Sabanci-Okan System at ImageClef 2011: Plant identification task

Berrin Yanikoglu[1], Erchan Aptoula[2], and Caglar Tirkaz[1]

[1] Sabanci University, Istanbul, Turkey 34956
[2] Okan University, Istanbul, Turkey, 34959
berrin@sabanciuniv.edu
erchan.aptoula@okan.edu.tr
caglart@sabanciuniv.edu

Abstract. We describe our participation in the plant identification task of ImageClef 2011. Our approach employs a variety of texture, shape as well as color descriptors. Due to the morphometric properties of plants, mathematical morphology has been advocated as the main methodology for texture characterization, supported by a multitude of contour-based shape and color features. We submitted a single run, where the focus has been almost exclusively on scan and scan-like images, due primarily to lack of time. Moreover, special care has been taken to obtain a fully automatic system, operating only on image data. While our photo results are low, we consider our submission successful, since besides being our first attempt, our accuracy is the highest when considering the average of the scan and scan-like results, upon which we had concentrated our efforts.

Keywords: Plant identification, mathematical morphology, morphological covariance, Fourier descriptors, Support Vector machines.

1 Introduction

The plant identification task in ImageCLEF 2011 consisted of labelling images of plants that were captured by different means (scans, scan-like photos called pseudo-scans and unrestricted photos). The details of the recognition task are described in [6]. A content-based image retrieval (CBIR) system for plants would be very useful for plant enthusiasts or botanists who would like to learn more about a plant they encounter. The goal of the competition was to benchmark state-of-the-art in this open problem where there are very few systems for identifying unconstrained whole or partial plant images [9, 13]. The existing research in this area is concentrated on isolated leaf identification [4, 3, 10, 12, 15, 14].

Content-based plant identification problem faces many challenges such as color, illumination, size variations that are also common in other CBIR problems, as well as some specific problems such as the variations in the composition of the leaves that makes the plant shape variable. In addition, one can see that color is less identifying in the plant retrieval problem compared to many other retrieval

Fig. 1. Color variations of "Eurasian Smoketree".

problems, since most plants have green tones as their main color with subtle differences. In the rare cases that color is discriminative for a certain plant, that is when the plant has an unusual color, then it may be the case that the leaves of that plant may also have other colors, due to individual plant or seasonal variations (e.g. Gingko, Eurasian smoketree), as shown in Fig. 1. Another issue with color in plant identification is due to the challenges posed by the color of the flowers: a flowering plant should be matched despite differences in flower colors.

While shape is quite discriminative in identifying isolated leafs, it is not as useful in identifying full plant images, since the global shape of a plant is affected from its leaf composition [9]. In that regard, isolated leaf identification can be said to be a simpler problem compared to the unconstrained images of full or partial plants. One method to address this problem could be to extract an individual leaf image by segmenting the overall plant image. Texture on the other hand seems to be a more robust and useful feature category for plant identification, and is widely used in plant identification.

2   Overall architecture

Upon examination of the training images that were categorized according to the capture type, we observed that the scanned and scan-like categories were similar in difficulty and both seemed significantly easier compared to the photo category that included a larger variation in scale. Due to shortage of time, we decided to concentrate our efforts on the scan categories, while the photo category was tackled during the last week – which was insufficient for such a difficult problem.

The final system is designed as two separate sub-systems, one for scan and scan-like images and another one for photos. Since the meta-data included the acquisition type, an input image is automatically sent to the correct subsystem. The acquisition method was the only meta-data used in the overall system.

Based on our previous work on plant identification [9], we had some experience with the usefulness of different feature groups. For handling photographs, which was the problem addressed in our previous work, we found that global shape descriptors and many of the descriptors considered in the present work, would not be useful if the photo consisted of an overlapping set of leaves, rather than a single leaf. On the other hand for scan and scan-like categories, all three main feature categories are useful: color, texture and shape.

