Journal of Mobile Multimedia, Vol. 9, No. 1&2 (2013) 145–154

A MOBILE APPLICATION FOR ROBUST FEATURE EXTRACTION AND CULTIVAR CLASSIFICATION OF LEAVES[a]

DOMINIK L. MICHELS[b]
Multimedia, Simulation and Virtual Reality Group
University of Bonn, Germany
michels@uni-bonn.de

GERRIT A. SOBOTTKA
Multimedia, Simulation and Virtual Reality Group
University of Bonn, Germany
sobottka@cs.uni-bonn.de

We illustrate the development of an application for cultivar classification of leaf images based on the extraction of the network of its main veins that runs on mobile devices like smart phones or tablets. Such mobile devices can be docked to farming robots in order to support the farming process. Our application uses an efficient Gabor filter-based tracing algorithm which is able to perform a robust network extraction. The results are used as input data for the classification with a support vector machine.
In order to demonstrate the advantageous behavior and the robustness of this method, we perform an evaluation on a test set consisting of 150 light transmitted images of different vine leaves.

Keywords: Mobile Application, Leaf Classification, Feature-Based Classification, Feature Extraction, Gabor Filter, Edge Tracing, Support Vector Machines

1 Introduction

Population growth, climate change, and the shortage of resources have caused an increased global interest of the agricultural community in intelligent farming methods. As a result, the integration of agricultural concepts and modern IT has paved the way for tremendous crop yield increases over the last decade. Dedicated robots for farm working, usually four-wheeled vehicles with robot manipulators, have been developed as part of the smart farming process. Equipped with recent satellite and sensor technologies they are able to autonomously navigate through vineyards, corn- or strawberry fields and at the same time take over the sowing and harvesting work of the agricultural laborer. Exhausted, depleted, and pesticide contaminated soils, on the other hand, reveal the disastrous consequences of the long-term use of classical monocrops and force modern agriculture more and more to embark on whole system approaches like sustainable agriculture, integrated farming, and permacultures, i.e. well-designed agriculturally productive ecosystems with the diversity, long-term stability, and

[a]Extended version of Dominik Ludewig Michels and Gerrit Alexander Sobottka, "A Gabor Filter-Based Approach to Leaf Vein Extraction and Cultivar Classification", Lecture Notes in Computer Science, Proceedings of the International Conference on Computational Science and Its Applications, ICCSA 2013, June 24-27, 2013, Ho Chi Minh City, Vietnam. Pages 150-159.
[b]The authors are affiliated with the Multimedia, Simulation and Virtual Reality Group at the Institute of Computer Science II of the University of Bonn, Friedrich-Ebert-Allee 144, 53113 Bonn, Germany.

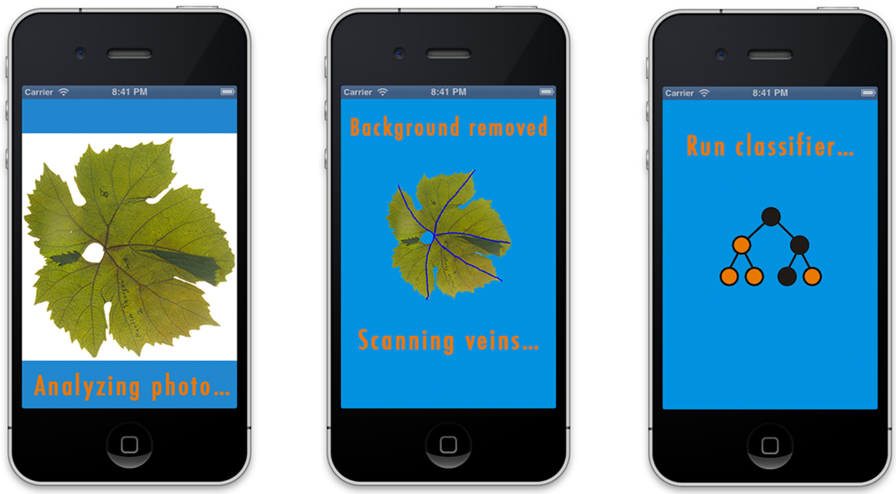

Fig. 1. Illustration of our leaf classification system running on a mobile device emulator.

