Real-time target detection and steerable spray for vegetable crops*

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

Abstract— **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, wherever blanket spraying can be replaced with targeted spray, in applications including thinning, micro-nutrient provision and weeding. The Ladybird platform is introduced and a pipeline from image-based seedling detection, geometry and transformation including camera-arm calibration, inverse-kinematics and target sequence optimisation.**

I. INTRODUCTION

Robotics for vegetable farming has many potential benefits, from digitising the farm by acquiring timely, accurate and ubiquitous information about every plant, to closing feedback loops to micro-manage their specific requirements [1]. A promising area is the targeted application of fluids to the crop, where blanket spraying has been the only economically viable alternative. Applications include thinning, site-specific micro-nutrition and weed control [1]. Targeted spray for weeding is an important area for robotics, due to the cost of herbicide, the development of resistant weed strains and environmental impact [2], as it has been shown that microsprayers may reduce herbicide use to as little as 0.01% compared to blanket spraying [3].

In this paper, we present a software pipeline for automated micro-spray delivery using the Ladybird, which is a flexible robotic research platform for the vegetable industry, designed and built at The University of Sydney.

II. BACKGROUND

The desire for selective spraying pre-dates the availability of technology to implement it. A prescient paper by Thompson et al. in 1991 [4] outlines the motivation for selective weed control in terms of non-uniformity in both spatial density and competitiveness; weeds are generally not evenly distributed and certain species compete more strongly with crop plants (thus decreasing yield) than others. Although not technically feasible at the time, Thompson identified the relevant system components: on-board weed detection and localisation and controllable spray nozzles.

A good survey of modern methods to 2008 is by Slaughter et al. [5], including a review of weed detection algorithms using visible range and multi-spectral imagery. Our system uses monocular vision for detection and stereo vision to locate plants relative to the robot in 3D.

Much of the work in selective spraying focuses on micro-drift suppression through physical fluid properties, the effect of surfactants on micro-drift and nozzle design [2]. Nozzles are typically platform-fixed, in an array to increase the sprayable area [6]. Such systems have been proposed for seed-line weeds [7] and thermal (hot oil) control in organic crops [8]. Commercial systems that integrate detection with spray control are available [9].

Nozzle arrays are generally not steerable and are typically designed for area coverage. Precise targeting is possible using tiny nozzles, as demonstrated using an inkjet printer head [3]. The limitation of tiny fixed nozzles is that their workspace is also small. There is initial work in targeting individual plants with a manipulator arm, but at a slow rate (10-13 s per weed) [10]. We instead propose a fast, steerable nozzle for precision spraying. The use of a simple steering controller avoids the computational complexity to plan for multi-link manipulators [11, 12]. Steerability dramatically increases the sprayable area of a single nozzle. Although in this work we demonstrate the simple case where we target all plants, our system is designed for extension to selectively target individual species.

Chemical thinning is another common application of targeted spray systems that has a long history [13] and is now realised commercially [14]. We focus on less structured scenarios where the spatial density of targets is non-uniform.

III. LADYBIRD

The Ladybird robot (Fig. 1) was designed and built in 2014 at the Australian Centre for Field Robotics (ACFR) at The University of Sydney as a flexible tool to support research and development of robotics applications for the vegetable industry. The system was designed for *flexibility* and *modularity*. For efficient research, we require one system to cater for a wide variety of applications, however, modularity enables decomposition, to rapidly design lower cost bespoke systems for individual commercial applications. The following is a list of key features:

- **Drive mechanism**: features four identical, modular, electric drive units. Each has two mechanically decoupled axes, to rotate the wheel orientation and to drive the wheel. Decoupling minimises soil shear and power consumption.
- **Power system**: a bank of Lithium Iron Phosphate (LiFePO4) batteries store and provide power. The system is charged from the mains and topped up by solar power in the field. Solar provides enough energy to

This work is supported by the Australian Centre for Field Robotics at the University of Sydney and Horticulture Australia Limited through project AH11009 Autonomous Perception Systems for Horticulture Tree Crops.

Authors are with ACFR, School of Aerospace, Mechanical and Mechatronic Engineering, The University of Sydney NSW 2006, Australia j.underwood@acfr.usyd.edu.au

run the system at low speed on flat terrain, with a net power drain at top speed over rough terrain. Longevity is determined by the weather and the driving duty cycle and mains charging has not been required for typical week-long field-trials.

