Investigation into the classification of diseases of sugar beet leaves using multispectral images

S.D. Bauer, F. Korč and W. Förstner
University of Bonn, Institute of Geodesy and Geoinformation, Department of Photogrammetry, Nussallee 15, D-53115 Bonn, Germany; sabine.bauer@uni-bonn

Abstract

This paper reports on methods for the automatic detection and classification of leaf diseases based on high resolution multispectral images. Leaf diseases are economically important as they could cause a yield loss. Early and reliable detection of leaf diseases therefore is of utmost practical relevance - especially in the context of precision agriculture for localized treatment with fungicides. Our interest is the analysis of sugar beet due to their economical impact. Leaves of sugar beet may be infected by several diseases, such as rust (*Uromyces betae*), powdery mildew (*Erysiphe betae*) and other leaf spot diseases (*Cercospora beticola* and *Ramularia beticola*). In order to obtain best classification results we apply conditional random fields. In contrast to pixel based classifiers we are able to model the local context and contrary to object centred classifiers we simultaneously segment and classify the image. In a first investigation we analyse multispectral images of single leaves taken in a lab under well controlled illumination conditions. The photographed sugar beet leaves are healthy or either infected with the leaf spot pathogen *Cercospora beticola* or with the rust fungus *Uromyces betae*. We compare the classification methods pixelwise maximum posterior classification (MAP), objectwise MAP as soon as global MAP and global maximum posterior marginal classification using the spatial context within a conditional random field model.

Keywords: pattern recognition, leaf diseases, maximum posterior classification, conditional random fields, watershed algorithm

Introduction

This paper presents an investigation into three classification methods to detect leaf diseases on sugar beet plants. This is a prerequisite for precision farming, in order to obtain complete information about the spatial distribution of the infected areas in the field. This replaces sparsely sampled tests. Identifying leaf diseases normally is destructive. The leaves are cut off the plants and scanned or photographed in the lab (cf. Boissard, 2008, Pydipati, 2006). Non-destructive approaches exist to obtain the 3D-structure of plants. In this approaches are adopted X-ray (Stuppy, 2003) or laser scanning systems (Hanan, 2004). X-ray cannot be applied in the field and laser scanning is costly and in case of wind, due to the motion of the leaved, cannot yield spatially consistent information. Pan (2004) used stereo photos for obtaining the 3D-structure of leaves, but there the user must select matching-points in stereo images.
In contrast, our approach to identify leaf diseases is non-destructive and applicable in the field. We take several images in order to enable stereo analysis. So we can integrate observations of different cameras, in this case RGB and Infrared, as well as investigate the evolution of a disease over time on the same leaf. Both the stereo reconstruction and the classification are fully automatic, thus enable real time analysis, a prerequisite for precision farming.
The investigation focuses on leaf disease of sugar beet plants due to their economical impact in Germany. But the methods can be transferred to other species. Leaves of sugar beet may be infected by several diseases, such as rust (*Uromyces betae*), powdery mildew (*Erysiphe betae*) and other

leaf spot diseases (*Cercospora beticola* and *Ramularia beticola*). In this investigation we restrict us to the leaf spot pathogen *Cercospora beticola* and the rust fungus *Uromyces betae.*
The investigation will show:

- The non-destructive and automatic classification allows the analysis of the development of a disease.
- Pixelwise classification can be improved by taking learning the spatial distribution of disease patterns and using this knowledge for the classification.
- *Cercospora beticola* can be automatically identified with a probability of 85% and *Uromyces betae* with a probability of 22%.

Materials and methods

In the following we introduce first the experimental set-up (Figure 1). Then we explain how we fuse the images from the RGB- and the Infrared-camera. Finally we define our classification strategy.

Experiment set-up
The multispectral stereo images were taken in a lab under well controlled illumination conditions. We restrict us in this investigation of single leaves. For each leaf we take four RGB images ('FujiFilm FinePix S5600', 2,592×1,944; 1 px equates 0.0967 mm) and one infrared image ('Tetracam ADC', 1,280×1,024; 1 px are 0.2123 mm) from an altitude of 30 cm and different positions with an approximate mutual distance of 10 cm. The infrared camera has the channels RED, GREEN and NIR (700 - 950 nm). The photographed sugar beet leaves are healthy or either infected with the leaf spot pathogen *Cercospora beticola* or with the rust fungus *Uromyces betae*.

