LettuceThink : A open and versatile robotic platform for weeding and crop monitoring on microfarms

David Colliaux[1], Peter Hanappe[1]
[1]P2P Food Lab, Sony Computer Science Laboratories
6, rue Amyot, 75005 Paris, surname.name@csl.sony.fr

ABSTRACT

Precision farming and in-field phenotyping benefit from developments in robotics. However, the technological advances are usually dedicated to large scale facilities, either monoculture farms or research facilities. We introduce the LettuceThink robotic platform as a versatile tool to ease the work on market farms. The design and software are modular and open source so that it can be built with standard parts and modified to fit specific needs. We demonstrate its use with two applications: precise mechanical weeding and 3D reconstruction of plants.

Keywords: Robotics, Weeding, Phenotyping, Microfarms..

1. INTRODUCTION

We are seeing a recent trend in Europe, in which young farmers are settling on small-surface, organic farms to grow a wide variety of vegetable crops. A major challenge on these microfarms is to overcome the workload as the smaller surfaces are compensated for by applying more spatially-dense growing techniques that require a large amount of manual work. The traditional machines used in agriculture are optimized for speed and power, and are not adapted to dense cultures (Bechar, C 2016). The LettuceThink platform is a lightweight wheeled robot equipped with precision sensors and actuators controlled with open source software. We present below two applications relevant to market farms. The weeding of crop beds is a task that demands the most amount of work and it is therefore the first application we wish to address. As a second application, we present the precise characterization of plant growth and plant structure. This analysis provides useful information to prevent the spread of diseases and to plan harvests.

The LettuceThink robot is currently being deployed on experimental fields to demonstrate its usability and efficiency. The acquired images will also complement annotated data sets used for plant recognition (Goëau et al, 2016) with market farm data. We wish to collaborate with related projects and we will release the hardware design and the software on a repository under free license.

The following sections provide details of the current state of our work. The open source robotic system is discussed in Section 2. Details on the weeding application, LettuceHoe, can be found in Section 3. Finally, an overview of the tool and application for 3D scanning is presented in Section 4.

D Colliaux and P Hanappe. “LettuceThink : A open and versatile robotic platform for weeding and crop monitoring on microfarms” .
EFITA WCCA 2017 Conference, Montpellier Supagro, Montpellier, France, July 2-6, 2017.
Paper number: 109

2. AN OPEN PLATFORM

2.1 Hardware design

We aim to build a platform that is inexpensive and easy to build and modify by farmers and other users. An open and accessible solution is a requisite to enable bottom-up, community-driven innovation for farming tools that use robotics and artificial intelligence. As an example of this bottom-up dynamics, we point to the stunning growth of 3D printing, laser cutting, and CNC routing tools over the past decade. Linked to this notion of open design is the idea of distributed production in which an object's description is available globally but produced locally, for example, in FabLabs (Gershenfeld, 2012, von Hippel, 2017). We use well-known, off-the-shelf components and simple assembly techniques to lower the costs, to benefit from the available online documentation, and to facilitate reuse.

The main frame of the robot consists of aluminium tubes (or wooden battens) that are assembled together using standards nuts and bolts. The total size is 1.67m x 1.26m x 1.64m (LxWxH). The robot bridges over a vegetable row, with two wheels on either side of the row, so it can work on the vegetables in the workspace underneath the robot. The soil-to-robot clearance is 80 cm, which allows it to pass over many commercial crops.

Four independent wheel modules are fixed to the main frame. Each wheel module can be controlled independently and has a commercial, electrical wheel designed for small scooters for the traction and a stepper motor with rotary encoder for the direction (Fig. 1). The robot is currently controlled remotely using a standard radio control because we have not yet included an automatic navigation system. The electronics consists mostly of a Raspberry Pi and Arduino compatible boards connected over Ethernet for communication (Fig. 2).

The robot carries a CNC machine with three degrees of freedom to place tools such as the weeding hoe or the 3D camera precisely in the workarea beneath it. We currently use the X-Carve's CNC kit. Our solution resembles the FarmBot project, which similarly uses a rail system to manage a small vegetable plot. However, by using a mobile system, LettuceThink can cover much larger areas.

The robot runs on two 12V lead batteries and carries a 280W solar panel to recharge the batteries. We estimate that the panel can provide enough energy to make the robot autonomous in terms of energy, but field experiments will have to confirm this.

2.2 Open source software

The software is also released under a free license and is available at http://github.com/p2pfoodlab/LettuceThink/. It is split into code for handling the general operations of the robot (navigation, power management, CNC control, topcam) and code for specific applications (LettuceHoe, LettuceScan). Each "app" provides its own user-interface, work tool, and control logic. We are in the process of porting the main modules to ROS1.