After experimenting with a large number of descriptors, we selected a 115-dimensional feature vector for the scan/scan-like sub-system and its 91-dimensional subset for the photo category. The features used in our system are explained in Section 3. For training the system, we trained a classifier combination using Support Vector Machines (SVMs), as explained in Section 4.

3   Feature extraction

As we are dealing with objects characterized mainly by their morphometric properties, whenever possible we attributed special preference to using morphological solutions. Since mathematical morphology, a nonlinear image processing framework, excels at shape based image analysis as well as at exploiting the spatial relationships of pixels.

An additional motivation in this regard has been to test our recently conceived morphological texture descriptors [1] in the context of a real-world application. Although a rich variety of content descriptors has been investigated, we present in this section primarily those that were included in the final system.

3.1   Texture Features

As far as scan and scan-like images are concerned, one can easily remark that we are dealing with relatively low scale variations, that can be easily countered with some form of normalization, while plant alignment is also not a major issue. Consequently, scale invariance set aside, from a texture description point of view, we require descriptors possessing a) a high discriminatory potential, b) illumination invariance, and unless we apply some form of angle normalization, then also c) rotation invariance.

When it comes to photos however, global texture characterization methods are bound to fail, since besides requiring all kinds of invariances, the background varies extremely in terms of complexity, thus presenting a considerable challenge. Hence, in order to apply any global morphological texture operators, a successful segmentation isolating the plant is necessary.

A set of novel morphological grayscale texture descriptors, possessing the aforementioned qualities has been recently introduced [1], leading to the highest classification scores among grayscale approaches, with a variety of texture benchmark collections, including Outex, CUReT and ALOT. They have been formulated as extensions for morphological covariance, that equip it with rotation and illumination invariance. Among them, we focused particularly on *circular covariance histograms* (CCH) and *rotation invariant points* (RIT). In summary, these two features achieve rotation invariance straightforwardly by replacing the point pairs of standard morphological covariance, with a circular structuring element (SE), possessing its center. Although any isotropic SE would suffice, this particular shape has the advantage of preserving the principle of covariance, consisting of comparing pixels at various distances.

As to illumination invariance, they take advantage of the complete lattice foundation of mathematical morphology. More precisely, morphological operators operate on pixel extrema, and not on their linear combinations. In other words, even if in a set of pixels the overall intensity levels change, as long as the relative order of pixels with respect to their intensity remains the same, the morphological operator under consideration, be it erosion, dilation or a combination thereof, will be unaffected, and will still pick as extremum the same pixel, albeit with a modified intensity value. That is why, conversely to granulometries and covariance, where a Lebesgue measure is used in order to quantify the morphological series, CCH and RIT rely on using directly the characteristic scale of each pixel. While the entire input image is described by means of its histogram of characteristic scales.

As to the difference of RIT from CCH, it is computed similarly, with the exception of first decomposing the circular SE into anti-diametrical point triplets. Thus there is an additional step of computing a label image, by means of a pixel based fusion. In particular, a rotation invariant measure is used to this end (e. g. minimum, maximum), upon all the intermediately filtered images by point triplets of various orientations.

3.2   Color Features

Since the previously chosen texture descriptors have not yet been extended to color data, it was decided to employ the *parallel* color texture description strategy, where color is described independently from texture. Among the investigated methods we can mention multi-resolution histograms based on morphological scale-spaces [2], in both polar and perceptual color spaces, non-uniformly sub-quantized saturation weighted LSH histograms, as well as color invariants and color moments. Yet, following an experimental evaluation of these color descriptors, only color moments have been included in the final system.

