

Thistle Detection

Søren I. Olsen^(✉), Jon Nielsen, and Jesper Rasmussen

University of Copenhagen, Copenhagen, Denmark

ingvor@di.ku.dk

Abstract. Detecting green thistles in yellow mature cereals seems an easy task. In practice however, the two colors may be close and changing over time and place. The total thistle area may be negligible, illumination effects at low sun angles etc. may make a stable detection less easy. This paper describes a method *Weeddetect* for detecting thistles in drone images of fields with mature wheat. Preliminary experiments indicate a classification accuracy of about 95%.

1 Introduction

Pesticide regulations and a relatively new EU directive [5] on integrated pest management create strong incentives to limit herbicide applications. In Denmark, several pesticide action plans have been launched since the late 1980s with the aim to reduce herbicide use [12]. One way to reduce the herbicide use is to apply site-specific weed management, which is an option when weeds are located in patches, rather than spread uniformly over the field. Site-specific weed management can effectively reduce herbicide use, since herbicides are only applied to parts of the field [2]. This requires reliable remote sensing and sprayers with individual controllable boom sections or a series of controllable nozzles that enable spatially variable applications of herbicides. Preliminary analyses [14] indicate that the amount of herbicide use for pre-harvest thistle (*Cirsium arvense*) control with glyphosate can be reduced by at least 60% and that a reduction of 80% is within reach. The problem is to generate reliable and cost-effective maps of the weed patches. One approach is to use user-friendly drones equipped with RGB cameras as the basis for image analysis and mapping.

The use of drones as acquisition platform has the advantage of being cheap, hence allowing the farmers to invest in the technology. Also, images of sufficiently high resolution may be obtained from an altitude allowing a complete coverage of a normal sized Danish field in one flight. Importantly, thistles shorter than the cereals may not be visible when viewed obliquely from 2–3 m height, but is clearly visible from above.

To differentiate between different types of weed in drone images, pure color based analysis probably will be insufficient. Here low altitude (high resolution) images allowing leaf shape analysis will be needed. However, since the dominating weed type is thistles, and since these may be detected from larger altitude allowing mapping of a larger area, we will here use the terms *weed* and *thistles* interchangeably and will not try to classify the detections into the different types of weed.

The objective of this paper is to demonstrate that single-image weed- (and in particular thistle-) detection in drone images of cereal fields is possible. The paper includes a short review of previous work, a description of the proposed method, dubbed *Weeddetect*, results of preliminary experiments and suggestions for further work.

2 Previous Work

The amount of previous work on color image processing is enormous. A few relevant textbooks include [6, 7, 10]. Less work exist on drone image analysis of field images. The most relevant contributions include [4, 8, 9, 11, 14]. The latter is directly relevant for the present work. Here each image is split into 64×64 pixel patches, that are converted to a scalar value using the *Excess green*-projection [13] $ExG = 2G - R - B$. Then, the average of the 5% largest ExG-values within the patch is computed and thresholded by a fixed parameter value learned from a set of annotated images. It is shown [14] that this method may achieve a classification accuracy above 95% on many images, but also that a number of problems exist. Most importantly, optimal performance could not be reached without adjusting the threshold parameter to each image capture campaign. This is because of different crop and light conditions makes automatic usage problematic. For some images, far too many or far too few patches were classified correctly as weed. To a large degree this seems to be caused by lack of preprocessing, but also the ExG-projection and the usage of a fixed threshold plays a role.

In [3] texture analysis was attempted in addition to color analysis. Texture analysis may seem useful because cereals and the thistle leaves are different in shape. However, texture based detection is inherently difficult no matter the approach, because even for small sized image patches, the distribution of weed features and cereal features mix. To obtain stable statistics, larger image patches is required but here the fraction of visible weed may be very small. Experiments [3] showed that detection was possible but unstable and less accurate compared to the color-based procedure reviewed above. Also, texture analysis was far more computationally demanding.

In the present work both pre- and post processing is addressed, but the major contribution of this work is a procedure for robustly estimating an optimal threshold for possibly unimodal distributions composed of a cereal and a weed component. Also, a new projection vector, significantly outperforming previous choices are presented.