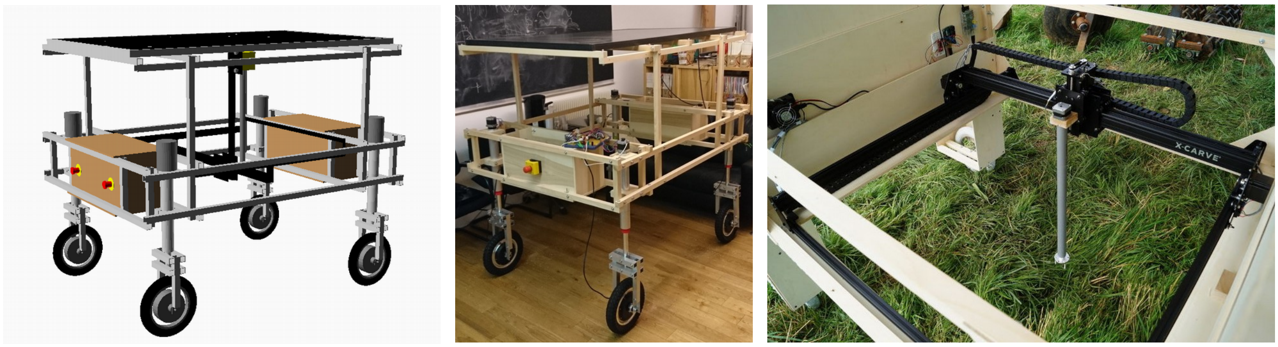

Figure 1. Robot design: (Left) The 3D design in OpenScad. Middle: The prototype in construction. (Right) The X-Carve CNC with an earlier version of the LettuceHoe tool.

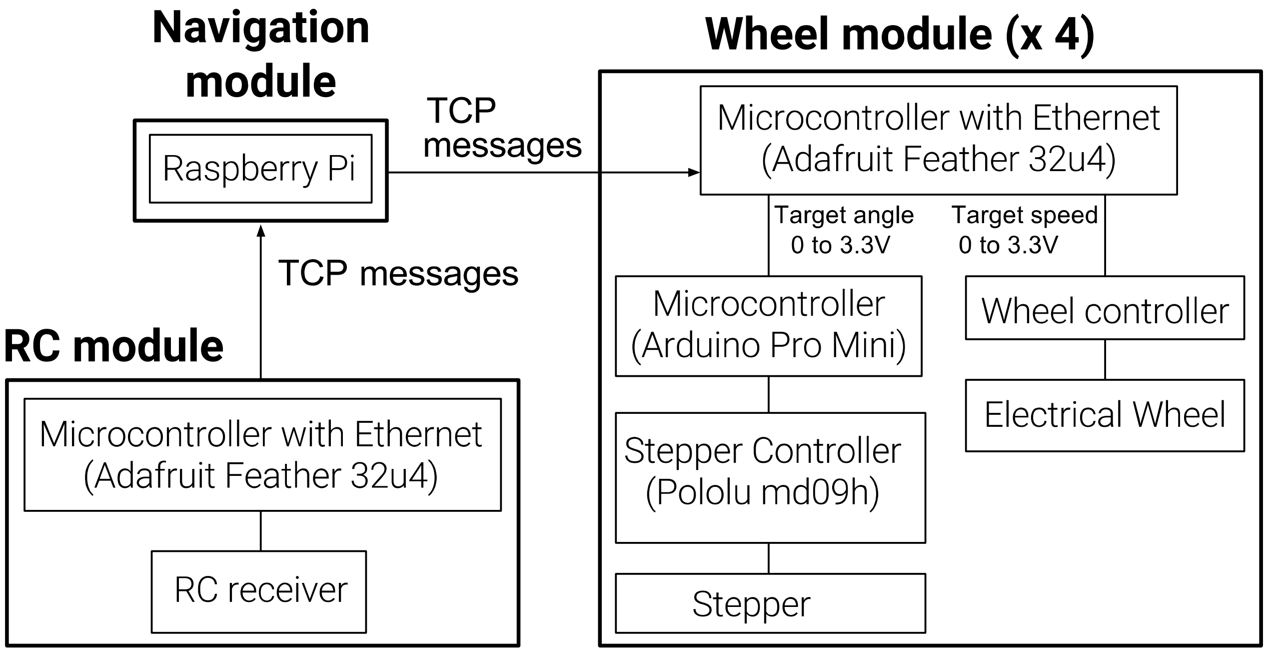

Figure 2. Robot design: The hardware modules for the navigation.

3. LETTUCEHOE: A TOOL AND AN APPLICATION TO WEEDING

Among the applications of agriculture robotics, weeding is especially interesting (Slaughter et al 2008). Several methods have been proposed for weeding based either on chemical, electrical, thermal or mechanical perturbations of the weeds. Some of these systems rely on the discrimination between weeds and plants of interest which, while based on color, hyperspectral, shape or location cues, may be hard problem to solve and require a large database to train classifiers. We present our solution to circumvent this problem and show early results from the exploratory experiment that started recently.

