Hybrid Consensus Learning for Legume Species and Cultivars Classification

Mónica G. Larese(✉) and Pablo M. Granitto

CIFASIS, French Argentine International Center for Information and Systems Sciences, UAM (France) / UNR-CONICET (Argentina), Bv. 27 de Febrero 210 Bis, 2000 Rosario, Argentina
{larese,granitto}@cifasis-conicet.gov.ar

Abstract. In this work we propose an automatic method aimed at classifying five legume species and varieties using leaf venation features. Firstly, we segment the leaf veins and measure several multiscale morphological features on the vein segments and the areoles. Next, we build a hybrid consensus of experts formed by five different automatic classifiers to perform the classification using the extracted features. We propose to use two strategies in order to assign the importance to the votes of the algorithms in the consensus. The first one is considering all the algorithms equally important. The second one is based on the accuracy of the standalone classifiers. The performance of both consensus classifiers show to outperform the standalone classification algorithms in the five class recognition task.

Keywords: Legume and variety classification · Venation images · Consensus learning

1 Introduction

In the current literature, many approaches have been proposed aimed at performing automatic plant classification via leaf image analysis. The shape of the leaves [1,5–7,13], the color and the texture [20] or the combination of the three former traits [2,10] are the most common analyzed features.

Recently, when attempting to classify plant species by means of their leaves, several works have pointed out the importance of incorporating vein information among other features [8,18], or using them solely [16,17]. This last case is specially useful when the leaves share a similar appearance in shape, size, color and texture, as it happens when dealing with several varieties within the same species [15].

Previous related works in the literature which investigated the classification of leaves using exclusively the venation features focused on identifying only different species [16,17], but did not deal with the varieties recognition problem. Varieties identification is challenging since the differences in venation between

---

Author to whom all correspondence should be addressed.

L. Agapito et al. (Eds.): ECCV 2014 Workshops, Part IV, LNCS 8928, pp. 201–214, 2015.
DOI: 10.1007/978-3-319-16220-1_15

the classes are not identifiable by humans. However, it is reasonable to expect some venation differences caused by variety adaptation (e.g., draught tolerance). In a recent work, Larese *et al.* [15] proposed a method based on the classification of multiscale vein features in order to also evaluate the automatic recognition of soybean varieties. However, the species identification and the varieties identification were treated as two separated problems. In addition, only single classifiers were proposed to perform the classification task.

In contrast, in this paper we focus on the problem of classifying leaves which can belong to different species, or come from different varieties from one common species. Specifically, we experimented with 3 different legumes, namely red and white beans (*Phaseolus Vulgaris*), and soybean (*Glycine max (L) Merr*). For the latter case we have available specimens from 3 possible varieties. In consequence, we are dealing with a five-class multiclass problem. This problem is more difficult than the two 3-class problems considered in previous work [15], since species and varieties are now mixed up. We pursue a single procedure able to perform the leaf identification both for very different leaves, such as from different species, as well as leaves with similar appearance (leaf shape, size, color, texture), such as from different varieties from the same species.

Following the approach in Larese *et al.* [15], we characterize the leaves by means of multiscale morphological vein and areole features. Firstly, we process the leaf images by means of computing the Unconstrained Hit-or-Miss Transform (UHMT) [22] at different scales in order to highlight the veins. Next, we apply contrast enhancement and adaptive thresholding in order to segment the venation. After that, we measure different vein and areole multiscale morphological features in order to describe the venation morphology.

Our method is focused on developing automatic classification based exclusively on the characterization of the leaf venation system. For species recognition, these differences are recognizable in plain sight and can be approximately described by the human experts (e.g. differences in veins orientation). However, for the varieties recognition, the experts cannot establish the differences between the veins of each cultivar. Our method provides with an exploratory procedure or knowledge discovery tool which allows to determine if there exist different venation patterns for the different species and varieties, even though humans cannot identify them in plain sight.

Larese *et al.* [15] showed that it is possible to find some differences between the varieties, as well as how to analyze the relevance of the traits for the classes characterization. However, the accuracies reported using single classifiers in the recognition were not very high. This may be due to the existence of subtle differences between the cultivars, showing the need to improve the automatic classifiers. In this paper we investigate how to perform species and varieties classification of the features by means of consensus theory and multiple classifiers fusion, which have shown to improve accuracies in the literature [3,14,24]. Better classification methods could lead to an improved identification of the most informative distinctive features, which could be related to genotype differences. Within this context, we build a heterogeneous committee formed by five different automatic classifiers. The classifiers are aggregated through a majority vote rule in order to obtain a single prediction for the class. The five individual classifiers that we employ are of very diverse nature, consisting of Support Vector Machines [23], Penalized Discriminant Analysis [11], Näive Bayes [12] and two ensemble algorithms, namely Random Forests [4] and AdaBoost [9].