synergistic properties of natural ecosystems. This poses new requirements for the autonomous harvesting vehicles: it is mandatory to furnish them with robust identification mechanisms that ensure a correct plant recognition by means of their phenotypic characteristics. Consequently, non-destructive approaches to the problem of computer aided analysis and screening of plant phenotypes have experienced a growing interest in the crop science community over the last decade. Matured computer vision techniques are employed for an automated plant recognition by means of the vein networks of their leaves which act as a unique classifier–a kind of fingerprint–for a specific cultivar.

In our contribution we tackle the challenging problem of extracting the main veins of vine leaves by means of a Gabor filter approach. The extracted vein data is a unique classifier for a certain cultivar and can thus be used as input for support vector machines to perform the final classification. We evaluate the robustness of our method on a test set consisting of 150 images of vine leaves with different color patterns. Since mobile devices like smart phones and tablet computers have become ubiqituous and are available at extremely low prices, it is a convenient and cost-effective way to use such hardware as on-board equipment for farming robots instead of special hardware produced on demand. We show that today it is already possible to run on such mobile devices—without any limitation—complex computation software like our classification ansatz, cf. Fig. (1).

2 Related Work

Leaf vein extraction and cultivar classification is a well-established field in the computer vision and machine learning community. In this section we give a brief overview over the recent achievements.

In [9] the authors present a leaf vein extraction method based on the gray-scale morphology. An independent component analysis is used in [5] to realize a robust vein extraction method.

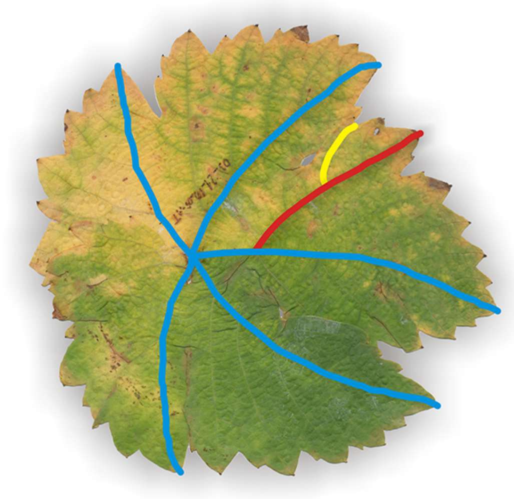

Fig. 2. Illustration of different level veins colored in blue (corresponds to level 0), red (corresponds to level 1), and yellow (corresponds to level 2).

A leaf recognition algorithm based upon probabilistic neural networks is presented in [8]. Their approach allows for a robust plant classification. In contrast, the method described in [4] makes use of region-based features. A rather exotic approach is presented in [1] where the authors discuss an ant colony algorithm.

In contrast to these methods we aim to exploit knowledge from the position space on the first hand and extend it with information from the frequency domain. This is achieved on the firm basis of appropriate Gabor filters that have originally been developed by Dennis Gabor in the last century.

In order to find the point of intersection of the five main veins, the configuration of a starting template is determined by means of a principal component analysis (cf. [6]) and a subsequent optimization step based on the simulated annealing method that has been introduced in [3].

2.1 Problem Setting

Our goal is to automatically extract the network of main veins in vine leaves, the so called veins of level zero. A vein emanating from a level n vein is of level $n+1$, cf. Fig. (2). It is important to note that we make intensive use of the fact that the phenotypical appearance of vine leaves always shows five level zero veins. These veins have a common start point–the so called central point of the leaf–and a well-defined endpoint. The extraction shall be performed automatically on a input image. The desired output is a vectorized representation of the veins of level zero as a pixel sequence or as a so called *chain code*.

We demonstrate the efficiency of our approach on a test set containing 150 images of vine leaves of different cultivars, specifically Kerner (35 images), Müller-Thurgau respectively Rivaner (38 images), Riesling (40 images), and Scheurebe (37 images). The images show diverse characteristics like overlapping parts, handwritten labels, and severe discolorations, e.g. due to leaf diseases, cf. Fig. (3).