- **Computation**: is provided by a single Nuvo-3005E-I7QC computer, with an Intel Core i7-3610QE CPU and 16GB RAM, running Ubuntu Linux. We use a light-weight, open-source software architecture [15] based on simple data streams, which facilitates multi-language support with C++ for core elements, Python at the higher level and Bash scripting at the top.
- **Sensing**: a multi-modal sensor suite provides information for autonomy and crop perception (see Fig. 1(b).) Forward and rear facing lidar and a spherical camera support obstacle avoidance and crop row detection, while RTK GPS/INS allows for map-based farm traversal. Crop sensing is provided by hyperspectral line-scanning, stereo vision (with strobe) and thermal infrared cameras.
- **Manipulation**: is provided by a six degree-of-freedom Universal Robotics UR5 arm. Their proprietary *urscript* language is used to interface to the arm. The arm can be configured for direct manipulation (such as tilling) and Ladybird also carries a spray unit with a nozzle end-effector, for targeted spray. The configuration is shown in Fig. 1(c).

IV. REAL-TIME DETECTION AND SPRAY

The real-time target detection and spray application is performed for a batch of targets while the Ladybird is stationary. The stereo camera is used for target detection and locating, then the coordinates are transformed to the arm frame and the inverse kinematics are solved to point the spray nozzle at each target. The sequence is optimised and the targets are sprayed. Ladybird moves to the next patch[1] and the sequence repeats.

A. Detection

Target detection is performed using monocular colour image processing with the right camera of the stereo pair. Although stereo 3D geometry is used subsequently for aiming, this separation allows any pure image based detection or classification algorithm to be adopted. For simplicity, we apply vegetation detection with no crop classification, but image based weed classifiers can be used interchangeably in the pipeline.

We adapt the "Excess Green minus Excess Red" (E_{g-r}) method as it outperforms other competing algorithms for vegetation detection in visual images [16]. The original formulation is expressed as:

$$\begin{aligned} E_g &= 2G - R - B \\ E_r &= 1.4R - G \\ E_{g-r} &= E_g - E_r \end{aligned} \quad (1)$$

where vegetation is detected with $E_{g-r} > 0$. An advantages is claimed to be the "fixed, built-in zero threshold" that does not require tuning [16], however, the constants (2 and 1.4) are tunable and may be optimised for a given application. We further subtract excess blue ($E_b = 1.4B - G$) as we found empirically that blue plastic features on our robot (sometimes visible as in Fig. 1(d)) could cause misclassification. The result is "Excess Green minus Excess Red minus Excess Blue $E_{g-r} - b$, which is better expressed as *green dominance*:

$$G > k(R + B) \quad (2)$$

with tunable k. The zero crossing threshold $E_{g-r-b} > 0$ corresponds to $k = 0.6$, which is equivalent to the original Excess Green formulation [17]. A canonical condition exists for this equation (and also E_{g-r}) when pixel noise dominates due to under exposure (shadows), thus we also specify a minimum intensity for classification:

$$(R + B) > t \quad (3)$$

When both inequalities in Equations 2 and 3 are met, a pixel is considered to be vegetation. Vegetation pixels are then filtered with erosion (by e pixels) to remove small noisy clusters and dilation (by d pixels) to combine multiple leaves to represent individual plants. The centroids of these regions provide the precise locations of the targets. The stages of the detection pipeline are illustrated in Fig. 2(a)-(d).

The pipeline is parametrised by $\Theta = \{k, t, e, d\}$, which were learnt from the data. A total of 676 seedlings were hand-labelled and divided in half to form an optimisation and test batch. Exhaustive search over Θ maximised the f1-score against the labels with $\Theta = \{0.65, 20, 2, 10\}$ for $f1 = 0.900$.[2] Imperfections were typically caused by adjacent plants being falsely combined into one and by human error in labelling, both of which are apparent in Fig. 2(d). Finally, the depths of all centroids are calculated from stereo, using semi-global block matching [18] to provide the Cartesian location in the camera sensor frame.

B. Reference frames and calibration

Before a point detected by the camera can be targeted by the robotic arm, the transformation from the camera to the arm's base must be found. A calibration process for a similar system [19] was modified for this purpose.

A chess board was attached to the arm's end-effector and moved through 40 different poses (see Fig. 1(d)). At each location, the arm pose is recorded and the corresponding

[1]For the experiments in this paper, Ladybird was manually driven to each test patch and the system was engaged by a button on the remote control.

