

Neural Networks and Low-Cost Optical Filters for Plant Segmentation

Floris De Smedt, Ive Billauws and Toon Goedemé

Lessius Mechelen Campus De Nayer, K.U.Leuven ESAT-PSI-VISICS

Toon.Goedeme@mechelen.lessius.eu

Abstract

In this paper, we describe a low-cost system for image segmentation in agricultural applications. In applications such as weed detection and classification, plants must be segmented from the image. Our novel approach combines low-cost NIR filters used on a standard camera with neural networks to achieve a significantly higher accuracy as compared to classic threshold techniques.

1. Introduction

The first step of a computer vision application is the selection of regions of interest. Computer vision in agricultural applications is no exception to this. The research of Sogaard et al., for instance, uses plant segmentation as the preprocessing step for camera guided agricultural equipment [5]. In this preprocessing step, distinction must be made between plants and background. In this paper, we compare different methods to do this segmentation and propose a novel technique with a better performance than the existing methods.

This paper is organized as follows. We begin with an overview of existing methods that use a linear combination of pixel values to detect the green of plants in the image (see section 2). In section 3, we describe our alternative neural network based method that makes the segmentation more robust to light variations. In section 4, the optical filters we used in our experiments are described. In order to make use of the NIR spectrum, we propose some less expensive methods without using an expensive industrial optical filter. We end with an overview of our experimental results (section 5), a conclusion and overview of further valorisation experiments.

2. Related work

Segmentation of plants is a common and important problem in computer vision in agricultural applications. Therefore much work has been done about this matter. Most of the techniques use a linear combination of the RGB-values of the pixels. The output of these techniques is a color index value which can be thresholded to make distinction between plant pixels and background.

Woebbecke et al. [8] used a linear combination that amplified the amount of green in images in their Excess Green Method (ExG):

$$ExG = 2 \times G - R - B$$

Plant pixels, with a high green value, will result in a higher ExG index and thus a bright pixel in the output image. The ExG method proved to work well, except in case of the presence of specific camera effects. An example of these effects is the higher saturation of pixel values on the border of leaves. Therefore Meyer et al. [4] added an extra constraint in their Modified Excess Green method. If the result of the following equation is true, the pixel value in the output image is set to the result of the excess green method function, else it is set to 0 (black):

$$G \geq R \wedge G \geq B \wedge G \geq 120$$

The Normalized Difference Vegetation Index (NDVI) is an index often used in remote sensing applications to measure the presence of vegetation. Healthy plants absorb most of the visible light, but they reflect a significant amount of Near Infrared (NIR). Based on these facts, the NDVI can be found using the following formula:

$$NDVI = \frac{NIR - R}{NIR + R}$$

This index could be used in other applications as well.

Based on the NDVI, Woebbecke et al. [7] defined an alternative in the visual spectrum: the NDI (normalized difference index), which uses the green and red channel, instead of NIR and red:

$$NDI = \frac{G - R}{G + R}$$

Another index is suggested by Marchant et al. [3]:

$$F = \frac{r_m}{g_m^A}$$

with

$$r_m = \frac{R}{B} \quad g_m = \frac{G}{B}$$

and A a constant based on the characteristics of the camera. Note that the (M)ExG, the NDVI and NDI indices represent plant pixels as bright pixels, while the F-index represents them as dark pixels

These methods have proven their functionality, but our experiments (see section 5) show that their performance is not optimal when changing light conditions are present. For this reason, we introduce a neural network based green detector in the following section

3. Neural network

In contrast to the methods described above, we chose to use a neural network for this segmentation task. With enough training material at hand, one can let a neural network train the optimal combination of the three color bands. The power of a neural network lies in the robustness to variations of conditions (i.e. change of light, variations of color). If the neural network is trained on a data set which contains images of all types of conditions, it will be able to face these conditions [2]. In analogous applications such as skin color segmentation, neural networks have shown their potential [1,6].




Figure 1: The neural network used in our experiment

We use neural networks to make a difference between plant pixels and background pixels, independently of the colors of the plants/background. We have to train the neural network once for each type of input-data (images in the visual spectrum, images of the NIR-spectrum and images on the border between these two). The neural network we used has 2 layers (fig. 1), with in the first layer, the 3 RGB-values of a pixel as input to each neuron. Each input value is multiplied by a weight value that is assigned to the input signal of a neuron. The outputs of the neurons in this layer are sent to the neuron in the output-layer. The activation-function of the neurons is a sigmoid function and the training technique is back-propagation. The back-propagation techniques uses the derive of the activation function to adjust the weight values of the neural network, so we have to use an activation function that is differentiable.

