Peer reviewed version

Link to published version (if available):
10.1109/ICIP.2015.7351537

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

TOWARDS AUTOMATING VISUAL IN-FIELD MONITORING OF CROP HEALTH

David Gibson, Tilo Burghardt, Neill Campbell, Nishan Canagarajah

Faculty of Engineering, University of Bristol
Bristol, UK

ABSTRACT

We present an application that demonstrates a proof of concept system for automated in-the-field monitoring of disease in wheat crops. Such in-situ applications are required to be robust in the presence of clutter, provide rapid and accurate analysis and are able to operate at scale. We propose a processing pipeline that detects key wheat diseases in cluttered field imagery. First, we describe and evaluate a high dimensional texture descriptor combined with a randomised forest approach for automated primary leaf recognition. Second, we show that a combined nearest neighbour classifier and voting system applied to segmented leaf regions can robustly determine the presence and type of disease. The system has been tested on a real-world database of images of wheat leaves captured in-the-field using a standard smart phone.

Index Terms— ecological informatics, log-Gabor filter, randomised forests, nearest neighbour voting.

1. INTRODUCTION

Automatically detecting the occurrence of disease in agricultural crops is a challenge of increasing importance [1] as global demands on food supplies grow. As the pressure on natural resources increases the need for more efficient farming practices is required to mitigate crop failure and optimise the use of pesticides and fertilizers. Early detection of crop disease onset plays a critical role in how a disease can be successfully treated; late detection often leads to significant crop failure. The Food and Environment Research Agency (FERA), in the United Kingdom, maintains an e-Agri-based service that provides online warnings, updates and status of disease prevalence in several crops across the country [2]. The service is maintained by human operators who travel around the country collecting samples of leaves from known fields and crops. Currently, the leaves are transported back to FERA head quarters where researchers visually analysis each leaf to asses the presence, type and extent of disease. This is an expensive process that does not scale well; increasing the level of automation of this process will have a significant impact of crop monitoring at scale. With the use of smart phones

Fig. 1: **System Overview.** An photograph of a crop leaf taken in the field under natural conditions (a) is presented to the system. Utilising a randomised forest over log-Gabor filter responses yields a heat map (b) that allows automated fitting of a primarily leaf model (c) to the image. Sampling from this segmentation (d), nearest neighbour classification over Haralik texture descriptors in a colour-difference space (e) is employed finally to produce a tag that flags disease occurrence and type.

is becoming common across the globe, we propose an automated, visually operating disease recognition system able to process smart phone images taken in the field. Figure 1 provides an overview of the proposed system. The application enables a non-expert user to capture images of crop leaves in a relatively unconstrained manner for uploading to a remote automatic disease detection and analysis system. The system output is informative to the user with regard to the type and levels of disease present, which could be enhanced by suggested actions to be taken to mitigate any risks from disease.

2. BACKGROUND

Various aspects relevant to automatic visual health monitoring of crops have been studied in previous work [3, 4, 5]. For various other plants researchers have investigated the suitability of different acquisition scenarios, feature spaces, normalisation procedures, and system designs along reports on visual detectors for diseases in plants other than crops.

Fig. 2: **Field Imagery and Ground Truth.** Examples of leaf images (resolved at 3264x2448 resized to 816x612 pixels with 8bit per RGB channel) as captured by an expert from FERA (top row) and the labelled ground truth (bottom row). Images (a) and (b) contain healthy primary leaves (green) and are labelled as 'green leaf' **(gl)**. Images (c) and (d) are examples of *Septoria tritici* **(st)** and images (e) and (f) are examples of yellow rust **(yr)**, primary leaves that are diseased are dark red. The labelling includes secondary leaves with and without disease (brown **(sg)**, blue **(sd)**) and leaves where it is hard to tell if disease is present (purple).

Many methods to detect plant disease in various environment have been investigated [6, 7]. In [8], segmentation and statistical methods were used to automate crop health monitoring. While being highly automated, such techniques rely on constrained environmental and crop layout to provide imagery that is well suited to standard segmentation approaches.

Advances in segmentation algorithms lead to the development of pixel based classification systems [9]. High dimensional features are used to train a classifier which then labels previous unseen pixels as belonging to a particular class type. High dimensional pixel feature have also been used to classify global image content [10]. In both these works Gabor or log-Gabor features were used [11]. Such complex features capture both local edge and blob information at multiple scales and rotations. While this can increase the dimensionality of a given feature vector it is a useful property if the scale and orientation of the subject is unknown or cannot be reliable predicted. While neural networks have been used in such classifier systems support vector machines (SVM) and AdaBoost [12] have become popular alternatives.