[2]The optimal value for k was found to be similar to the $E_{g-r-b} > 0$ condition, but this is not bound to be true for different sensors, fields, crops and weather conditions.

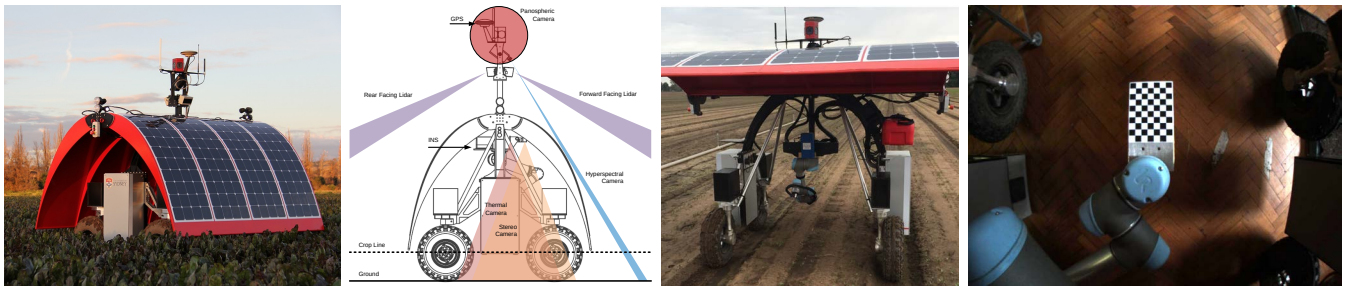

(a) (b) (c) (d)

Fig. 1. Ladybird: (a) scanning a field, (b) sensor configuration, (c) spot spraying in the field, (d) one of 40 images for camera-arm calibration

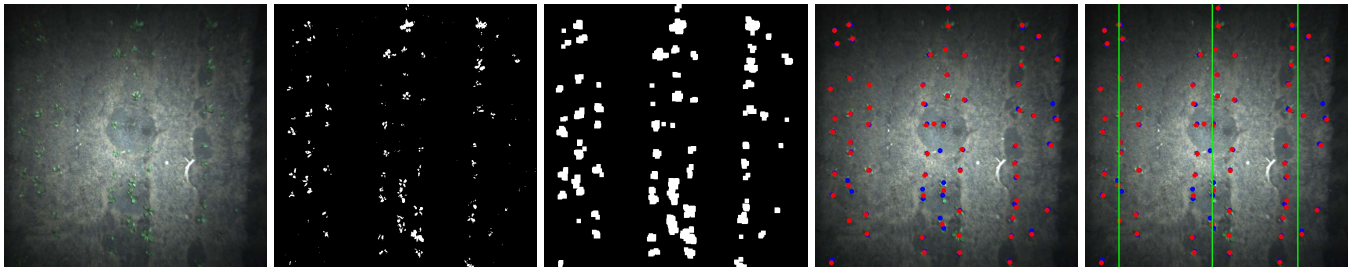

(a) raw image (b) green dominance (c) erosion and dilation (d) image and centroids (e) row detection

Fig. 2. The vegetation detection pipeline from left to right. The coloured dots in (d) show detected centroids, red=algorithm, blue=hand-labels.

corners of the chessboard are identified within the images. The two sources of information are related by the following transformation:

$$p_{cam} = K\ T_{base}^{cam}\ T_{end}^{base}\ T_{chess}^{end}\ p_{chess} \tag{4}$$

where

p_{chess} is a corner point on the chess board

T_a^b are the transformation matrices from frame a to b, linking the chess board, arm end-effector, arm base and camera.

K is the camera projection matrix

p_{cam} is the pixel position of the corner point in the camera's image

This equation contains two unknown transformations T_{chess}^{end} and T_{base}^{cam}, which we find using interior point optimization, minimizing the difference between the above projection and the location of the corners detected in the camera image. To increase the robustness of our method we use the trimmed mean, rejecting the worst 20% of points as outliers. The calibration variance is estimated by bootstrapping the corner point locations and rerunning the method 100 times.

C. Targeting

The spray nozzle is mounted to point along the positive y-axis of the arm's end-effector frame. With sufficient pressure, the flow of liquid travels in an approximately straight line, therefore the aim is to align the y-axis to point directly at the target. The arm has six joints, all of which are used to move the spray nozzle to a configurable "home" position, from which only two of the joints are required to point at any target within the workspace.