To create the training data, we use a mask which selects the plant-pixels and background-pixels. The input of the network are the 3 RGB-values, the output is dependent on the type of pixel. For a background-pixel we used a value of 0.1 for a plant-pixel we used a value of 0.9.

4. NIR imaging

Most cameras have CCD ICs which are sensitive to the entire visual spectrum (400nm-750nm) and the lower part of the Infrared spectrum, i.e. Near Infrared (NIR, 750nm-1400nm). Figure 2 shows the spectral sensitivity of our camera (ImagingSource DBK 21BF04-Z2). The NIR frequencies detected by the camera results in a RGB-value, dependent on the received frequency. This NIR sensitivity often results in unwanted effects in the resulting image, so usually an NIR filter is placed in front of the CCD chip. However, plants reflect a significant amount of NIR, while ground absorbs most of the NIR light (fig. 3). This indicates that cameras without NIR filter could potentially lead to a better plant-background segmentation.

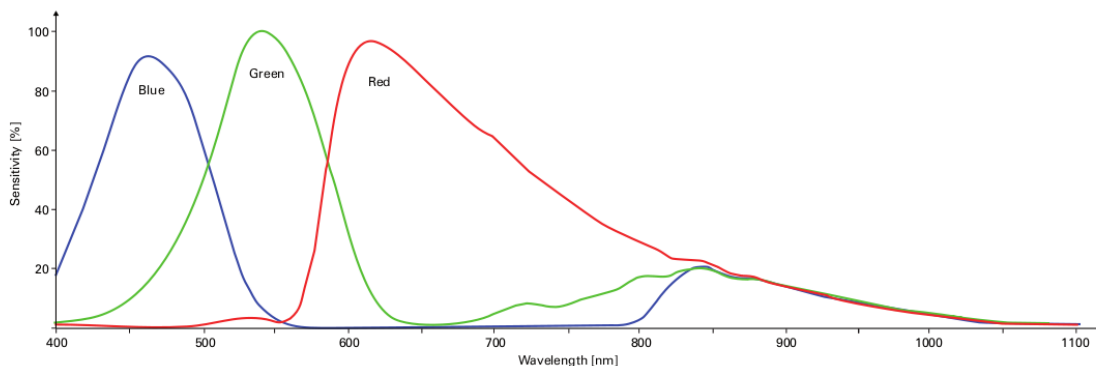


Figure 2: Spectral sensitivity diagram of Sony ICX098BQ chip

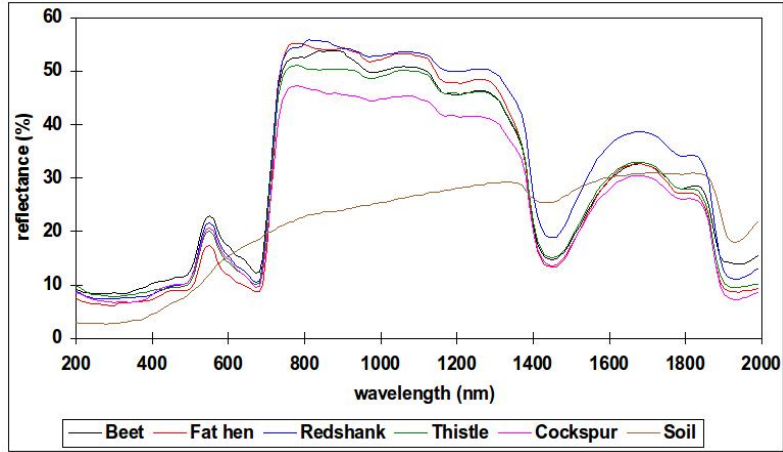


Figure 3: Spectral sensitivity of different plants and ground-types




Figure 4: Spectral Diagram of the Hoya RT-830 filter

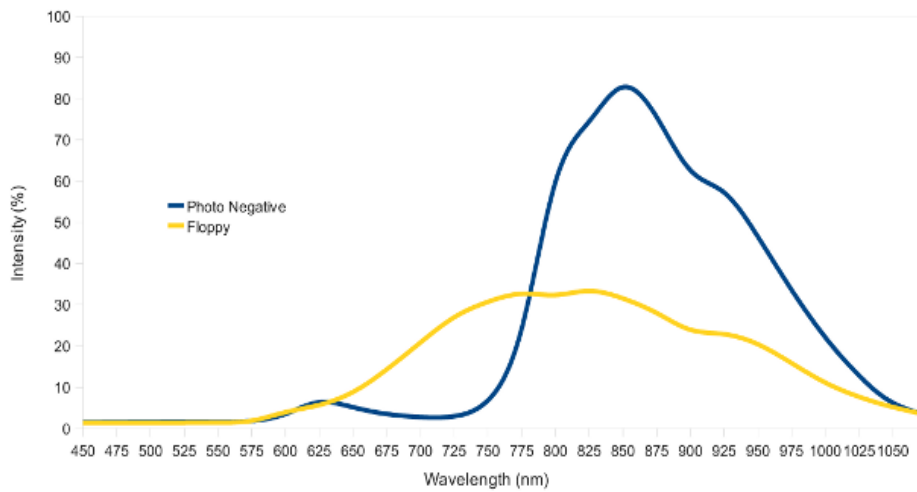


Figure 5: Spectral diagram of a photo negative and a floppy disk

For determining the effect of this NIR sensitivity, we have experimented with several filters. To compare the results of less expensive filters with an industrial solution, we used a bandpass filter of the NIR-spectrum. The filter used is a Hoya RT-830 with a center wavelength of 830 nm and a full width-half max wavelength of 260 nm (fig. 4).

Because this paper focuses on the development of a cost-effective method, we looked also into other NIR filter possibilities. A black photo film negative eliminates most of the visual spectrum and let a part of the NIR-frequencies pass. Also, the internal magnetic material of a 3 1/2" floppy disk has this property (fig. 5). The broader range of the floppy disk spectrum results in a higher presence of red color in the image, although the overall transparency is lower.

5. Experimental results

