The Extraction Of Venation From Leaf Images By Evolved Vein Classifiers And Ant Colony Algorithms

James S Cope, Paolo Remagnino, Sarah Barman, and Paul Wilkin

Digital Imaging Research Centre, Kingston University, London, UK
{j.cope,p.remagnino,s.barman}@kingston.ac.uk
Royal Botanic Gardens, Kew, London, UK
p.wilkin@kew.org,

Abstract. Leaf venation is an important source of data for research in comparative plant biology. This paper presents a method for evolving classifiers capable of extracting the venation from leaf images. Quantitative and qualitative analysis of the classifier produced is carried out. The results show that the method is capable of the extraction of near complete primary and secondary venations with relatively little noise. For comparison, a method using ant colony algorithms is also discussed.

1 Introduction

In the field of comparative biology, novel sources of data are continuously being sought to enable or enhance research varying from studies of evolution to generating tools for taxon identification. Leaves are especially important in this regard, because in many applied fields, such as studies of ecology or palaeontology, reproductive organs, which may often provide an easier form of identification, are unavailable or present for only a limited season. Leaves are present during all seasons when plants are in growth. There are also millions of dried specimens available in herbaria around the world, many of which have already been imaged. While these specimens may possess reproductive organs, the main character features are often concealed in images through being internal or poor preparation. However, almost all specimens possess well-preserved and relatively easily imaged leaf material.

Traditional methods employed by botanists for describing leaves rely on terminology and are wholly qualitative and open to some level of interpretation [3]. In recent decades plant science has begun to use a range of quantitative morphometric methods in comparative studies [12, 5]. However, such data currently exist for a small minority of plant taxa, largely due to the limitations imposed by manual data capture. Research such as that cited above has shown that the most useful features of leaves for use in comparative biology are usually the two-dimensional outline shape, characters of the margin and structure of the vein network (Figure 1). Thus a fully automated method of extracting consistent,

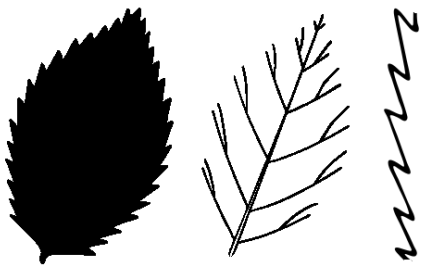

Fig. 1: Leaf Shape, Venation & Margin

mathematically sound information from images of leaves would be a great aid in plant comparative biology.

This paper presents a first step towards extracting the primary and secondary venation from leaf images, using a classifier that has been evolved to recognise those pixels which belong to veins. For comparison we also explore the use of ant colony algorithms for vein extraction.

2 Previous Work

In recent years, a number of techniques have been employed for extracting leaf venation. Clarke [1] compares the results from two simple methods, smoothing and edge detection, and a scale space algorithm, with the best results that they could achieve manually using Photoshop. Fu & Chi [4] used a two stage approach on leaves which had been photographed using a fluorescent light bank to enhance the venation. First, edge detection methods were used to determine a suitable greyscale threshold for removing most of the non-vein pixels. An artificial neural network classifier was then used to refine the results. Li & Chi [7] successfully extracted the venation from leaf sub-images using Independent Component Analysis (ICA) [2], though when used on whole leaves, the results were only comparable to the Prewitt edge detection operator. Artificial ant swarms were also used by Mullen [9] to trace venation and outlines in leaves via an edge detection method. Kirchgeßner [6] describes a method of tracking vein structures on leaves, and representing them using b-splines which contain the hierarchical venation information. This method, however, required some manual interaction to initialise a search.

The method presented in this paper produces robust vein extractions from whole leaf images without backlighting the leaves. Section 3.1 outlines how the pixels are classified. The features used are specified in section 3.2, whilst section 3.3 describes how the classifiers are evolved. Results for this method are given in section 3.5. Section 4.1 describes how ant colony algorithms can be used for vein extraction, with section 4.2 containing the results for this and a comparison of the two methods, with discussion of how they might be combined to further improve the results.