3 Weed Detection

Previous work [3, 4] shows that simple approaches may work remarkably well. However, too often such approaches fail due to lack of robustness. One practical problem is that change of camera or use of auto-white balancing at image capture

changes the color balance. This makes the images more grayish and reduces the difference between yellow (cereals) and green (thistles). In general, the colors change with the lighting conditions. For color based weed detection the more saturated colors, the easier the detection will be. Another problem arises when the angle between the sun and the horizon is below about 60° . In this case, because the view is approximately ortho, the image part opposite to the sun azimuthal angle will show the illuminated side of the flora, while the image part towards the sun will be darker. This effect clearly depend on the shape of the crop leaves, the viewing direction etc. If not corrected for, too many thistles in the affected image area will not be detected.

Depending on the resolution of the spraying equipment, the detections need to be resampled. For an altitude of 50 m and using a standard commercial camera with fixed standard optics, the pixels side length may correspond to a few centimeters or less at ground. This resolution is far too high to be used for sprayers. Typically, the controllable spraying resolution is about 3 m. The resampling should be made accordingly taking altitude and effective focal length into account. Finally, the farmer/user should be able to specify the trade off between minimizing the risk of false positives (causing too much spraying) and false negatives (causing too little spraying).

Previous attempts [3, 4] within thistle detection have used a fixed threshold to separate green from yellow. Despite all preprocessing such as color balancing, such non-adaptive approach may not be sufficiently robust and an adaptive procedure for threshold selection is needed.

3.1 Preprocessing

To remedy an undesirable color balance we follow the approach of Cheng et al. [1]. However, instead of adjusting to a white-balance, i.e. the maximum vector $(1, 1, 1)$ we use the vector $(1, 0.9, 0.7)$ reflecting that the images show yellow and green colors and very little blue. Colors in images obtained using auto white balance are restored reasonably well (in the sense that the resulting color distribution becomes similar to what is seen at other image captures). To a large extent, this procedure replaces a more tedious camera color calibration prior to image capture.

To eliminate the effect of uneven illumination obtained when the angle between the horizon and the sun is not large, we first transform the representation to HSV. Here, the effect is concentrated both in the saturation component and the value component, whereas the hue component is hardly affected. The shape and position of the effect depends on the crop type, the sun height and azimuthal position, the camera tilt, and would probably be difficult to estimate. We choose to correct both saturation S and value V using a linear fit to each image, i.e.:

$$S(x, y) \longrightarrow \frac{\bar{S}}{a_s + b_s x + c_s y} S(x, y) \quad \text{and} \quad V(x, y) \longrightarrow \frac{\bar{V}}{a_v + b_v x + c_v y} V(x, y) \quad (1)$$




Fig. 1. (a) An original image with uneven illumination. (b) Compensated illumination.

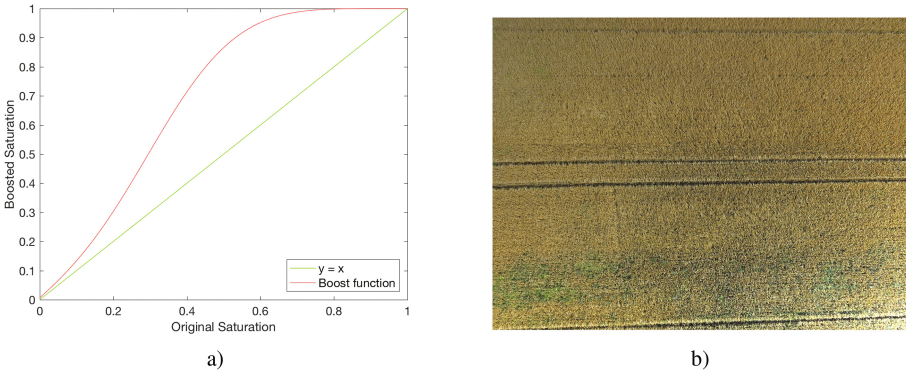


Fig. 2. (a) Saturation boosting function. (b) Corrected and saturation-boosted image (Color figure online)

where \bar{S} and \bar{V} are global averages of saturation and value, and where the denominators show planar fits to the local averages of saturation and value. This simple approach is robust to situations where spots (e.g. of flowers) make the local average deviate significantly from the fit. More complex models (than a planar fit) have shown too sensitive to such cases. An example of a correction is shown in Fig. 1.