3.1 Principles

In the LettuceThink robot, a CNC machine is augmented with a rotating hoe and is moved through the working area. The task is to move the hoe along a path that covers the ground in the working area while avoiding the plants. First, we consider situations where the culture bed is initially free of weeds: young plants are grown in a greenhouse and then planted out in the prepared and weed-free culture beds. This context makes the weeding process more straightforward as it can be executed in 3 steps:

- Detect regions of the workspace occupied by plants of interest.
- Generate a tool path covering the workspace except for the regions occupied by plants.
- Run the tool through the generated path.

Our hypothesis is that a weekly application of this method prevents most of the weeds from developing.

3.2 Implementation

3.2.1 Plant detection

A large collection of color indices have been proposed for the detection of plants. The most robust index was selected by testing on images, see Fig. 3], from a wide variety of external conditions (with different lighting conditions and types of soil). The index used in practice is an excess green index, as described in Hamuda (2016), with slight difference in the normalization compared. It is based on the rescaled values of the image channels, at pixel (i,j):

$$\bar{\mathbf{R}}_{i,j}=\frac{\mathbf{R}_{i,j}}{\max\mathbf{R}},\bar{\mathbf{G}}_{i,j}=\frac{\mathbf{G}_{i,j}}{\max\mathbf{G}},\bar{\mathbf{B}}_{i,j}=\frac{\mathbf{B}_{i,j}}{\max\mathbf{B}}$$

where $\mathbf{R}$, $\mathbf{G}$ and $\mathbf{B}$ are the matrices corresponding to red, green and blue channels of the image. The excess green index is then defined, at pixel (i,j), as

$$\mathbf{C}_{i,j}=\frac{3\bar{\mathbf{G}}_{i,j}}{\bar{\mathbf{R}}_{i,j}+\bar{\mathbf{G}}_{i,j}+\bar{\mathbf{B}}_{i,j}}-1$$

