Multispectral detection of floral buds for automated thinning of pear

Niels Wouters[a] • Bart De Ketelaere[a] • Tom Deckers[b] • Josse De Baerdemaeker[a] • Wouter Saeys[a]

[a] *Department of Biosystems, KU Leuven,*

Kasteelpark Arenberg 30, B-3001 Leuven, Belgium

Phone: +32 16 372403

Fax: +32 16 321994

E-mail: niels.wouters@biw.kuleuven.be

[b] *Research Station for fruit growing (pcfruit)*

Fruittuinweg 1, B-3800 Sint-Truiden (Kerkom), Belgium

Abstract Thinning of pome and stone fruit involves the reduction of tree crop load in order to regulate fruit set and quality. As it is typically carried out through manual labor, thinning comprises a large part of a grower's production costs. Mechanized thinning has been shown to be a cost-effective alternative but the performance of existing thinning devices needs to be further improved by taking the variation in bearing capacity of the individual trees into account.

In this work, a multispectral camera system is developed to detect the floral buds of pear (cv. *Conference*) during the growth stages prior to bloom. During a two-year field trial, the multispectral system was used to measure orchard scenes in six distinct optical wavebands under controlled illumination. These wavebands are situated in the visible and near infrared region of the spectrum and were selected based on hyperspectral laboratory measurements described in previous work.

The recorded multispectral images were converted to a database containing the spatial-spectral signatures of the objects present in the orchard. Subsequently, canonical correlation analysis was applied to create a spectral discriminant model that detects pixels originating from floral buds. This model was then applied to the recorded data after which an image analysis algorithm was designed and optimized to predict the number of floral buds. In total, approximately 87% of the visible floral buds were detected correctly with a low false discovery rate (<16 %). Therefore, it is expected that the multispectral sensor can be used to improve the efficiency of existing thinning devices. Additionally, it could as well be used as a stand-alone sensor for early-season yield estimation.

Keywords: multispectral sensing, pear, mechanical thinning, feature recognition

1. Introduction

Horticulture involves many tedious and labor-intensive tasks which require the employment of expensive, trained personnel. As it is becoming increasingly difficult for growers to hire a sufficiently large work force (Maas and van der Steeg, 2011), an increasing amount of research is targeted to automate or augment the operation of cultivation techniques. Many of these horticultural practices require some form of feedback to either assess the state of the canopy (e.g. disease detection) or determine the location of certain objects (e.g. harvesting). As humans typically rely on their sight to perform these tasks, robotic systems are often equipped with a vision system to match or even improve on the performance of their human counterparts. In this work, we focus on the development of a vision system in the context of automated thinning in fruit orchards.

As fruit trees have a natural tendency to produce heavy crop loads, the sugars produced in the leaves (sources) need to be distributed over too many fruits (sinks). This often results in the production of many small fruits which are not suited for fresh market sale. Thinning decreases the competition for photosynthetic products by removing the excess buds, flowers or fruitlets. This not only allows the remaining fruits to reach commercially interesting sizes, but also increases fruit quality, tree vigor and yield regularity (Lopez, 2011; Theron, 2010 and Meland, 2009). Research has shown that early thinning – at or even prior to bloom – leads to stronger positive effects than the traditional late season thinning, because it minimizes the investment of the trees in fruits which will not be harvested (Theron, 2010; Meland, 2009; Link, 2000; Bertschinger et al., 1998). Together with pruning and harvesting, thinning is one of the most labor-intensive – and thus expensive – cultivation measures as these are still typically performed by hand. Consequently, a strong demand exists among growers for alternative thinning methods.

Over the years, the potential of chemical thinning has been extensively studied. Though it can be considered a practical and cost-effective method, it cannot completely and reliably replace hand thinning (Miller and Tworkoski, 2010). Generally speaking, chemical thinning suffers from two main drawbacks. Firstly, the efficacy of the currently available thinning agents is strongly related to cultivar and environmental conditions (Kviklys and Robinson 2010; Peck and Merwin 2009). Secondly, chemical thinning often has detrimental effects on the environment, tree vigor and human health (e.g. laborers). It is for this reason that many chemical thinning agents have been withdrawn from the market (Hong, 2010). However, even under perfect conditions, growers still have to await the actual response of the trees as chemical thinning offers no direct feedback.

Mechanical thinning machines developed in recent years have demonstrated that automated thinning can be a viable alternative for the traditional methods and can yield economic savings. String thinners realize apple and peach blossom thinning by means of fast rotating flexible strings (Hehnen et al., 2012; Martin-Gorriz et al., 2012; Martin-Gorriz et al., 2011; Baugher et al., 2010). Spiked drum-shakers were used for peach fruitlet thinning by using rotating drums to transfer shaking energy to the canopy branches (Miller et al., 2011; Schupp et al., 2008). Wouters et al. (2014) removed floral pear buds by pulses of compressed air. Finally, Yang (2012) and Nielsen et al. (2012) developed a prototype robotic manipulator and clamplike end effector for brushing off peach blossoms. Other techniques such as trunk shaking (Gloser and Hasey, 2006) or limb shaking (Martin-Gorriz et al., 2010; Rosa et al., 2008) have been investigated as well, but were found less effective.

Although positive results were realized by these automated techniques, their thinning speed and efficiency need to be further improved by taking into account the tree-to-tree variability. As the floral bud distribution is non-uniform throughout an orchard, certain trees – or regions on a tree – will benefit from less or more severe thinning. Since most of the existing techniques often cause injuries to shoots, leaves and bark, thinning in a way tailored to the needs of each individual tree would prevent unnecessary tree damage. This maintains tree vigor and reduces the risk of disease spread (Kon et al., 2013; Ngugi and Schupp, 2009; Schupp et al., 2008; Bertschinger et al., 1998). Furthermore, it would allow to prevent overthinning of high-value crops.

In recent years, several researchers have investigated vision systems to detect and quantify fruit blossoms with the goal to provide this information as feedback to a thinning machine. Gebbers et al. (2013) introduced a shock absorbing stereo camera platform to map the flower density on apple trees. They used this information to control the rotation speed of a string thinner and thereby the thinning intensity. Nielsen et al. (2012) achieved good peach blossom detection by means of a trinocular stereo color camera. They were able to locate the three dimensional (3D) position of the blossoms with a spatial accuracy of less than 1 cm. Emery et al. (2010) developed a scanning laser range imaging system to measure the 3D shape of peach trees with a spatial accuracy of 1.2 cm.

These detection techniques all rely on the sharp color contrast between the blossoms and their environment as quantified using standard RGB cameras. However, this approach is not suitable for detecting floral buds prior to bloom as the brightly colored petal leaves are still contained within the buds. To our knowledge, no attempt has been made to develop a sensor to detect floral buds prior to bloom.

Previous work has shown that multispectral imaging can be successfully applied for object recognition in many agricultural applications (e.g. Bas et. al., 2013; Bulanon et al., 2010; Okamoto and Lee, 2009; Wallays et al., 2009). This technique produces images with a higher contrast between objects of interest by combining more and narrower wavebands than the red, green or blue regions of the spectrum.

In previous work (Wouters et al., 2013), we determined the optimal wavebands for building a multispectral vision system which is able to detect floral pear buds in the phenological stages before bloom (*Pyrus communis* cv. *Conference*). Using these wavebands, a discrimination model was built that already showed good pixel classification under laboratory conditions (i.e. 95 % correct pixel classification). However, additional steps are required to make this technique suitable for floral bud detection under field conditions. In this we work, we deal with the following three challenges: (1) going from pixel to object recognition, (2) taking into account the presence of additional objects which are not included in the original discriminant model and (3) performing the detection at faster, more realistic speeds. First, a new multispectral setup is elaborated which was tested during a two-year field trial. Hereafter, details are provided on the construction of a new pixel classification model and the image analysis used to realize object detection. Finally, conclusions are presented regarding the potential of the detection system and suggestions are made for future research.

2. Materials and methods

2.1. Image acquisition setup

A low-cost custom movable camera platform was built to perform multispectral measurements in field conditions [Fig. 1(a)], similar to the setup used by Bulanon et al. (2010). The setup consists of a 12 bit monochrome CCD camera (TXG14NIR, Baumer, Frauenfeld, Switzerland) with a resolution of 1392 x 1040 pixels and a 16 mm monofocal manual iris lens (C1614A, Pentax, Tokyo, Japan). In front of the lens a fast rotating multispectral filterwheel (FW103H/M, Thorlabs Inc, Newton, NJ, USA) is placed which houses six optical bandpass filters in the range 400-1000 nm with a diameter of 25 mm. These filters are rotated sequentially in front of the lens with a change time of approximately 55 ms between adjacent filters. This operation enables to perform fast multispectral measurements (< 1 s) with no or very limited distortions between the different filter images, e.g. motion blurring caused by wind. The filters are commercially available bandpass filters which were selected to have bandpass regions that match as closely as possible to the desired optimal wavebands to discriminate between floral buds and their environment (Wouters et al., 2013). Both the actual and optimal transmission bands of the filters are displayed in Table 1.