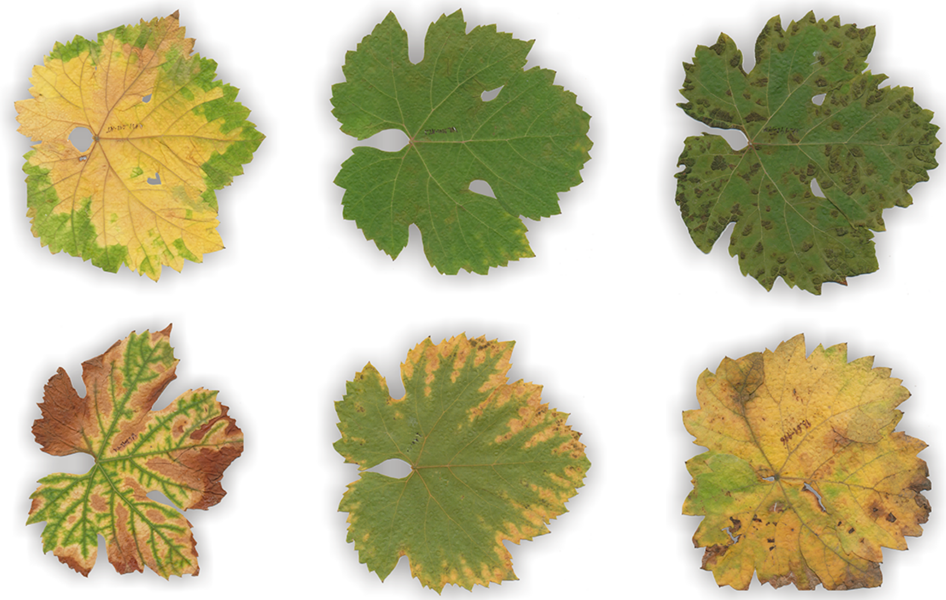

Fig. 3. Some leaves from our test set. Diverse characteristics, like overlapping parts, handwritten labels, and severe discolorations can be identified.

One of the major problems we have to deal with is that the difference in the intensities of the level zero veins and their higher order branches is not large enough to prevent a standard vein detection algorithm from spuriously changing its tracing direction from the principal to a secondary direction at the branch-offs. Beside this level zero veins typically tend to become smaller with increasing distance from the center point. As a consequence, it is difficult to distinguish between veins of different levels. This in turn promotes erroneous classifications of the vein levels. Here, our adaptive Gabor filter-based approach comes into play.

The choice to take veins of level zero is well founded due to the fact, that it has become an established phenotypic feature. For example it is recorded under the number OIV 070-1 by the International Plant Genetic Resources Institute (cf. [2]).

3 A Gabor Filter-Based Approach

In order to overcome the aforementioned problems we aim at combining knowledge from the position space as well as the frequency domain. This is achieved with so called Gabor filters that are closely related to short-time Fourier transforms. In order to allow the reader to acquire a deeper understanding of how our approach works we give a brief overview of the mathematical concepts behind Gabor filters and explain how they can be effectively applied in the context of leaf vein detection.

3.1 Gabor Filter

Appropriate transformations to the frequency domain allow us to study the frequency information of a time dependent phenomenon. In the one-dimensional, continuous case such a transformation is given by the Fourier transform

$$\mathcal{F}[f](\omega) = \int_{\mathbb{R}} f(t)\,\mathsf{exp}(-2\pi i\omega t)\,\mathsf{d}t,$$

which maps the signal f in the time domain to the Fourier transformed version $\mathcal{F}[f]$ in the frequency domain. This transformation is represented by a scalar product in the functional space with the arguments f and a complex exponential with frequency ω. Hence, $\mathcal{F}[f](\omega)$ can be considered as the complex amplitude of the occurrence of the fundamental oscillation with frequency ω in the signal f. In other words, the graph of $\mathcal{F}[f]$ shows how much of the signal f lies within each given frequency.