Figure 6 shows some results of the different plant segmentation algorithms described in section 2. On randomly selected pixels of his visual spectrum plant image, we trained our neural network, which resulted in the far right segmentation results. The actual purpose of the training process of the neural network is to search for the most optimal linear combination of the RGB-values of a pixel. The resulting linear combination depends on what we select as plant pixels and what we select as background pixels, so this process is independent of the objects or source (NIR frequencies or visual spectrum frequencies) we want to segment. It is clear that our NN approach outperforms the other methods. The high detection rate of (M)ExG and NDI in these images is caused by the (almost) perfect green of the leaves and the high difference between plant pixels and background. We also can see that the red color of the stalk is not detected by the segmentation methods based on green pixel values. Figure 7 shows ROC curves for the different approaches we derived based on manually labelled ground truth pixels. This curve also shows clearly the superior performance of our NN as compared to the other approaches. The choice of the used threshold is based on this ROC curve (fig 7), as the value resulting in the smallest Euclidean distance to the most perfect point (0,1), i.e. all pixels correctly classified.

With real collected weed plants, we created a data set of different types of weeds and different types of ground. The presence of different types of weed is important so the neural network is not trained for a single type of weed, but can be used on fields where different types of weed are present. It is also important that different types of soil are present in the images, because a segmentation is based

on the difference in light reflection between ground and




Figure 6: Segmentation methods, from top to bottom: Modified ExG, NDI, Marchant and Neural Networks

plants. The more the dataset varies, the more complex it is to train a neural network to do a correct segmentation. The images of the dataset are taken on a ground of stone. The absence of extra plants makes it easier to create a correct mask that defines the plant pixels and ground pixels in the dataset. The presence of the stone ground will not enhance the results of this paper since the presence of a extra type of soil will make the image more difficult to segment. A set of 8000 pixels (approximately 2.6% of the number of pixels in the image), randomly chosen from the odd pixels in the image is used for training. The same image is used to test the correctness of the learned segmentation method.

Table 1. Comparison between existing segmentation methods and neural network

	Correct classified pixels (%)	Plant pixels found (%)	Ground pixels found (%)
Modified ExG	96.2210	90.9946	98.6880
NDI	92.4784	88.1333	94.5251
Marchant	93.7070	89.8840	95.5116
Neural network	97.2912	96.6627	97.5879