Fusion of the images from the RGB- and the Infrared-camera
We calibrate the camera using a 3D-test field to obtain the inner orientation, especially to eliminate lens distortion. Afterwards we automatically determine the six parameters of the pose (position and rotation) of the cameras using the program AURELO (Läbe, 2006). Based on the inner and exterior orientation of the cameras we determine the surface model, i.e. the 3D-structure of the leaves with the INPHO-Software MATCH-T (Lemaire, 2008). This fusion of the five images per leaf allows the rectification of the leaf w.r.t a reference plane. For each pixel of this rectified image we obtain seven values: (1) blue from the RGB-camera, (2)-(5) 2× green and 2×red from the RGB- and the Infrared cameras, and (6) infrared.

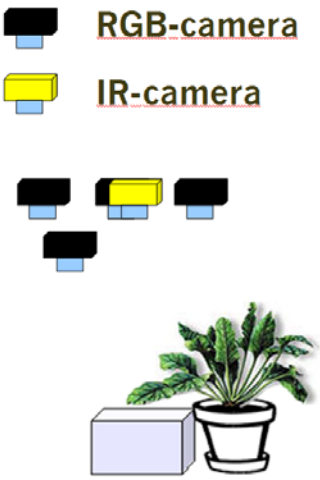

Figure 1. Experiment set-up.

Classification methods

We investigate three versions of a maximum posterior (MAP) classification. MAP classification works in such a way, that the joint probability or the product of the likelihood function $P(e|\omega_i)$ and the *a-priori* probability $P(\omega_i)$ is being maximised:

$$P(\omega_i, \boldsymbol{e}) = P(\omega_i|\boldsymbol{e})P(\boldsymbol{e}) = P(\boldsymbol{e}|\omega_i)P(\omega_i) \quad (1)$$

The vector $e=[e_j]$ stands for the feature vectors, which are in our case the different colour information. The class $\omega_i \in \Omega$ are one of the four classes: background, healthy leaf areas, with *Cercospora beticola* or *Uromyces betae* infected areas. This is equivalent to maximizing the posterior probability $P(\omega_i|e)$ as the probability for the feature vector $P(e)$ can be unconsidered, as it is constant.

The three different methods for classification are first a *pixelwise* MAP classification, further referred to as PMAP. Second an *objectwise* MAP classification, with watershed regions as objects, further referred to as OMAP and third a global MAP classification by considering the spatial context in *a conditional random field*, further referred to as CRF. They refer to different objects, pixels for PMAP and CRF or regions for OMAP, and use different priors, local ones for the PMAP and OMAP and global ones for the CRF.

First we classify binary leaf and background. In a second step we discriminate healthy and infected leaf areas. Therefore we make two separated classifications for leaf against *Cercospora beticola* and leaf against *Uromyces betae*.

Pixelwise maximum posterior classification PMAP: The *a-priori* probability $P(\omega_i)$ can easily computed from the frequency of the occurrence of the different classes. To find the likelihood function $P(e|\omega_i)$ we use the expectation maximization (EM) algorithm, where we assume the density to be a mixture of two Gaussian distributions per class. As result from the EM-algorithm we get for each class a matrix with the mean vector and the covariance matrix.

$$P(e|\omega_i) = \sum_{j=1}^{k} p(k_j) * P(e|\mu_{k_j}, \Sigma_{k_j}) \quad (2)$$

In the classification process we must now compute the posterior probability, i.e. the density for each class Equation 2 multiplied with the *a-priori* probability, and obtain that class, whose probability, is the greatest one, cf. Equation 1.

A pixel-wise classification has the disadvantage, that all pixels are being classified separately. So it can happen, that in a region most of the pixels are classified for example as background and one or two as leaf. Such classification errors can be reduced, if the neighbourhood of the pixels is taken into account. One solution is an objectwise classification another one the use of conditional random fields, which will be presented in the next chapter.

Objectwise classification OMAP: For an objectwise classification the images must be segmented. In our study we use the watershed algorithm (Vincent, 1991). This algorithm derives a partitioning of the image area into a set of disjunct regions based on a gradient image as input. We use the RED, GREEN and NIR channel to compute the gradient image.

