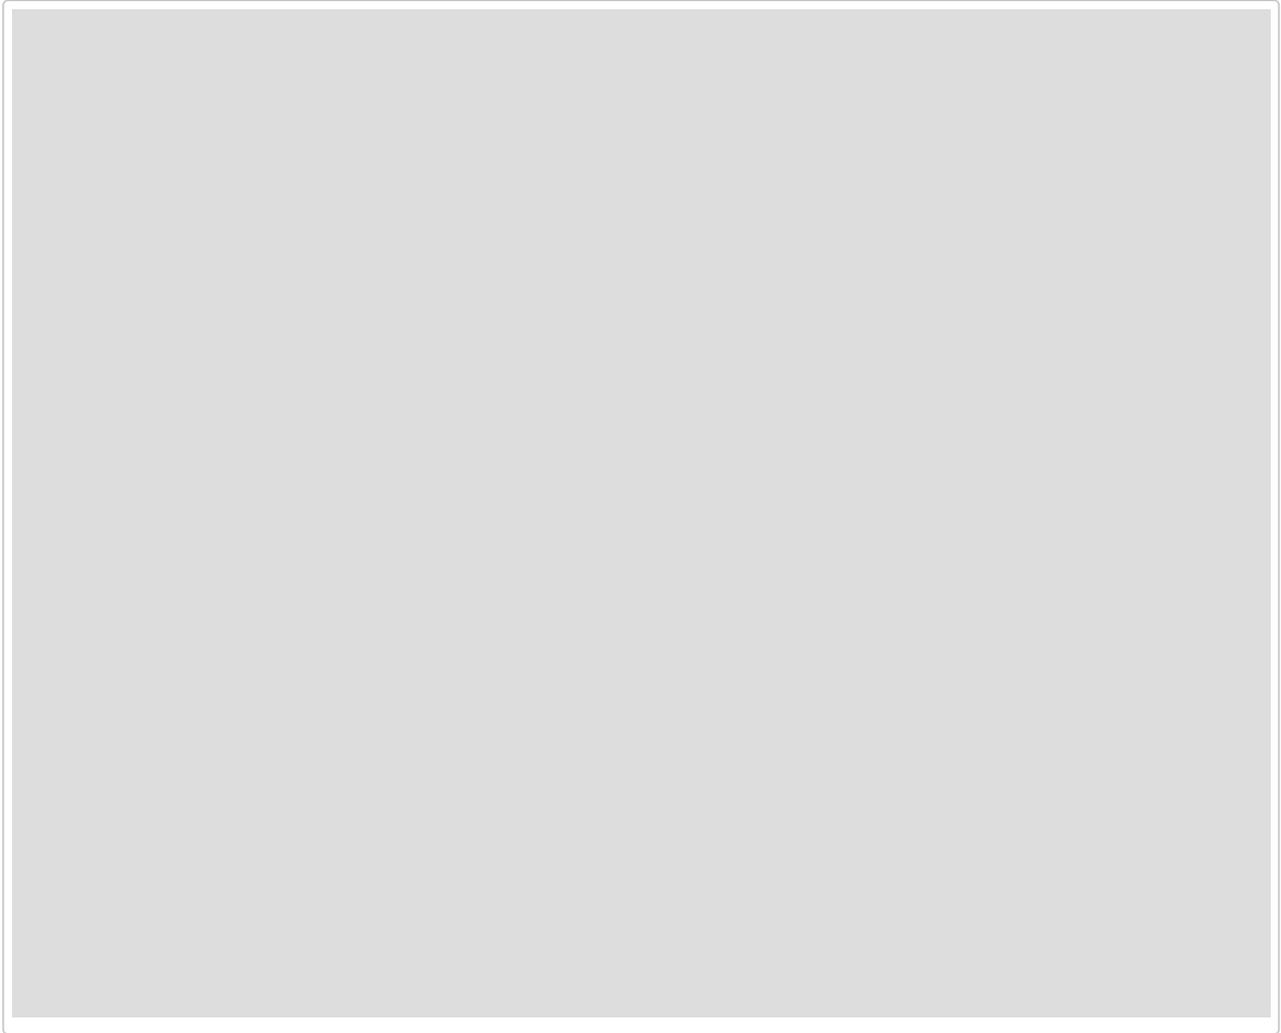
This paper demonstrates a generalised multi-scale feature learning approach to multi-class segmentation, applied to the estimation of fruit yield on tree-crops. The learning approach makes the algorithm flexible and adaptable to different classification problems, and hence applicable to a wide variety of tree-crop applications. Extensive experiments were performed on a dataset consisting of 8000 colour images collected in an apple orchard. This paper shows that the algorithm was able to segment apples with different sizes and colours in an outdoor environment with natural lighting conditions, with a single model obtained from images captured using a monocular colour camera. The segmentation results are applied to the problem of fruit counting and the results are compared against manual counting. The results show a squared correlation coefficient of $R^2 = 0.81$.