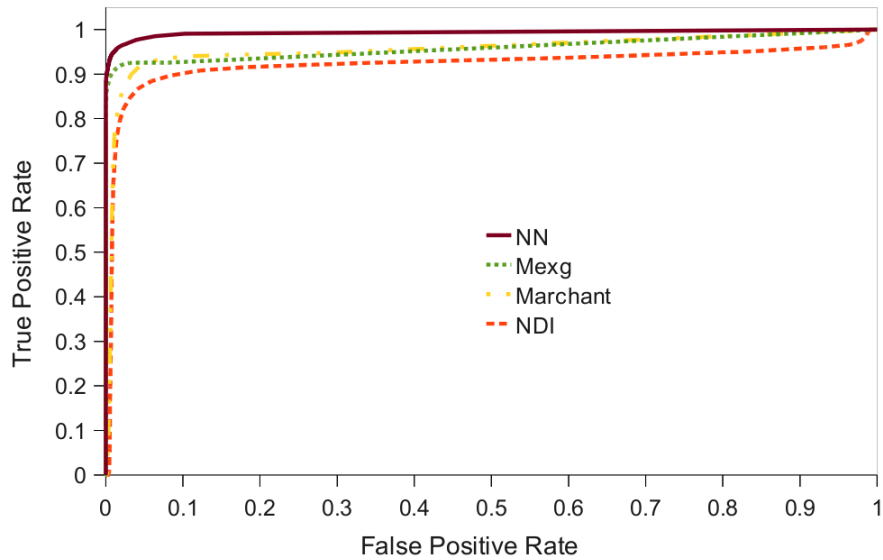


Figure 7: ROC-curve of vegetation indexes and Neural Network (NN)


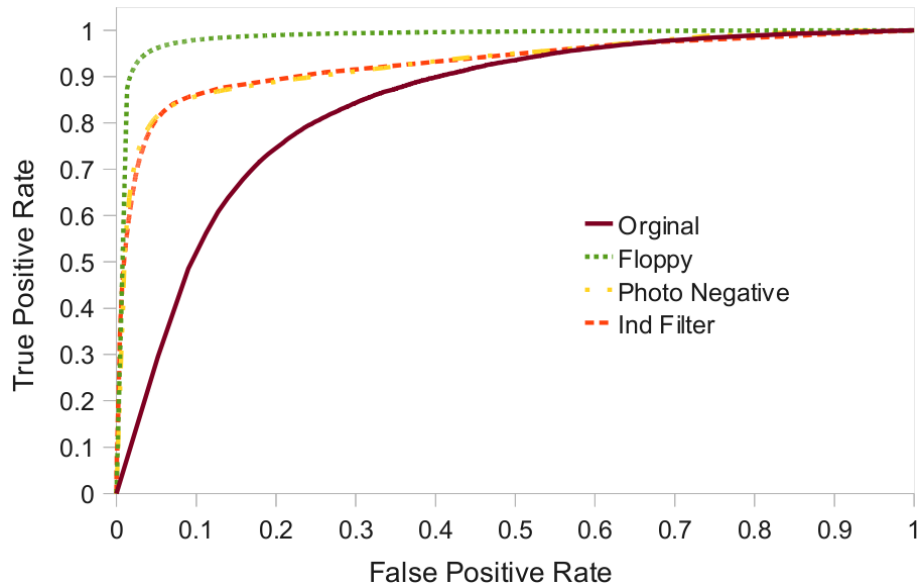


Figure 8: Results neural networks, on top the used image, on the bottom the segmentation result. From left to right: Original (no filter), industrial filter, photo negative and floppy.

Table 2. Comparative results for the different filters

	Correct classified pixels (%)	Plant pixels found (%)	Ground pixels found (%)
Original	87.1318	30.1099	96.2178
Industrial filter	93.2943	71.2531	97.0237
Photo negative	93.5182	80.1003	95.8016
Floppy	97.0697	91.3971	98.0355

**Figure 9: ROC-curve of Neural Networks with different filtered input images: original (no filter), industrial filter, photo negative and floppy disk**

The results of the Neural Network-based segmentations are shown in figure 8. Also here we have chosen a threshold value based on an ROC curve (fig 9). Table 2 summarises the results of these experiments as compared to hand-labelled ground truth data. The ROC curve and the example segmentation results show that surprisingly the floppy disk as filter yields the best recognition rate. This can be explained by the fact that it is still a bit transparent to visible light, and therefore combines the properties of NIR imaging and visible light.

6. Conclusion

In this paper, we proposed a novel method for the reliable segmentation of plant images. We compared the traditional methods, which are based on the green color of plants, with trained neural networks, which are color independent, and we concluded that the use of neural networks makes an improvement in detection rate. Next, we showed that including the NIR-spectrum can be very useful in the segmentation of plants and background. We

compared different types of filters of the NIR-spectrum and we can conclude that the use of a partial NIR filter, selecting a combination of NIR and a piece of the visual spectrum (red), gives the highest correct segmentation percentage. Moreover, our experiments with less expensive optical filter materials prove that something as simple as floppy disk plastic yields surprisingly good results, even better than expensive industrial-grade optical filter material.

7. Future work and experiments

The method described in this paper defines a solid base for further research on vision-based agricultural applications. In this project, we applied the segmentation technique described above for two cases: the growth measurement of lettuce and weed detection in Brussels chicory seedlings.