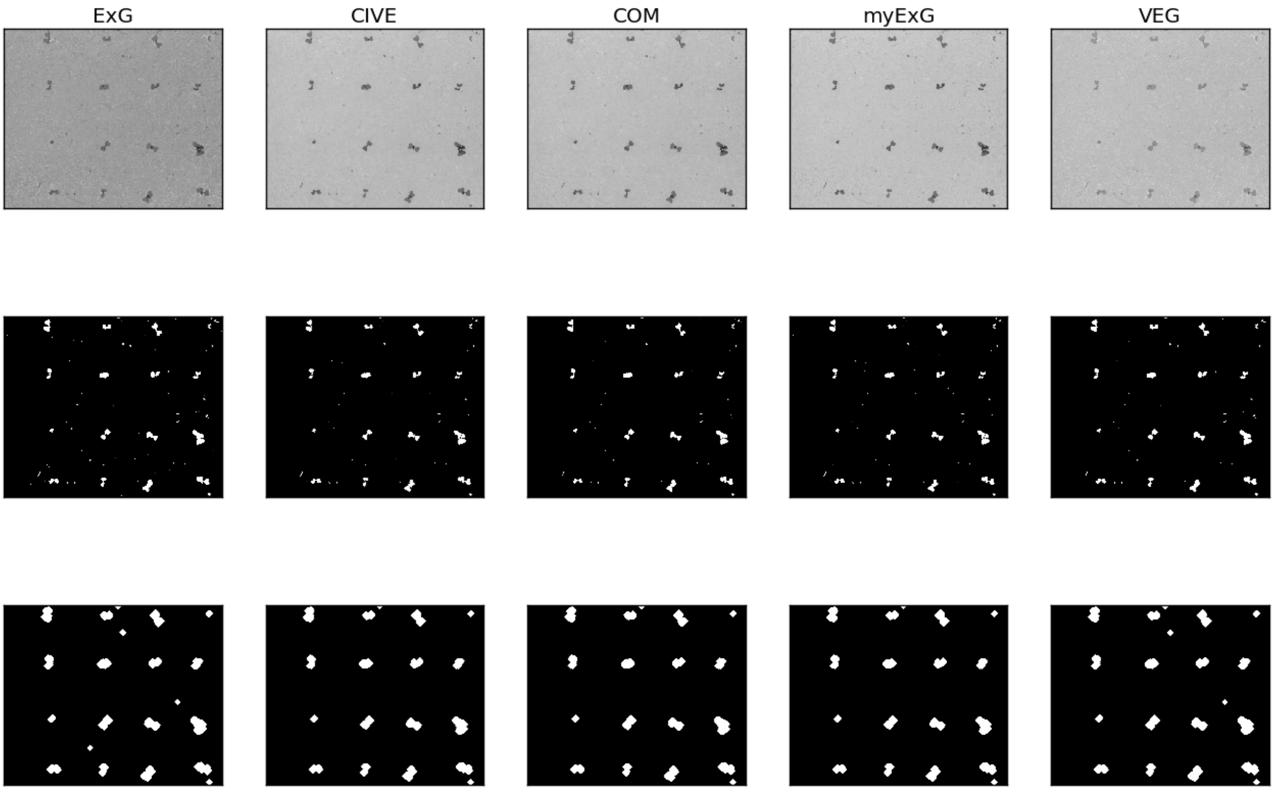

Figure 3. Comparaison of common color indices, listed in Hamuda (2016), for the computation of the plant masks. Several of the indices have similar performances on most images.

On images of young plants on ground, the histogram for this index is bimodal with one main peak for the ground and a secondary peak for the plants. Clustering of the image can be performed using K-means or Gaussian mixture models to determine the threshold for the separation the parts of the histogram related to the ground and those related to plants (an implementation of automatic thresholding using Otsu's method is available in OpenCV performs well on most images).

The thresholded image may still contain outlier pixels, with high excess green index, from the ground. To ensure that the mask describing the regions occupied by plants of interest does not include those pixels successive erosions and dilations are applied. The parameters of the morphological operations (size and shape of the kernel and number of iterations) are used to choose the minimum size of the regions identified as plants (see Fig. 4).

3.2.1 Path planning

Based on the mask of the segmented plants, the area to be avoided by the center of the hoe is taken as all points in plant area and all points located at less than the radius of the tool from the plant/ground border. The shortest solution to the 2 dimensional covering path problem on a convex domain with no obstacles is a simple zig-zag path, commonly known as the boustrophedon and reviewed in Galceran (2013).

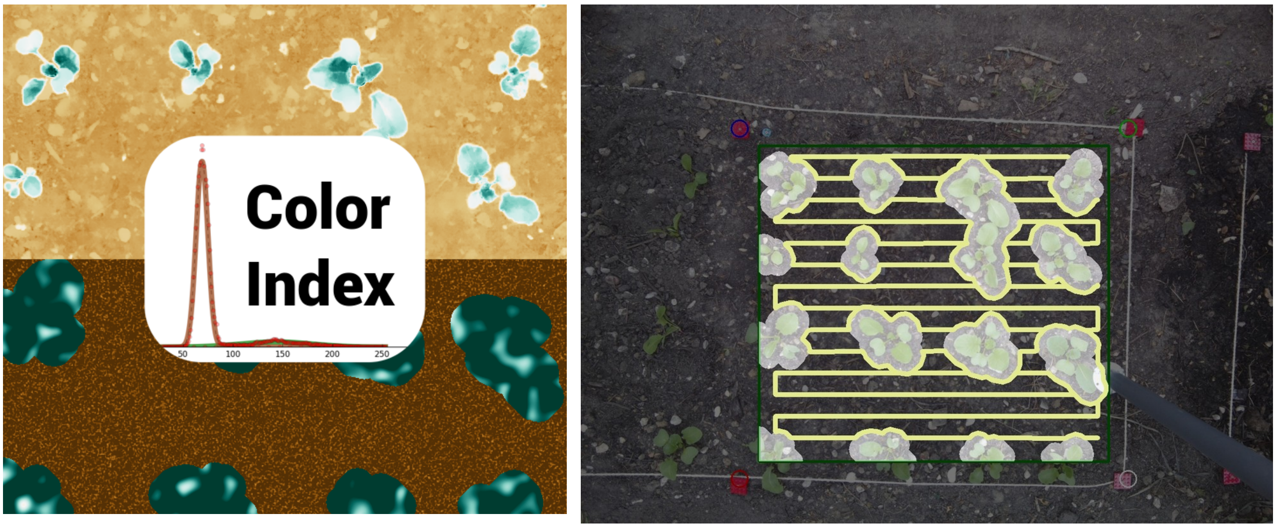

Figure 4. Weeding application algorithms: (Left) Excess green color index and mask defining the domains occupied by plants. (Right) Tool path generated using the modified boustrophedon algorithm.

As a first path generation algorithm, the boustrophedon path for the workspace with forbidden areas is modified so that parts of the path that go through the forbidden regions are substituted with valid regions. At each portion of the boustrophedon which crosses a forbidden region, entry and exit points are identified. To have pixel accuracy in the detection of those points, both the initial boustrophedon path and forbidden areas contours are densified using linear interpolation. For each pair of entry/exit points, the shortest path along the border of the corresponding forbidden zone connecting those point is substituted to the straight line of the original boustrophedon. Finally, the number of points composing the path is reduced using the Douglas-Peucker algorithm.