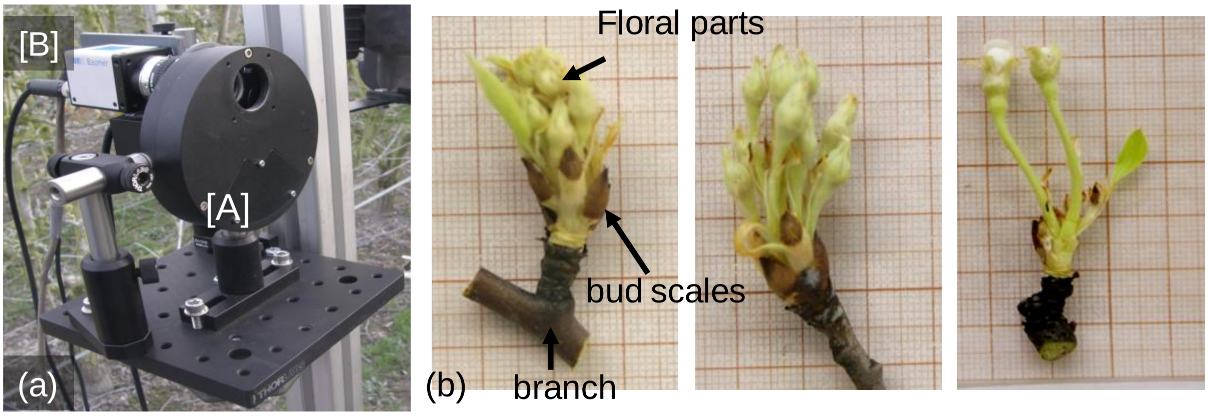

Fig. 1 (a) Camera platform used during the field measurements. Main components are: [A] a fast rotating filter wheel containing six optical bandpass filters and [B] a monochrome camera. **(b)** Appearance of the floral buds during the examined phenological stages. The three main constituents are indicated. Stages are displayed chronologically, from left to right: "Green cluster", "Green bud" and "White bud". The buds are displayed in front of graph paper to give a measure of scale (1 square = 1 mm²).

Table 1 Comparison between the optimal wavebands found by Wouters et al. (2013) and the actual wavebands of the filters used during the field experiments.

Order of importance[x]	1	2	3	4	5	6
Optimal waveband [nm]	595 – 610	925 – 975	440 – 490	685 – 700	755 – 805	535 – 565
Actual waveband[y] [nm]	589 – 625	925 – 975	430 – 490	672 – 712	752 – 798	532 – 554
Filter name	NT84-102[a]	NT86-072[a]	MF460-60[b]	NT67-038[a]	NT84-106[a]	NT67-032[a]

a: manufactured by Edmund Optics, Barrington, NJ, USA

b: manufactured by Thorlabs Inc, Newton, NJ, USA

x: as determined by Wouters et al. (2013)

y: the bandwidth of the actual wavebands is determined by their "full-width at half maximum", i.e. the width between the points of the passband wavelengths where the transmittance is 50% of that of the central wavelength of the filter.

To check the effect of the difference between the optimal and actual wavebands, the methodology and dataset used to select the optimal wavebands (Wouters et al., 2013) were again used to predict the pixel classification performance of the actual filters. It was found that difference between the actual and optimal wavebands reduced the predicted correct pixel classification by less than 1 %. This is attributed to the typical high correlation between information gathered from (partly) overlapping wavelengths (Table 1). Therefore, the effect of choosing the commercially available filters instead of the optimal wavebands can be considered negligible.

Data acquisition and control of the setup was realized by means of a laptop running a custom software written in Labview 2009 (National Instruments, Austin, Texas, USA).

2.2. Orchard description and phenology

During the growing seasons of 2012 and 2013, field measurements were conducted in a commercial pear orchard situated in Bierbeek, Belgium (50°49’36.35”N, 4°47’40.35”E). Trees of the pear cultivar *Conference* were trained in an intensive V-hedge system with four main fruiting branches on one central stem (Quince C rootstock, planted in 1992). The trees possessed an average height of 2.5 m and were spaced at 3.5 m x 1.3 m (1978 trees.ha^{-1}).

Multispectral images were acquired during three early phenological stages occurring before bloom [Fig. 1(b)], i.e. “Green cluster”, “Green Bud” and “White bud”. In these stadia, the main constituents of the trees are branches, bud scales and developing floral parts. Most of the canopy leaves are still contained inside the leaf buds. In Fig 2. the appearance of the orchard at the time of the field trials is illustrated. Note from this figure that the branches of the trees are typically covered by green-colored algae.

2.3. Experimental procedure

Measurements were conducted at nighttime with use of artificial illumination. This approach had two advantages: firstly, it allowed to control the quality of the illumination of each scene and reduced the variability between measurements which result from variations in the illumination due to clouds and solar movements at daytime. Secondly, measuring at night simplified the observed scenes as the visibility of background objects was greatly reduced. This simplified image analysis as less background objects needed to be filtered out.

Multispectral images were recorded at random locations throughout the orchard. Before each measurement, the setup was placed at a distance of approximately 1 m from the canopy which resulted in a field of view of roughly 410 by 550 mm. The height at which the setup was placed was chosen randomly as well. Illumination of a scene was provided by a 500 W halogen lamp which was of the same type as the light source used during the hyperspectral measurements in the laboratory (Wouters et al., 2013). The power of this light source was sufficient to illuminate the line of trees in front of the camera, but not high enough that trees in the background are visible. In the first season, the lamp was held stationary by hand at a fixed position relative to the camera setup (Fig. 2). As this method can give rise to less even illumination, the lamp was but was mounted on a fixed support in the second season in order to provide a more stable illumination.

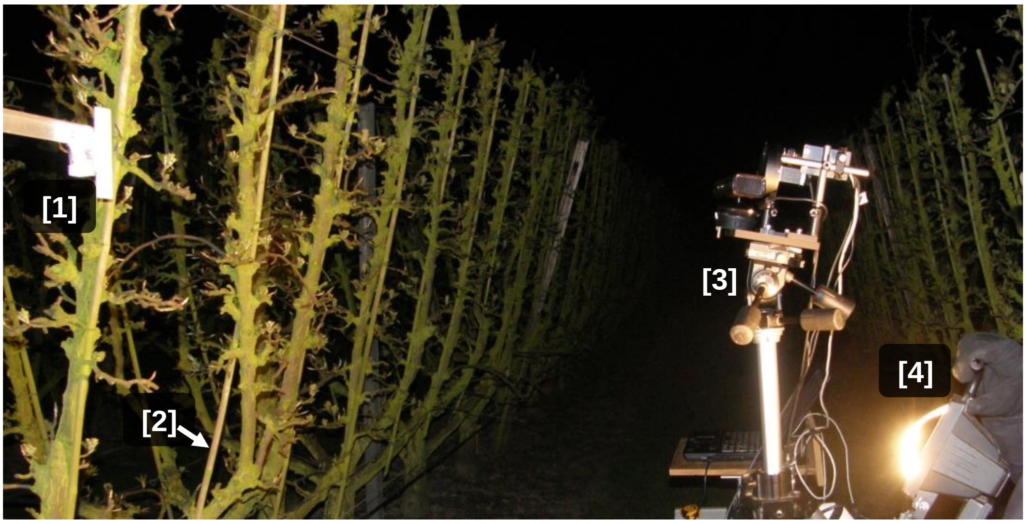

Fig. 2 Picture of the setup, taken during the field trials. Numbers refer to: [1] optical reference made out of PTFE, [2] bamboo support, [3] camera platform and [4] handheld illumination.

Since the multispectral camera applied optical filters from both the visible and near infrared region of the spectrum, undesirable chromatic aberration effects might occur. In order to minimize this effect, a small aperture of the lens was chosen (f/8). This increased depth of field (DOF) at the expense of an increase in the required exposure time. From each scene, multiple multispectral images were recorded at various exposure times. In this way, the highest quality images (maximum dynamic range with little to none saturated pixels) could be selected for each scene for further analysis. All relevant information related to each experiment (exposure times, duration of a measurement, file names, etc.) was automatically recorded in a log-file. Finally, a classical RGB picture of every scene was taken with a standard RGB camera (SP-55OUZ, Olympus Corporation, Tokyo, Japan). The number of floral buds in each scene was counted by hand.