Some important diseases in lettuce are shown to be correlated to the growth speed of the plant. When the plant grows too quickly, the amount of absorbed water is

too large for the plant cells to contain. A too fast growth rate has therefore the risk of forming grey spots on the leaves, the result of burst open cells. Figure 10 illustrates this.



Figure 10: Rind disease in lettuce due to an overly fast growth rate

Because critical consumers favour spotless vegetables, the forming of the above described rind disease must be avoided. Therefore, we carried out experiments at the *Proefstation voor de Groenteteelt* in Sint-Katelijne-Waver, Belgium. An overhead placed camera is mounted so that it can be moved over the entire lettuce hydroculture plantage collecting images of the lettuce plants (figure 11). In the growth season, we acquired these image data sets twice a week.




Figure 11: Example collected lettuce image

In figure 11, the result of our segmentation approach presented in this paper is shown. The lettuce plants are clearly segmented from the background. On these segmentation results, we measured the size of each lettuce

plant at each camera run. The resulting values were collected and averaged per plant zone in order to deal with measurement noise. Our agricultural research partners are examining these growth results further in order to make conclusions.




Figure 12: Segmentation result of fig. 11

In a second test case, we focused on the detection of weed by combining plant-detection (and classification) and row detection. The purpose is to only use herbicides on the weeds and not on the desired crop plants. The crop plants are on rows, so in theory each plant that deviates from this row can be considered to be weed.



Figure 13: Test field for Belgian chicory, showing our camera mounted on an agricultural machine

We performed experiments on the test field of the *Nationale proeftuin voor de Witloofteelt* in Herent, Belgium (fig. 13). The test was timed two weeks after seeding of the Belgian chicory plants. Most of the plants were grown out by that time to two-leaf or four-leaf seedlings, along with spontaneously upcoming weed with about the same size. Our camera including floppy disk filter was mounted on an agricultural machine in such a

way that it had a approximately constant distance to the plant bed. We acquired a large set of images over a distance of 400 m on different places on the field.

On this dataset, the first step is applying the segmentation technique as described above. Then, chicory plants and weed are distinguished on the fact that the real crop plants are seeded on straight rows, even with more-or-less constant interdistances. Weed plants pop up at random positions.



Figure 14: Illustration of the Belgian chicory results. Above: segmentation results of one camera image. Below: original camera image with superimposed the two estimated crop row positions (red).

In the field of view of our camera, two parallel rows of chicory seedlings are visible. In our segmentation results (figure 14), these rows are not visible in every image because the seeds did not germinate all. Therefore, we applied a Kalman filter to estimate the position of the rows in the sequence of images. With these crop row positions known, it is easy to tell the difference of real crop plants and weed plants.

8. Acknowledgement

This work is supported by the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen). The work is part of a one-year research project SIVOL, “Snelle Implementatie van Vormgebaseerde Objectherkenning in Landbouwttoepassingen” (Fast Implementation of Shape-based Object recognition in Agricultural applications). Aim of this project is the introduction of state-of-the art computer vision techniques in the agricultural sector. This paper presents a novel robust preprocessing step, necessary for all agricultural camera applications.

9. References

- [1] Decker et al., 2007. Lighting Independent Skin Tone Detection Using Neural Networks. IEICE - Trans. Inf. Syst. Vol. E90-D, No. 8, pp 1195-1198
- [2] Haykin S., 1999. Neural networks. a comprehensive foundation, Pearson Education
- [3] Marchant J. A. and Onyango C. M., 2000. Shadow-invariant Classification for Scenes Illuminated by Daylight. Journal of the Optical Society of America. Vol. 17, No. 11, pp 1952–1961.
- [4] Meyer, G.E., T. Mehta, M.F. Kocher, D.A. Mortensen, A. Samal. 1998. Textural imaging and discriminant analysis for distinguishing weeds for spot spraying. Transactions of the ASAE 41(4): 1189-1197.
- [5] Sjøgaard H.T. and Olsen H.J., 2002. Determination of crop rows by image analysis without segmentation, Computers and Electronics in Agriculture, Volume 38, Issue 2, February 2003, Pages 141-158
- [6] Seow M.-J. et al., 2003. Neural Network Based Skin Color Model for Face Detection. Applied Image Pattern Recognition Workshop. Vol. 0, pp 141.
- [7] Woebbecke D. M. et al, 1992. Plant Species Identification, Size and Enumeration Using Machine Vision Techniques on Near-Binary Images. SPIE Optics in Agriculture and Forestry. Vol. 1836, pp 208–219.
- [8] Woebbecke D. M. et al, 1995. Color Indices for Weed Identification Under Various Soil, Residue, and Lighting Conditions. Transactions of the ASAE. Vol. 38, No. 1, pp 259–269.