3 Extraction By Evolved Vein Classifiers

3.1 Classifying The Vein Pixels

A genetic algorithm is used to evolve a set of classifiers for detecting veins. Each classifier consists of a pair of bounds for each of the features used. If the values of all the features for a pixel fall within all the bounds for a classifier, then it is classified as vein. The vein pixels found by all the classifiers in the set are combined, and all other pixels are classified as non-vein. These classifiers are similar those used by Liu & Tang [8]. More specifically, the set of vein pixels, V, is determined as follows:

$$V = \{(x,y)|0 \leq x < w, 0 \leq y < h, \\ \exists c \in C s.t. \forall f_i \in F_{xy}(c_{i0} \leq f_i \leq c_{i1})\}$$

where

- w,h are the image width and height respectively
- C is the set of all classifiers
- c_{i0} is the lower bound for the i^{th} feature for the classifier c
- c_{i1} is the upper bound for the i^{th} feature for the classifier c
- F_{xy} is the set of feature values for the pixel at (x,y)
- f_i is the value for the i^{th} feature

3.2 Feature Extraction

A set of 9 features are extracted for each pixel for use in classification. The features used are as follows:

1. Pixel greyscale value $f_1 = I(x,y)$
2. Edge gradient magnitude (from Sobel)
3. Average of greyscale values in a 7×7 neighbourhood

$$f_3 = \frac{1}{49} \sum_{\substack{x-3 \leq i \leq x+3 \\ y-3 \leq j \leq y+3}} I(i,j)$$

4. Greyscale value minus neighbourhood average

$$f_4 = I(x,y) - \frac{1}{49} \sum_{\substack{x-3 \leq i \leq x+3 \\ y-3 \leq j \leq y+3}} I(i,j)$$

5. Greyscale value minus leaf lamina average

$$f_5 = I(x,y) - \frac{1}{|lamina|} \sum_{\substack{0 \leq i < width \\ 0 \leq j < height \\ (i,j) \in lamina}} I(i,j)$$

Where *lamina* is the set of all pixels which are part of the leaf's lamina (surface), found by using Otsu's thresholding [10] to remove the leaf from the background.

The average local gradient direction of pixels in a 11×11 neighbourhood around the current pixel is calculated. This size neighbourhood was chosen because for most vein pixels this will include both sides of the vein. The greyscale values of the points 5 pixels from the current one in both directions along the gradient and perpendicular to the gradient are calculated. If the current pixel is part of a vein, the pixels perpendicular to the gradient direction are likely to also be vein pixels, and so similar to the current pixel, whilst the pixels along the gradient direction are likely to be non-vein, and therefore quite different.

$$i_1 = I(x + 5 * Sin(\alpha), y + 5 * Cos(\alpha))$$

$$i_2 = I(x - 5 * Sin(\alpha), y - 5 * Cos(\alpha))$$

$$j_1 = I(x + 5 * Sin(\alpha + \frac{\pi}{2}), y + 5 * Cos(\alpha + \frac{\pi}{2}))$$

$$j_2 = I(x - 5 * Sin(\alpha + \frac{\pi}{2}), y - 5 * Cos(\alpha + \frac{\pi}{2}))$$

where α is the gradient direction.
The remaining features are then:

6. The difference between pixels either side of potential vein $f_6 = |i_1 - i_2|$
7. The difference between pixels along potential vein

$$f_7 = |j_1 - j_2|$$

8. Greyscale value minus average value of the two pixels either side of the potential vein

$$f_8 = I(x, y) - \frac{i_1 + i_2}{2}$$

9. Greyscale value minus average value of the two pixels along the potential vein

$$f_9 = I(x, y) - \frac{j_1 + j_2}{2}$$

To allow the same genetic operators to be used on features with very varied distributions, the feature values for the training data are mapped to a uniform distribution. This mapping is recorded and applied to any data being subsequently classified.

3.3 Evolving The Classifiers

Classifiers are evolved one after another using a genetic algorithm, and added to the classifier set until no more classifiers with a fitness above a certain threshold can be generated. The only genetic operators used are mutations, as crossover operations are likely to combine classifiers that work on different types of vein

pixels, thereby having a negative effect. For example, a classifier that finds thin sections of vein may require higher edge gradient values and lower greyscale values than a classifier finding the pixels in the middle of thicker veins. Crossing over these two classifiers would result in ones which classified neither of these vein pixel types. Bounds are mutated with probability 0.3 by adding or subtracting an amount randomly drawn from the range [0,0.01]. The population is re-initialised after each classifier is added to the set. Each individual is initialised by centring the bounds around the feature values for a vein pixel randomly selected from the training data, with the width of the bounds drawn from a Gaussian distribution. This increases the likelihood of the classifier being effective, as one vein pixel will always be correctly classified by it, along with any similar vein pixels.
The fitness function used is as follows:

$$fitness_i = \frac{|T_i \setminus \bigcup_{j \in C} T_j|}{|F_i \setminus \bigcup_{j \in C} F_j| + k}$$

where:

- T_i is the set vein pixels correctly classified by classifier i (true positives).
- F_i is the set non-vein pixels incorrectly classified by classifier i (false positives).
- C is the current set of classifiers selected in previous iterations and k is a constant.