The modified boustrophedon is well suited for young plants but as plants grow, although the ground region to be covered by the tool shrinks, the resulting path length increases (see Fig. 5) due to multiple passages on the contours of the forbidden regions. As an alternative to this simple algorithm, a more computationally demanding algorithm is designed using the following steps:

- Generate cells of size comparable with the tool on the ground regions,
- Find the shortest path passing through all centers of the cells once and only once,
- Reduce the number of points along the path using the Douglas-Peucker algorithm.

There are several possible algorithm to generate ground cells, e.g. with the SLIC superpixel algorithm. For the second step, a traveling salesman problem (TSP), there are also multiple methods available. Elastic net, a kind of self-organizing maps, reaches approximate solutions with good performance when the problem dimension is reasonable (Durbin et al 1987) and we found that it perform well on this problem. The centers of superpixels ($\mathbf{S}_i$ with $i \in 1 \cdots N$) are associated with cities in the TSP so that the shortest path along those points will cover the ground. Starting with a number of nodes (or neurons) ($\mathbf{P}_j$ with $j \in 1 \cdots M$) denser than number ($N < M$) and organized along a path, we solve the optimization problem with a cost function of 2 terms:

- $C_1(\mathbf{P}_j, K) = -\alpha K \sum_i log \sum_j e^{-|\mathbf{S}_j - \mathbf{P}_i|^2/(2K^2)}$ ¿ is the term accounting for the path passing through the centers of cells modelled as a attraction potential by centers of cells on path nodes.
- $C_2(\mathbf{P}_j) = \beta \sum_j (\mathbf{P}_j - \mathbf{P}_{j+1})^2$ is the term accounting for the path length modelled as an elastic interaction among path nodes.

The update rule for moving path nodes is then:

$$\Delta \mathbf{P}_j = \alpha \sum_i w_{ij} (\mathbf{S_i} - \mathbf{P}_j) + \beta K (\mathbf{P}_{j+1} + \mathbf{P}_{j-1} - 2\mathbf{P}_j)$$

with

$$w_{ij} = \frac{e^{-(\mathbf{S_i} - \mathbf{P}_j)^2/(2K^2)}}{\sum_k e^{-(\mathbf{S_i} - \mathbf{P}_k)^2/(2K^2)}}$$

Here the parameter K is a regularizing factor, varied along iterations of the algorithm, so that the initial steps are not trapped in a local minimum. At the begin of the algorithm, K is large and each path node feels the potential of several centers around and as the algorithm goes on interactions are frozen so that only interactions with closest centers remain.
We tested path generation with both modified boustrophedon and elastic net on a workspace including plants at various (simulated) growth stages. The path is shorter with the modified boustrophedon at early stages showing that the elastic net solution is not globally optimal. As plants grow, the elastic net solution gets shorter than the modified boustrophedon, reflecting the inefficiencies mentioned above for the modified boustrophedon. Paths obtained with the elastic net are decreasing in length as plants grow, as expected, and at some stage of the growth, paths generated by the elastic net are shorter than the ones obtained by the modified boustrophedon. At this point, elastic net solutions are better although at higher computational cost. Another possible algorithm, yet to be tested, is based on a cellular decomposition of the ground domains and running a boustrophedon path on each cell (Chosset, 2000).

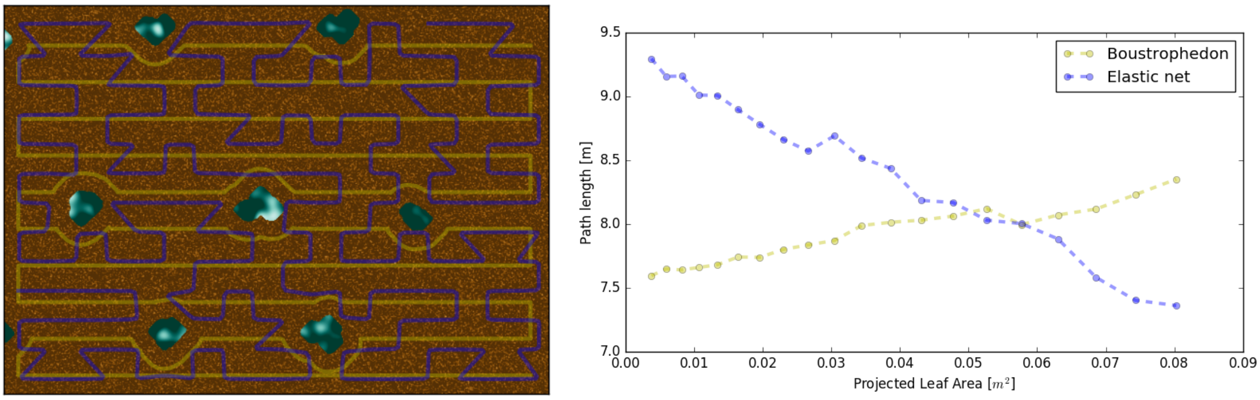

Figure 5. Comparaison of paths generated by the modified boustrophedon (yellow) and the elastic net (purple): (Left) Sample paths when plants are small. (Right) Path length depending on the size of plants (simulated through iterated dilation of the small plants). Curves crossing suggest a plant size for which elastic net generates shorter paths, but at the cost of longer computation time.