The algorithm works well, if a good contrast between the different regions exists and the regions relate to a unique class. In our case, the contrast between the healthy leaf and the *Cercospora beticola* leaf spots isn't good, how it is seen in Figure 2. The white ring indicates the border of the leaf spot. One region must be allocated to one class even though the region may contain pixels from different classes. In such a case the region is being related to the majority class in that region. Figure 3 shows watershed regions on a leaf, which is infected with *Uromyces betae*. Obviously the watershed algorithm does not find edges around the small dark spots. Therefore all this regions will be allocated to the class 'healthy leaf areas'. For this reason leaves infected with *Uromyces betae* cannot be classified with an object-wise classifier.

Global MAP with conditional random fields (CRF): In the context of classification, we are generally interested in estimating the posterior distribution over labels given the observations. In this work, we use a Conditional Random Field (CRF) to model the posterior distribution as a Markov Random

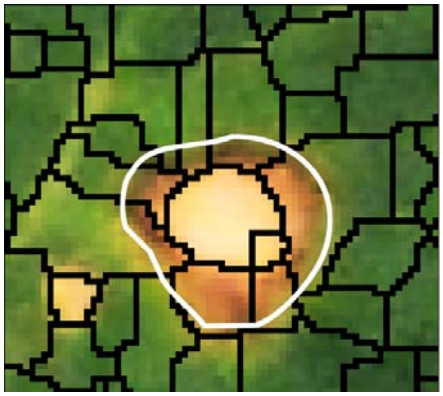

Figure 2. Watershed regions by *Cercospora beticola.*

Figure 3. Watershed regions by *Uromyces betae.*

Field (MRF), globally conditioned on the observed data. MRF allows probabilistic modelling of local contextual constraints in labelling problems and is the most commonly used model for modelling spatial interactions in image analysis (Li, 2001). A CRF provides an approach for combining local classifiers A that allow the use of arbitrary overlapping features, with adaptive data-dependent label interaction B. As MRF, it is a global model, however directly modelling the posterior probability distribution

$$P(\omega|e) \propto \prod_i A(\omega_i|e) \prod_{ij} B(\omega_i, \omega_j|e) \tag{3}$$

Such formulation gives us freedom to choose any domain specific method to identify class labels and neighbourhood dependencies and provides several advantages compared to the traditional MRF model. For discussion on the formulation, on the differences between a traditional MRF formulation and the CRF model, and comparison of the performance of the two models and an independent sitewise classifier, see references in Korč (2008).

Evaluation

The classification methods PMAP and OMAP are tested on 243 data sets, from which 98 data sets are infected with *Cercospora beticola* and 145 with *Uromyces betae* in different development stages. The photos are taken distributed about three weeks after the inoculation. These data sets are divided in subsets for a 5-fold cross validation separately performed for the *Cercospora beticola* data sets and the *Uromyces betae* data sets. We use each of the subsets of the data sets for training and the other four for testing. For memory reasons, from the training data sets we choose randomly 250,000 pixels for each class for training by the pixelwise MAP classification. For the objectwise MAP classification we use up to 250,000 regions for training.

In our CRF experiments we adopt the multiclass formulation from Kumar (2004) and learn the CRF model parameters automatically from training data as it is described in Korč (2008). Final labelling is computed using the max-product belief propagation (Tappen, 2003), where we employ the software accompanying (Szeliski, 2008).
To compare the formulations, we applied an independent sitewise classifier and the CRF model to our pixelwise class segmentation task. The aim of these experiments is to assign each image pixel with one of the three class labels. For this purpose, we prepared two datasets, the *Uromyces* dataset with 871 images and the *Cercospora* dataset with 867 images. Both datasets contain imagepatches of size 64×64 pixels. Each pixel is assigned with a feature vector composed of the pixel RGB colour vector extended with the information from the infrared channel. In other words, each pixel is associated with a 4D feature vector. We use the 2-norm of the vector difference for modelling the pairwise label interaction B in (3). In each experiment, 10 randomly chosen images are used as the training set for parameter learning and 400 images are randomly chosen from the rest of the images for testing.

Results and discussion

In this chapter we give representative results of our classification experiments using different feature vectors and methods. First we illustrate the results of the pixel-wise classification PMAP, second the results of the object-wise OMAP and finally the results of the conditional random field CRF. In all tables on the left side are the ground truth labels and on the top the test classes.