To explain, a color image corresponds to a function I defining RGB triplets for image positions $(x, y) : I : (x, y) \mapsto (R(x, y), G(x, y), B(x, y))$. By regarding RGB triplets as data points coming from a distribution, it is possible to define moments. Mindru et al. [11] have defined *generalized color moments* M_{pq}^{abc}:

$$M_{pq}^{abc} = \int \int x^p y^q [I_R(x, y)]^a [I_G(x, y)]^b [I_B(x, y)]^c dx dy \qquad (1)$$

M_{pq}^{abc} is referred to as a generalized color moment of order $p+q$ and degree $a+b+c$. This descriptor uses all generalized color moments up to the second degree and the first order, which leads to nine possible combinations for the degree: M_{pq}^{100}, M_{pq}^{010}, M_{pq}^{001}, M_{pq}^{200}, M_{pq}^{110}, M_{pq}^{020}, M_{pq}^{011}, M_{pq}^{002} and M_{pq}^{101}. These are combined with three possible combinations for the order, M_{00}^{abc}, M_{10}^{abc} and M_{01}^{abc}, which makes a 27-dimensional feature vector, possessing additionally shift-invariance, if the average is subtracted from all input channels.

3.3 Shape Features

Undoubtedly, shape plays a major role in plant identification and a plethora of shape descriptors, usually categorized as *region-* and *contour-based*, are available. In our case we employed a variety of shape descriptors from both categories.

Fourier descriptors We used the Fourier descriptors that are widely used to describe shape boundaries, as the main shape feature in our system. The Fourier Transform coefficients of a discrete signal $f(t)$ of length N is defined as:

$$C_k = \frac{1}{N} \sum_{t=0}^{N-1} f(t) e^{-j2\pi tk/N} \qquad k = 0, 1..., N-1 \tag{2}$$

In our case, $f(t)$ is the 8-directional chaincode of the plant, N is the number of points in the chaincode, and C_k is the k-th Fourier coefficient.

The coefficients computed on the chaincode is invariant to translation since the chaincode is invariant to translation. Rotation invariance is achieved by using only the magnitude of the coefficients and ignoring the phase information. Scale invariance is achieved by dividing all the coefficients by the magnitude of the DC component. We used the first 50 coefficients to obtain a fixed-length feature and to eliminate the noise in the leaf contour.

Width length/volume factor: These two descriptors are slight variations of the *leaf width factor* (LWF) introduced by Hossain and Amin [8]. Specifically, given an isolated leaf image (f), their method consists in dividing it into n strips, perpendicular to its major axis (Fig. 2). For the final n-dimensional feature, they compute the length of each strip (l_i), divided by the length of the entire leaf (l):

$$LWF_n = \{l_i/l\}_{1 \leq i \leq n} \tag{3}$$
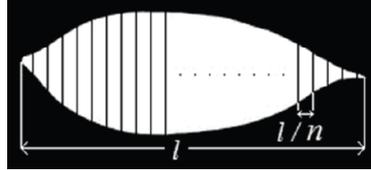
Fig. 2. Example of leaf decomposition [8].

We derived two new features from this. The *Width length factor* normalizes the lengths of each strip by the maximum width of the leaf. This is necessary as we normalize the length of each leaf into a fixed size during preprocessing, leaving the width as variable. The second derived feature is obtained by integrating into LWF the grayscale variations of each strip (f_i), thus obtaining *Width volume factor* (WVF). Specifically, we employ the ratio of volumes (i.e. sum of pixel values) instead of lengths:

$$WVF_n = \{\mathrm{Vol}(f_i)/\mathrm{Vol}(f)\}_{1 \leq i \leq n} \tag{4}$$

Convexity: This mono-dimensional feature aims to describe the overall contour smoothness of its binary input, which is assumed to be consisting of a single connected component. To explain, after isolating and binarizing the plant image, we compute its convex hull (CH) and then trivially derive its convexity:

$$Convexity(f) = \frac{\text{Area}(\text{CH}(f)) - \text{Area}(f)}{\text{Area}(f)} \tag{5}$$