To further ease the following separation of the green and the yellow, the saturation values are nonlinearly increased. This is illustrated in Fig. 2. Low values are left almost unchanged whereas medium and in particular large values are increased towards the maximal value.

3.2 Detection

Based on a large number of annotated areas, the preprocessed images were examined and verified to split the color space linearly into two slightly overlapping

subspaces. This is illustrated in Fig. 3(a) for a random subset of training points. The best fitting plane separating the two populations had a unit surface normal of $(-0.609, 0.773, -0.178)$. The best separating plane vary from image to image but all had approximately the same normal. Compared to [14] this vector is significantly different from the normalized ExG-vector, putting much less weight on the blue color component. By using the normal to the best separating plane as projection vector, the following threshold estimation is eased.

Separation of the projected (smaller) yellow cereal values from the larger green thistle values is difficult because the amount of the latter is usually only a tiny fraction. Thus, for almost all images the distribution of projected values is unimodal. Often a visual inspection will not reveal the thistle component. Bimodality exceptionally shows up when large image areas are covered with weeds, grass or trees.

If the cereal pixel projection values are modelled as i.i.d. random values, then the central limit theorem suggests that the yellow cereal distribution component should be Normal. Visually, this is confirmed. However, both the mean value and the variance vary from image to image. Also for images with large amount of green plants, the corresponding distribution component interferes with the right part of the normal cereal distribution. This suggest an approach where first the parameters of the Normal cereal distribution is estimated, then the weed distribution is defined as the right residual, and the threshold is selected as the one minimizing the sum of false positives and false negatives. This is illustrated in Fig. 3(b). The procedure is fast because the binning of the histograms may be chosen to limit the number of possible threshold values allowing an exhaustive search.

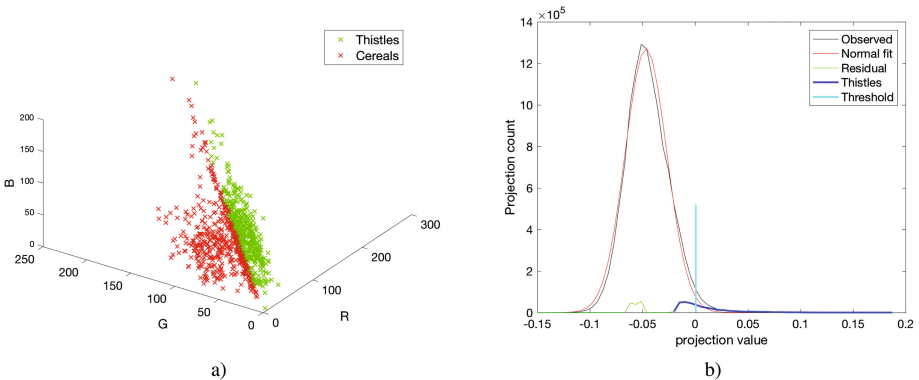


Fig. 3. (a) Projection in red of cereal- and in green of thistle- RGB-values. The two classes are almost linearly separable. (b) A typical distribution of projected values. The black curve is the measured distribution. The red curve is the estimated cereal Normal distribution. The curve in green and blue show the estimated left and right residual weed distribution. The vertical cyan line marks the optimal threshold. (Color figure online)

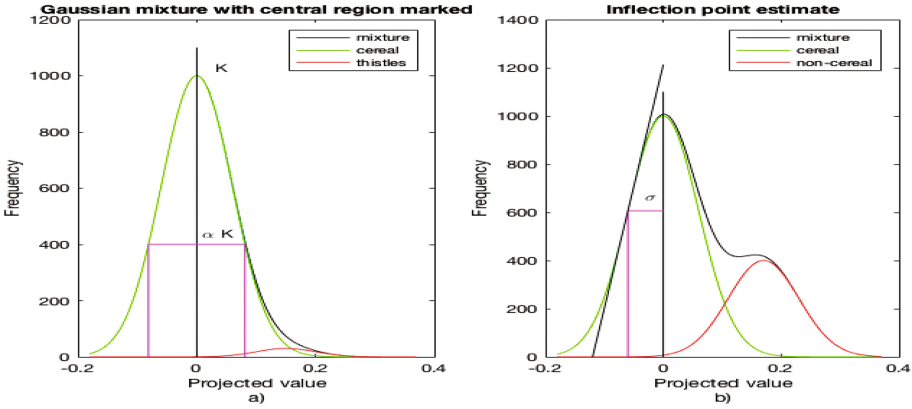


Fig. 4. (a) Estimation based on the central area marked in magenta. Used when the non-cereal distribution shown in red only affects the right tail of the distribution. (b) Estimation based on the position of the left inflection point and the tangent here. Used when the middle and right part of the distribution seriously deviates from a normal distribution.