For the purpose of data normalization – discussed in section 2.4 – an optical reference was placed in the field of view of the camera (Fig. 2). This reference was placed at the same location in each scene. In the first season this was a small white polytetrafluoroethylene (PTFE) plate. Since the luminosity of this reference was relatively high in comparison to the other objects in a scene, the full dynamic range of the camera could not be used to measure these objects as this resulted in a saturation of the pixels of the reference. For this reason, the white reference was replaced in the second season by a grey-colored reference made from polyvinylchloride (PVC) which possessed a luminosity similar to that of the trees in a scene, resolving the higher mentioned issue. Both PTFE and PVC display stable optical behavior in the 400-1000 nm range without clear absorption peaks.

Field trials season 1. Experiments took place from March 26th until March 30th. Respectively 13, 15 and 15 scenes were recorded during the stadia "Green Bud", "green Cluster" and "White bud". In total, 48 fruiting branches carrying 353 floral buds were imaged. The observed number of floral buds was distributed approximately equal over all three phenological stadia.

Field trials season 2. Field tests were conducted starting on April 16th and lasting until April 24th. Respectively 15, 15 and 14 scenes were recorded during the stadia "Green Bud", "green Cluster" and "White bud". The total number of observed floral buds was 315, spread over 44 fruiting branches. The majority of these floral buds (47%) were observed during the "White bud" stadium. The images of the stadia "Green cluster" and "Green bud" contained, respectively 14% and 39% of the floral buds.

2.4. Pre-processing

For each scene, only the image for each optical filter λ_i with the highest dynamic range was retained for further analysis. To be able to compare images taken at varying exposure times *t*, the raw images $S(\lambda_i,t_j)$ were converted to reflectance images $r(\lambda_i)$ by normalizing them with respect to the average intensity of the optical reference $I(\lambda_i,t_i)$ which was recorded in the same image*:*

$$r(\lambda_i) = \frac{S(\lambda_i, t_j) - D(\lambda_i, t_j)}{I(\lambda_i, t_j) - D(\lambda_i, t_j)} \quad , \; i = 1 \ldots 6 \tag{1}$$

Both the raw image $S(\lambda_i,t_j)$ and the reference signal $I(\lambda_i,t_j)$ were corrected for the dark current image $D(\lambda_i,t_j)$. The latter is related to the noise caused by the electronics and was measured by capturing images with the sensor shielded from incident light by means of the lens cap.

As discussed in section 2.3, a different optical reference was used during each growing season. Since both references possessed a different relative luminosity this had an effect on the attained reflectance values $r(\lambda_i)$, as can be seen from formula (1). To be able to compare reflectance values across the two measuring seasons, a correction of the reflectances $r(\lambda_i)$ obtained during the second season was carried out to estimate the corresponding reflectance value of the PTFE reference used during the first season. To this end, both references were measured simultaneously with the multispectral setup at various exposure times. This resulted in a linear calibration curve which related the observed intensities $I(\lambda_i, t_j)$ for the PVC reference to those for the PTFE reference (r^2=99.9 %). A separate calibration curve was fitted for each filter.

The resulting reflectance images acquired for each filter were concatenated to create a multispectral image cube of each measured scene. Hereafter, a spatial-spectral object database was constructed by manually indicating in each multispectral image to which type of object the foreground pixels belonged. The four most common types (or classes) of objects occurring in the multispectral images are termed as the “main components”. These were bud scales, developing flower parts, branches and bamboo supports (Fig. 2). The latter are a part of a scaffolding that is required to support the weight of the branches. Besides these main components, other types of objects were marked in the images as well. Most of these objects were a part of the supporting scaffolding for the trees. Examples include plastic and metal wires and wooden and concrete posts.

2.5. Image analysis

In this section the methodology to translate the information contained in the multispectral images to the recognition of floral buds is described. The operation of the detection algorithm can be divided into three main parts (Fig. 3). First, a statistical model creates a probability image P of each scene. The latter is an image in which each pixel is assigned a likelihood that it belongs to a floral bud. In the two subsequent steps, morphological image processing is applied to the probability images in order to identify the floral buds as objects (segmentation) and remove noise. Finally, the performance of the detection algorithm presented here was optimized by means of a desirability index. The parameters of the detection algorithm subjected to this optimization are denoted as "χ_i", with *i* representing a number assigned to each parameter. All analyses were performed in Matlab, version 7.5.0 (MathWorks Inc., MA, USA) on an Intel® Core™ i7 CPU Q720 @1.60 GHz with 8GB RAM.

2.5.1. Pixel classification model

Object detection in images can be greatly facilitated if the majority of the pixels in an image can be correctly attributed to a certain class. For this reason, a pixel discriminant model was built by means of canonical correlation analysis (CCA), as described by Sharma (1995) and applied in previous work (Wouters et al., 2013) [Fig. 3 – step 1A]. CCA is a multivariate analysis technique which produces orthogonal discriminant functions that have maximum separation between groups. Three discriminant functions are required to discriminate between the four main components. In the discriminant space spanned by these functions, pixels are classified based on their Bayesian posterior probability. A Box's M test showed that the covariance matrices of the different groups were unequal ($p < 0.001$). Therefore, Quadratic discriminant analysis (QDA) was used. Finally, a 1392 x 1040 pixels probability image P was made for each scene by assigning to each pixel the posterior probability of it belonging to the group "flower parts" [Fig. 3 – step 1B].