This function grants high fitnesses to individuals which, if added to the classifier set, would significantly increase the number of true positives, but not the number of false positives. The constant k is used to adjust the balance between a high true positive/false positive ratio, and a high total number of true positives. If k is set too low the ratio will be very high, but the final classifier set may over-fit the training data. If k is set too high it will result in a high number of false positives. A value of $k = 5$ was found to be appropriate.

3.4 Redundancy

Classifiers can potentially be made redundant by other classifiers added to set later. In other words, a classifier may no longer uniquely classify many vein pixels whilst still incorrectly classifying some non-vein pixels. It is beneficial to remove such classifiers as this may greatly reduce the number of false positives whilst only slightly reducing the number of true positives.

Redundant classifiers are identified by removing candidates from the set and measuring any improvement in overall classification quality. The classifier whose removal produces the largest increase in quality is permanently removed from the set. This process is repeated until no more classifiers are found to be redundant.

3.5 Results

The classifier was trained using 8000 pixels randomly selected from 14 leaf images, 2 from each of 7 species. These pixels were then manually labelled as either

vein or non-vein. The resulting classifier was then tested on 7 new leaf images, one from each of the species used for training. The ROC curve in figure 2 shows the results (solid line). With a false positive rate of 0.0166, a true positive rate of 0.853 was achieved. The curve was produced by varying the value of k, and the size of the initial classifier bounds.

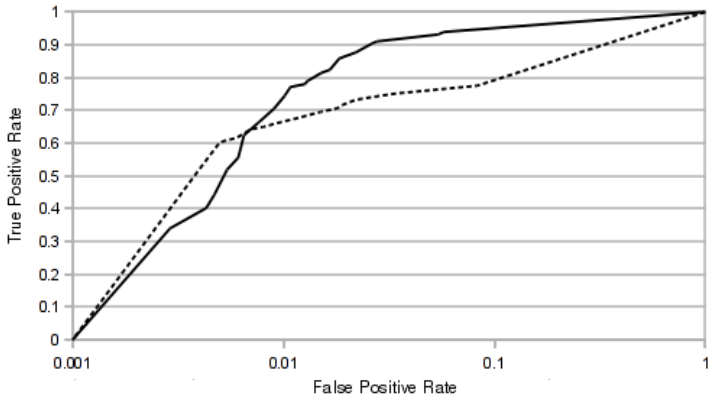

Fig. 2: ROC Curve. Solid line - evolved classifiers. Dashed line - ant algorithm

The classifier was also used on the full leaf images from the test set, in order to extract the full venation pattern. Examples of these results are shown in figure 3.

4 Extraction By Ant Colonies

4.1 The Algorithm

Our second approach to vein extraction is to use an ant colony algorithm. A population of ant-like agents are placed at random across the image. These "ants" then move across the image, moving from pixel to pixel based upon some heuristic evaluation of that pixel, known as the pixel's visibility, and also based on the level of "pheromone" at that pixel. The pheromones are an indicator deposited by ants to signal to other ants the value of a particular pixel. As time progresses, the pheromone levels build up to create a pheromone map for the image, with high levels in desirable regions, and low levels in undesirable regions. In our case we use the edge magnitude as the measure of visibility, to encourage the ants to traverse along the veins, thereby extracting continuous sections of venation. The probability, P_{ij}, of an ant at pixel i moving to pixel j is calculated