But there is no time domain information available in $\mathcal{F}[f]$. Since we are interested in a combination of time and frequency information, it is a common method to add a τ-shifted time domain window function g to the first argument of the scalar product in the sense of a multiplication with the signal f. Therefore, the resulting so called windowed or short-time Fourier transform

$$\mathcal{F}_g[f](\omega,\tau) = \int_{\mathbb{R}} f(t)\,g(t-\tau)\,\mathsf{exp}(-2\pi i\omega t)\,\mathsf{d}t,$$

describes the frequency behavior of the signal f in the time domain neighborhood of τ.

The transformation $\mathcal{G}[f] := \mathcal{F}_{g_\sigma}[f]$ with the Gaussian window function

$$g_\sigma(t) = \frac{1}{\sqrt{2\pi\sigma^2}}\,\mathsf{exp}\left(-\frac{t^2}{2\sigma^2}\right),$$

with variance σ is known as the Gabor transformation of the signal f and shows how much of the signal f limited to $t \in [\tau-\sigma, \tau+\sigma]$ matches a given frequency. An arbitrarily close resolution in the time as well as in the frequency domain is not possible at the same point in time and limited by Heisenberg's Uncertainty principle. One can show that the Gabor transformation is a Fourier transformation with minimal uncertainty.

The Gabor filter G_θ with orientation θ used in this approach is formally given by

$$= G_\theta(\mathbf{x},\lambda,\sigma,\psi,\gamma) = \mathsf{exp}\left(-\frac{\hat{x}_1^2+\gamma^2\hat{x}_2^2}{2\sigma^2}\right)\,\mathsf{exp}\left(i\left(\frac{2\pi}{\lambda}\hat{x}_1+\psi\right)\right), \tag{1}$$

at the point $\mathbf{x} = (x_1, x_2)^\mathsf{T}$. It can be regarded as an oriented two-dimensional discrete version of the Gabor transformation $\mathcal{G}$.

The left factor in Eqn. (1) describes a two-dimensional elliptic Gaussian in which the spatial aspect ratio γ denotes the ellipticity. The variance is again given by σ and the direction is determined by the angle θ of the normal direction that influences the vector $\hat{\mathbf{x}} = \mathbf{R}_\theta^\mathsf{T} \cdot \mathbf{x}$. The matrix $\mathbf{R}_\theta$ describes the two-dimensional mathematically positive rotation by the angle θ, which is given by

$$= \mathbf{R}_\theta = \begin{pmatrix} \mathsf{cos}\theta & -\mathsf{sin}\theta \\ \mathsf{sin}\theta & \mathsf{cos}\theta \end{pmatrix}.$$

The right factor in Eqn. (1) analogously describes a complex exponential with phase shift ψ and frequency $\omega = 1/\lambda$, in which λ denotes the wavelength.

With the use of Euler's formula we can split the complex Gabor filter Eqn. (1) in a real and an imaginary part. Only the real part is evaluated in our algorithm and given by

$$= G_{\mathrm{real}_\theta}(\mathbf{x}, \lambda, \sigma, \psi, \gamma) = \mathsf{exp}\left(-\frac{\hat{x}_1^2 + \gamma^2 \hat{x}_2^2}{2\sigma^2}\right) \mathsf{cos}\left(\frac{2\pi}{\lambda}\hat{x}_1 + \psi\right). \tag{2}$$

The Gabor filter G_{real_θ} is applied to an image F by using the two-dimensional convolution $F \otimes G_{\mathrm{real}_\theta}$, pointwise defined by

$$= F(x_1, x_2) \otimes G(x_1, x_2) = \sum_{\tau_1} \sum_{\tau_2} F(\tau_1, \tau_2) G(x_1 - \tau_1, x_2 - \tau_2).$$

The main idea behind our approach is that for a given angle θ a θ-oriented vein of order zero will always dominate the Gabor filter response in contrast to the higher order veins that branch off. We exploit this fact by tracing the course of the veins on $F \otimes G_{\mathrm{real}_\theta}$.

3.2 The Algorithm

This motivates the following algorithm: In a precomputing step the central point of the leaf is detected with a template-based matching strategy as illustrated in Fig. (4).