In this work, we consider two different strategies in order to weight the vote of each classifier in the committee. First of all, we consider all the algorithms equally important, *i.e.*, the vote of each algorithm is 1/5. The other strategy consists on weighting each classifier according to its accuracy obtained in a validation step. We show that we obtain better classification results by the aggregation of the algorithms into a committee of experts, compared to the ones achieved by the standalone classifiers. This is possible since our algorithm can take advantage of the contribution that each standalone algorithm provides to the consensus. When compared to manual classification, the hybrid automatic consensus is the only algorithm which outperforms the human experts for all the classes under consideration simultaneously.

The rest of the paper is organized as follows. In Section 2 we describe the dataset used in this work. We explain the segmentation and feature extraction steps in Section 3. We detail the individual classification algorithms, as well as the hybrid consensus strategy, in Section 4. We present and discuss the experimental results as long as the comparison of the performances for the different algorithms in Section 5. Finally, we draw some conclusions in Section 6.

2 Species and Varieties Dataset

The image dataset is composed by 866 color leaf images provided by Instituto Nacional de Tecnología Agropecuaria (INTA, Oliveros, Argentina). It consists of 272 images of red bean leaves, 172 images of white bean leaves (*Phaseolus Vulgaris*) and 422 images of soybean leaves (*Glycine max (L) Merr*). The soybean images are divided into three cultivars: 198 from cultivar #1, 176 from cultivar #2 and 48 from cultivar #3. They correspond to the images of the two first foliage leaves (pre-formed in the seed) of each specimen after 12 days of seedling grow. First foliage leaves were selected since their characteristics are less influenced by the environment. The leaves were acquired using a standard flatbed scanner (Hewlett Packard Scanjet-G 3110) at a resolution of 200 pixels per inch, and the images were stored as 24-bit uncompressed TIFF images. The abaxial surfaces of the leaves were scanned since veins appear stronger on this side. All the leaves lay in the same vertical position, thus avoiding significant rotation influences.

Figure 1 shows some exemplars from each one of the species and varieties which compose the dataset. The reader should notice that the differences between individuals from some of the classes do not compensate for the high variability also present between individuals within the same class. Thus, this application problem is characterized by relatively low inter-class and high intra-class variabilities.

3 Segmentation and Feature Extraction

In this paper we followed the segmentation and multiscale feature extraction procedures described by Larese *et al.* [15]. We worked with the gray scale image for each leaf. The leaf segmentation is based on the computation of the Unconstrained Hit-or-Miss Transform (UHMT)[22] on five leaf scale images (at 60%, 70%, 80%, 90% and 100% of the original image size). Each UHMT image highlights a different level of vein detail. Next, we added the five UHMTs (resized to the original size) to form the combined UHMT. This combined image highlights both small and large veins. In addition, we also picked up the UHMTs at scales 100% (the original image size), 80% and 60% for further processing.

On the four resulting UHMT images we applied contrast enhancement techniques (adaptive histogram equalization) and umbralization (adaptive thresholding) in order to obtain the segmented veins.

In order to discard the leaf shape contour, we cropped a centered patch of 100×100 pixels from each one of the four previously described segmented images. On these four patches, we measured the 208 vein and areole multiscale morphological features per leaf image (52 features $\times 4$ patches) described in the work by Larese *et al.* [15]. The 52 features are described in Appendix A for completeness.

4 Classification Algorithms

We considered 5 different classifiers, namely Support Vector Machines with Gaussian kernel, Penalized Discriminant Analysis, Näive Bayes, Random Forests and AdaBoost. We briefly describe them in the following. In the last subsection we present the consensus learning proposed approach, formed by the aggregation of the previously mentioned standalone classifiers.

We used inner 5-fold cross validation to optimize the parameters of the classifiers where necessary.