2013 - Orchard Fruit Segmentation using Multi-spectral Feature Learning, *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*

Calvin Hung, Juan Nieto, Zachary Taylor, James Underwood, Salah Sukkarieh

This paper presents a multi-class image segmentation approach to automate fruit segmentation. A feature learning algorithm combined with a conditional random field is applied to multi-spectral image data. Current classification methods used in agriculture scenarios tend to use hand crafted application-based features. In contrast, our approach uses unsupervised feature learning to automatically capture most relevant features from the data. This property makes our approach robust against variance in canopy trees and therefore has the potential to be applied to different domains. The proposed algorithm is applied to a fruit segmentation problem for a robotic agricultural surveillance mission, aiming to provide yield estimation with high accuracy and robustness against fruit variance. Experimental results with data collected in an almond farm are shown. The segmentation is performed with features extracted from multi-spectral (colour and infrared) data. We achieve a global classification accuracy of 88%.



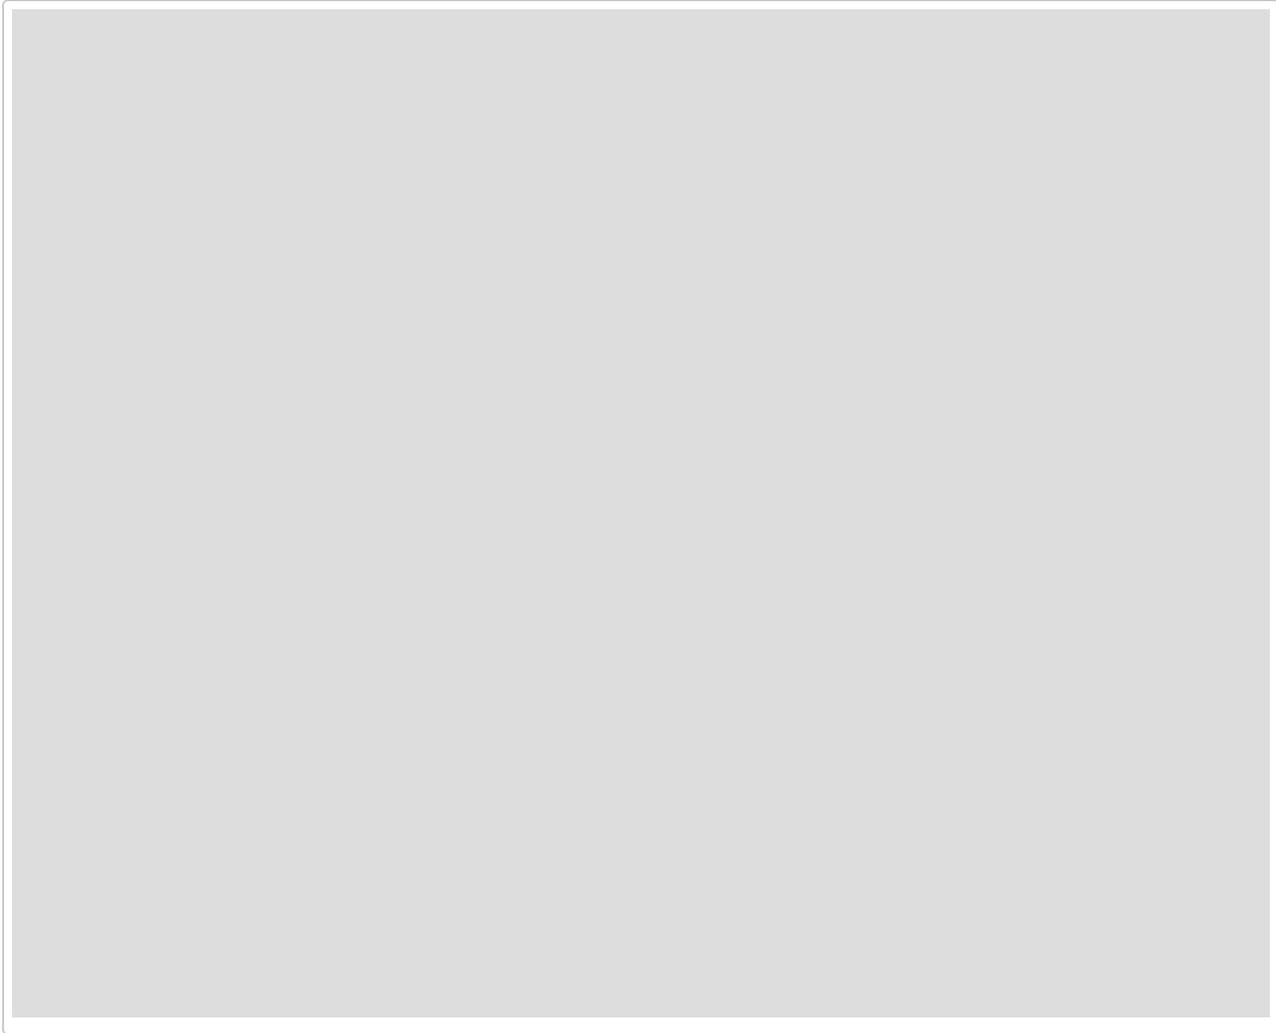
2013

2013 - A robot amongst the herd: A pilot investigation regarding the behavioural response of dairy cows, *Spatially Enabled Livestock Management Symposium*

Cameron Clark, Sergio Garcia, Kendra Kerrisk, James Underwood, Juan Nieto, Mark Calleija, Salah Sukkarieh, Greg Cronin

It is widely recognised that technology has a vast role to play in helping Australia's farmers reduce the time spent on repetitive tasks, increasing the attraction and retention of employed labour in the industry and to provide and act on data to increase farm productivity to sustainable levels. The continuous monitoring and movement of livestock between areas defined for grazing or from these areas to procedural locations (i.e. yards, dairy facility) is a repetitive task that is ideally suited to automation. A pilot study conducted at the University of Sydney's dairy research farm in Camden determined the behavioural response of dairy cows to an unmanned ground vehicle (UGV) across time (Figure 1). The ability to use the UGV in an operationally relevant way and the ability of the sensors and perception algorithms on-board the UGV to automatically identify and track the motion of the dairy cows are covered by Underwood et al. in the current proceedings. Following the morning milking, the first 20 cows to be milked were separated from the main herd at 0830 h and offered 0.5 ha of an ad-libitum kikuyu pasture allocation (50 kg DM/cow to ground level). A pre-defined figure eight route was determined for the UGV within this 0.5 ha. After allowing time for the cows to settle, the robot entered the pasture allocation at 0900 h and traversed the figure eight route at a speed of 2.5 km/h (average traverse time was 7 min). Between traverses the robot was parked outside the allocation until the process was repeated a further four times at 15 min intervals. The 0.5 ha was virtually split into four sectors for observation purposes, with four observers covering one sector each. To determine the interaction between the UGV and cows, the number of cows exiting or entering each sector when the UGV was in or out of the given sector was recorded. Data were analysed by GLMM within REML. The model was as follows: Cows out = Fixed (Robot (presence/absence) * Traverse number) + Random (Cow).

There was a significant effect of Robot ($P = 0.02$) and Traverse ($P < 0.01$) on the number of cows exiting a sector, however, there was no interaction between these fixed effects. Twice as many cows exited a sector when the robot was present (8%) as compared with absent (4%). More cows exited a sector in traverse 1 (14%) as compared with all other traverses (mean = 4%). The greater number of cows exiting a sector in the first traverse was likely associated with an initial period of increased cow movement as cows foraged. These results also indicate that dairy cows habituate to the moving UGV quickly. Future work will aim to fully automate the process of herding and integrate this process with other data requirements such as ground cover and soil moisture levels. The authors would like to acknowledge the staff of both Corstorphine dairy and the Australian Centre of Field Robotics for their excellent advice and support.