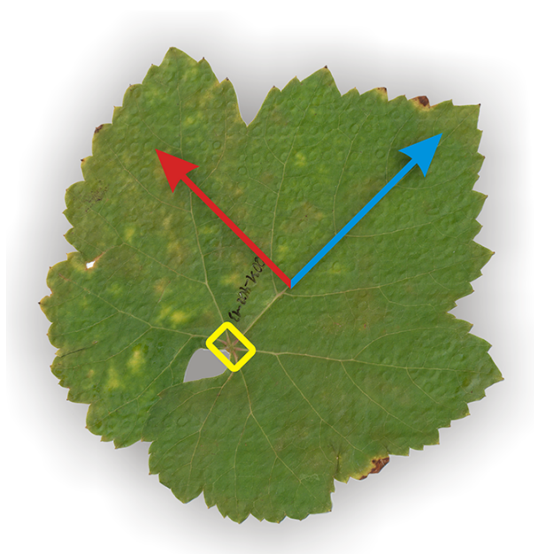

Fig. 4. Illustration of the template-based strategy to detect the central point. The orientation is found by a principal component analysis. The directions correspond to the maximal and minimal variance respectively. Because the orientation is determined in such a way, the simulated annealing algorithm has to minimize a scalar fitness function of only two variables, which depends on the shift of the template image.

The image template is rotated so that its direction fits the orientation of the leaf which in turn is found by means of a principal component analysis (cf. [6]) on the binarized image. The image template is then shifted along the main directions until the color distance to the leaf pixels reaches a minimum. This process is controlled by a simulated annealing strategy, cf. [3]. In order to reduce the influence of minor variations, the average of the squared color distances is used as fitness function. The barycenter is a good initial guess to start from. The image template is created from the average image of the intensity normalized center regions of twelve leaves from our test set.

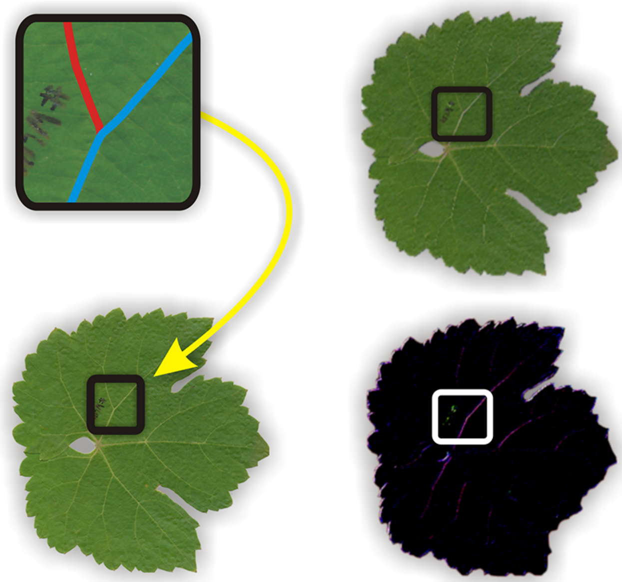

Fig. 5. Illustration of a single filter step of our vein tracing method. It is not possible to decide which one of the two colored parts belongs to the level zero vein (left) by just comparing pixel intensities. After applying G_{real_θ} to the red channel F of the image with angle θ computed from prior curvature values the obtained filter response $F \otimes G_{\mathrm{real}_\theta}$ sheds a light on the actual course of the vein because it masks all other directions (right).

Starting at the central point found by the procedure described above, the algorithm traces the course of the veins by comparing the intensity values on the red channel. Here the contrast between the veins and the other parts of the leaf is significantly improved, cf. Fig. (5). The reason for this is that the leaf pigments consist of the green chlorophyll and the different reddishbrown carotenes.