4.1 Support Vector Machines

Support Vector Machines (SVM) [23] is a state-of-the-art classifier which assumes that applying an appropriate nonlinear mapping of the data into a sufficiently high dimensional space, two classes can be separated by an optimum hyperplane. This decision hyperplane is chosen in such a way that the distance between the nearest patterns of different classes (i.e., the margin) is maximized. SVM depends on a regularization parameter, C, which controls the trade-off between the complexity of the classifier and the number of allowed misclassifications. In this work we used inner validation during the training phase in order to set this parameter.

The decision surface may be linear or nonlinear. In the latter case, a kernel function can be used to map the patterns into a high dimensional space. In this work, we considered SVMs with a Gaussian kernel (SVMG). We optimized the Gaussian standard deviation in a validation step during the training.

SVM is a binary classifier in nature. In order to extend its use to the present multiclass problem, we used the one-vs-one strategy.

4.2 Penalized Discriminant Analysis

Fisher's Linear Discriminant Analysis (LDA) [12] is a classical classifier and dimension reduction tool which searches for linear combinations of the features in such a way that the class means of the linear combinations are maximally separated relative to the intra-class variance. The classification of new observations is then performed by assigning them to the closest centroid according to a distance metric (typically the Mahalanobis distance) in the transformed space.

In order to improve LDA, Hastie *et al.* [11] proposed Penalized Discriminant Analysis (PDA). PDA is a regularized version of LDA, which adds a penalty term to the intra-class covariance matrix. PDA is useful for image classification problems with large number of highly correlated features.

In this work, we used standard Ridge Regression (GenRidge) [12], which has the ridge constant λ as the only free parameter. This constant penalizes high values of the fitted variables, and is similar to the C parameter in SVM. We automatically selected this parameter using a validation set in the training phase.

4.3 Näive Bayes

Näive Bayes (NB) [12] is a simple probabilistic classifier which assumes independence among the features, specially useful when dealing with a high-dimensional feature space. In spite of this unrealistic assumption, NB has shown to perform very well in real world applications.

The method applies the Bayes' Theorem in order to compute the *a posteriori* probabilities (class conditional probabilities) for a test observation. These probabilities are calculated as the product of the individual distributions for each feature (since they are assumed to be independent). In this work we assumed Gaussian distributions.

Once the probabilistic model is constructed, the classification is performed according to the Maximum a Posteriori decision rule (the most probable class is chosen).

4.4 Random Forests

Random Forests (RF) [4] is a state-of-the-art ensemble algorithm where the individual classifiers are a set of de-correlated trees. They perform comparably well to other state-of-the-art classifiers and are also very fast.

The algorithm constructs a set of unpruned trees from B random samples with replacement (bootstrap versions) of the original training dataset. For each random forest tree, a random sample of m variables from the full set of p variables ($m \leq p$) is selected to split the data at each node and grow the decision tree. The final classification result is the class corresponding to the majority vote of the ensemble of trees. In this work, we used 500 trees and a standard value of $m = \sqrt{p}$.

4.5 AdaBoost

Boosting classifiers are based on the idea that if many "weak" classifiers (slightly better than chance) are combined into a "strong" classifier, the overall performance will be highly improved [21]. AdaBoost [9] creates a sequence of weak classifiers aimed at discriminating the training observations. Initially, all the observations are assigned a unique weight. This distribution of weights is modified along with the number of iterations (rounds), *i.e.*, observations which are badly classified (more difficult to learn) are given higher weights. The algorithm attempts to find an optimum classifier at each round. Each weak classifier is weighted according to its performance on the current distribution of weights on the observations. At the end, the final strong classifier is the weighted linear combination of the weak classifiers. The algorithm minimizes the expectation of the exponential loss. In this work, we used 500 rounds and stumps as weak learners.

4.6 Hybrid Consensus of Experts

Consensus theory [3] combines multiple single probability distributions in order to build a unique predictive model from the opinion of several experts, assuming that the individual judgments are based on Bayesian decision theory.

Let $D = \{(\mathbf{x}_i, y_i)\}$, with $i = \{1, ..., n\}$, be a training dataset of n pairs of feature vectors $\mathbf{x}_i \in \mathbb{R}^p$ and class labels $y_i \in \{1, 2, ..., k\}$. The M probability distributions from each one of the M experts are combined using a so-called *consensus rule*. The consensus rule P_k for each pattern $\mathbf{x}_i$ is calculated for each one of the K classes. The simplest and most common consensus rule is the *Linear Opinion rule* (LOP), which is computed as a weighted linear combination of the posterior probabilities P_j from each expert (Eq. 1), with $j = \{1, ..., M\}$, and α_j denoting the weights associated to each expert. The coefficients α_j are non-negative and $\sum_{j=1}^{M} \alpha_j = 1$. The simplest approach of the weighting scheme consists in assigning the same weights to all the experts.