The challenge is to estimate the mean and variance of the (left) cereal distribution not knowing how severely the right part is corrupted by other components. First, the contribution of the latter distribution component is assumed to be minor. In this case, the mean is estimated from the central data around the mode and with value larger than α (say 0.4) times the mode value. This support area is illustrated in Fig. 4(a). If the mean deviates much from the mode, α is increased and the mean reestimated. This procedure is iterated up to 3 times, each time reducing the support area for the estimation. If the estimated value still deviates significantly from the mode an alternative estimation is used (see below).

As for the mean estimation, the variance is estimated by gradually reducing the (central) support area around the mode up to 3 times. The estimation is accepted if the absolute mean difference of the estimated Gaussian and the data is below a threshold. One problem is that the accuracy of the estimated variance is reduced when the tails of the distribution is discarded—The variance is systematically underestimated. It is easy to show that the variance compensation factor will be:

$$\gamma(\alpha) = \sqrt{2\pi} \int_{-\beta}^{\beta} t^2 e^{-t^2/2} dt$$

where $\beta = \sqrt{-2 \ln(\alpha)}$. Thus, γ is independent on the standard deviation of the true distribution. Other expressions for the variance of asymmetric or one-sided truncated normals may be found in [15]. The expression for γ has no closed form solution, but is (for a limited number of values of α) estimated from simulations.

If the estimation of either the mean value or the standard deviation fails, an alternative estimation procedure is attempted. This is applied when a large fraction of the (right part of the) normal cereal distribution is mixed with the distributions of non-cereals. As illustrated in Fig. 4(b), here, only the left part of the empirical distribution is used. After smoothing, the left inflection point I is localized. From the value $v(I)$ in this point and an estimate of the derivative $d(I)$, the variance, and then the mean value may easily be computed by $\hat{\sigma} = v(I)/d(I)$ and $\hat{\mu} = I + v(I)/d(I)$. Since the latter procedure is based on far fewer data, it is invoked as a last resort. Preliminary experiments indicate a robust, but not very accurate estimation. In the experiments reported later the alternative estimation was not in use because all images (with ground truth) only showed cereals and thistles.

3.3 Postprocessing

The result of the pixel-based classification is not appropriate for thistle mapping. A large number of tiny (few pixels large) green areas often may be incorrectly registered as thistles. Mostly, the tiny areas are caused by ground vegetation (most often within the central image area). To eliminate the small detections a morphological opening is applied. This significantly reduces the number of segments even though the structure element is chosen small.

Finally, pixels based detection is not really of use for glyphosate spraying because the image resolution is far higher than the resolution of the sprayer. First, the image patch size in pixels corresponding to the desired resolution in meters is computed. Then, to avoid aliasing and missing detections, a windowed sum of pixel detections within the final resolution map is computed and a final acceptance threshold T applied. This threshold is not determined based on computer science or math, but on agricultural objectives. It specifies the farmers' trade-off between risking not to spray thistles versus unnecessary spraying. Figure 5 shows an example of an original image, the pixel-based classification and the final thistle map.




Fig. 5. (a) An original image. (b) Initial pixel based detection. (c) Final 1 m²-resolution detection with thistles areas marked as white.

4 Experimental Results

To measure the performance of the system, dubbed *Weeddetect*, 97 field images of mature wheat and 101 field images of mature barley was taken from altitudes of 20, 30 and 50 m. Based on a previous program called *ThistleTool* [14] a number of image patches were extracted. All patches showed a field area of 3×3 m (corresponding to side lengths of 450, 300 and 180 pixels). Approximately half the patches showed crop, the remaining thistles. For each patch the position of the central 1×1 m sub-patch, used for performance measurement, was recorded. The patches were presented to an agricultural expert and classified as showing either weed or only crop. This forms the basic ground truth. See Table 1 for details.