Results of the PMAP classification
In a first step we make a binary classification for separating the leaf from the background. Therefore we uses the features red, green and blue from the RGB camera, and additional the near infrared (NIR) channel from the infrared camera. The results are shown in the Table 1.
The rate of correct classified pixels lies by 99.53%. For the most application areas this would be a good result. After separating the leaf from the background we detect the healthy and infected areas on the leaf. First we investigate the PMAP classifier on the 98 leaves infected with *Cercospora beticola*. There we use feature vectors only with RGB information and then the gain of additional use of the NIR channel. Table 2 shows the median of the confusion tables for the two cases over all leaves. In both cases the classification rate due to the large number of healthy pixels adds up to 99%. Without the NIR information the percentage of true positives, i.e. correctly classifying *Cercospora* leaf spots, called C-detection rate, is 42%, whereas the adding NIR causes some improvement about 12% to 54%. In order to analyse the low C-detection rates we show the distribution of the C-detection rate of *Cercospora* leaf spots over the infected leaves in Figure 4. The C-detection rate ranges from 0% to nearly 100%. I.e. on 12 infected leaves no leaf spot were detected.
As we are not interested in classification of the pixels but the classification of the leaves, we investigate the probability of finding a truly infected leaf as infected. Especially we are interested

Table 1. Leaf against background in *Cercospora beticola* by pixelwise MAP classification.

	RGB plus NIR classification rate 99.53%	
	Leaf	Background
Leaf	99.59%	0.41%
Background	0.50%	99.50%

Table 2. PMAP: Median of the confusion tables. Healthy leaf areas against *Cercospora beticola*.

	Only RGB classification rate 99.33%		RGB plus NIR classification rate 99.27%	
	C. leaf spots	Healthy areas	C. leaf spots	Healthy areas
C. leaf spots	41.52%	58.48%	54.15%	45.85%
Healthy areas	0.14%	99.86%	0.24%	99.76%

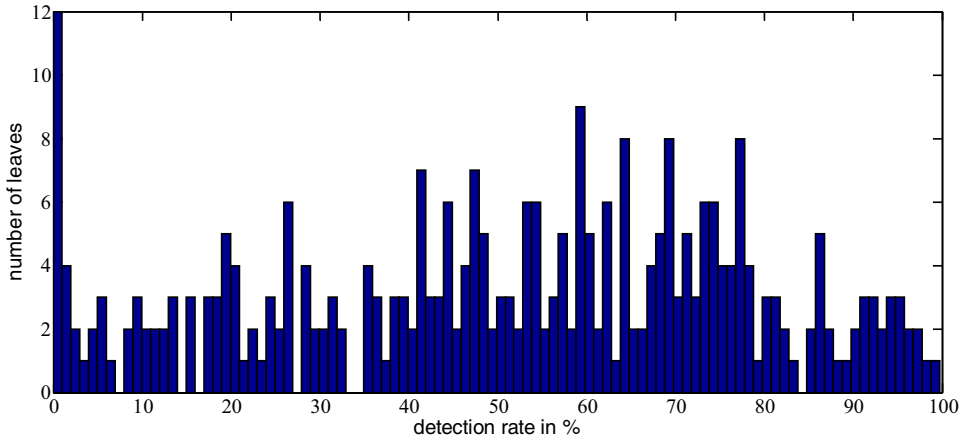

Figure 4. C-detection rate of Cercospora beticola: used channels are RGB and NIR.

in an early detection after inoculation. In case we treat a leaf as infected if at least more than 0% of the leaf area is infected, i.e. at least one pixel is classified as infected. We find the C-detection rate as a function of the day after inoculation being 100% from day 4 on. How it is seen in Figure 5 right, almost all healthy leaves are also classified as infected.

If we use a threshold for the infected pixels of 0.2% of the leaf area the detection rate from the healthy leaves increases but also the detection rate of the infected leaves decreases. After the tenth day of inoculation all leaves are infected, so the detection rate of healthy leaves must be there zero.