The following equation describes the transformation of a point in the arm base frame to the end-effector frame, including the full chain of six joints [20], where each joint transformation is parametrised by a single movable joint angle $T_{n+1}^n = f(\theta_n)$:

$$\begin{aligned} p_{end} &= T_{base}^{end}\ p_{base} \\ &= T_2^{end}\ T_3^2\ T_4^3\ T_5^4\ T_6^5\ T_{base}^6\ p_{base} \end{aligned} \tag{5}$$

The y-axis alignment condition is expressed by the following squared error objective cost function:

$$cost = p_{end,x}^2 + p_{end,z}^2 \tag{6}$$

Targets are identified by the camera, transformed to the arm base frame using Equation 4 and to the end-effector frame with Equation 5, resulting in an alignment cost from Equation 6 that is a function of the two joint angles we wish to move. Combined, they give an optimisation problem:

$$\underset{\theta_1,\theta_3}{\operatorname{argmin}}\ cost \mid p_{end,y} > 0 \tag{7}$$

The additional constraint selects the solution with the nozzle pointing towards the target rather than away. The constrained optimisation problem is solved using the Nelder Mead simplex algorithm [21, 22], adopting the barrier method with a step function as the barrier, though the initial conditions (nozzle pointing straight down) are sufficiently close to not rely upon the barrier.

TABLE I

SPRAY SEQUENCE TIMING

	spray	total(s)	per image (s)	per hectare (h)
back-to-front	no	323.8	32.4	44.7
greedy	no	221.2	22.1	30.5
row-raster	no	224.8	22.5	31.0
back-to-front	yes	453.3	45.3	62.6
greedy	yes	351.9	35.2	48.5
row-raster	yes	348.9	34.9	48.2

D. Sequence Optimisation

A single image produces a list of targets (roughly 68 seedlings per image in our field tests). The joint angles are calculated as above for all targets prior to spraying any, allowing the sequence to be optimised for the shortest total time. We tested three methods: *back-to-front*, where targets are ordered from the rear to the front of the vehicle, *greedy*, where the next nearest target (in joint space) is scheduled repeatedly and *row-raster*, where the seed rows are detected using k-means (in the left/right image axis, independently for each image) and targets are sprayed in a raster scan, up and back each row (see Fig. 2(e)). The time taken to spray the 676 seedlings in the labelled set (comprising 10 images) is reported in Table I, both with and without the spray, because the arm pauses at each target proportional to the spary volume, which is application dependent). The reported times include all pipeline computation, which is negligible compared to the motion and spray. The times do not include the motion of the Ladybird between patches, which took approximately two to four seconds between each scan.

The results show that the greedy sort and raster methods are substantially faster than back-to-front scheduling, taking only 68% to 77% of the time, resulting in an average of 35 seconds to spray the 68 seedlings in each image. There was no significant difference between the greedy and raster methods. Further optimisation may be possible using state-of-the-art solutions to the travelling salesman problem, but it is likely that the additional computation time will exceed the savings.

V. CONCLUSIONS

This paper described a system for autonomously delivering a quantum of fluid to individual seedlings in a vegetable crop field. The pipeline performs detection, transformation (via calibration), inverse kinematics and optimisation for targeted spray delivery. Any image-based detection methods can be used, to allow specific applications such as targeted weeding or thinning. The time to spray every seedling in a hectare was estimated to be upwards of 30.5 hours, however, typical weeding and thinning applications require only a small fraction of the plants to be sprayed, reducing the total time proportionally. Furthermore, the time can be decreased by using multiple spray nozzles and multiple platforms for parallelisation and the optimal configuration is likely to depend on the specifics of the application and the crop.

ACKNOWLEDGMENT

Thanks to Ed Fagan and Mulyan Farms for their ongoing support. Thanks to Andrey Sokolov, Vinny Do and Vsevolod Vlaskine for the software development and to Akash Arora for help with the arm. Further information and videos available at: `http://sydney.edu.au/acfr/agriculture`.

REFERENCES

[1] H. J. Heege, "Precision in crop farming," 2013.

[2] S. L. Young and D. K. Giles, "Targeted and microdose chemical applications," in *Automation: The Future of Weed Control in Cropping Systems*, S. L. Young and F. J. Pierce, Eds. Netherlands: Springer, 2014, pp. 139–147.