In Fig. 6 two patches classified as crop and two patches classified as weed are shown. Two of the patches are easy to classify (expert or not), while the remaining two less clearly belong to either of the classes. Even for an expert in agriculture there is a significant amount of patches where the ground truth classification is debatable. In addition to the data specified above, the extracted patches for the 19 images of Wheat observed using an altitude of 50 m was reclassified 3 times by the same expert from agriculture. For the first reclassification the expert was told to be as liberal in his classification into weed as for the basic classification. For the second and third reclassification he was told to classify

Table 1. Amount of data used for testing including the manual classification of the automatically extracted image patches

| Type | Height | Images | Patches | Weed | Crop |
|--------|--------|--------|---------|------|------|
| Wheat | 50 | 19 | 950 | 475 | 475 |
| Wheat | 30 | 39 | 1574 | 599 | 975 |
| Wheat | 20 | 39 | 1638 | 663 | 975 |
| Barley | 50 | 18 | 900 | 450 | 450 |
| Barley | 30 | 38 | 1900 | 950 | 950 |
| Barley | 20 | 45 | 2103 | 978 | 1125 |




Fig. 6. The leftmost two images patches are classified as crop where the rightmost two are classified as weed. The classification of the middle two patches is debatable. The central area used for performance evaluation is indicated by the small magenta markers. (Color figure online)

Table 2. Patch classification distribution for initial classification and 3 reclassifications of the image patches extracted from the images of Wheat from 50 m altitude.

| Classification | Weed | Crop |
|-----------------------|------|------|
| Basic Liberal | 481 | 419 |
| Liberal 2 | 494 | 406 |
| Slightly Conservative | 317 | 583 |
| Conservative | 320 | 580 |

Table 3. Performance on two types of crop and three altitudes. The columns marked *Within sample* mark the raw results. The columns marked *Estimated total* show the performance after normalization with the true amount of weed seen in the images.

| Crop | Height | Estimated total | | | Within sample | | |
|--------|--------|-----------------|------|------|---------------|------|------|
| | | TPR | SPC | ACC | TPR | SPC | ACC |
| Wheat | 20 | .954 | .956 | .955 | .984 | .890 | .933 |
| Wheat | 30 | .992 | .929 | .943 | .997 | .948 | .912 |
| Wheat | 50 | .938 | .965 | .960 | .990 | .866 | .928 |
| Barley | 20 | .871 | .890 | .878 | .923 | .787 | .888 |
| Barley | 30 | .920 | .855 | .888 | .942 | .777 | .874 |
| Barley | 50 | .887 | .945 | .914 | .920 | .919 | .920 |

slightly conservative and conservative. In Table 2 below the distribution of the initial and the three reclassifications are shown. These data serves as supplementary ground truth used to evaluate the sensitivity of *Weeddetect* as a function of ground truth bias.

Next, *Weeddetect* was applied to all images for a number of different values of the final threshold value T and the final binary classification (with 1 m resolution) was compared to the ground truth. For each altitude, crop type and value of T , the true positive rate (TPR) or sensitivity, and the true negative rate or specificity (SPC), as well as the accuracy (ACC) was computed. In Table 3, these measurements are marked as *Within sample*. Because the fraction of classified weed patches is far larger than observed in general, an estimate of the performance on the total image area is made. These measures marked *Estimated total* more correctly shows the capability of the approach.

Although the results for the different altitudes were slightly different (30 m giving the best results for Wheat and 50 m for Barley), the ranking was constant as a functions of the threshold T . Thus to summarize data, the TPR and SPC were averaged across altitude. Figure 7 shows the ROC-curves for Weed and Barley as functions of the threshold parameter T . The figure shows that *Weeddetect* performs well on Wheat where high values of both TPR and SPC can be achieved simultaneously. For Barley, the performance is significantly worse. This is in agreement with agricultural experience.

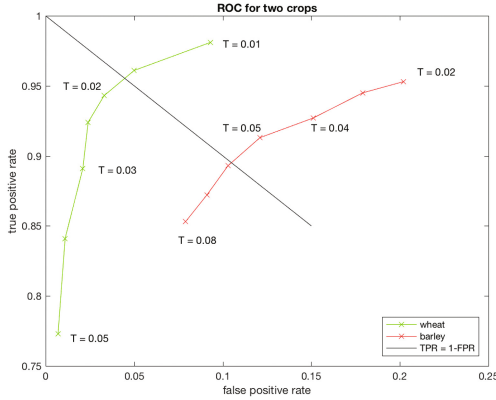


Fig. 7. ROC-curves for “Estimated total” measures for *Weeddetect*. The left curve is for Wheat and the right for Barley. It is clearly more difficult to obtain ideal results for Barley. The curves are obtained by averaging across altitude.