When objects exhibit large variations in orientation and textural structure model based approaches can become overly complex and difficult to learn. Ensembles of randomised forest have been shown to compete with state-of-the-art learning algorithms, while avoiding the need for explicit models [13]. For the segmentation of leaves under challenging visual conditions this is an attractive route to investigate. Given a patch based classification of a scene containing a primary leaf, it is desirable to extract a leaf object and further classify the leaf object regarding the presence and type of disease.

We show that a secondary sub-leaf patch based nearest neighbours approach combined with a leaf based disease voting system results in an accurate level of leaf disease diagnosis. These system and its component approaches are described in the following sections.

3. DATA SPECIFICATION

The data set consists of 153 annotated images captured in the field by a researcher from FERA using a standard smart phone (see Figure 2, top row). The researcher is an expert at manually classifying leaves; he captured the data at arm's length fully independently from the computer vision team to avoid biased data. He was asked to choose a framing and quality level which farmers could capture themselves, and that allows for a reliable determination and analysis of the primary leaf.

In order to generate a reliable ground truth, the images were labelled in a manner similar to that of the PASCAL VOC challenge [14]. The bottom row of Figure 2 shows several examples and explains the annotation rationale. Note, that additional subclasses have been added in an attempt to disambiguate the primary target from very similar secondary leaves and other clutter (see Figure 2 for details).

The amount of visual variation that is generally present in these images of leaves captured in the field is significant; leaf shapes can vary dramatically even within species, leaf colours and textural properties vary over time, disease patterns are non-uniformly distributed and in-field environmental conditions constantly change.

4. METHOD

The proposed processing pipeline operates in a staged fashion; starting with the characterisation of primary leaf features, followed by primary leaf detection, disease feature extraction over the leaf region, and finally disease classification. The following section motivates and details these various steps.

Primary Leaf Descriptor. First, the crop leaf in view is modelled by its texture properties utilising dense, local features coupled with a global geometric model. In order to re-

a.   b.   c.   d. (partial)   e.   f. (failed)

Fig. 3: **Primary Leaf Extraction.** Top row shows the per patch based posterior $[P(p|D_{xy})]$ of primary leaf presence as output of the randomised forest for the examples in Figure 2. The bottom row shows the cutout leaf segments as found by primary leaf detection according to Equation 1. Example (d) shows how the proposed approach fails to separate exactly two very similar adjacent leaves. The particularly challenging example (f) fails due to primary leaf ambiguity.

solve local appearance properties of this primary leaf against other content, we filter the image using a subset of the complex 2D log-Gabor filter family $g_{\sigma,\theta}$, selecting four scales σ and six orientations θ, as suggested by Kovesi [11].

Individual locations are then represented by a raw feature vector $([g]_1, ..., [g]_{256})$ that concatenates the filter responses $[g]$ over a local spatial 16x16 neighbourhood. Accordingly, each location is then described by 256 jets of 24 dimensional *real* or 48 dimensional *real* and *imaginary* entries, which yields a 6144 or 12288 dimensional uncompressed feature vector for each location, respectively. Principal component analysis (PCA) is finally applied to compact the dimensionality producing a compressed n-dimensional descriptor D whose dimensionality n is set such that the model covers 1-4 standard deviations of the variance. This allows for a choice regarding the level of description detail that is coupled to its variance and, thus, limits the impact of potential overfitting.

Random Forest Classifier. The dimensions of D span a pattern space over which a spatially localised estimate of primary leaf p presence can be defined, that is a posterior $P(p|D)$. In order to implement this practically, a random forest is trained on example descriptors of known class, i.e. samples where $P(D|p) \in \{0, 1\}$. Training results in the construction of 100 trees using the implementation by [15] over 2-folds of the data (three-fold cross validation). Note that the classes of primary disease and primary-no-disease are combined against all other classes to form the positive set. Once trained, for a given descriptor the m leaf nodes are activated in the forest and give their individual posterior estimates for primary leaf presence. They can be collected over the entire forest to provide a probabilistic estimate $P(p|D_{xy})$ per descriptor D_{xy} representing the textural structure of the image location (x, y).