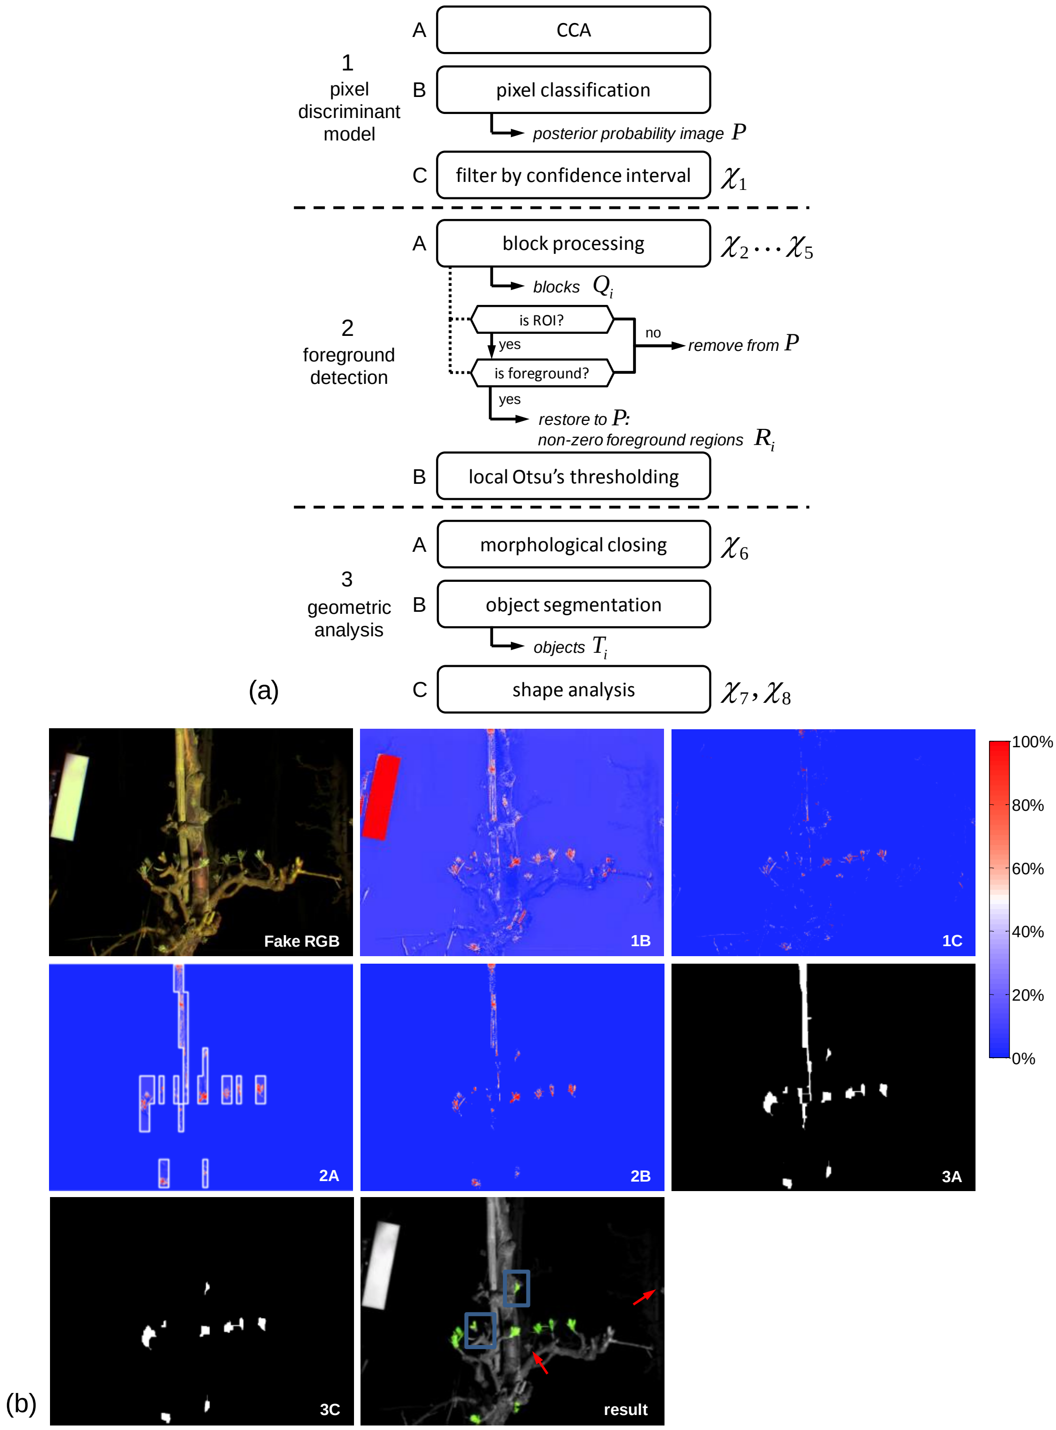

Fig. 3 (a) Schematic overview of the floral bud detection algorithm. Parameters subjected to optimization are denoted as χ_i next to the relevant step. **(b)** Step-by-step illustration of the workings of the detection algorithm. Numbers in the lower right corner of each sub-image refer to the corresponding step in (a). In the lower right sub-image, detected buds are displayed in a green overlay. The red arrows point indicate undetected buds, whereas the blue squares mark false detections.

Compared to the hyperspectral laboratory measurements (Wouters et al., 2013), a new discriminant model was required in order to account for the difference between the optimal wavebands and the passbands of the actual optical filters used during the field trials (Table 1). Furthermore, next to bud scales, floral parts and branches, an additional component (i.e. the bamboo supports) needed to be included in the new model since it was a prevalent feature in the captured scenes.

To quantify the effect of the difference between the optimal and actual wavebands, the analysis described by Wouters et al. (2013) was repeated using the properties of the actual wavebands as input.

Since the CCA procedure was only applied on the four main components, pixels belonging to other types of objects are necessarily assigned to one of these groups as well. This decreased the quality of P. As a remedy, pixel observations are filtered based on the confidence intervals of each group in the discriminant space spanned by the first two discriminant functions (Fig. 4). The confidence intervals are calculated based on the covariance matrix of each group. The level of the confidence intervals χ_1 was considered as tuning parameter in the optimization (see section 2.5.3). All observations that do not belong exclusively to the "flower parts" confidence interval (green ellipse in Fig. 4) are rejected from P, i.e. their pixel value is set to zero. [Fig. 3 – step 1C]. Those pixels are colored black in Fig. 4, whereas the pixels considered to be floral parts are indicated in blue. The confidence intervals of the other main components shown in Fig. 4 (red ellipses) overlap heavily, because they were separated by the third discriminant function which is not shown.

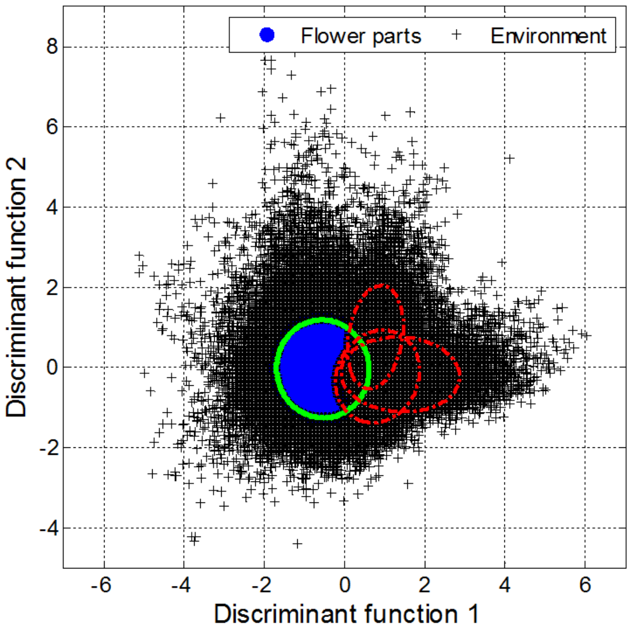

Fig. 4 Illustration of filtering by confidence interval. The graph shows the projection of all multispectral pixels of one scene into the discriminant space spanned by the first two discriminant functions. Illustrative confidence intervals for each main component are plotted on top as ellipses. Only data points exclusively inside the “Flower parts”-ellipse are retained for further analysis.

2.5.2. Morphological Image processing and segmentation

Subsequently, morphological image processing techniques are applied to P in order to interpret the information contained in the separate pixels and identify the floral buds as objects. As a reminder: pixel values close to 1 indicate a high likelihood of belonging to a flower part, whereas pixel values close to zero are considered not of interest. The algorithm carries out the following steps:

First, P is divided into equal-sized blocks Q_i of size χ_2 x χ_3 pixels [Fig. 3 – step 2A & Fig. 4] which are then processed individually. The goal of step 2A is to retain only the blocks Q_i with a high chance of containing (a part) of a floral bud. This yields two possible outcomes:

1. Q_i that are not considered of interest are not retained for further analysis and the value of all their pixels is set to zero.
2. Q_i of interest are retained. Furthermore, pixel values in these blocks that are set to zero in step 1C (filtering by confidence interval, i.e. removal of pixels outside of the green ellipse in Fig. 4) are restored to their original (non-zero) value. This is done because it improves the performance of the next steps of the algorithm by restoring some of the pixel information which was mistakenly removed in that filtering step 1C.

To determine which Q_i should be retained, the following procedure is carried out: For each block Q_i is checked whether the percentage of non-zero pixels is greater than the tunable value χ_4. If this is the case, Q_i is considered a region of interest (ROI). Otherwise Q_i is rejected from further analysis. In a next step, all retained Q_i are subjected to Otsu's adaptive thresholding (Otsu, 1975) to determine their local foreground. If the median value of the foreground pixels in Q_i is greater than χ_5, the block Q_i is retained for further analysis and its pixel values are restored as described above.

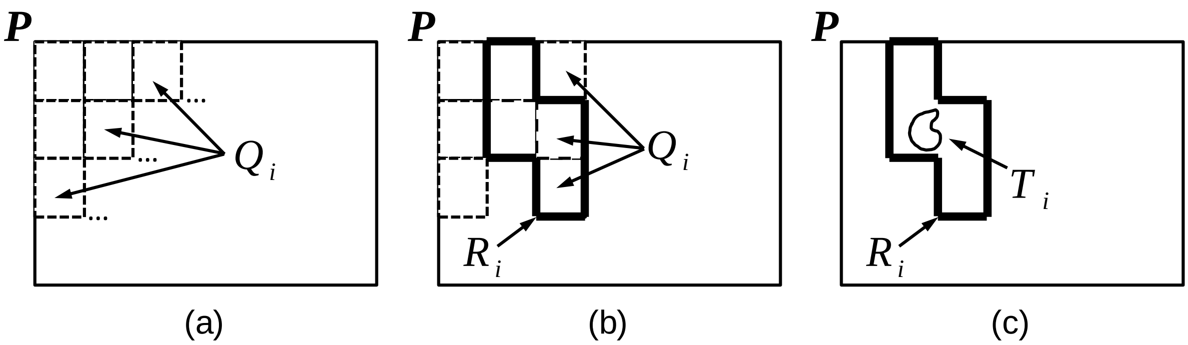

Fig. 4 Relation between the parameters Q_i., R_i and T_i. (a) Division of the probability image P into equal-sized blocks Q_i. (b) Adjacent Q_i which are considered ROI are grouped into a region R_i. (c) Inside R_i, an object T_i is located.

Summarizing, the procedure described above reduces the image P to a number of discrete non-zero regions $Q_{i,}$. Groups of adjacent Q_i that are retained are termed as regions R_i (Fig. 5). For the purpose of illustration, the edges of these R_i. have been marked in white in sub-image 2A of Fig. 3(b).

In a next step [Fig. 3 – step 2B], Otsu's thresholding is repeated separately in each R_i,. Since the R_i contain the information from a larger part of the recorded scene, a better separation between foreground and background can be realized than for each of the Q_i separately. After this operation, P is converted to a binary image, representing the foreground of each region R_i.

Finally, the geometric properties of the non-zero pixels/regions in P are analyzed. To merge (small) adjacent foreground regions, a closing operation is performed by means of a square χ_6 x χ_6 pixels structuring element [Fig. 3 – step 3A]. Regions of connected pixels are labeled as T_i (segmentation) [Fig. 3 – step 3B].