$$P_k(\mathbf{x}_i) = \sum_{j=1}^{M} \alpha_j P_j(\mathbf{x}_i) \tag{1}$$

Pattern $\mathbf{x}_i$ is then classified to belong to class $\hat{y}_i$, which is the class with the highest probability, as below.

$$\hat{y}_i = \max_{k=1}^{K} P_k(\mathbf{x}_i) \tag{2}$$

In this work, we propose to constitute a *Hybrid Consensus of Experts (HCE)* formed by one instance of each one of the five classifiers described in the previous subsections. This approach is also known in the literature as multiple classifiers fusion or combination of multiple classifiers [14,24].

In this context, we set P_j to be the probability predicted by each one of the $M = 5$ classifiers. We propose to use two different strategies in order to assign the importance to the opinion of each classifier. The first strategy considers them equally important, and therefore $\alpha_j = 1/5$. The second strategy takes into account the accuracy obtained by each classifier in a validation step. In this case we set α_j to be this accuracy value (a real value between 0 and 1 normalized by the sum of the α's). We call this second strategy *Weighted Hybrid Consensus of Experts (WHCE)*.

5 Experimental Results

In Fig. 2 we show one example leaf for each class, *i.e.*, 3 varieties of soybean, white and red beans. Below each leaf we show the corresponding segmented vein images. From this figure it can be noticed that there are many morphological differences between the veins of soybean and white or red beans. White and red beans present some visual differences, even though they are less strong than with respect to soybean. On the contrary, the three varieties of soybean look like very much the same. Moreover, we have encountered that in this case there is a high intra-class variability within each variety.

We filtered out the features having near zero variance. We scaled the remaining features in order to have zero mean and variance 1. For all the classification experiments, we computed the average classification accuracy after performing 10 runs of 5-fold cross validation. We also used 5-fold cross validation for inner validation to automatically optimize the parameters of the classifiers where necessary.

In Table 1 we show the average total accuracy computed for each one of the standalone algorithms as well as for the two hybrid consensus strategies. According to this table, it can be noticed that the most accurate results can be achieved by PDA, HCE and WHCE, with over 72%. The lowest value is obtained by NB (59%), whereas SVM, RF and AB show a similar performance around 71%.

In order to analyze the results with more detail, we present in Table 2 the average per class accuracies. These results are important since the five classes are not balanced, and we want to ensure that none of the implemented algorithms tends to favour majority classes. From this table it is noticeable that the easiest leaves to be recognized for all the algorithms are red and white beans. The distinction among the three soybean cultivars is the most difficult problem, as revealed by the general low accuracies obtained by all the methods. In the case of NB, the third cultivar of soybean is also well predicted. However, in this case it is to the detriment of cultivars #1 and #2 recognition, leading to a low overall accuracy as previously described in Table 1. With this respect, NB has a tendency to prefer soybean cultivar #3 over the other two soybean cultivars.

Table 1. Total accuracy (mean±S_E) for the five-class species and varieties classification problem

Classification Algorithm	Total accuracy (mean±S_E%)
SVM	71.05 ± 0.41
PDA	73.64 ± 0.39
NB	58.65 ± 0.42
RF	71.03 ± 0.40
AB	71.30 ± 0.39
HCE	72.45 ± 0.42
WHCE	72.97 ± 0.44

Table 2. Accuracy (mean±S_E) for the five-class species and varieties classification problem

Classification Algorithm	Per class Accuracy (mean±S_E%)				
	RBean (272 images)	WBean (172 images)	SBean#1 (198 images)	SBean#2 (176 images)	SBean#3 (48 images)
SVM	89.86 ± 0.60	82.96 ± 0.91	67.31 ± 1.16	51.41 ± 1.32	9.16 ± 1.37
PDA	91.54 ± 0.45	89.53 ± 0.72	62.73 ± 1.06	52.66 ± 1.12	37.00 ± 1.86
NB	84.60 ± 0.64	76.85 ± 1.11	22.21 ± 1.00	36.70 ± 1.25	77.13 ± 2.06
RF	89.27 ± 0.63	81.79 ± 0.93	66.60 ± 1.04	53.01 ± 0.93	13.38 ± 1.34
AB	89.75 ± 0.65	83.07 ± 0.85	64.29 ± 1.04	53.37 ± 1.26	19.40 ± 1.57
HCE	89.89 ± 0.55	86.97 ± 0.90	59.79 ± 1.05	53.13 ± 1.15	44.56 ± 2.33
WHCE	90.26 ± 0.54	87.26 ± 0.85	61.86 ± 1.07	53.58 ± 1.16	40.60 ± 2.08
Manual classification	83.28 ± 3.71	70.82 ± 13.15	44.95 ± 2.00	42.78 ± 5.37	43.98 ± 6.97