Also shown in Fig. 7 is the line $TPR = 1 - FPR$ corresponding to simultaneously minimal values of false positives and false negatives. This trade-off has been chosen by experts in agriculture to represent the ideal case. The two types of crop clearly require different threshold values to obtain this trade-off. In the following we therefore fix the threshold parameter values to 0.060 for Barley and to 0.015 for Wheat. Important to notice is that this choice does not necessarily correspond to the threshold values for which ACC is maximized. Table 3 show the detailed performance measures for the chosen threshold values. For the chosen values of the threshold T for the two types of crop (but otherwise fixed and identical parameter settings), the accuracy does not change significantly across altitudes, but differs significantly between Wheat and Barley.

Compared with the previous approach *ThistleTool* [14], Table 4 shows the over-all performance when averaging over altitude. The table shows that for Wheat, the TPR is slightly better and SPC is slightly worse for *Weeddetect*. For Barley, the TPR is slightly better for *Weeddetect* while at the same time being significantly better in specificity. Important here is that the numbers for

Table 4. Comparison of *Weeddetect* and *ThistleTool* using Within sample measures and averaging across altitude. Notice that for *ThistleTool* threshold tuning to each field was applied.

| Method | Crop | TPR | SPC | ACC |
|--------------------|--------|------|------|------|
| <i>Weeddetect</i> | Wheat | .961 | .950 | .953 |
| | Barley | .887 | .945 | .914 |
| <i>ThistleTool</i> | Wheat | .940 | .982 | .963 |
| | Barley | .872 | .724 | .823 |

Table 5. Performance on 900 image patches of Wheat fields viewed from 50 m altitude. The columns marked *Within sample* mark the raw results. The columns marked *Estimated total* show the performance after normalization with the true amount of weed seen in the images. The rows correspond to different ground truth definitions.

| Ground truth | Estimated total | | | Within sample | | |
|-----------------------|-----------------|------|------|---------------|------|------|
| | TPR | SPC | ACC | TPR | SPC | ACC |
| Basic liberal | .938 | .965 | .960 | .990 | .866 | .928 |
| Liberal 2 | .920 | .972 | .963 | .986 | .879 | .935 |
| Slightly conservative | 1.000 | .906 | .915 | 1.000 | .629 | .753 |
| Conservative | 1.000 | .912 | .922 | 1.000 | .641 | .762 |

Weeddetect are obtained for a fixed set of parameters while for *ThistleTool* tuning to each specific field was allowed. Thus, the preprocessing made in *Weeddetect* seems to have eliminated the problems previously observed with *ThistleTool*.

Finally, still keeping all parameters fixed, including the threshold values for the two types of crop, the performance on the three supplementary data sets were measured. Table 5 below show, for comparison, in the first line the results on the initial data set. In the next three lines the result obtained on the supplementary data sets (see Table 2) are shown. Table 5 shows, that bias in the expert classification of patches has a significant impact on the performance measures. If the classification is more conservative (fewer patches are marked to show thistles) then it's easy for *Weeddetect* to reach a perfect TPR. However, for such classification more patches marked as crop will be detected as thistles and the accuracy will be reduced. When the expert classifies patches more eagerly into thistles the performance is better and more balanced although the TPR is lower. The difference between the results using the two liberal classifications indicate the ground truth accuracy here may be about 1%.

Weeddetect is currently implemented in Matlab and it takes about 8s to process one $3K \times 4K$ image on a high-end laptop. If ported to say C++, analysis of all images covering a field may be reduced to a few minutes facilitating application in practice.

5 Conclusions

We have presented a fully automatic system *Weeddetect* for detecting weeds, in particular thistles, in drone images of mature wheat. The system automatically compensates for uneven illumination and unbalanced colors. The preliminary experiments indicate that the system performs well on all altitudes from 20 to 50 m but that the performance depends on the type of crop. For Barley the user specified threshold on the minimal necessary amount of detections within one square meter should be increased significantly with respect to the parameter value for Wheat. Although the method is developed for detecting thistles in wheat it may be possible to apply it on other cereals such as Barley. In such

cases it is likely that the threshold has to be tuned to obtain the desired tradeoff between false positives and false negatives.