During the tracing procedure the algorithm keeps track of the curvature of the vein. The curvature is defined as slope of the secant of the last five percent of the main vein with respect to the estimated total length of the detected vein. We use a simple heuristic to detect when the algorithm is about to leave the exact course of the main vein: if the ratio of the difference of the current and the last curvature and that of the last three curvature differences is greater than ten percent the algorithm jumps back to the location where the curvature has changed and applies the real Gabor filter Eqn. (2) to the red channel F of the image. Since we feed the Gabor filter with the angle θ–computed from the curvature values–of the direction of the current subsection of the vein it is guaranteed that all subsections of veins with this particular direction will be visible in the filter response while other directions are masked. Therefore, the course of the vein is traced on the filter response $F \otimes G_{\mathrm{real}_\theta}$ and not on the original image. It should be clear that when the curvature and therefore the angle of the current subsection of the vein under consideration changes "too much" the filter response must be recomputed. Our tracing algorithm stops when the border of the leaf is reached. This is easily detected because of the high contrast between the background and the leaf image.

The main veins typically become thinner as their distance from the central point increases. We account for this fact by adapting the wavelength and the variance parameters λ and σ of the Gabor filter through interpolation from the intensity values of the original image in the area where–based on the last curvature values–the next part of the vein is expected. In contrast to the wavelength and the variance, the influence of the spatial aspect ratio γ and the phase shift ψ can be regarded as global. Thus, we work with a constant aspect ratio and ignore the phase shift.

4 Cultivar Classification

Our Gabor filter-based algorithm has been implemented in a vine leaf classification system written in Java. The different steps carried out for the final classification of a given vine leaf are illustrated in Fig. (6). The curvature values of the main veins are stored in a vector and fed into a support vector machine. In our implementation we embark on the WEKA SVM (cf. [7]) and use a simple linear kernel function. The training set contains 20 images of the test set: five leaves of each of the four white vine cultivars.

5 Evaluation and Results

We apply the devised Gabor filter-based algorithm to the test set consisting of 150 images of vine leaves of different cultivars with a broad spectrum of color characteristics, cf. Sec. (2.1). The set therefore contains 750 different zero order veins. The vein extraction algorithm failed in only 13 of the 750 cases ($< 2\%$). This was due to the fact that there has been either a hole in the outer leaf lobes or a shadow cast prevented the recognition of the correct course of the vein.

The cultivars of the untrained 130 images in the test set have been classified with the following success rates: Kerner 93% (28 of 30), Müller-Thurgau respectively Rivaner 97% (32 of 33), Riesling 94% (33 of 35), and Scheurebe 100% (32 of 32). Hence, all classifications were performed with usable success rates significantly above the 93% threshold. The classification of a single leaf took about 5 sec on a 3.20 GHz Intel Core i7-3930K, 32 GB RAM, under Microsoft Windows 7. It comprises all the processing steps described in Fig. (6). The input images have been used in a downscaled resolution of 400×400 pixels, because there was no significant difference in the success rate compared to the larger scaled images.

An illustration of our leaf classification system running on a mobile device emulator is shown in Fig. (1).

5.1 Conclusion and Future Work

We have demonstrated that working on appropriate Gabor filter responses allows us to accurately extract the main veins of vine leaves. Our approach finds the correct paths in almost all test cases. At the same time it is able to handle challenging problems like discolorations and even in the case of equally oriented handwritten labels it stays on the correct path because