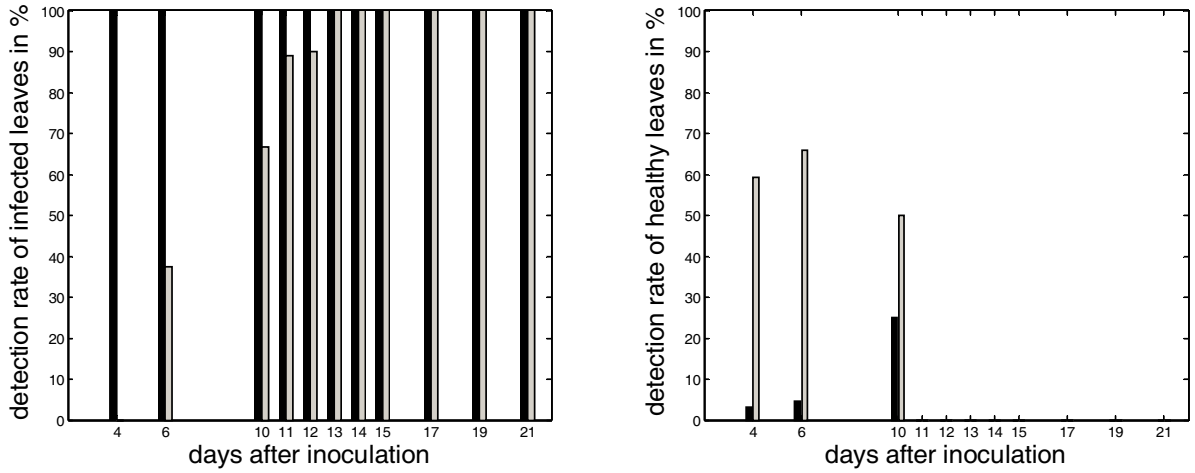

Figure 5. C-Detection rate of infected (left) and healthy leaves (right) in *Cercospora beticola.* Black: Threshold 0%. Gray: Threshold 0.2%.

Figure 6 shows the infected area over time after inoculation for ground truth (solid) and the result of the classification (dashed). The percentage of the infected areas on the leaves increase up to 10% on the 21st day after the inoculation. The classification systematically overestimates the infected area. The classification results of *Uromyces betae* is shown in Table 3. The overall classification rate is 99.95%, but the correctly classified rust spots are only 13%, which indicates a weak classi-fication result. The number of *Uromyces* leaf spots is so marginal, that the *a-priori* probability for this class is only 0.06%. So this bad result hasn't a great effect of the whole classification rate.

Results of the OMAP classification

As outlined in the chapter 'Classification methods', *Uromyces betae* cannot be classified with an object-wise classifier based on watershed segments. So we are only executing a pixel-wise classification with *Cercospora beticola*. We chose the features RGB plus NIR, as they have shown the best results in the PMAP classification of *Cercospora beticola*. The overall classification rate is with 99% just as well as in the PMAP classification (Table 4). But the detection rate of *Cercospora beticola* decreases considerably. The reason for this is that the watershed algorithm cut all borders of the leaf spots, as shown in Figure 2.

Results of the CRF classification

Detailed classification results on the *Cercospora* dataset and the *Uromyces* dataset using the CRF model are shown in Table 5 and Table 6, respectively.

The advantage of using the CRF is illustrated in Figure 7, where on the left side it can be observed that the classification result of the full CRF is considerably smoother compared to the pixelwise classification. Further, we note that solution yielded by CRF in general cannot be reached through simple data independent smoothing of a pixelwise classification result. CRF solution is smooth

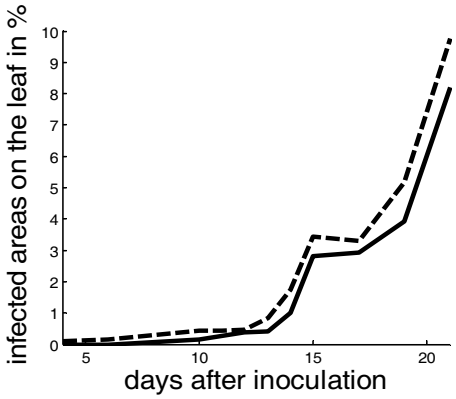

Figure 6. Infected areas on the leaf in% in *Cercospora beticola*. dashed: detected areas. solid: groundtruth.