Primary Leaf Extraction. Applied over the entire (downsampled 816x612) input image this yields a heat map that, once normalised and under the assumption of pixel independence, can be interpreted as an estimate of the overall spatial distribution $[P(p|D_{xy})]$ of primary leaf presence. The top row of Figure 3 exemplifies such maps.

A rectangular leaf region R of learned constant mean scale and aspect ratio is fitted by maximising:

$$\sum_{(x,y) \in R} P(p|D_{xy}) \tag{1}$$

Disease Classification. In order to discriminate healthy from diseased primary leaf segments, sample regions are drawn from the extracted region and represented in an R-G-colour difference space as illustrated in Figure 4. This procedure compactly captures chlorophyllic changes in the leaf indicative of disease. Each sample patch is then characterised by five Haralick features [16], that is; contrast, correlation, energy, homogeneity and entropy of the co-occurence function. A Nearest Neighbour classifier is finally applied in an image-to-class fashion, following [17], to yield the disease tag of the system associated to a photograph.

5. RESULTS

There are 51 examples of each class; no disease, septoria tritici and yellow rust, a 151 images in total consisting of over 1.2 million patches, see Figure 2. Three fold cross validation has been used to train and test the patch classifier. For leaf segmentation both no disease and diseased leaves are considered as primary leaves, the positive class, while the background and all other secondary leaves constitute the negative class. The rectangular leaf region R is used to segment the primary leaf and the bounding box of R is used to compute the object classification scores.


Fig. 4: Five patches from an automatically segmented leaf (left). The diseased areas are highly contrasting in red-green space (right).

Fig. 5: A plot of AUC for primary leaf segmentation versus percentage of retained variance in the patch feature set. Values are the mean on the three validation sets.

Results for Primary Leaf Segmentation. The initial goal is to segment the primary leaf from the background and other clutter in the scene. See Figure 3, top row for an example of the output of the randomised forest. After segmentation the mean shape is applied using a local greedy search so as to maximise the the likelihood under the area. The search is initialised with the largest components of a fully covariant gaussian mixture model that is fitted to the likelihood output. Figure 3 bottom row for examples of the output of the mean shape fitting process.

Performance of patch classification and object segmentation is shown in Figure 5. The full patch based feature representations for the real (REAL) and real plus imaginary (REALIMAG) log-Gabor components are reduced in dimensionality using PCA in the standard manner. The feature sets are reduced so that one to four standard deviations of the feature variance is retain for each of the sets of REAL and REAL-IMG. From Figure 5 it can be seen that accuracy of patch classification levels out after two standard deviations, around 60 dimensions. However, the object classification peaks at around two standard deviations, this is attributed to a smoother output from the patch classifier assisting the shape fitting search in the presence of noise and dramatic lighting effects.

Results for Disease Analysis. The top 70% output of the leaf segmentation system is taken as input for disease presence and type analysis. The output is noisy and contains overlapping boundaries, occlusion and shadows. The segmented leaves where divided into the central 5 equally spaced, 64x64 pixel, patches, Figure 6 shows an example. The 750 patches were labelled according to there class; no disease, septoria tritici, yellow rust with an additional outlier class. While the variation in colour over the image set is large, areas of disease always exhibit a brownish yellow tint compared to the relatively uniform greenness of leaf areas without disease. A red minus green colour space was used to emphasise areas of disease and Haralick features [16, 18] were used to de-

scribe each patch. The features include; contrast, correlation, energy, homogeneity and entropy. Figure 6 top, shows the ROCs for each of the patch type as classified used a $k = 4$ nearest neighbours approach, with a $50{:}50$ train/test ratio.

In Figure 6 bottom, the final leaf classification is shown. For each segmented test leaf, the output of the nearest neighbour patch classification is subjected to winner takes all voting.
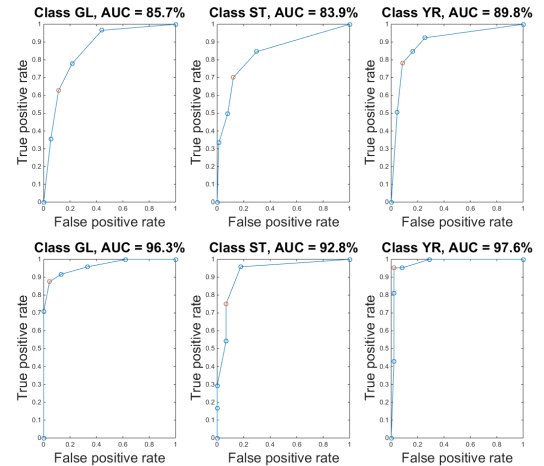