Basic shape statistics: This descriptor (BSS) on the other hand, operates on the contour profile of its binary input. Specifically, we start by computing the center of mass of a given binary plant image (f), assumed to be consisting of a single connected component. Then we obtain its morphological internal gradient, computed by means of 3×3 square SE:

$$g_i(f) = f - \varepsilon(f) \tag{6}$$

Next, we calculate the Euclidean distances from the aforementioned center to each of the border pixels, which leads to a discrete series $S(f)$; in case of rotation of the input image, $S(f)$ is only shifted horizontally. Thus the final feature is obtained by means of 4 simple statistical measures on S:

$$BSS(f) = \{\max(S(f)), \min(S(f)), \text{med}(S(f)), \text{var}(S(f))\} \tag{7}$$

using its maximum (max), minimum (min), median (med) and variance (var). And since they are horizontally translation invariant w.r.t $S(f)$, they lead to a simple yet effective and rotation invariant description.

Border Covariance: Similarly to basic shape statistics, border covariance (BK) also operates on the contour profile $S(f)$ of its binary input, under the same assumptions. This time however, instead of computing simple statistical measures, we aim to capture contour regularity. To this end we employ morphological covariance, along with a horizontal pair of points. In other words, we treat the contour profile as a mono-dimensional texture.

We modified the standard morphological covariance operator so as to employ openings and closings instead of erosions, in order to capture respectively both bright details on dark background, as well as dark details on bright background:

$$BK_T^n(f) = \left\{ \frac{\text{Vol}(T_{P_{2,v}}(S(f)))}{\text{Vol}(S(f))} \right\}_v \tag{8}$$

where T denotes the morphological operator (either an opening γ or a closing ϕ), Vol the sum of pixel values of f and $P_{2,v}$ a pair of points separated by a vector **v**. Moreover, it should be noted that since we use a horizontal pair of points (i. e. translation invariant w.r.t the contour profile), the resulting feature is thus rotation invariant.

4 Classification

Scan and Scan-like Images

Mainly due to the collaboration of two universities, we trained two separate classifiers using two different sets of features. For the first classifier (Classifier$_1$), we used a 67-dimensional shape feature consisting of the 50 Fourier descriptors, the width length factor, eccentricity and solidity. For the second classifier (Classifier$_2$), we used a 115-dimensional feature consisting of all the contour, texture and shape features described in Section 3, excluding the Fourier descriptors. For both classifiers, we used an SVM using the radial basis function kernel.

The outputs of these classifiers are distances of the test instance to each plant class. We then trained a third classifier to learn how to combine these two classifiers at score level. Hence, the feature vector used in training the combiner is of length $2 \times K$, where K is the number of classes in the problem.

In the final stage of classification the most probable 5 classes are selected and a multi-class SVM is trained specifically for those classes for disambiguation. This was done because we found it beneficial to train classifiers that would learn to distinguish similar classes (e.g. different kinds of maples which are very similar amongs themselves, compared to the other plants). While the original idea was to train one such classifier according to the number of lobes in the leaves, the difficulty in assessing this information and the remaining complexity of this task led us to train a new classifier on the fly, using only the training instances from the 5 most probable classes and all of the 182 (=115+67) features. We use the outcome of this stage as the final classification decision. Cross-validation accuracies obtained for each classifier using the training data set are summarized in Table 1.

Table 1. Cross-validation accuracies of each classifier used for the scan and scan-like categories.

Stage	Accuracy (%)
Classifier$_1$ using only shape features	71.46
Classifier$_2$ using all features (except Fourier descriptors)	89.69
Classifier combination	90.10
After resolving ambiguities	93.64

Photographs

As far as photos are concerned, due to time constraints, neither their feature extraction nor their classification received the attention they deserved. Since shape features were not used, we trained only a single SVM classifier, using a 91-dimensional feature vector. For this classifier, default parameters (cost = 25, first degree polynomial kernel) of the Weka SVM software [7] and the Sequential Minimal Optimization (SMO) algorithm were used.

5   Experimental design

In this section we describe the implementational choices that have been made, as well as the experiments that have been carried out, while designing & optimizing our plant identification system. The majority of experimentations has been realized using Weka (v.3.6.4) [7]. In all stages of our experiments and in the results given in throughout this paper, we used cross-validation on the training data and measured the average accuracy obtained across images, rather than the average user-based accuracy used in the competition.