3.3 Exploratory experiments in outdoor environments

We are testing the weeding application on radishes cultivated in plain ground. We continue to refine the application to account for unexpected aspects of an open-field experiment. For example, red markers were added on the field to make the identification of the workspace easier on images. One of the challenges that we encountered is the high contrast in the image generated by shade, which may result in plants or markers detection to fail. As plants grow, markers may be hidden by plants or no ground may be left to hoe. Although when plants are big weeding may be unnecessary since the growth of weeds is slowed down due to the foliage, the path generation algorithm could be adapted to generate covering path of the connected components so that it can be used when plants are overlapping leaving isolated islands of ground to be hoed. We also noticed that, for radish, when the plant grows leaves are few centimeters above the ground so that using the contour of leaves as border for the forbidden zone may leave a large portion of the ground untouched. In this case it may be better to take finer details of the plant architecture into account in the definition of the forbidden zone.

Figure 6 Test of the weeding application on radishes cultivated in open-field: (Left) The hoe is passed before transplanting and every week. (Middle) The hoe is passed before transplanting only. (Right) The ground is left untouched.

Preliminary evaluation shows that a pass every week is enough to maintain a low density of weeds compared to a bed where the hoe was not used during plant growth and that had many more weeds. In both case, the ground was hoed before transplanting the radishes. A single pass before transplanting is quite efficient in delaying the growth of weeds by about a week.

4. LETTUCESCAN: TOOLS AND APPLICATIONS TO CROP MONITORING AND 3D SCANNING

The second application that we present in this paper, at the crossroad of crop monitoring and plant phenotyping, is the characterization of the plant structure and plant growth. This app aims at providing information to the farmers about the health and maturity of its crops.