The equivalent diameter and the aspect ratio of each region T_i are calculated [Fig. 3 – step 3C]. The former is the diameter (in pixels) of a circle containing the same number of non-zero pixels and gives a measure of a region's size. The latter is the ratio of the minor axis over the major axis of the ellipse that encloses each region. All T_i with an equivalent diameter smaller than χ_7 are considered noise and consequently removed from P. Likewise, T_i with an aspect ratio smaller than χ_8 are deemed too long and narrow to be a floral bud and thus are removed from P as well. All remaining T_i are considered a floral bud.

2.5.3. Parameter optimization

In order to obtain a well performing detection algorithm, the parameters χ_i were tuned by an optimization procedure. Each unique combination of χ_i-values was called a set *{k},* in which k represented a unique index assigned to that set.

For any set $\{k\}$, the performance of the algorithm was described by means of quality assessment scores originating from information retrieval statistics (Manning et al., 2008). This was done because no true negatives could be defined, as is the case for many object detection problems. The recall ρ (or true positive rate) was defined as the fraction of floral buds which were correctly detected (completeness of detection). The precision π (or positive predictive value) represented the fraction of detections that were in fact real floral buds (purity of detection). Both ratios are given by the following equations:

$$\rho = \frac{TP}{TP + FN} \tag{2}$$

$$\pi = \frac{TP}{TP + FP} \tag{3}$$

TP represents the number of correctly detected buds (true positives), while FP represents the number of false detections (false positives). FN is the number of undetected buds (false negatives). A high recall indicates that most of the floral buds were detected, while a high precision indicates a low false discovery rate. In order to have an efficient algorithm, it is clear that both ρ and π should have a value close to 1. However, this was not straightforward as maximizing one ratio typically tends to reduce the other. For this reason, a desirability index $D_{\{k\}}$ was introduced as a measure that combines the scores from both recall and precision (Derringer and Suich, 1980):

$$D_{\{k\}} = \rho_{\{k\}}^{w} \times \pi_{\{k\}}^{1-w} \tag{4}$$

$D_{\{k\}}$ was defined as the product between recall and precision and ranged between 0 (no detection) and 1 (perfect detection). The weight w ($\in [0,1]$) modifies the relative importance of either ratio.

Since it was not feasible to test the performance of the detection algorithm for all possible sets $\{k\}$, a multi-parameter optimization was conducted by means of an iterative procedure using a subset of the multispectral images as training set (see section 2.6). Each optimization step simultaneously targeted all the χ_i-parameters belonging to one of the three parts of the image analysis algorithm as shown in Fig. 3(b), being "pixel discriminant model", "foreground detection" and "geometric analysis". A range of possible values was chosen for each of the χ_i-parameters of the part under consideration, e.g. χ_2 to χ_5 for the part "foreground detection". The algorithm was then run for all combinations of these parameter values, while keeping the values of the parameters of the other two parts constant. The combination of parameter values which yielded the largest $D_{\{k\}}$ was assigned to the relevant parameters. Then, the procedure was repeated for the parameters of the next part of the detection algorithm. This process continued cyclically until no further improvement of $D_{\{k\}}$ could be realized, resulting in the optimal set *{k}*.

Initial estimates and admissible ranges of the χ_i-parameters were chosen within reasonable bounds that were assessed based on the properties of the objects contained in the spatial-spectral database described in section 2.4. At the start of the optimization procedure, the range of the parameters was chosen quite broad as to encompass a large range of possible settings. As the procedure converged closer to the largest $D_{\{k\}}$, the range was chosen incrementally finer. The weight *w* in formula (4) was assigned a value of 0.5, as both recall and precision were considered equally important.

2.6. Training and validation

The performance of the detection algorithm was validated in two ways. First, the multispectral images captured during the first growing season were subjected to a three-fold holdout cross-validation (*type A*). In this analysis, the multispectral images recorded in the first season were divided by phenological stadium into three groups of approximately 15 images each. For each of the three iterations of the cross-validation, two of the groups were used (together) as the training set for the optimization of the detection algorithm, whereas the remaining group was used for validation.

The second way of validation (*type B*) was conducted in a similar fashion, but now all data of the first and second growing season were used as training and validation set, respectively.

For each training set, the CCA procedure [Fig. 3 – step 1A] was based on a subset of the spatial-spectral database. This subset was created by sampling 200 random (multispectral) pixels per object of the main components included in the training set. This resulted in datasets which included about 7 times more pixels of floral buds than of each of the other three main components. This ratio is representative for the relative occurrence of each main component in the recorded scenes. Finally, the bud recognition performance of the algorithm was further investigated to find out the cause for false detections and undetected floral buds.

3. Results

3.1. Pixel classification model

Comparison between the actual and optimal wavebands showed that the use of the 'non-optimal' wavebands only reduces the pixel classification accuracy by approximately 0.5 %. So, the effect of choosing the commercially available filters instead of the optimal wavebands can be considered negligible.

3.2. Parameter optimization

For all training sets, the optimization procedure gave very similar optimal parameter sets (values not shown). Only the selected values of the parameters χ_2 and χ_3 [see Fig. 3(a)] varied slightly between training sets, i.e. with relative differences of less than 10%.

As an illustration, the results of the optimization procedure conducted on the training set of the type B (inter-season) validation are shown in Fig. 6. The obtained precision π and recall ρ were plotted for each unique parameter set *{k}* investigated during optimization. The point {ρ =84.14% , π= 85.34%} is shown as a purple cross in Fig. 6 and was found to have the largest desirability index according to formula (4), i.e. its position was closest to the optimum {ρ =100% , π= 100%}. For the type A validation, similar graphs were obtained.

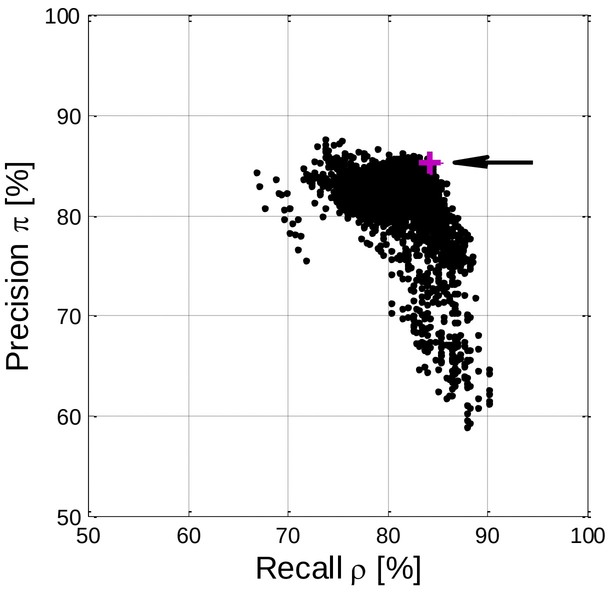

Fig. 6 Outcome of the optimization procedure for the type B validation. Each point on the graph represents the classification results for a unique combination of χ_i-values. In total, 3935 points are plotted. The arrow and purple cross mark the location of the point with the highest desirability index.

3.3. Floral bud detection

The results obtained for the two types of validation are summarized in Table 2. The results are grouped per phenological stadium. For clarity, the false discovery rate (*FDR*) is listed as well. This number represents the number of false detections relative to the number of actual floral buds.

Table 2 Overview of the classification performance of the detection algorithm.

		Training				**Validation**			
validation	phenological Stadium	# floral buds	recall ρ [%]	precision π [%]	*FDR*[x] [%]	# floral buds	recall ρ [%]	precision π [%]	*FDR*[x] [%]
Type A[a]	green cluster	250	80.00	86.21	12.80	125	82.40	78.63	22.40
	green bud	196	82.65	72.00	32.14	98	81.63	76.92	24.49
	white bud	260	84.62	78.57	23.08	130	84.62	82.71	17.69
	Total	706	82.58	79.14	21.77	353	83.00	79.62	21.25
Type B[b]	green cluster	125	84.00	88.24	11.20	44	68.18	100.00	0.00
	green bud	98	83.67	80.39	20.41	122	71.31	94.57	4.10
	white bud	130	84.62	86.61	13.08	149	85.91	76.65	26.17
	Total	353	84.14	85.34	14.45	315	77.78	84.78	13.97

a: three-fold cross-validation: for each row, the stadium shown in the 2nd column was used for validation

b: training was performed on all stadia of season 1 combined, validation was done on all stadia of season 2

x: false discovery rate is defined as the ratio of false detections to the total number of real buds, i.e. $FP.(TP+FN)^{-1}$

For the type A validation, similar results were obtained for both training and validation. The detection algorithm was able to correctly recognize approximately 83% of the floral buds. The average *FDR* was 22%. For the type B validation, all the data of the first season were included in the training set improving the classification results – for the training set – slightly, i.e. a recall of 84% and an *FDR* of 14%. For the validation set, the recall value was somewhat lower at 78%, but still a low *FDR* of 14% was realized. Most of these false detections originated from the scenes recorded during the

"White bud" stadium due to the occurrence of opening leaf buds. In all, similar classification results were obtained for all corresponding training and validation sets. This indicates that no overfitting was present.