For practical reasons, we chose to first optimize our descriptor combination, then the preprocessing scheme, followed finally by the classification step. Given their visual similarity, we handled scan and scan-like images identically, while treating photos separately. Special care has been taken to obtain a fully automatic system.

For the sake of simplicity, and in order to minimize the number of variables, initial feature selection experiments have been carried out with a nearest neighbour (1NN) classifier along with the χ^2-distance, while subsequently we switched to Support Vector Machines (SVM). Whenever necessary, conversion to grayscale has been realized using the weighted combination of RGB channels, as $0.299 \times R + 0.587 \times G + 0.114 \times B$, while binarization has been achieved using Otsu's threshold.

5.1   Feature Selection

Preliminary experimentation with various features has been realized by handling shape, color and texture descriptors separately, in an effort to determine the most suitable among them for the problem under consideration. After this step, we experimented with their combination and parameterization.

Scan and scan-like images: At this stage a relatively simple preprocessing step was realized, consisting of first extracting the bounding rectangle of the plant, followed by scale normalization resulting in a fixed height of 600 pixels. This was applied to all 71 classes containing a total of 3066 scan and scan-like images. Then a series of cross-validation experiments took place, in an effort to determine the most suitable descriptors for distinguishing among these classes. The results are shown in Table 2, along with their arguments.

The accuracy scores have been obtained by dividing the available data randomly into train (1444 samples) and test (1622 samples) sets, using the aforementioned classification settings. Interestingly, one can observe that texture exhibits the highest discriminatory potential, followed by color and shape.

In addition to measuring the individual discriminatory potential of each feature, we experimented with many of their combinations. The resulting scores are given in Table 3, where we present the classification accuracies obtained with various combinations of feature sets.

Table 2. The descriptors used for scan and scan-like images along with their accuracies.

ID	Descriptor	Applied in	Size	Accuracy (%)
1	$CCH_{\gamma,12,\#1} + CCH_{\phi,12,\#1}$	Grayscale	12+12	72.32
2	$RIT_{\gamma,12,min} + RIT_{\phi,12,min}$	Grayscale	12+12	53.45
3	WVF_{11}	Grayscale	11	44.94
4	Convexity	Binary	1	7.89
5	BSS	Binary	4	37.30
6	$BK_{\gamma}^{n} + BK_{\phi}^{n}$	Binary	12+12	21.33
7	Color moments	RGB	27	49.94
	Total		115	

Table 3. Accuracies (%) for the combinations of the descriptors given in Table 2, for scan and scan-like data.

Descriptors	1NN	SVM/SMO
1,3	79.59	80.27
1,2,3	85.02	86.93
1,2,3,4	85.45	87.11
1,2,3,4,5	87.36	89.15
1,2,3,4,5,6	88.59	90.20
1,2,3,4,5,6,7	88.41	**90.57**

Photographs: Conversely to scan and scan-like images, the main challenge presented by photos lies in isolating the plant from its often very complicated background. Due mainly to lack of time, we hardly had any chance of constructing an optimized feature set for this image category, as done for the other images. Instead, we transferred almost directly our descriptor choices for scan and scan-like data, with no additional preprocessing whatsoever. Nonetheless, one of the very few experiments with photos that has been carried out, consisted of simply testing the combination of the features given in Table 2, with the end of adapting them to the new content.

In particular, we joined the scan and scan-like images with photos, thus obtaining a total of 3996 samples, and divided them equally and randomly into training and test sets. The classification accuracies are provided in Table 4. Considering the background complexity of photos, contour-based shape descriptors suffered a significant performance loss; which was expected, since they rely heavily on correct border extraction. Border covariance in particular has been unable to contribute any longer, so it has been removed from the set of descriptors used to characterize photos. Consequently the length of the feature vector used with photos is $115 - 24 = 91$. In summary, the addition of photographs has decreased the overall classification performance considerably, with shape descriptors being affected the most.