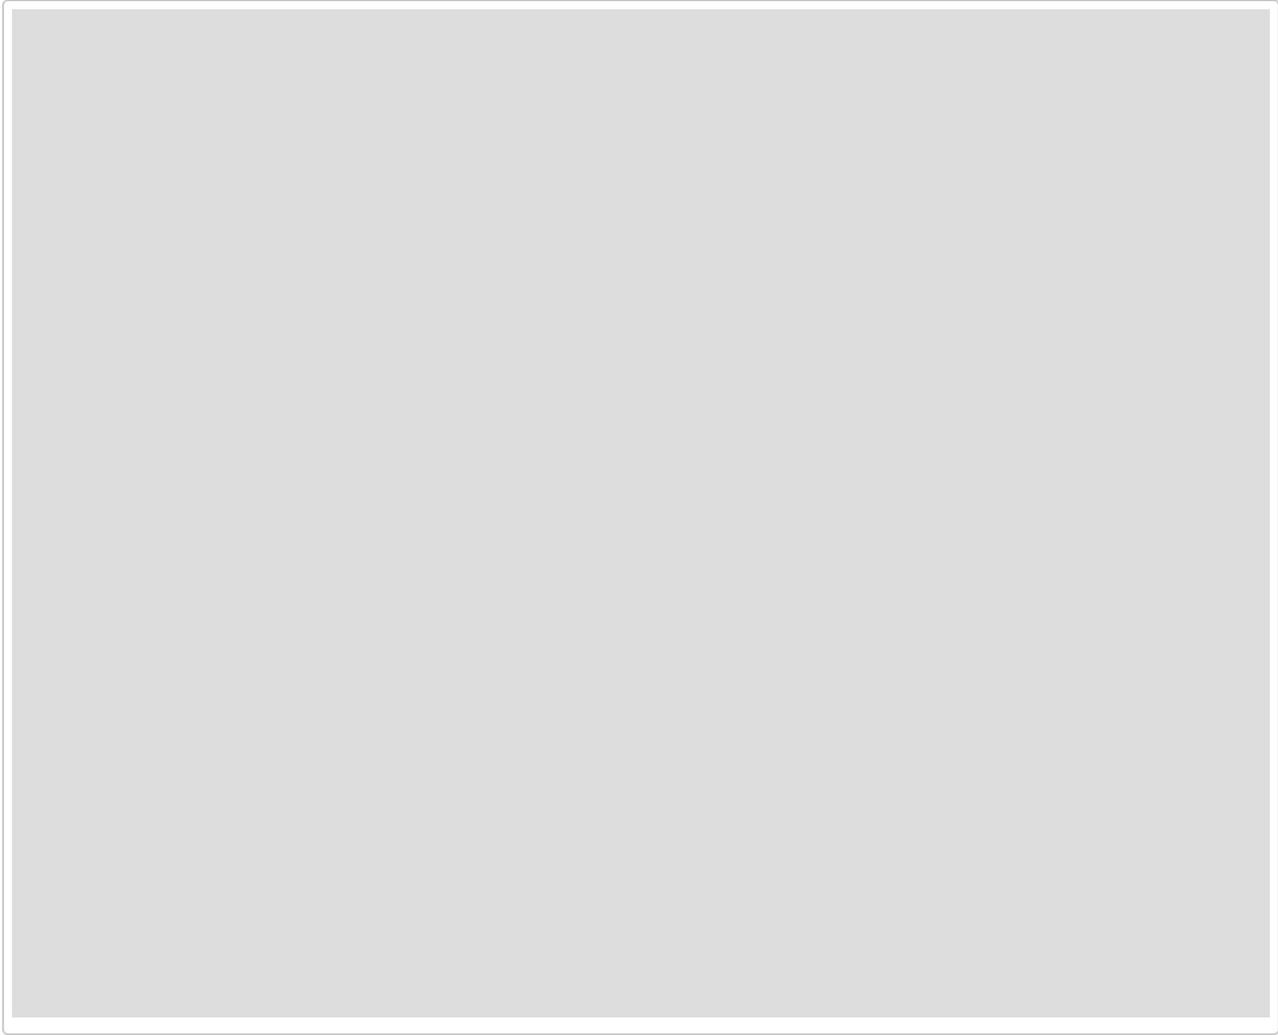


2013 - A robot amongst the herd: Remote detection and tracking of cows, *Spatially Enabled Livestock Management Symposium*