In Fig. 7 and Fig. 8 the performance of the detection algorithm is illustrated for both growing seasons. These figures show a fake color image of each scene which was obtained by combining the multispectral images of the wavebands 925-975 nm, 685-700 nm and 755-805 nm. For reference purposes, the classical RGB image is shown as well. Next to this, the (posterior) probability images *P* resulting from the CCA procedure are displayed. Finally, the floral buds detected by the algorithm are shown as a green overlay on a grayscale image of each scene. In these figures, false detections (red arrows) and undetected buds (blue squares) are marked as well. It should be noted that these figures do not show the 'best' results (i.e. perfect detection), but were chosen to illustrate the different types of detection errors (see section 3.3).

In the scene shown in Fig. 7, 8 out of 13 buds were correctly detected next to 2 false detections. In the scene depicted in Fig. 8, 7 out of 10 buds were found with no false detections. The time required by the algorithm for processing one multispectral image was 5 to 6 seconds.

3.4. Detection errors

Since the algorithm was not able to provide perfect classification, the causes for both false detections and undetected buds were further investigated. The main reason for not detecting buds was because they were completely or partially hidden from the camera behind other tree components or were located at the edge of the recorded images (partial occlusion, arrow 1 in Fig. 7 and Fig. 8). Consequently, these buds were not or badly visible in the multispectral images. Some buds were badly illuminated (shadowed) and therefore produced reflectance values which could not be

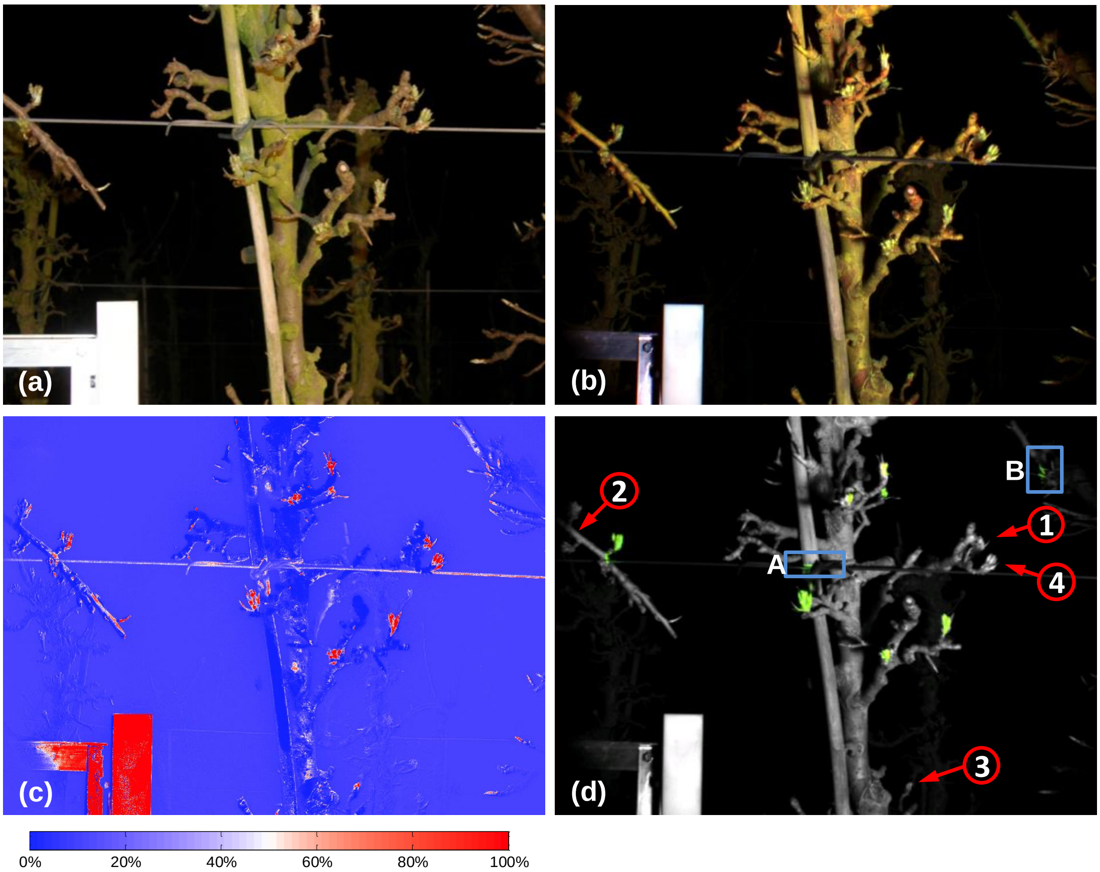

Fig. 7 Bud recognition in a scene captured on March 26th, 2012 during the "Green cluster" stadium. **(a)** Ground truth RGB image, **(b)** Fake color image created from the multispectral data, **(c)** Posterior probability image, **(d)** Floral bud detection: recognized buds are shown as a green overlay. The red arrows mark undetected buds. The blue squares mark false detections.

classified by means of the CCA procedure (arrow 2 in Fig 7. and Fig. 8). Buds that were too small (arrow 3 in Fig. 7) or were situated in a too noisy region of the probability image P were filtered out by the detection algorithm. Undetected buds which could not be classified into any of these categories were termed as an "artifact" (arrow 4 in Fig. 7).

Most of the false detections could be attributed to the occurrence of (large) leaf buds (square B in Fig. 7), especially during the "White bud" stadium when these buds started to open. Other false detections were caused by noise in the probability image P. Finally, in a few cases, parts of the white PTFE reference, plastic wires (square A in Fig 6.) or floral buds located in the background were falsely

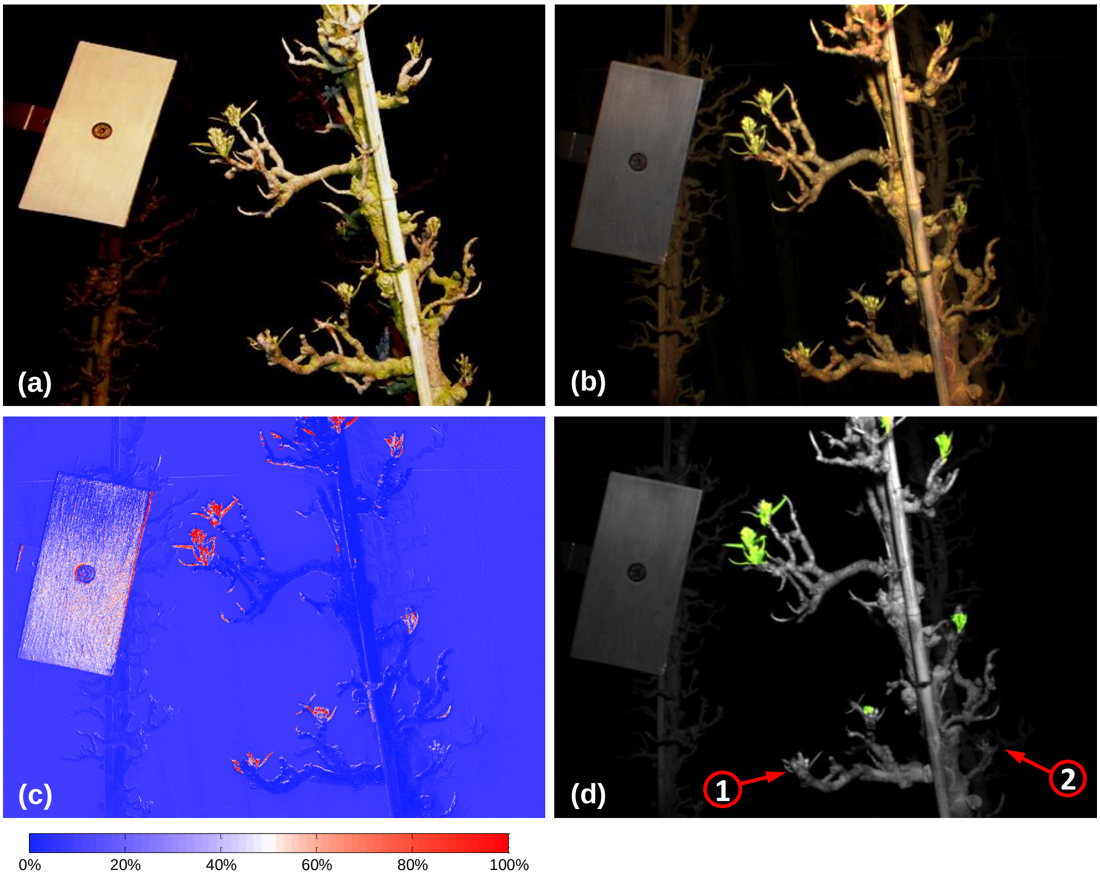

Fig. 8 Bud recognition in a scene captured on April 18th, 2013 during the "Green bud" stadium. **(a)** Ground truth RGB image, **(b)** Fake color image created from the multispectral data, **(c)** Posterior probability image, **(d)** Floral bud detection: recognized buds are shown as a green overlay. The red arrows mark undetected buds.

classified as foreground buds.

