Leaf Classification from Boundary Analysis

Anne Jorstad
AMSC 663 Project Proposal
Fall 2007

Advisor: David Jacobs

1 Abstract

In this project, I will analyze the boundary curves of leaf images to recognize and classify leaves by species. The input leaf boundary curve will be decomposed into a wavelet bases, which will then be compared to the decomposition of all samples in the database, and the closest matches will be returned. The inner-distance measure will also be computed and compared at multiple scales. This project is constructed to complement an existing leaf classification system which uses general leaf geometry, and not pure boundary information, to recognize leaf species.

2 Background

There is an ongoing project between members of the University of Maryland, Columbia University, and National Museum of Natural History Smithsonian Institution to create an electronic field guide for plants [1, 5]. The ultimate goal of this project is to develop a system where a user in the field can take a picture of an unknown plant, feed it to the system carried on a portable computer, and have the system classify the species and display sample images of the closest matches within a couple of seconds. A basic system has been implemented for the leaves of woody plants of the Baltimore-Washington, DC area, and manages to provide reasonable guesses of species classification for most examples. Currently, the classification is done by comparing the overall shape of each leaf to each specimen in the system. I propose to extend the classification algorithm to examine the variation in leaf boundaries. Leaf edges can be smooth, serrated or lobed, and although these characteristics do affect the overall shape description of the leaves, the specific amount of variation in leaf edges has not been studied for this system. I will develop a software module to classify leaves based exclusively on information extracted from the leaf edges. We hypothesize that this will provide information independent from the current classification scheme, so that when combined with the current system, the classifier will be able to perform with higher accuracy.

3 Method

The current system examines the global shape of each leaf, and no specific consideration is given to the variation in boundary patterns. It is possible for two leaves with similar global shape to have very distinct features at a small scale over local edge features. There are a variety of ways to look at edge variation.

3.1 Harmonic Analysis

The first class of methods that will be considered will involve harmonic analysis. The leaf boundary is a curve which can be decomposed into a wavelet basis, and the boundary constructions can then be compared in the frequency domain. The boundaries of the leaves are already provided with good accuracy, so boundary detection will not be part of my project. Each boundary is represented as a set of discrete 2D points, which can be handled as a 1D complex function, $x+iy$, allowing for the use of 1D techniques. There are many algorithms in Computer Vision dealing with texture analysis, some of which involving the representation of a texture as a combination of fundamental textures [4]. This is very similar to a wavelet basis for functions. These techniques can be exploited, along with traditional harmonic analysis, to create a representation for each leaf boundary from basis functions, which will provide information that can be assessed quantitatively for leaf classification.

3.2 Inner-Distance

Distinct from harmonic analysis techniques, there are other measures of local boundary variation that should be considered. Recent work by Professor Jacobs has developed an idea of “Inner-Distance” [3]. This is a measure of the shortest distance between two points on an object that passes entirely through the object itself (similar to a geodesic on a general manifold, but instead of a manifold we consider closed but not necessarily convex shapes). For a smooth-edged leaf, the inner-distance between two globally nearby points is small, but if the edge is serrated, the local inner-distance will be relatively much greater. The inner-distance is already successfully used in the current classification system, but it is only considered over one set of evenly spaced points around the leaf, comparing all the relative distances to the corresponding points of other leaves. To extract a better understanding of local variations, sets of points at several smaller scales should be considered. Such a hierarchy of resolutions on the boundary has been shown to successfully classify shapes in previously explored settings [2]. This measure will be more or less equivalent to measuring the amount of change in edge gradients that exist over short distances.

3.3 Convexity

Another measure to be considered is convexity. Serrated leaves are much less convex than those with smooth edges, so various convexity measures can be explored to help distinguish these cases. There are several variations on convexity measures which might be useful [6], for example:

$$Convexity = \frac{Area(object)}{Area(ConvexHull(object))}$$

or

$$Convexity = \frac{Perimeter(object)}{Perimeter(ConvexHull(object))}\ .$$

However, it is possible that measuring the convexity of individual leaves will not provide enough new information to classify a leaf any further that the system already manages without this measure. Convexity measures cannot provide local information on a very small scale, as convexity is a global concept, but measuring convexity can certainly help classify leaves into larger families of species, which might help in the first stages of classification.

4 Verification

To verify the correctness of my work, I will generate artificial “leaves” with known properties. For example, an ellipse has no high frequency oscillations, and all its inner-distance measures will be straight lines, but an appropriately constructed jagged oval can have the desired detectable features. Once the algorithm is proved on these simple cases, I will then test my code on the leaf datasets. Here I should be able to demonstrate *reasonable* classification accuracy on relevant examples, but because I am not considering any global shape descriptions, it is expected that not all distinguishing features should be recognized. The goal is that the classification offered by the new schemes provide information independent from the current system. If this is the case, the decisions of the new code will in many cases be distinct from those of the existing system, but when the two implementations are combined, more accurate overall classification results will be achieved. It is possible that the new recognition code will not add significant improvement to the system as a whole, but it is believed that the new information will improve the results enough to validate the extra time required to run the new code.

5 Specifications

The existing Electronic Field Guide system is written in MATLAB and C, and my project will be continuing along these lines. Most of my work will be written in MATLAB, making use of the Image Processing and Wavelet Toolboxes. It is not expected that the proposed algorithms be so slow that they would require recoding in C for noticeably improved performance, but it is possible that some of the code will eventually be converted to C. The end product is to run on a portable computer, so all computations must be quick and efficient. I will therefore be writing and testing my work on a PC.

6 Schedule

I will start my work by engaging in significant background reading, and familiarizing myself with the current system as currently implemented. By early November, I intend to have an outline of the functions I will need to implement for the wavelet-based algorithm. By the end of the first semester, I will have basic code for this algorithm, and this code will be tested on the artificial examples I will construct. The first part of second semester, I will implement the algorithm considering inner-distances on several scales, and hope to be testing this by spring break. I will then work on making both algorithms more robust, and will combine the results of the two methods, and possibly also that of some convexity measures if they demonstrate promise of being helpful, in an effective manner. The final step is to run my code on the existing datasets. I will have started work on my final report by mid-April.

References

[1] Gaurav Agarwal, Haibin Ling, David Jacobs, Sameer Shirdhonkar, W. John Kress, Rusty Russell, Peter Belhumeur, Nandan Dixit, Steve Feiner, Dhruv Mahajan, Kalyan Sunkavalli, Ravi Ramamoorthi, Sean White. "First Steps Toward an Electronic Field Guide for Plants". Taxon, vol. 55, no. 3, Aug. 2006.

[2] Pedro F. Felzenszwalb, Jushua D. Schwartz. "Hierarchical Matching of Deformable Shapes". IEEE Conference on Computer Vision and Pattern Recognition, 2007.

[3] Haibin Ling, David W. Jacobs. "Using the Inner-Distance for Classification of Articulated Shapes". IEEE Conference on Computer Vision and Pattern Recognition, vol. II, June 2005.

[4] Jitendra Malik, Serge Belongie, thomas Leung, Jainbo Shi. "Contour and Texture Analysis for Image Segmentation". International Journal of Computer Vision, vol. 34, no. 1, July 2001.

[5] Sean White, Dominic Marino, Steven Feiner. "Designing a Mobile User Interface for Automated Species Identification". Proceedings of the SIGCHI, April 2007.

[6] Jovisa Zunic, Paul L. Rosin. "A New Convexity Measure for Polygons". IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 26, no. 7, July 2004.