James Underwood, Mark Calleija, Juan Nieto, Salah Sukkarieh, Cameron Clark, Sergio Garcia, Kendra Kerrisk, Greg Cronin

Recent advances in sensing, automation and information technology in Australia and globally have resulted in commercially successful field-robotic applications and the timing is appropriate to consider what additional roles these systems may serve in the dairy industry during the next decade, to decrease the cost of milk production. A first test of existing unmanned ground vehicle (UGV) technology was conducted to assess current capability in this domain. Three aspects were studied: 1) the response of dairy cows to the presence of a UGV (see Clark et al. in current proceedings), 2) the ability to use the UGV in an operationally relevant way (remote controlled herding) and 3) the ability of the on-board sensors and algorithms to automatically detect and track dairy cows. Success in all three is a pre-requisite for UGVs to have a role in automated herding of dairy cows. All three aspects were addressed using 3D LiDAR data. Raw data were displayed remotely for operator situational awareness and control from a distance of up to 200 m. The data were georeferenced and processed through a perception pipeline, to estimate global cow trajectories (with no instrumentation on the cows), including instantaneous position, velocity and a record of the paths they followed (Figure 1). This enables quantitative analysis and modelling of their response to the UGV, and provides real-time feedback required for future development of autonomous herding. A pre-set route was driven five times, amongst 20 cows in a 0.5 ha paddock. 3D LiDAR data showed mean cow velocities away from the UGV of [0.06, 0.04, 0.02, 0.01, 0.01] m/s, indicating that dairy cows habituate quickly to UGV movement as per human observations (Clark et al. in current proceedings). The cows were then herded three times from the same 0.5 ha paddock by remotely controlling the UGV. For each experiment, every cow was successfully herded through the gate without human intervention. Summary statistics are shown in Table 1. The mean velocity of the cows was at most 0.1 m/s, which was considered a 'calm herding' from human observation, with potential animal welfare benefits such as reduced lameness. In conclusion, this experiment showed remote herding to be possible and that 3D LiDAR sensing and existing perception algorithms are able to detect and track cows. Further work is required to build on these findings and automate the process of herding.