The small changes made to the detection setup (discussed in section 2.1) had a positive effect on the data recorded during the second season. For this reason, the performance of the detection algorithm was reevaluated by no longer taking into account those false detections observed during the first season which were caused by effects easily avoided in the second season, e.g. false detections because of the PTFE reference. Additionally, the buds which were obscured from view (occluded) were omitted from the algorithm's performance reassessment as well. In this way, the recalculated results reflect the detection algorithm's capability to identify unobscured buds. In Table 3, it is illustrated that under the aforementioned assumptions, a recall value around 87% was achieved for both training and validation.

Table 3 Reevaluation of the classification results displayed in Table 2. The results displayed here do no longer take into account occluded buds and false detections related to causes which could be easily avoided, as was done in the second season.

	Training			**Validation**		
	recall ρ [%]	precision π [%]	*FDR*[x] [%]	recall ρ [%]	precision π [%]	*FDR*[x] [%]
Season 1[a]	86.72	83.29	17.40	87.20	85.42	14.88
Season 2[b]	88.39	90.83	8.92	86.88	84.78	15.60

a: type A validation: results were averaged over the three training-validations

b: type B validation

x: false discovery rate is defined as the ratio of false detections to the total number of real buds, i.e. $FP.(TP+FN)^{-1}$

4. Discussion

The detection algorithm was able to detect a high percentage of the floral buds visible to the camera with a low number of false detections. The results displayed in Table 2 and Table 3 attest to the robustness of the detection since good results were obtained for both type A (inter stadium) and type B (inter season) validation.

Although work on the development of other camera systems for flower detection has been reported, a performance comparison was not possible as no clear performance results have been reported for these systems. Moreover, these systems all target blossoms which look very different from floral buds prior to bloom.

As can be observed in the RGB images (Fig. 7 and Fig. 8) floral bud detection from these images is absolutely not obvious. This is due to the presence of a large amount of green in the canopy (algae) and the irregular shape of the trees. In this context, the superiority of multispectral imaging can already be seen in the fake color RGB images (Fig. 7 and Fig. 8) where the green algae on the branches are no longer visible.

A disadvantage of the technique described in this work is the need for an optical reference for image normalization. This might be resolved by means of an optical power meter that measures the average light intensity coming from a scene. The signals obtained can then be used for image normalization.

4.1. Image analysis

The algorithm was designed to first process (multispectral) color information, only then followed by shape analysis. This was done because typically color based segmentation is easier than shape based segmentation (Nielsen et al., 2012). As expected, measuring at nighttime was beneficial for the image quality as the visibility of background objects reduced rapidly with increasing distance. However, due to the characteristic shape of the trees, it was difficult to provide a good illumination of all floral buds without shadows with a single lamp. Pixels that are positioned in a shadowed region typically were assigned a lower posterior probability than well-lit pixels. Applying stronger illumination may partially resolve this issue. However, this might unfavorably increase the rate of floral buds detected in the background. A better approach would be to design a dedicated illumination unit which provides a more uniform illumination of the tree. Nevertheless, most of the shadowed floral buds still contrasted clearly with the background in the probability image *P*. For this reason – as explained in section 2.5.2 – the detection algorithm included Otsu's thresholding to

separate the local foreground and background by means of block processing [Fig. 3 – Step 2A]. This technique is particularly well suited for dealing with local variations in illumination quality.
In an earlier version of the detection algorithm, the shape analysis was conducted directly after step 2A. However, it was found that implementing a second step (2B) of local floral bud/environment separation drastically improved bud recognition. This second step takes into account larger regions of connected blocks of the image instead of separate, independent blocks [Fig. 3(b) – step 2A]. Therefore, a more accurate threshold could be calculated, retaining a higher number of floral bud pixels.

Currently, the time required to process a single multispectral image is too long for a real-time implementation of the algorithm. This was expected, as no efforts have been made with respect to the computational efficiency of the algorithm. Therefore, it is expected that the processing speed can be significantly increased by applying a more speed-efficient code combined with imaging at a lower resolution. It is expected that the number of pixels can be safely reduced by a factor of 4 without losing performance.

4.2. Detection errors

The general quality of the recorded images was lower in the first season which resulted in noisier probability images. This was mainly due to the handheld illumination which caused slight variations in illumination between measurements at different exposure times. In addition, exposure times were set manually in the first season. This made the recorded images more susceptible to motion blur because of wind. These issues were avoided in the second season by fixing the position of the light source and implementing an auto exposure function in the software. This improved image quality and compatibility between measurements recorded at different exposure times.

Notwithstanding, the lower image quality of the first season had an effect on the type B validation. Because the detection algorithm was trained on the noisier dataset of season 1, a set of χ_i-parameters was selected which best dealt with this noise. More specifically, the parameter χ_2 relates to the noise level allowed in the probability images. Since the images of the second season contained less noise, using a less strict threshold value for χ_2 allowed to increase the recall rate shown in Table 2 up to 82% with only a small reduction in precision (less than 4%).

4.2.1. Undetected buds

The main reason for not detecting floral buds was occlusion, especially during the second season. In total, about 6% of the captured buds were occluded. The algorithm developed by Nielsen et al. (2012) filtered out approximately 18% of blossoms due to occlusion. Presumably, this higher percentage might be caused by the larger spatial volume occupied by blossoms compared to floral buds. This makes it more likely for a part of the blossoms to be obscured from view by others.
Irrespective of the detection system used or better control of the tree shapes by pruning methods (Schupp and Baugher, 2011), a part of the floral buds is expected to be occluded at any rate due to their semi-random location on the trees. These buds should be thinned by additional manual follow up thinning, which is recommended even after mechanical thinning (Schupp et al., 2008).

4.2.2. False detections

The predominant reason for false detections originated from the presence of leaf buds in the orchards. When leaf buds start to open they look similar to floral buds. Most of these false detections occurred during the “White bud” stadium when the leaf buds were largest. This effect was strongest during the second season, because measurements were conducted until very late in the “White bud” stadium, close to bloom. Besides, it should be noted that the detection algorithm still filtered out most of the leaf buds by either applying the confidence interval in the discriminant space [Fig. 3 – step 1C] or the shape analysis [step 3C]. Further reduction of leaf bud detection might be obtained by including these as a separate group in the CCA procedure.

4.3. Potential of the detection system

It is expected that this detection system will provide a valuable tool to improve the performance and selectivity of mechanical thinners. For instance, measuring the floral bud distribution before and after thinning immediately provides feedback to the growers about the efficacy of the thinning procedure and tells them where and to which degree manual follow up thinning might be required.

Some additional improvements are still required to transform the research setup described in this work to a stand-alone sensor platform. The two most prominent issues are the implementation of a faster, real-time version of the algorithm and a faster camera system. The latter could be achieved by implementing a multi-CCD camera or a camera chip equipped with a Bayer-like filter in which the pixels are sensitive to the specific wavebands used in this research. In this context, it would also be interesting to investigate the performance of the sensor when less wavebands are applied. Reducing the number of wavebands not only reduces the complexity, but also the cost price and required computing power.

It is expected that the detection algorithm can also be used during the stadium “Mouse Ear” which takes place prior to the stadia investigated in this research. In this phenological stadium the green floral parts become visible for the first time and have a similar outlook as during the “Green cluster” stadium. In total, this would give the multispectral floral bud detection a usable time period of approximately two weeks.

Finally, the potential of the multispectral sensor is not limited to automated thinning alone. It could also be used for other applications such as variable rate spraying or early yield detection. The latter is typically done by manually counting the number of floral buds on a few sample trees and extrapolating this result over the entire orchard. As this often gives a false sense of the crop load, more correct information would aid growers in deciding which horticultural measures are required. However, more research is required to study the actual relation between the number of floral buds and the final yield, because this relation depends strongly on seasonal conditions. Carbon balance models that link environmental conditions to tree status and final yield may be a useful tool in this context (Robinson and Lakso, 2011).