However, the poor performance of NB for cultivars #1 and #2 is compensated by the other 4 classification algorithms in the HCE and WHCE. Moreover, NB contribution helps the consensus to reinforce soybean cultivar #3 recognition, since this class is badly detected by all the standalone classifiers except for NB.

In our experiments, the best results for red and white beans are provided by PDA (92% and 90%, respectively), followed by WHCE and HCE, which are slightly lower. In the case of soybean cultivar #1 all the standalone algorithms (except for NB, as explained before) obtain better performance, being SVM and RF the best accuracies (67% for both of them). On the other hand, WHCE gets 62% and HCE 60%. Regarding the second cultivar, WHCE reaches an accuracy of 54%, closely followed by AB, HCE, RF and PDA. In the case of the third cultivar, WHCE and HCE obtain the highest accuracies. As previously explained, this value is only beated by NB. The closest accuracy value is 37% obtained by PDA, whereas the rest of the classifiers obtain less than 20% of accuracy.

When comparing HCE and WHCE to each other, we find that, as previously shown in Table 1, the average accuracy of WHCE is slightly higher. From Table 2 we can notice that the usage of the validation accuracies as weights in the consensus reduces the impact of the NB vote, slightly augmenting the accuracy of 4 classes, but resulting in a lower recognition of soybean cultivar #3.

For reference purposes only, we show the average classification results obtained by five human experts who manually classified the leaves. It is worth noticing that the experts solved two easier problems. Instead of classifying the leaves into 5 classes, they performed the classification into 3 classes, namely red bean, white bean and soybean. In a second independent experiment, they classified only soybean leaves into the three possible cultivars. From Table 2 it is evident that HCE obtained a better performance than the manual classification for all the five classes under consideration in the context of a more difficult problem. None of the five standalone classification algorithms could achieve this goal. WHCE obtains a lower performance only for cultivar #3.

The implementation of the consensus allows to improve the classification of the more difficult classes, while obtaining a good performance for the rest of them. Additionally, there is not need to choose between several different classifiers, since the consensus provides accuracies which are, at least, as good as the best standalone classifier in the committee, taking advantage of the goodness of each algorithm.

6 Concluding Remarks

In this work we propose to use a hybrid consensus approach in order to classify five different species and varieties leaves. We use multiscale morphological features extracted from the segmented leaf veins. The hybrid consensus is formed by five standalone classifiers, namely Support Vector Machines, Penalized Discriminant Analysis, Näive Bayes, Random Forests and AdaBoost. These classifiers are very different in nature, allowing to introduce diversity in the committee. This diversity helps to compensate the classification difficulties found by the algorithms in the complex task of plant recognition.

We implemented two different strategies to weight the vote of the classifiers in the consensus. The first one assigns the same importance to all the standalone classifiers. The other one assigns a different weight to the classification algorithms, which is based on the accuracy that they obtained during a validation step. Both strategies help to improve the recognition of the more difficult classes, and the second one provides with a slightly higher overall result.

The problem under analysis is difficult in the sense that within each class the leaf differences are high, and specially for the cultivars classification, even humans find it difficult to perform recognition since in this case the characteristics for each cultivar are not clear. The hybrid consensus is the only algorithm able to overcome human recognition for all the classes under analysis simultaneously. The usage of an automatic algorithm aimed at performing the classification provides with a reliable and repetible process, removing the tediousness of the task. Moreover, the classification procedure can help to highlight distinctive vein features which could be related to differences in the genotype of the species and cultivars under analysis.

We are currently working on the research of new features which can better characterize the veins. These features may include semantic information and relations between the vein branches. Additionally, we are considering adding patches from other locations of the leaf apart from the central patch used in this work.

Acknowledgments. MGL and PMG acknowledge grant support from ANPCyT PICT 2012-0181. We also acknowledge technical support from R. Craviotto, M. Arango and C. Gallo at Instituto Nacional de Tecnología Agropecuaria (INTA Oliveros, Argentina).