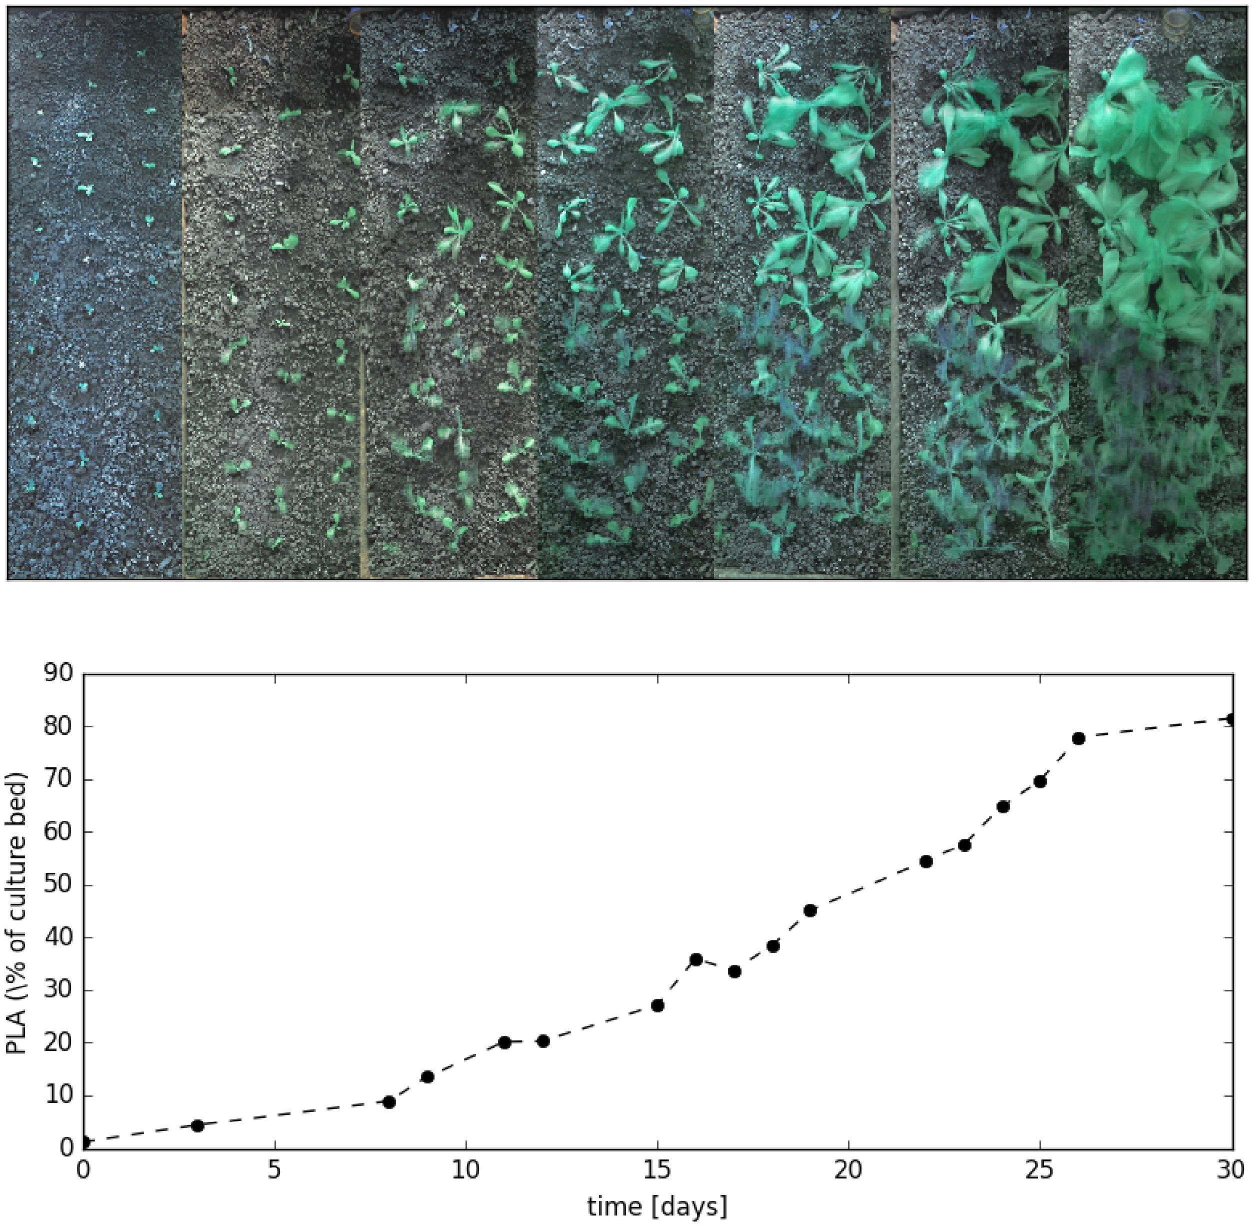

Figure 7 Crop monitoring based on 2d images: (Top) Snapshots of the culture bed along a month. (Bottom) Estimation of the plants size from the projected leaf area.

The pictures from the top-camera of the robot give a good approximation of the volume of lettuces but for a finer description or for other plants this set-up is lacking. An RGB+Depth camera mounted on the robotic arm and traveling around the plants provides much more information. We show that a good 3d reconstruction can be obtained with commercial time-of-flight cameras. The characteristics of the plants (leaf area, leaf count, leaf orientation relative to the shoot and to each other) can then be estimated by processing the acquired point clouds (Li ,2014, Chaudhury, 2017).

A first component, 2D crop monitoring, is easily implemented from the robot operations in LettuceHoe since the weeding application generates images of the plant bed. The surface occupied by segmented plants is measured on each of these images and the resulting projected leaf area is an estimation of the size of the plant. As seen on Fig. 6, plant growth dynamics are monitored from the projected leaf area (PLA).

We expect the accuracy of this estimation of plant volume from 2D images to decrease as plant grow but as can be seen on Fig. 6, it is still be very useful in detecting growth stage transitions. It would also be able to characterize variations of growth dynamics across the field.

Ultimately, the accuracy of PLA in estimating plant volume should be evaluated using stronger methods, for example, relying on 3D imaging. We thus started using a commercial 3D camera (Softkinetic DepthSense) based on RGB and TOF sensors. The camera is attached to the robotic arm and it can be controlled in translation on 3 axes and in rotation around 2 axes. This 3D scanning system aims at the precise evaluation of the volume of plants as well as their architecture.