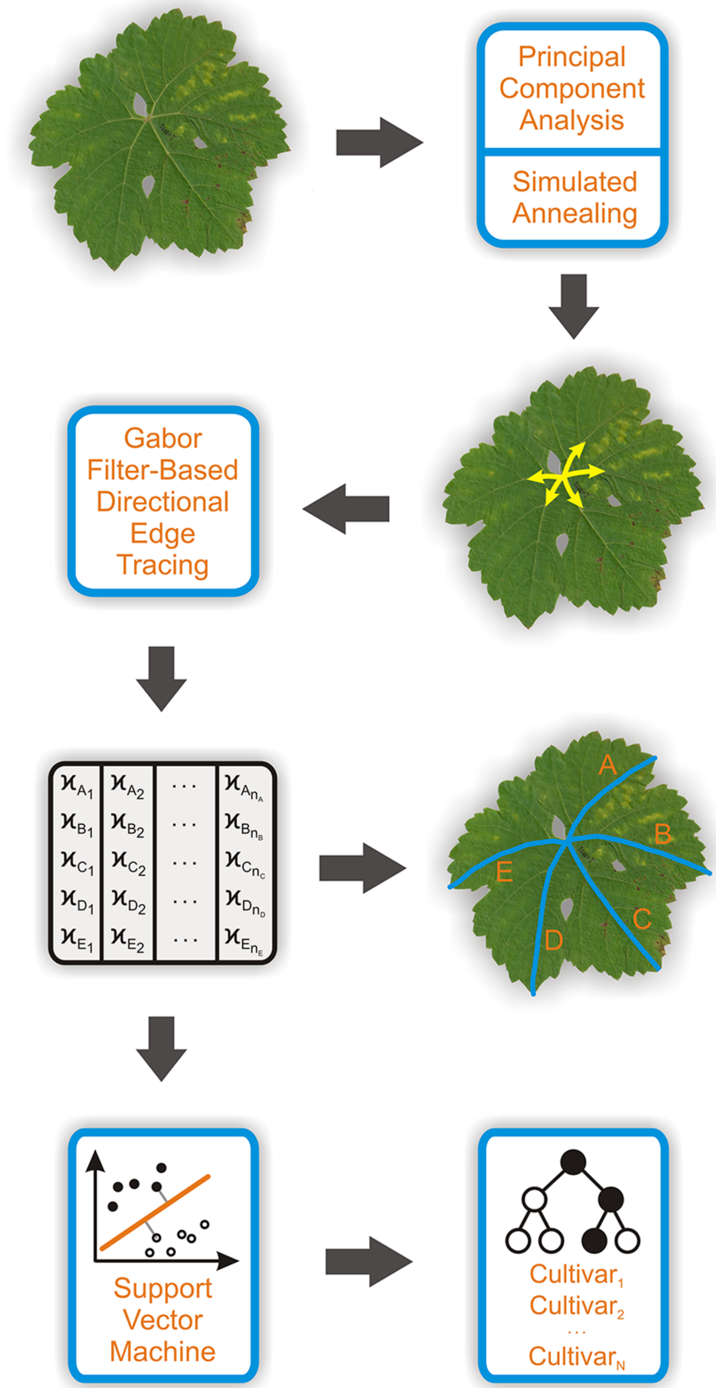

Fig. 6. Illustration of our full-fledged vine leaf classification system: the final classification is based on the curvature vectors that act as a kind of fingerprint using a support vector machine.

the Gabor filter masks all differently oriented parts of the label. However, problems can still occur in the case of holes caused by pest infestation or shadow casts. Our future work focuses on a robust method that detects such holes and applies a gap filling algorithm, e.g. as part of the precomputing step of our leaf classification system. Moreover, we aim at generalizing the presented approach to other cultivars with different phenotypical characteristics in the leaves to overcome the current restriction to five veined vine leaves.

References

1. J. S. Cope, P. Remagnino, S. Barman, and P. Wilkin. The extraction of venation from leaf images by evolved vein classifiers and ant colony algorithms. In *ACIVS (1)*, pages 135–144, 2010.
2. I. P. G. R. Institute. *Descriptors for Grapevine: (Vitis Spp.)*. Descriptors IBPGRI. International Plant Genetic Resources Institute, 1997.
3. S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
4. C.-L. Lee and S.-Y. Chen. Classification of leaf images. *International Journal of Imaging Systems and Technology*, 16(1):15–23, 2006.
5. Y. Li, Z. Chi, and D. D. Feng. Leaf vein extraction using independent component analysis. In *Systems, Man and Cybernetics, 2006. SMC '06. IEEE International Conference on*, volume 5, pages 3890–3894, 2006.
6. K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2:559–572, 1901.
7. WEKA 3. Data Mining with Open Source Machine Learning Software in Java.
8. S. Wu, F. Bao, E. Xu, Y.-X. Wang, Y.-F. Chang, and Q.-L. Xiang. A leaf recognition algorithm for plant classification using probabilistic neural network. In *Signal Processing and Information Technology, 2007 IEEE International Symposium on*, pages 11 –16, dec. 2007.
9. X. Zheng and X. Wang. Leaf vein extraction based on gray-scale morphology. *International Journal of Image, Graphics and Signal Processing*, 2:25–31, 2010.