Appendix A

Following, the 52 morphological features computed on the veins and areoles are described. These features are an adaptation for classification purposes of the individual features proposed by Price *et al.* [19].

Feature #1: Total number of edges, *i.e.*, estimated veins.
Feature #2: Total number of nodes. The number of connecting nodes between edges.
Feature #3: Total network length. Total distance (in mm) along the skeleton of the vein image patch.
Features #4, #5, #6: Median/min/max edge length. The edge length (in mm) is the distance along the skeleton of a vein.
Features #7, #8, #9: Median/min/max edge width. The edge width (in mm) is the mean of the doubled distances between each skeleton pixel of the current edge and the nearest non-vein pixel, i.e., areole pixel.
Features #10, #11, #12: Median/min/max edge 2D area. The edge 2D area (in mm^2) is the sum of the widths at every skeleton pixel of the current edge times the length of one pixel.
Features #13, #14, #15:. Median/min/max edge surface area. The surface area (in mm^2) of the cylinder centered at the edge skeleton is computed as the sum of the individual surface areas for each skeleton pixel of the current edge, as $\sum_i SA_i = 2\pi(d_i/2)l_i$, where d_i is the diameter (width) and l_i is the length for a skeleton pixel i.
Features #16, #17, #18:. Median/min/max edge volume. The edge volume (in mm^3) corresponds to the volume of the same cylinder as in surface area, and is computed as $\sum_i V_i = \pi(d_i/2)^2 l_i$.
Features #19, #20, #21: Median/min/max edge orientation. The orientation is the angle in degrees between the x-axis and the major axis of the ellipse with the same second moments as the vein.
Feature #22: Total number of areoles in the image patch.
Features #23, #24, #25: Median/min/max areole perimeter. The perimeter (in mm) is the distance along the pixels of the border of the areole.
Features #26, #27, #28: Median/min/max areole area. The areole area (in mm^2) is the number of pixels in each areole times the area of one pixel.

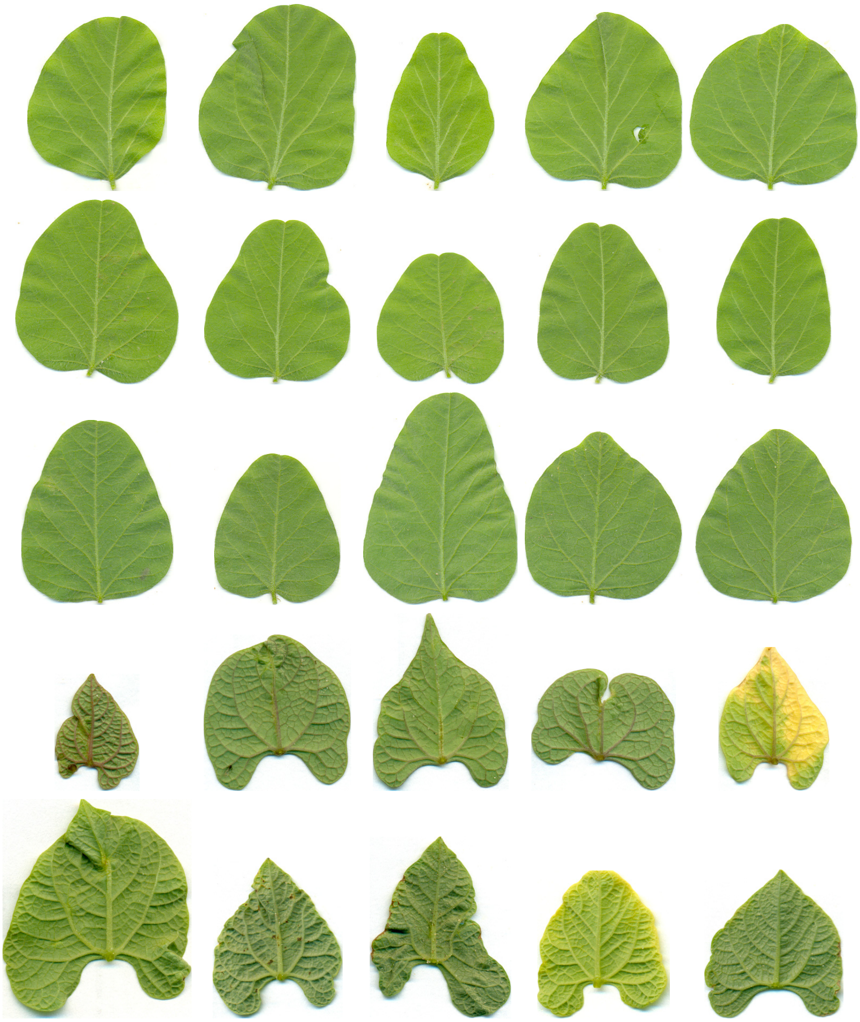

Fig. 1. Sample leaves from each class. First row: Soybean (cultivar #1). Second row: Soybean (cultivar #2). Third row: Soybean (cultivar #3). Fourth row: Red bean. Fifth row: White bean.