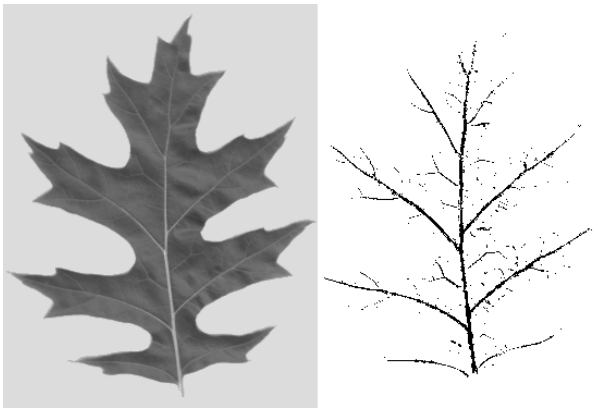

(a) Quercus Shumardii

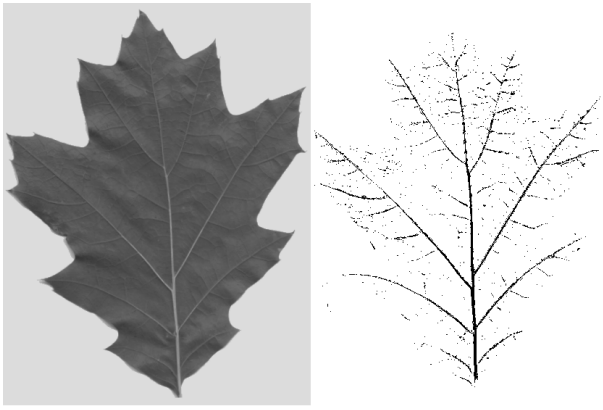

(b) Quercus Rubra

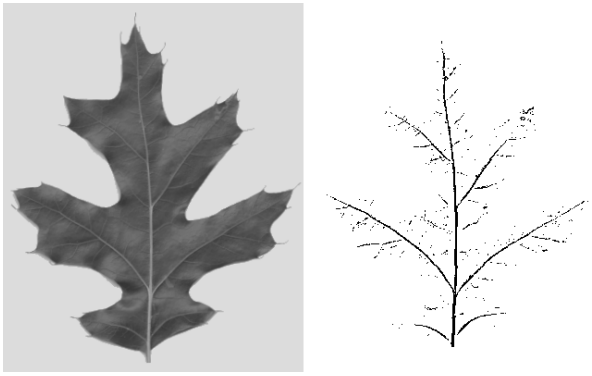

(c) Quercus Ellipsoidalis

Fig. 3: Results for extraction by evolved classifers

as follows:

$$P_{ij} = \begin{cases} \dfrac{\tau_j^\alpha \eta_j^\beta}{\sum_{k \in K_i} \tau_k^\alpha \eta_k^\beta} & \text{if } j \in K_i \\ 0 & \text{otherwise} \end{cases}$$

where τ_j and η_j are the pheromone level and visibility respectively at pixel j, α and β are the weightings for these two components, and K_i is the set of pixels neighbouring pixel i. To prevent the ants converging on the strong edges outlining the leaf instead of the venation, the visibility for all background pixels (again calculated using Otsu's method) and all pixels within a short distance of the background (in this case, a distance of 10 pixels) is set to 0. After all the ants have performed one move, the pheromone levels are updated:

$$\tau_i = (1 - \rho)\tau_i + \delta a_i \eta_i$$

where ρ is the rate at which pheromones evaporate, δ is the update rate, and a_i is the number of ants at pixel i. There is a risk that ants will simply move between the same small set of pixels, building up pheromone levels until it is highly unlikely for them to escape. This is prevented by keeping a list of the last 10 pixels visited by each ant, and forbidding the ant from re-visiting any of these pixels. After a set number of moves have taken place, the pheromone map is thresholded to produce a binary vein classification. This method is based on the method described in [9].

4.2 Results And Comparison Of Methods

Figure 4 contains examples of typical results obtained using this method. For each leaf the algorithm was run for 500 steps, using 2000 ants. The pheromone map was then thresholded at 2% of the maximum pheromone level. These values were chosen as they appeared to give the best qualitative results. The results differ from those obtained using the evolved classifiers in a number of ways. Firstly, due to the use of only the edge gradients to guide the ants across the image, the results contained only the hollow outline of the venation, whereas the other method extracts the full vein. One advantage of using ants is that it helps in extracting continous venation, whilst the evolved classifiers extract veins with many small gaps in them. On the downside, when a vein contains a section with only a low edge magnitude, the ants are unable to continue to extract the rest of that vein as the pixel-by-pixel evolved classifiers are able to do. The effects of this can be seen near the top of the first image in figure 4, where a large section of venation is completely absent. Furthermore, whereas much of the false positive results from the first method are isolated pixels that can be easily removed, the ants produce larger, connected areas of noise, that may be harder to distinguish from the actual venation.

By applying morphological closing, the hollow vein centres can be filled in (figure 5). From these, quantitative results can be calculated, as shown in figure 2

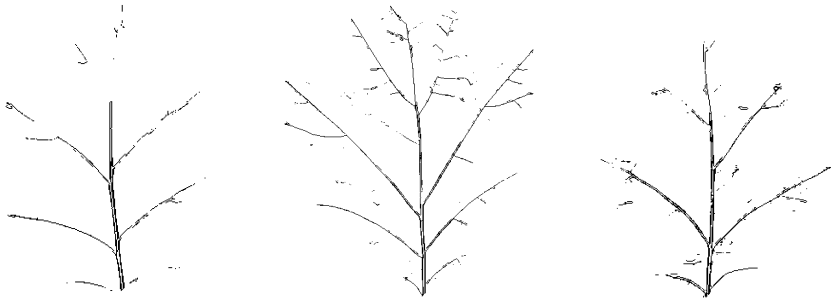

Fig. 4: Results using the ant colony algorithm

(dashed line). It can be see that the ant algorithm still performs worse than the evolved classifiers, except when the true positive rate falls below approximately 0.63.