5.2 Preprocessing

Having determined a set of features for describing the plant collection, we focused on optimizing the preprocessing stage, in order to further improve performance.

Table 4. Cross-validation accuracies (%) for the combinations of the descriptors of Table 2 with all images, using 1NN and SMO classifiers.

Descriptors	1NN	SMO
1,3	70.23	71.03
1,2,3	74.62	74.98
1,2,3,4	75.59	76.77
1,2,3,4,5	77.02	78.67
1,2,3,4,5,6	75.01	76.44
1,2,3,4,5,7	79.57	**81.51**

Besides the already applied scale normalization, we considered illumination normalization as well, through histogram equalization as proposed in Ref. [5]. However this did not contribute in any substantial way, probably due to the fact that our primary descriptors are already illumination invariant. Moreover, the removal of the leaf petiole has also been tackled, with the end of obtaining a more accurate plant border for contour-based shape descriptors, but unfortunately its effect has been negligible, probably due to the partial success of the removal procedure.

Furthermore, although most of our operators are rotation invariant, WVF and color moments are not. That is why, we chose to apply an additional step, that would align the plant vertically along its major axis. However, the small gains were hindered by mistakes in the angle estimate; thus this normalization did not improve our classification rates. Consequently, as far as scan and scan-like images are concerned their preprocessing consists of extracting the bounding rectangle of a given plant image, followed by its scale normalization to a fixed height of 600 pixels.

Photographs on the other hand, given their content variation, should benefit from a background removal step. Due to time constraints, we experimented with relatively simple and automatic hue based background removal; but it did not contribute significantly to classification performance, and therefore were not included in the final system's operation. In short, photographs were not preprocessed in any way.

5.3 Classifier Optimization

Having determined the descriptors and preprocessing operations for each sub-problem (scan and photo), we optimized the parameters of the classifiers. For the two base classifiers used in scan/scan-like categories, the cost and spread parameters (C and γ) of the SVM are learned using 5-fold cross validation and grid search. As for the combiner, we first divided the training set into half. Then, we trained the two classifiers with their optimum parameters on the first half and used the second half to produce distances. In the next step, we found the optimum parameters for the combiner using 5-fold cross validation and grid search on the parameters using the produced distances. Furthermore, as we analyzed the errors of the system on training data, we decided to add a final classifier which is trained with only the closely scoring (top-5 classes) classes' instances,

as explained in Section 4. For the single classifier used in photos, we used the default SVM parameters due to shortness of time.

6   Results and discussion

According to the official results (Table 5) our run achieved 6^{th} place in overall classification, 2^{nd} place with scan type images, 6^{th} place with scan-like images and 16^{th} place with photographs. However, we did achieve the best score when considering the average of the scan and scan-like results, upon which we had concentrated our efforts.

Nevertheless, although the placements came as no surprise, the scores on the other hand have been lower than our expectations, especially when compared with the values obtained during our cross-validation tests with the training data (Tables 3 and 4). As a matter of fact, there is a very significant drop of approximately 40% overall. This important difference may be due to several factors, one of which is overfitting of the classifiers. Our descriptor optimization stage may have very well led to excessively powerful features capable of distinguishing the training data quite effectively, yet when faced with the test dataset, containing distinct plants of the same genus, the same features failed to generalize their performance. Furthermore, the scoring function is no longer the average accuracy across images, but instead the average accuracy obtained per user; and of course there is always the possibility of implementation errors.

In conclusion, given that this is our first participation, we consider our attempt satisfactory and even successful, in the sense that we accomplished our main goal; which consisted of identifying effectively the plants in scan and scan-like images. All the same, our score is far from being perfect. Future work in this category will include testing non-morphological descriptors, in an attempt to harness the advantages of both non-linear and linear image analysis methodologies. As far as photographs are concerned, our main focus will be on the preprocessing stage, with the end goal of isolating the plant effectively from its background, so as to be able to employ the same optimized features as with the other image types.