Features #29, #30, #31: Median/min/max areole convex area. The convex area (in mm^2) is the area of the convex hull for the areole.
Features #32, #33, #34: Median/min/max areole solidity. The solidity is a dimensionless parameter between 0 and 1 which measures the proportion of the pixels in the convex hull that are also in the area (ratio between the areole area and the convex area).
Features #35, #36, #37: Median/min/max areole major axis. The major axis (in mm) corresponds to the ellipse with the same normalized second

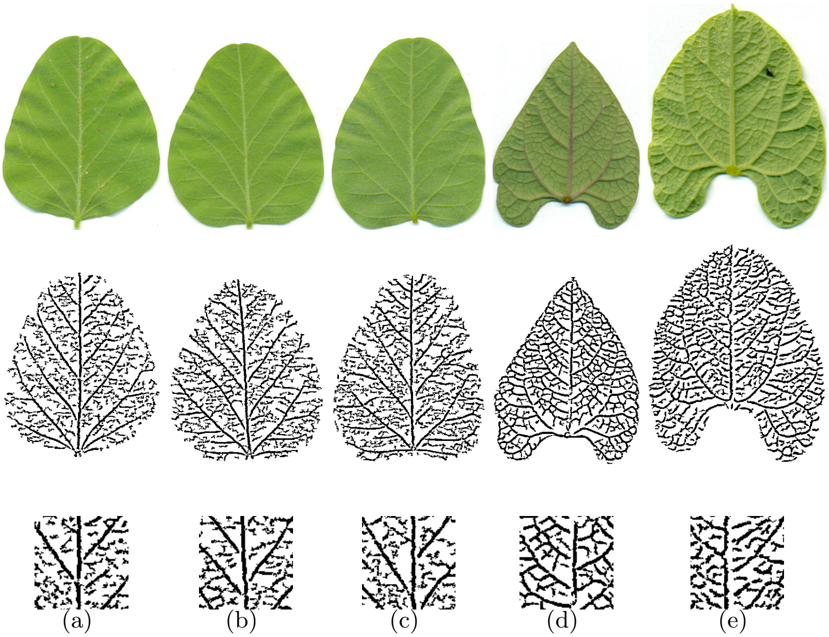

Fig. 2. Example leaves with their corresponding segmented veins and (amplified) cropped centered patches (only segmentation for the combined UHMT is shown). (a) Soybean (cultivar #1). (b) Soybean (cultivar #2). (b) Soybean (cultivar #3). (d) Red bean. (e) White bean.

moments as the areole.

Features #38, #39, #40: Median/min/max areole minor axis. The minor axis (in *mm*) corresponds to the ellipse with the same normalized second moments as the areole.

Features #41, #42, #43: Median/min/max areole eccentricity. The eccentricity is a dimensionless parameter between 0 (a circle) and 1 (a line), which measures the ratio of the distance between the foci of the ellipse having the same normalized second moments as the areole and its major axis.

Features #44, #45, #46: Median/min/max areole equivalent diameter. The equivalent diameter (in *mm*) is the diameter of a circle having the same area as the areole.

Features #47, #48, #49: Median/min/max areole mean distance. The mean distance (in *mm*) is the mean value of the Euclidean distances between each areole pixel and the nearest vein pixel.

Features #50, #51, #52: Median/min/max areole variance distance. The variance distance (in *mm*) is the variance of the Euclidean distances between each areole pixel and the nearest vein pixel.