The experiments revealed that a slight bias in the expert classification of ground truth, had a significant impact on the obtained results. This both point at the difficulty of the problem, but also raises fundamental questions on how to test systems when noise and bias in ground truth may be present.

In future work we plan to combine the detections obtained at single images into one field map. Such map will be easy to geo-reference and convert to a spraying-program. Also consistent detection in multiple images may be used to increase the quality of the final map.

Acknowledgements. This work was supported by the Danish Environmental Protection Agency project MST-667-00141.

References

1. Cheng, D.L., Prasad, D., Brown, M.S.: Illuminant estimation for color constancy: why spatial domain methods work and the role of the color distribution. *J. Opt. Soc. Am. - A (JOSA)* **31**(5), 1049–1058 (2014)
2. Christensen, S., Rasmussen, J., Pedersen, S.M., Dorado, J., Fernandez-Quintanilla, C.: Prospects for site specific weed management. In: *Second International Conference on Robotics and Associated High-Technologies and Equipment for Agriculture and Forestry. New trends in Mobile Robotics, Perception and Actuation for Agriculture and Forestry*, Madrid, pp. 541–549 (2014)
3. Egilsson G.J.: Detecting weed on images of cereal fields acquired by drones. Master thesis project, Department of Computer Science, University of Copenhagen (2014)
4. Egilsson, G.J., Pedersen, K.S., Olsen, S.I., Nielsen, J., Ntakos, G., Rasmussen, J.: Pre-harvest assessment of perennial weeds in cereals based on images from unmanned aerial systems (UAS). In: *The 17th European Weed Research Society Symposium*, Montpellier, France (2015)
5. The Directive on Sustainable Use of pesticides EU 128/EC Directive of the European parliament and of the council establishing a framework for Community action to achieve the sustainable use of pesticides. OJL 309/71, 24 November 2009
6. Fernandez-Maloigne, C. (ed.): *Advanced Color Image Processing and Analysis*. Springer, New York (2012)
7. Forsyth, D., Ponce, J.: *Computer Vision: A modern Approach*, 2 edn. Pearson (2012)
8. Franco, C.A., Pedersen, S.M., Papharalampos, H., Ørum, J.E.: An image-based decision support methodology for weed management. *Precis. Agric.* **15**, 595–601 (2015)
9. Garcia-Ruiz, F.: UAV based imaging for crop, weed and disease monitoring. Ph.D. dissertation, Department of Plant and Environmental Sciences, Faculty of Science, University of Copenhagen (2014)
10. Gonzales, R.C., Woods, R.E.: *Digital Image Processing*, 3rd edn. Prentice Hall, Upper Saddle River (2008)
11. Hansen, K.D., Garcia-Ruiz, F., Kazmi, W., Bisgaard, M., la Cour-Harbo, A., Rasmussen, J., Andersen, H.J.: An autonomous robotic system for mapping weeds in fields. In: *Proceedings of the 8th IFAC Symposium on Intelligent Autonomous Vehicles*, Australia, pp. 217–224 (2013)

12. Jørgensen, L.N., Kudsk, P.: Twenty years' experience with reduced agrochemical inputs. In: Proceedings of the Home Grown Cereal Authority, Research and Development Conference. Arable Crop Protection in the Balance Profit and the Environment, pp. 16.1–16.10 (2006)
13. Rasmussen, J., Nielsen, J., Garcia-Ruiz, F., Christensen, S., Streibig, J.C.: Potential uses of small unmanned aircraft systems (UAS) in weed research. *Weed Res.* **53**, 242–248 (2013)
14. Rasmussen, J., Nielsen, J., Olsen, S.I., Petersen, K.S., Jensen, J.E., Streibig, J.C.: Droner til monitorering af flerårigt ukrudt i korn; (In Danish). Miljøstyrelsen, Miljøministeriet. <http://mst.dk/service/publikationer/>. Accessed Jan 2017
15. Barr, D.R., Sherill, E.T.: Mean and variance of truncated normal distributions. *Am. Stat.* **53**(4), 357–361 (1999)