[3] H. S. Midtiby, S. K. Mathiassen, K. J. Andersson, and R. N. Jørgensen, "Performance evaluation of a crop/weed discriminating microsprayer," *Comput. Electron. Agr.*, vol. 77, no. 1, pp. 35 – 40, 2011.

[4] J. F. Thompson, J. V. Stafford, and P. C. H. Miller, "Potential for automatic weed detection and selective herbicide application," *Crop. Prot.*, vol. 10, no. 4, pp. 254 – 259, 1991.

[5] D. C. Slaughter, D. K. Giles, and D. Downey, "Autonomous robotic weed control systems: A review," *Comput. Electron. Agr.*, vol. 61, no. 1, pp. 63 – 78, 2008.

[6] Y. Song, H. Sun, M. Li, and Q. Zhang, "Technology application of smart spray in agriculture: A review," *Intell. Autom. Soft Co.*, 2015. [Online]. Available: http://dx.doi.org/10.1080/10798587.2015.1015781

[7] D. K. Giles, D. C. Slaughter, D. Downey, J. C. Brevis-Acuna, and W. T. Lanini, "Application design for machine vision guided selective spraying of weeds in high value crops," *Asp. Appl. Biol. Pestic. Appl.*, vol. 71, pp. 75 – 82, 2004.

[8] Y. Zhang, E. S. Staab, D. C. Slaughter, D. K. Giles, and D. Downey, "Automated weed control in organic row crops using hyperspectral species identification and thermal micro-dosing," *Crop. Prot.*, vol. 41, pp. 96 – 105, 2012.

[9] WeedSeeker Spot Spray System, "http://www.trimble.com/Agriculture/weedseeker.aspx WebSite, March 2015.

[10] H. Y. Jeon and L. F. Tian, "Direct application end effector for a precise weed control robot," *Biosyst. Eng.*, vol. 104, no. 4, pp. 458 – 464, 2009.

[11] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," in *Proc. of IEEE ICRA*, 2009, pp. 489–494.

[12] J. J. H. Lee, K. Frey, R. Fitch, and S. Sukkarieh, "Fast path planning for precision weeding," in *Proc. of ARAA ACRA*, 2014.

[13] S. W. R. Cox and K. A. McLean, "Electro-chemical thinning of sugar beet," *J. Agric. Eng. Res.*, vol. 14, no. 4, pp. 332 – 343, 1969.

[14] Blue River Technology, "http://www.bluerivert.com/," WebSite, March 2015.

[15] A. C. for Field Robotics (ACFR), "Comma and Snark: generic C++ libraries and utilities for robotics," [Accessed 2015-03-18]. [Online]. Available: https://github.com/acfr/

[16] G. E. Meyer and J. C. Neto, "Verification of color vegetation indices for automated crop imaging applications," *Computers and Electronics in Agriculture*, vol. 63, no. 2, pp. 282–293, Oct. 2008.

[17] D. Woebbecke, G. Meyer, K. Von Bargen, and D. Mortensen, "Color indices for weed identification under various soil, residue, and lighting conditions," *Transactions of the ASAE*, vol. 38, no. 1, pp. 259–269, 1995.

[18] G. Bradski, "OpenCV: Camera calibration and 3d reconstruction," *Dr. Dobb's Journal of Software Tools*, 2015, [Accessed 2015-03-23]. [Online]. Available: http://docs.opencv.org/modules/calib3d/doc/camera_calibration_and_3d_reconstruction.html

[19] O. Nielsen and E. Pedersen, "Calibration of a robotic arm," 2010, [Accessed 2015-03-18]. [Online]. Available: http://robolabwiki.sdu.dk/mediawiki/index.php/Calibration_of_a_robotic_arm

[20] K. P. Hawkins, "Analytic inverse kinematics for the universal robots UR-5/UR-10 arms," dec 2013, [Accessed 2015-03-18]. [Online]. Available: https://smartech.gatech.edu/handle/1853/50782

[21] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal*, vol. 7, no. 4, pp. 308–313, 1965.

[22] E. Jones, T. Oliphant, P. Peterson *et al.*, "SciPy: Open source scientific tools for Python," 2001–, [Accessed 2015-03-18]. [Online]. Available: http://www.scipy.org/