References

1. Agarwal, G., Ling, H., Jacobs, D., Shirdhonkar, S., Kress, W., Russell, R., Belhumeur, P., Dixit, N., Feiner, S., Mahajan, D., Sunkavalli, K., White, S.: First Steps Toward an Electronic Field Guide for Plants. Taxon, Journal of the International Association for Plant Taxonomy **55**, 597–610 (2006)
2. Bama, B.S., Valli, S.M., Raju, S., Kumar, V.A.: Content based leaf image retrieval (CBLIR) using shape, color and texture features. Indian Journal of Computer Science and Engineering **2**(2), 202–211 (2011)
3. Benediktsson, J.A., Swain, P.H.: Consensus theoretic classification methods. IEEE Trans. Systems, Man Cybernet. **22**, 688–704 (1992)
4. Breiman, L.: Random Forests. Machine Learning **45**, 5–32 (2001)
5. Neto Camargo, J., Meyer, G.E., Jones, D.D., Samal, A.K.: Plant species identification using Elliptic Fourier leaf shape analysis. Computers and Electronics in Agriculture **50**, 121–134 (2006)
6. Chaki, J., Parekh, R.: Designing an automated system for plant leaf recognition. International Journal of Advances in Engineering & Technology **2**(1), 149–158 (2012)
7. Du, J.X., Wang, X.F., Zhang, G.J.: Leaf shape based plant species recognition. Applied Mathematics and Computation 185(2), 883–893 (2007), special Issue on Intelligent Computing Theory and Methodology
8. Du, J.X., Zhai, C.M., Wang, Q.P.: Recognition of plant leaf image based on fractal dimension features. Neurocomputing **116**, 150–156 (2013)
9. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. J. Comput. Syst. Sci. **55**(1), 119–139 (1997)
10. Golzarian, M.R., Frick, R.A.: Classification of images of wheat, ryegrass and brome grass species at early growth stages using principal component analysis. Plant Methods **7**, 28 (2011)
11. Hastie, T., Buja, A., Tibshirani, R.: Penalized discriminant analysis. Ann. Statist. **23**(1), 73–102 (1995)
12. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning, 2nd edn. Springer (2009)
13. Im, C., Nishida, H., Kunii, T.L.: Recognizing Plant Species by Leaf Shapes-A Case Study of the Acer Family. In: International Conference on Pattern Recognition, vol. 2, p. 1171 (1998)
14. Kuncheva, L.I., Bezdek, J.C., Duin, R.P.: Decision templates for multiple classifier fusion: an experimental comparison. Pattern Recognition 34(2), 299–314 (2001). http://www.sciencedirect.com/science/article/pii/S003132039900223X
15. Larese, M.G., Bayá, A.E., Craviotto, R.M., Arango, M.R., Gallo, C., Granitto, P.M.: Multiscale recognition of legume varieties based on leaf venation images. Expert Systems with Applications 41(10), 4638–4647 (2014), http://www.sciencedirect.com/science/article/pii/S0957417414000529
16. Larese, M.G., Craviotto, R.M., Arango, M.R., Gallo, C., Granitto, P.M.: Legume identification by leaf vein images classification. In: Alvarez, L., Mejail, M., Gomez, L., Jacobo, J. (eds.) CIARP 2012. LNCS, vol. 7441, pp. 447–454. Springer, Heidelberg (2012)
17. Larese, M.G., Namías, R., Craviotto, R.M., Arango, M.R., Gallo, C., Granitto, P.M.: Automatic classification of legumes using leaf vein image features. Pattern Recognition **47**(1), 158–168 (2014)

18. Park, J., Hwang, E., Nam, Y.: Utilizing venation features for efficient leaf image retrieval. J. Syst. Softw. **81**(1), 71–82 (2008)
19. Price, C.A., Symonova, O., Mileyko, Y., Hilley, T., Weitz, J.S.: Leaf Extraction and Analysis Framework Graphical User Interface: Segmenting and Analyzing the Structure of Leaf Veins and Areoles. Plant Physiology **155**, 236–245 (2011)
20. Pydipati, R., Burks, T.F., Lee, W.S.: Identification of citrus disease using color texture features and discriminant analysis. Computers and Electronics in Agriculture **52**, 49–59 (2006)
21. Schapire, R.E.: The Strength of Weak Learnability. Mach. Learn. **5**(2), 197–227 (1990)
22. Soille, P.: Morphological Image Analysis: Principles and Applications. Springer-Verlag (1999)
23. Vapnik, V.: The nature of statistical learning theory. Springer-Verlag (1995)
24. Xu, L., Krzyzak, A., Suen, C.: Methods of combining multiple classifiers and their applications to handwriting recognition. IEEE Transactions on Systems, Man and Cybernetics **22**(3), 418–435 (1992)