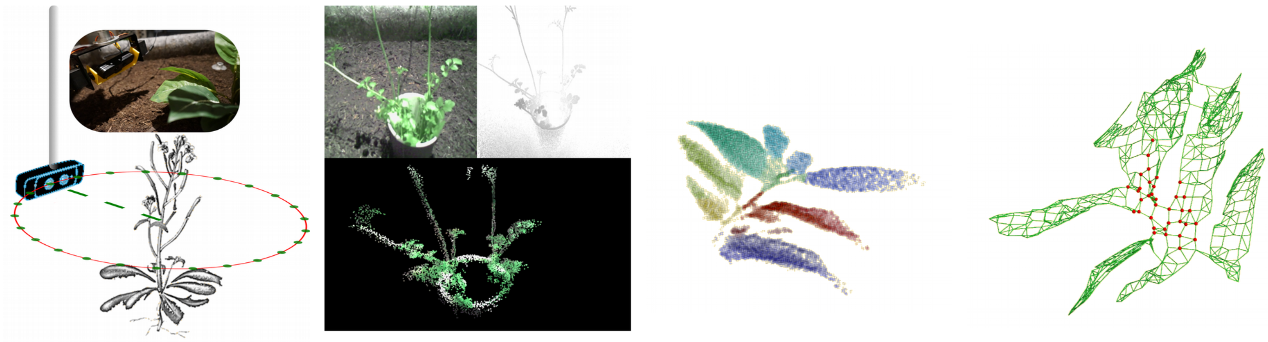

Figure 8 LettuceScan pipepeline: (Left) Data acqisition with a 5 DOF camera. (Center-Left) Combinaition of RGB and depth images into a coloured point cloud. (Center-Right) Segmention of the point cloud isolating leaves of a plant. (Right) Graph representation of a plant and identification of the shoot.

The 3D point clouds from multiple views are combined for the reconstruction of the plant. The camera moves around the plant along a circular path and a hundred images are sampled. At each point, a colored point cloud is generated from RGB and depth data and filtered to avoid outliers and smooth the point cloud. As the neighbouring images are close to together the transformation from one image to the other can be directly estimated on the dense point clouds using an iterative closest point algorithm (Chen et al, 2015). We found the dense transformation to perform well on data separated by as far as 12° on the circle although some specific configurations (like a leaf being perpendicular to the sensors) made the estimation difficult. After pairwise registration of point clouds, they are merged maintaining a constant density of points. The resulting point cloud is converted into more abstract representations like meshes, graphs or octree. For example, a sparse graph based on nearest neighbors in the point cloud is shown on Fig. 8 vertices with highest betweenness centrality index are highlighted in red, those are located on the shoot. As another post-processing step, the clustering of the point cloud by K-means result in segmentation of leaves. Further work is still needed to analyze other components of plant architecture like the characterization of individual leaves or phyllotactic patterns.

5. CONCLUSION

We showed in this paper that a lightweight robotic plateform is useful on small scale farm by relieving farmers of painful and time consumming tasks like weeding. We also highlighted other uses of the robot based on embedded imaging devices. Future work will be dedicated to show the potential usefulness of this application for farmers in planning harvest and detecting pests or stress in fields. This application may also drive new insights for researchers in plant biology, agronomy or ecology.

We are also hopeful that the modular, open source design and programming of the robot will bring interested researchers and farmers in proposing new, unexpected modules.

6. REFERENCES

Bechar, A. & Vigneault, C. (2016) ’Agricultural robots for field operations: Concepts and components’, *Biosystems Engineering*, Vol.149, pp.94–111

Chaudhury A., Ward C., Talasaz A., Ivanov A.G., Brophy M., Grodzinski B., Huner N.P.A., Patel R. V. & Barron J. L.(2017) ’Machine Vision System for 3D Plant Phenotyping’, *https://arxiv.org/abs/1705.00540*

Chen K., Lai Y.K. & Hu S.M. (2015) 3D indoor scene modeling from RGB-D data: a survey *Computational Visual Media*, vol.1, no. 4, pp 267-278

Choset H. (2000) ’Coverage of Known Spaces: The Boustrophedon Cellular Decomposition’ *Autonomous Robots* Volume 9, Issue 3, pp 247–253

Durbin R. & Willshaw D. (1987) ’An analogue approach to the travelling salesman problem using an elastic net method’ *Nature* 326, 689 - 691

Galceran E. & Carreras M. (2013) ’A survey on coverage path planning for robotics.’ *Robotics and Autonomous Systems* vol. 61, no. 12 , pp. 1258-1276

Gershenfeld N (2012) ’How to Make Almost Anything: The Digital Fabrication Revolution’, *Foreign Affairs,* 91:6, 2012.

Goëau H., Joly A. & Bonnet P. (2016) ’Plant Identification In An Open World, CLEF 2016 working notes.’ *LifeCLEF 2016 working notes, Evora, Portugal*, http://ceur-ws.org/Vol-1609/16090428.pdf

Hamuda E., Glavin M. & Jones E. (2016) ’A survey of image processing techniques for plant extraction and segmentation in the field.’ *Computers and Electronics in Agriculture*, vol.125, pp. 184-199

von Hippel, E (2017) ’Free Innovation’, Cambridge MA: MIT Press

Li L., Zhang Q. & Huang D. (2014) ’A review of imaging techniques for plant phenotyping’ *Sensors,* vol. 14, no. 11, 2014.

Slaughter D.C., Giles D.K., Downey D. (2008) ’Autonomous robotic weed control systems: A review*’ Computers and Electronics in Agriculture* Volume 61, Issue 1, Pages 63-78