Fig. 6: The ROCs for patch classification (top) and for leaf type classification after winner takes all voting (bottom). The ROCs for each case of no disease ('gl'), septoria tritici ('st') and yellow rust ('yr') are shown

6. CONCLUSION

In this paper we show that framed leaves can be segmented in highly cluttered and variable images captured in the field. High dimensional global texture features are used to successful distinguish between the foreground primary leaf and the background and other clutter. A learnt leaf shape is then applied to segment the most likely leaf area. A nearest neighbours approach classifies components of the segmented leaf as either; no disease, septoria tritici or yellow rust. Finally, a voting system discards outliers and generates a robust disease type classification or diagnosis.

Future work will involve collecting a much larger number of disease and non-disease examples across a range of the most important crop types. The classification of more variable shaped leaves and fruits will be investigated. The visual structure and extent of disease type within individual crops will be analysed.

7. ACKNOWLEDGEMENTS

8. REFERENCES

[1] WHO, "http://www.who.int/globalchange/ publications/climatechangechap7.pdf," .

[2] FERA, "http://www.cropmonitor.co.uk/," .

[3] J.K. Patil and R. Kumar, "Advances in image processing for detection of plant disease," in *Journal of Advanced Bioinformatics Applications and Research*, 2011, vol. 2.

[4] A. Meunkaewjinda, N. Ratchasima, P. Kumsawat, K. Attakitmongcol, and A. Srikaew, "Grape leaf disease detection from color imagery using hybrid intelligent system," in *Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, 2008.

[5] K. Wangab, S. Zhangab, Z. Wangab, Z. Liuab, and F. Yang, "Mobile smart device-based vegetable disease and insect pest recognition method," in *Intelligent Automation and Soft Computing, Special Issue on Intelligent Information Technology in Agriculture*, 2013, vol. 19.

[6] W. Luo, *Spatial/Temporal Modelling of Crop Disease Data using High-Dimensional Regression*, Ph.D. thesis, Department of Statistics, 2008.

[7] L. Wang, W. P. Sousa, and P. Gong, "Integration of object-based and pixel-based classification for mapping mangroves with ikonos imagery," in *International Journal of Remote Sensing*, 2004, vol. 25, pp. 5655–5668.

[8] B. Southall, B. F. Buxton, J. A. Marchant, and T. Hague, "On the performance characterisation of image segmentation algorithms: a case study," in *The Sixth European Conference on Computer Vision*. Lecture Notes in Computer Science, 2000, vol. 1843, pp. 351–365.

[9] N. Campbell and B. Thomas, "Automatic selection of gabor filters for pixel classification," in *Sixth International Conference on Image Processing and its Applications*, 1997, pp. 761–765.

[10] N. Rose, "Facial expression classification using gabor and log-gabor filters," in *IEEE Automatic Face and Gesture Recognition*. IEEE, 2006, pp. 346–350.

[11] P. Kovesi, "Edges are not just steps," in *The Fifth Asian Conference on Computer Vision*. Lecture Notes in Computer Science, 2002, pp. 822–827, http://www.csse.uwa.edu.au/ pk/research/matlabfns/.

[12] Y. Freund and R. E. Schapire, "A short introduction to boosting," in *Journal of Japanese Society for Artificial Intelligence*, 1999, vol. 14, pp. 771–780.

[13] F. Moosmann, E. Eric Nowak, and F. Jurie, "Randomized clustering forests for image classification," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2008, vol. 30, pp. 1632–1646.

[14] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," in *International Journal of Computer Vision*, 2010, vol. 88, pp. 303–338.

[15] Matlab2014, "Treebagger class," .

[16] R. M Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," in *IEEE Transactions on Systems, Man, and Cybernetics*, 1973, vol. 62, pp. 610–621.

[17] O. Boiman, E. Shechtman, and M. Irani, "In defense of nearest-neighbor based image classification," in *International Conference on Computer Vision and Pattern Recognition*, 2008.

[18] P. Siricharoen, B. Scotney, P. Morrow, G. Parr, D. Gibson, and N. Canagarajah, "Effects of different mixtures of features, colours and svm kernels on wheat disease classification," 2014.