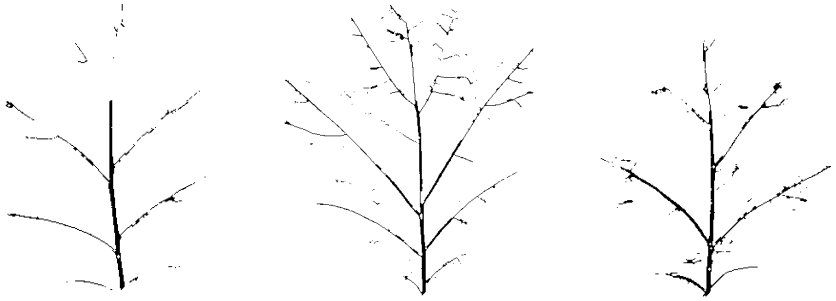

Fig. 5: Results after morphological closing

5 Conclusion

In this paper, a method of evolving classifiers capable of extracting the venation pattern from leaf images was presented. Qualitative results show that near complete primary and secondary venation patterns can extracted with relatively little noise. Extraction of tertiary venations is less reliable. This was compared to a vein extraction method using an ant colony algorithm. Whilst the ants were able to extract more continuous venation, the gaps which were present tended to be larger, as did the regions of noise. With this in mind, it may be possible to

combine the two methods to achieve even better results. The evolved classifier could be adapted to provide a probabilistic, rather than binary, classification, with a different ant colony using each classifier as its visibility measure to extract complete, continuous venation. This is one possible area for future work. Other future work will first involve finding a method of refining the extracted venation so that only the complete primary and secondary venation remains. A system for automatically generating a venation description, preserving the hierarchical vein structure, will then be developed. This will allow the venation from different leaves to be accurately compared.

References

1. Clarke, J., Barman, S., Remagnino, P., Bailey, K., Kirkup, D., Mayo, S., Wilkin, P.: Venation pattern analysis of leaf images. Lecture Notes In Computer Science 4292, 427–436 (2006)
2. Comon, P.: Independent component analysis, a new concept? Signal Processing 36(3), 287–314 (1994)
3. Ellis, B., Daly, D.C., Hickey, L.J., Johnson, K.R., Mitchell, J.D., Wilf, P., Wing, S.L.: Manual Of Leaf Architecture. Cornell University Press (2009)
4. Fu, H., Chi, Z.: Combined thresholding and neural network approach for vein pattern extraction from leaf images. In: IEE Proceedings. Vision Image And Signal Processing. vol. 153, pp. 881–892. Institution of Electrical Engineers (2006)
5. Jensen, R.J., Ciofani, K.M., Miramontes, L.C.: Lines, outlines, and landmarks: Morphometric analyses of leaves of acer rubrum, acer saccharinum (aceraceae) and their hybrid. Taxon 51(3), 475–492 (2002)
6. Kirchgessner, N., Scharr, H., Schurr, U.: Robust vein extraction on plant leaf images. In: 2nd IASTED International Conference Visualisation, Imaging And Image Processing (2002)
7. Li, Y., Chi, Z., Feng, D.D.: Leaf vein extraction using independent component analysis. In: IEEE International Conference On Systems, Man, And Cybernetics. pp. 3890–3984. IEEE (2006)
8. Liu, J., Tang, Y.Y.: Adaptive image segmentation with distributed behavior-based agents. IEEE Trans. Pattern Anal. Mach. Intell. 21(6), 544–551 (1999)
9. Mullen, R., Monekosso, D., Barman, S., Remagnino, P., Wilkin, P.: Artificial ants to extract leaf outlines and primary venation patterns. Lecture Notes In Computer Science 5217, 251–258 (2008)
10. Otsu, N.: A threshold selection method from gray level histograms. IEEE Trans. Systems, Man and Cybernetics 9, 62–66 (1979)
11. Park, J., Hwang, E., Nam, Y.: Utilizing venation features for efficient leaf image retrieval. Journal of Systems and Software 81(1), 71–82 (2008)
12. Plotze, R.d.O., Falvo, M., Padua, J.G., Bernacci, L.C., Vieira, M.L.C., Oliveira, G.C.X., Martinez, O.: Leaf shape analysis using the multiscale minkowski fractal dimension, a new morphometric method: A study with passiflora (passifloraceae). Canadian Journal Of Botany 83(3), 287–301 (2005)