5. Conclusions

In this work we demonstrated the feasibility of detecting floral pear buds during the early phenological stadia by means of a multispectral camera system. A custom image analysis algorithm was developed which was able to detect approximately 87 % of the (unoccluded) floral buds with a low false detection rate (< 16 %). It is expected that the detection algorithm's performance can still be further increased by tuning its parameters specifically for each phenological stadium. This especially applies to the "White bud" stadium where the development of the leaf buds resulted in an increased rate of false detections. Furthermore, it should be investigated whether good floral bud detection can also be achieved with the multispectral sensor during daytime conditions. This floral bud counting sensor could be used for automated thinning, variable rate spraying and yield estimation.

Acknowledgements This research has been funded by the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen, project 080497). Bart de Ketelaere is funded by the Industrial Research Fund. The authors gratefully acknowledge Proefcentrum Fruitteelt vzw and Fruitbedrijf Van der Velpen for their cooperation in this research. In addition, we also wish to thank Dirk Leroy, Frank Mathijs, Robbe Van Beers, Hannelore Bertels, Koen Rutten and Chyngyz Erkinbaev for their assistance during the field trials.

References

Bac, C., Hemming, J. and van Henten, E. (2013). Robust pixel-based classification of obstacles for robotic harvesting of sweet-pepper. *Computers and Electronics in Agriculture, Elsevier, 96*, 148-162.

Baugher, T. A., Ellis, K., Remcheck, J., Lesser, K., Schupp, J., Winzeler, E. and Reichard, K. (2010). Mechanical string thinner reduces crop load at variable stages of bloom development of peach and nectarine trees. *HortScience, American Society for Horticultural Science, 45*, 1327-1331.

Bertschinger, L., Stadler, W., Stadler, P., Weibel, F. and Schumacher, R. (1998). New methods of environmentally safe regulation of flower and fruit set and of alternate bearing of the apple crop. *Acta Hort., 466*, 65-70.

Bulanon, D., Burks, T., and Alchanatis, V. (2010). A multispectral imaging analysis for enhancing citrus fruit detection. *Environment Control in Biology, 48*, 81–91.

Derringer, G. and Suich, R. (1980). Simultaneous optimization of several response variables. *Journal of Quality Technology, 12*, 214-219.

Emery, K. G., Faubion, D. M., Walsh, C. S., and Tao., Y. (2010). Development of 3-D range imaging system to scan peach branches for selective robotic blossom thinning. ASABE Paper number 10-09202. St. Joseph, MI: The American Society of Agricultural and Biological Engineers.

Gebbers, R., Pflanz, M., Betz, A., Hille, B., Mattner, J., Rachow-Autrum, T., Özyurtlu, M., Schischmanow, A., Scheele, M., Schrenk, J. and Schrenk, L. (2013). OptiThin – Implementation of

Precision Horticulture by Tree-Specific Mechanical Thinning. *33. GIL-Jahrestagung Massendatenmanagement in der Agrar und Ernährungswirtschaft Erhebung Verarbeitung Nutzung.*

Glozer, K. and Hasey, J. (2006). Mechanical thinning in cling peach. *HortScience, American Society for Horticultural Science, 41*, 995.

Hehnen, D., Hanrahan, I., Lewis, K., McFerson, J. and Blanke, M. (2012). Mechanical flower thinning improves fruit quality of apples and promotes consistent bearing. *Scientia Horticulturae,134*, 241-244.

Hong, Z. (2010). Investigation of regulatory mechanisms of chemical-mediated fruit thinning in apple (Malus x domestica Borkh.). PhD thesis. Virginia Polytechnic Institute and State University.

Kon, T. M., Schupp, J. R., Winzeler, H. E. and Marini, R. P. (2013). Influence of Mechanical String Thinning Treatments on Vegetative and Reproductive Tissues, Fruit Set, Yield, and Fruit Quality of ‘Gala’Apple. *HortScience, American Society for Horticultural Science, 48*, 40-46.

Kviklys, D. and Robinson, T. (2010). Temperature before and after Application of Chemical Thinners Affects Thinning response of ’Empire’ Apple Trees. *Acta Hort. (ISHS) 884*, 525–530.

Link, H. (2000). Significance of flower and fruit thinning on fruit quality *Plant growth regulation, Springer, 31*, 17-26.

Lopez, G., Larrigaudière, C., Girona, J., Behboudian, M. and Marsal, J. (2011). Fruit thinning in 'Conference' pear grown under deficit irrigation: Implications for fruit quality at harvest and after cold storage. *Scientia Horticulturae, Elsevier, 129*, 64-70.

Maas, F. and van der Steeg, P. (2011). Crop Load Regulation in 'Conference' Pears. *Acta Hort. (ISHS), 909*, 367-379.

Manning, C.D., Raghavan, P. and Schütze, H. (2008). *Introduction to information retrieval.* Vol. 1. Cambridge University Press Cambridge.

Martin-Gorriz, B., Torregrosa, A. and Garcia Brunton, J. (2012). Post-bloom mechanical thinning for can peaches using a hand-held electrical device. *Scientia Horticulturae, 144*, 179-186.

Martin-Gorriz, B., Torregrosa, A. and Garcia Brunton, J. (2011). Feasibility of peach bloom thinning with hand-held mechanical devices. *Scientia Horticulturae, Elsevier, 129*, 91-97.

Martin-Gorriz, B., Torregrosa, A. and Garcia Brunton, J. (2010). Post-bloom thinning of peaches for canning with hand-held mechanical devices. *Scientia Horticulturae, Elsevier, 125*, 658-665.

Meland, M. (2009). Effects of different crop loads and thinning times on yield, fruit quality, and return bloom in Malus x domestica Borkh.'Elstar'. *J. Hortic. Sci. Biotech., Special Issue*, 117-121.

Miller, S., Schupp, J., Baugher, T. and Wolford, S. (2011). Performance of mechanical thinners for bloom or green fruit thinning in peaches. *Hortscience, 46,* 43-51.

Miller, S. and Tworkoski, T. (2010). Blossom thinning in apple and peach with an essential oil. *HortScience*, *45(8),* 1218-1225.

Ngugi, H.K. and Schupp, J.R. (2009). Evaluation of the risk of spreading fire blight in apple orchards with a mechanical string blossom thinner. *HortScience, American Society for Horticultural Science, 44*, 862-865.

Nielsen, M., Slaughter, D.C. and Gliever, C. (2012). Vision-based 3D peach tree reconstruction for automated blossom thinning. *Industrial Informatics, IEEE Transactions on, IEEE, 8(1)*, 188-196.

Okamoto, H., and Lee, W. (2009). Green citrus detection using hyperspectral imaging. *Computers and Electronics in Agriculture, 66,* 201–208.

Otsu, N. (1975). A threshold selection method from gray-level histograms. *Automatica, 11*, 23-27.

Peck, G. M. and Merwin, I. A. (2009). *A grower's guide to organic apples.* Cornell University Coop. Extension.

Robinson, T. and Lakso, A. (2011). Predicting chemical thinner response with a carbohydrate model. *Acta Hort. (ISHS)* 903, 743–750.

Rosa, U., Cheetancheri, K., Gliever, C., Lee, S., Thompson, J. and Slaughter, D. (2008). An electro-mechanical limb shaker for fruit thinning. *Computers and Electronics in Agriculture,* 1(2), 213-221.

Schupp, J.R. and Baugher, T.A. (2011). Peach blossom string thinner performance improved with selective pruning. *HortScience, American Society for Horticultural Science, 46*, 1486-1492.

Schupp, J., Baugher, T., Miller, S., Harsh, R. and Lesser, K. (2008). Mechanical thinning of peach and apple trees reduces labor input and increases fruit size. *HortTechnology, 18,* 660-670.

Sharma, S. (1995). *Applied multivariate techniques* (pp. 237–316). New York: John Wiley & Sons, Inc.

Stajnko, D. and Blanke, M. (2011). Yield prediction in fruit crops using image analysis. IX International Symposium on Integrating Canopy, Rootstock and Environmental Physiology in Orchard Systems, Geneva, New York

Theron, K. (2010). Size matters: Factors influencing fruit size in pear. *XI International Pear Symposium, 909*, 545-555.

Wallays, C., Missotten, B., De Baerdemaeker, J. and Saeys, W. (2009). Hyperspectral waveband selection for on-line measurement of grain cleanness. *Biosystems Engineering, Elsevier, 104*, 1-7.

Wouters, N., Van Beers, R., De Ketelaere, B., Deckers, T., De Baerdemaeker, J. and Saeys, W. (2014). *Quantification of detachment forces and single nozzle tests for pneumatic thinning of floral pear buds.* Manuscript submitted for publication.

Wouters, N., De Ketelaere, B., De Baerdemaeker, J. and Saeys, W. (2013). Hyperspectral waveband selection for automatic detection of floral pear buds. *Precision Agriculture, Springer, 14*, 86-98.

Yang, F.O. (2012). *Development of a table-top robot model for thinning of fruit “Master thesis”.* University of Illinois, Urbana, Illinois.