Table 3. Pixelwise MAP: Healthy leaf against *Uromyces betae.*

	RGB plus NIR classification rate 99.95%	
	U. leaf spots	Healthy areas
U. leaf spots	12.73%	87.27%
Healthy areas	0.04%	99.96%

Table 4. Objectwise MAP in *Cercospora.*

	RGB plus NIR classification rate 99.41%	
	C. leaf spots	Healthy areas
C. leaf spots	35.56%	64.44%
Healthy areas	0.11%	99.89%

Table 5. CRF classification in *Cercospora beticola.*

	RGB plus NIR classification rate 95.40%		
	C. leaf spots	Healthy areas	Background
C. leaf spots	84.77%	14.90%	0.33%
Healthy areas	1.91%	97.83%	0.26%
Background	2.21%	3.30%	94.49%

Table 6. CRF classification in *Uromyces betae.*

	RGB plus NIR classification rate 98.22%		
	U. leaf spots	Healthy areas	Background
U. leaf spots	22.03%	77.78%	0.19%
Healthy areas	0.21%	99.24%	0.55%
Background	1.48%	2.00%	96.52%

where data support it and does not impose smooth solution on object boundaries. The process of training the CRF model can be viewed as maximizing overall classification rate on the training data. However, classification rate of a particular class is dependent on the proportion of training examples of this class in the training data. For instance, the proportion of the *Uromyces* pixels in our training data is small and, hence, the classification rate of this class has small influence on the overall classification rate that is maximized. This explains decreased detection with respect to the reduced CRF shown in Figure 7 (Right).

To illustrate the gain of using the neighbourhood relations in the CRF, we compare the result of the full CRF model to a CRF without using the neighbourhood relations, thus without the factors B in (Equation 3). We observe that the overall classification rates of the full CRF model on these two datasets are respectively 95.40% and 98.22%, whereas the classification rates of the reduced CRF drop to respectively 93.82% and 97.65%. We stress that this pixelwise classifier is not identical to the PMAP classifier and is not meant to be compared to it in our experiments.

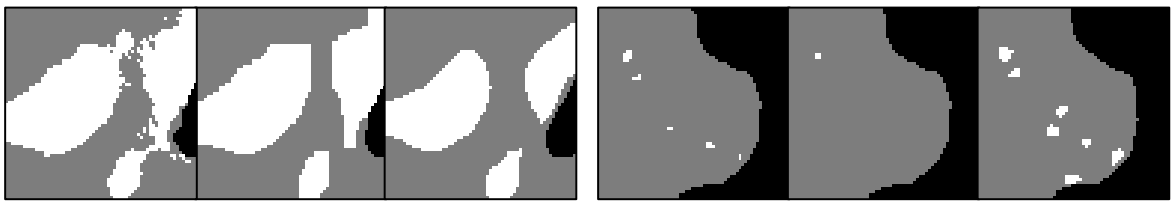

Figure 7. Left: *Cercospora beticola;* Right: *Uromyces betae*. In both cases: First image, CRF model reduced to a pixelwise classifier, then the result of the full CRF model and last the ground truth image. Black: Background, Gray: Leaf, White: Left *Cercospora beticola* and right *Uromyces betae.*

Conclusions and outlook

The overall classification rates of 99% seem promising for our application area. The detection rate of *Cercospora beticola* achieved by the CRF is 85%. In contrast, CRF classification rate of *Uromyces betae,* however, is only 22%. The reason for this is may be the small proportion of the *Uromyces* pixels in the training data.
In future work we will investigate whether the CRF classification rates of particular classes could be improved by weighting of the training examples.
In a next step, we aim to examine to what extent the two diseases can be distinguished.
Last, we want to compute the surface roughness using the existing 3D-structure of the leaves and exploit the potential of the surface roughness as feature for classification of leaf diseases. First results shown, that the surface roughness of a leaf three days after inoculation and the roughness of the same leaf 21 days after inoculation are different (Figure 8).
In future investigations on plant and canopy level the 3D-structure enables us also to compensate for illumination and occlusion effects as soon as we can refer to the true leaf area.