References

1. E. Aptoula. Extending morphological covariance. *Pattern Recognition*, May 2011. Submitted.
2. E. Aptoula and S. Lefèvre. Morphological description of color images for content-based retrieval. *IEEE Transactions on Image Processing*, 18(11):2505–2517, November 2009.
3. A. R. Backes and O. M. Bruno. Shape classification using complex network and multi-scale fractal dimension. *Pattern Recognition Letters*, 31(1):44 – 51, 2010.
4. O. M. Bruno, R. O. Plotze, M. Falvo, and M. Castro. Fractal dimension applied to plant identification. *Information Sciences*, 178(12):2722 – 2733, 2008.

Table 5. Partial plant retrieval results showing the top-performing systems ordered according to mean performances.

Run	Scan	Scan-like	Photo	Mean	Scan & Scan-like
IFSC USP_run2	0.562	0.402	**0.523**	**0.496**	0.482
inria_imedia_plantnet_run1	**0.685**	0.464	0.197	0.449	0.574
LIRIS_run3	0.546	0.513	0.251	0.437	0.529
LIRIS_run1	0.539	0.543	0.208	0.430	0.541
IFSC USP_run1	0.411	0.430	0.503	0.448	0.420
Sabanci_Okan-run1	0.682	0.476	0.053	0.404	**0.579**
LIRIS_run2	0.530	0.508	0.169	0.403	0.519
LIRIS_run4	0.537	0.538	0.121	0.399	0.537
inria_imedia_plantnet_run2	0.477	**0.554**	0.090	0.374	0.515
...	...	...	...	...	...

5. G. D. Finlayson, S. Hordley, G. Schaefer, and G. Y. Tian. Illuminant and device invariant colour using histogram equalisation. *Pattern Recognition*, 38(2):179–190, February 2005.

6. H. Goëau, P. Bonnet, A. Joly, N. Boujemaa, D. Barthelemy, J. Molino, P. Birnbaum, E. Mouysset, and M. Picard. The clef 2011 plant image classification task. In *CLEF 2011 working notes*, Amsterdam, The Netherlands, 2011.

7. M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. *SIGKDD Explorations Newsletter*, 11(1), June 2009.

8. J. Hossain and M. A. Amin. Leaf shape identification based plant biometrics. In *International Conference on Computer and Information Technology*, pages 458–463, Dhaka, Bangladesh, December 2010.

9. H. Kebapci, B. Yanikoglu, and G. Unal. Plant image retrieval using color, shape and texture features. *The Computer Journal*, 53(1):1–16, April 2010.

10. F. Lin, C. Zheng, X. Wang, and Q. Man. Multiple classification of plant leaves based on gabor transform and lbp operator. In *International Conference on Intelligent Computing*, pages 432–439, Shanghai, China, 2008.

11. F. Mindru, T. Tuytelaars, L. Van Gool, and T. Moons. Moment invariants for recognition under changing viewpoint and illumination. *Computer Vision and Image Understanding*, 94(1-3):3–27, April-June 2004.

12. Z. Wang, Z. Chi, and D. Feng. Shape based leaf image retrieval. *IEE Proceedings - Vision, Image and Signal Processing*, 150(1):34–43, Feb 2003.

13. D. M. Woebbecke, G. E. Meyer, K. Von Bargen, and D. A. Mortensen. Plant species identification, size, and enumeration using machine vision techniques on near-binary images. In *Optics in Agriculture and Forestry*, volume 1836, pages 208–219, Boston, USA, 1993. SPIE.

14. Itheri Yahiaoui, Nicolas Hervé, and Nozha Boujemaa. Shape-based image retrieval in botanical collections. In *PCM*, pages 357–364, 2006.

15. S. Yonekawa, N. Sakai, and O. Kitani. Identification of idealized leaf types using simple dimensionless shape factors by image analysis. *Transactions of the ASAE*, 39:1525–2533, 1996.