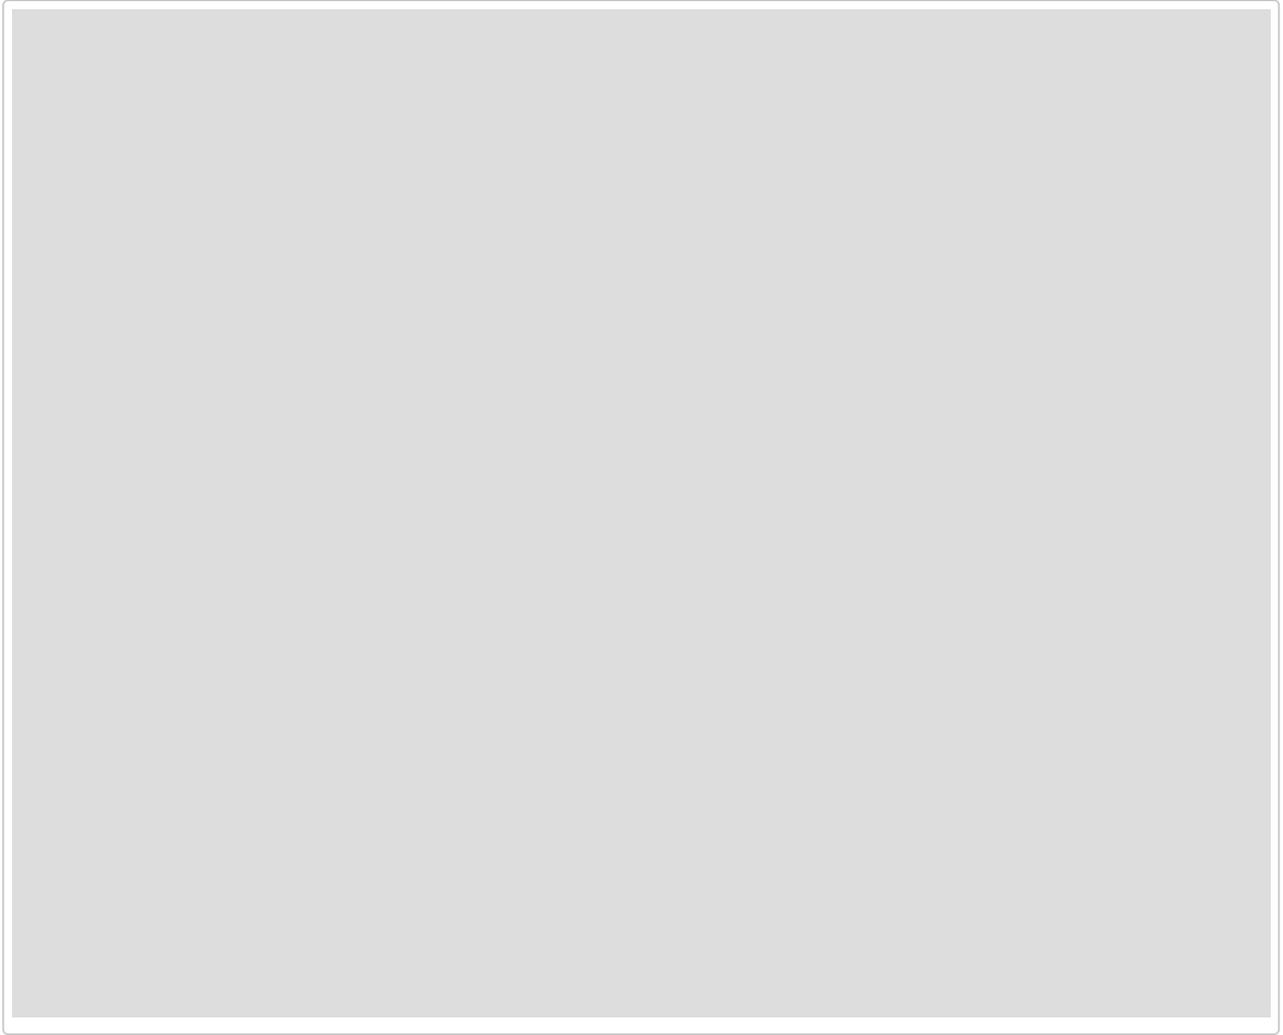
The authors would like to acknowledge the staff of both Corstorphine dairy and the Dairy Science group for their excellent advice and support.



2013 - Robotic Aircraft and Intelligent Surveillance Systems for Weed Detection, *Plant Protection Quarterly*

Calvin Hung, Salah Sukkarieh

This paper presents a summary of the autonomous weed detection R&D program at the Australian Centre for Field Robotics (ACFR) over the past seven years. The ACFR has used various aerial robots on various detection and mapping projects, targeting weeds including prickly acacia (*Acacia nilotica*), parkinsonia (*Parkinsonia aculeata*), mesquite (*Prosopis pallida*), wheel cacti (*Opuntia robusta*) and alligator weed (*Alternanthera philoxeroides*), and extending it to include pests such as detecting Red Imported Fire Ants (RIFA) mounds, in various parts of Australia. The algorithm research at ACFR leads to various intelligent detection and mapping software systems for accurate terrain mapping, vegetation segmentation and detection of different invasive species.



2012

2012 - Multi-class predictive template for tree crown detection, *ISPRS Journal of Photogrammetry and Remote Sensing*

Calvin Hung, Mitch Bryson, Salah Sukkarieh

This paper presents a novel approach for automatic segmentation and object detection of tree crowns in airborne images captured from a low-flying Unmanned Aerial Vehicle (UAV) in ecology monitoring applications. Cost effective monitoring in these applications necessitates the use of vision-band-only imaging on the UAV platform; the reduction in spectral resolution (compared to multi- or hyper-spectral imaging) is balanced by the high spatial resolution available (20 cm/pixel) from the low-flying UAV, when compared to existing satellite or manned-aerial survey data. Our approach to object detection thus uses both geometry and appearance information (through the use of tree shape and shadow information) in addition to spectral information to help accurately distinguish tree crowns within our application. A predictive geometric template for tree detection is constructed using on-board UAV navigation data, sun lighting information and information about the geometry of the target crown. A two-stage detection algorithm is then used to segment tree crowns based on spectral (colour) information convolved with information from the predictive template. Results of our approach are presented using airborne image data collected from a fixed-wing UAV during a weed monitoring and mapping mission over farmland in West Queensland, Australia.