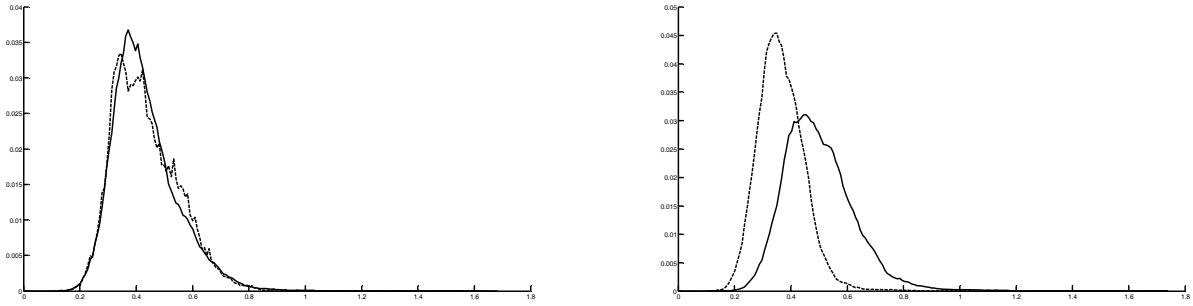

Figure 8. Roughness. Left: an infected leaf with *Cercospora beticola* 21 days after inoculation. Dashed: areas with *Cercospora* spots on the leaf. Continuous: healthy leaf areas. Right: Dashed: leaf infected with *Cercospora beticola* 21 days after inoculation. Continuous: healthy leaf.

Acknowledgements

This research is funded by the DFG Post Graduate Program 722 'Use of information technologies for precision crop protection' and partly by EU-STREP 027113 eTRIMS 'E-Training for Interpreting Images of Man-Made Scenes'. The authors are grateful to the Department of Phytomedicine at the University Bonn for their assistance.

References

Boissard, P., Martin, V. and Moisan, S., 2008. A cognitive vision approach to early pest detection in greenhouse crops. Computers and electronics in agriculture 62: 81-93.

Hanan, J.S., Loch, B. and McAleer, T., 2004. Processing laser scanner data to extract structural information. In: Godin, C., Hanan, J., Kurth, W., Lacointe, A., Takenaka, A., Prusinkiewicz, P., DeJong, T., Beveridge, C. and Andrieu, B. (Eds.), Proceedings of the 4th International Workshop on Functional–Structural Plant Models. Proceedings—FSPM04, Montpellier, pp. 9-12.

Korč, F. and Förstner, W., 2008. Approximate parameter learning in Conditional Random Fields: An empirical investigation. In G. Rigoll, editor, Pattern Recognition, number 5096 in LNCS, Springer. pp. 1120.

Kumar, S. and Hebert, M., 2004. Multiclass discriminative fields for partsbased object detection. In: Snowbird Learning Workshop.

Läbe, T. and Förstner, W., 2006. Automatic relative orientation of images. In: Proceedings of the 5th Turkish-German Joint Geodetic Days, March 29th – 31st, 2006, Berlin, ISBN 3-9809030-4-4.

Lemaire, C., 2008. Aspects of the DSM production with high resolution images. In: Proceedings of the XXIst ISPRS congress, July 3th – 11th, 2008, Beijing, China.

Pan, Z., Hu, W., Guo, X. and Zhao, C., 2004. An efficient image-based 3D reconstruction algorithm for plants. In: Kanade, T., Kittler, J., Kleinberg, J.M., Mattern, F., Mitchell, J.C., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M.Y. and Weikum, G. (Eds.), Lecture Notes in Computer Science. Springer, Heidelberg, pp. 751-760.

Pydipati, R., Burks, T.F., Lee, W.S., 2006. Identification of citrus disease using color texture features and discriminant analysis. In: Computers and electronics in agriculture, 52: 49-59.

Li, Stan Z., 2001. Markov random field modeling in image analysis. Springer-Verlag New York, Inc.

Stuppy, W.H., Maisano, J.A., Colbert, M.W., Rudall, P.J., Rowe, T.B., 2003. Three-dimensional analysis of plant structure using high-resolution X-ray computed tomography. Trends Plant Sci. 8: 2-6.

Szeliski, R. *et al.,* 2008. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. In: IEEE Trans. on Pattern Analysis and Machine Intelligence 30(6):1068-1080.

Tappen, M. F., Freeman, W. T., 2003. Comparison of graph cuts with belief propagation for stereo, using identical mrf parameters. In: IEEE International Conference on Computer Vision.

Vincent, L., and Soille, P., 1991. Watersheds in digital spaces: an ancient algorithm based on immersion simulations. IEEE Trans. Patt. Anal. Mach. Intell. 13: 583-598.