

Searching the World's Herbaria: A System for Visual Identification of Plant Species

Peter N. Belhumeur¹, Daozheng Chen², Steven Feiner¹, David W. Jacobs², W. John Kress³, Haibin Ling⁴, Ida Lopez³, Ravi Ramamoorthi¹, Sameer Sheorey², Sean White¹, and Ling Zhang³

¹ Dept. Of Computer Science, Columbia University, New York NY

² Dept. Of Computer Science, University of Maryland, College Park MD

³ Department Of Botany, National Museum of Natural History, Smithsonian Institution, Washington, DC

⁴ Information Science and Technology Center, Temple University, Philadelphia PA

Abstract. We describe a working computer vision system that aids in the identification of plant species. A user photographs an isolated leaf on a blank background, and the system extracts the leaf shape and matches it to the shape of leaves of known species. In a few seconds, the system displays the top matching species, along with textual descriptions and additional images. This system is currently in use by botanists at the Smithsonian Institution National Museum of Natural History. The primary contributions of this paper are: a description of a working computer vision system and its user interface for an important new application area; the introduction of three new datasets containing thousands of single leaf images, each labeled by species and verified by botanists at the US National Herbarium; recognition results for two of the three leaf datasets; and descriptions throughout of practical lessons learned in constructing this system.

1 Introduction

We have built a hand-held botanical identification system for use by botanists at the Smithsonian Institution. Employing customized computer vision algorithms, our system significantly speeds up the process of plant species identification. The system requires only that the user photograph a leaf specimen, returning within seconds images of the top matching species, along with supporting data such as textual descriptions and high resolution type specimen images. By using our system, a botanist in the field can quickly search entire collections of plant species—a process that previously took hours can now be done in seconds.

To date, we have created three datasets for the system: one that provides complete coverage of the flora of Plummerville Island (an island in the Potomac River owned by the National Park Service); a second that covers all woody plants in published flora of the Baltimore-Washington, DC area; and a nearly



Fig. 1. Left: A computer vision system for identifying temperate plants on the botanically well-studied Plummers Island, Maryland, USA. Right: Congressman John Tanner tries an augmented reality version of the system.

complete third dataset that covers all the trees of Central Park in NYC. The system is currently being used by botanists at the Smithsonian to help catalogue and monitor plant species. Figure 1 shows the system and various versions of the user interface (UI). Although a great deal of work remains to be done in this ongoing collaboration between computer vision researchers and scientists at the US National Herbarium, we hope that our system will serve as a model and possible stepping stone for future mobile systems that use computer vision-based recognition modules as one of their key components.

1.1 Motivation

Botanists in the field are racing to capture the complexity of the Earth's flora before climate change and development erase their living record. To greatly speed up the process of plant species identification, collection, and monitoring, botanists need to have the world's herbaria at their fingertips. Tools are needed to make the botanical information from the world's herbaria accessible to anyone with a laptop or cell phone, whether in a remote jungle or in NYC's Central Park.

Only recently has the data required to produce these tools been made available. Volumes of biological information are just now going on-line: natural history museums have recently provided on-line access to hundreds of thousands of images of specimens, including our own work in helping to digitize the complete Type Specimen Collection of the US National Herbarium. These massive digitization efforts could make species data accessible to all sorts of people including non-specialists, anywhere in the world.

Yet there is a critical shortfall in all these types of natural databases: finding a species quickly requires that the searcher know in advance the name of the species. Computer vision algorithms can remove this obstacle, allowing a user to search through this data using algorithms that match images of newly collected

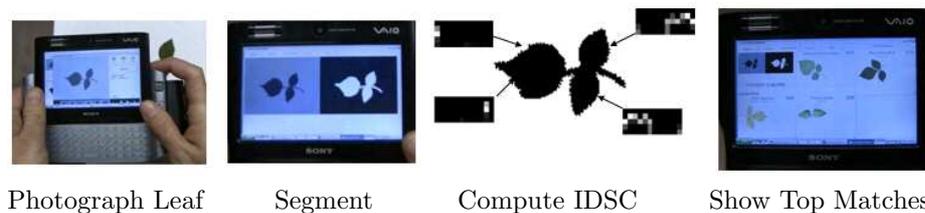


Fig. 2. A flow diagram of our plant identification system. A leaf from an unknown species of plant is photographed by the user. The system then segments the leaf image from its background, computes the IDSC shape representation used for matching, and then displays the top matches, as they are computed.

specimens with images of those previously discovered and described. Without such tools, a dichotomous key must be painfully navigated to search the many branches and seemingly endless nodes of the taxonomic tree. The process of identifying a single species using keys may take hours or days, even for specialists, and is exceedingly difficult to impossible for non-scientists.

1.2 System Design and Contributions

Using our system, a botanist in the field can choose a leaf and photograph it against a plain background. The leaf image is then compared to all possible matches, and in a matter of seconds the botanist is shown information about the best matching species. Figure 2 illustrates the process, with photographs of our system in action. Figure 4 shows the performance of our system. On the woody plants of the Baltimore-Washington, DC area (245 species) the system returns the correct species in the top ten matches more than 97% of the time.

This paper makes several contributions. First and foremost, we describe a complete working system for an important application that has received little attention from the computer vision community. We hope the reader will take from this paper an appreciation for the possible impact that computer vision can have on the study of biodiversity. Also, while many individual components of our system build on existing work, we have gained valuable experience getting these pieces to work effectively together, and we want to pass these lessons on to others in the field. Second, we describe several new datasets. Each dataset contains thousands of images of isolated leaves, along with segmentation information that extracts their shape. These each include leaves of about 150–250 different species of plants, with about 30 different leaves per species. These are by far the largest publicly available sets of leaf images and provide a unique challenge set for researchers on shape understanding. Third, we demonstrate recognition results for shape matching on two of these datasets (Figure 4). This can be viewed as a high-performance baseline system for shape matching. In this context, we pose a challenge problem to the computer vision community. We describe a set of

performance criteria and offer to include in our deployed system code for any algorithm that can meet these criteria.

After describing prior work in Section 2, we describe in Section 3 extensive datasets that we have collected for this project, which we are now making publicly available. In Section 4, we address a number of practical considerations needed to get a color-based EM algorithm to effectively segment images of isolated leaves. In Section 5, we summarize the shape comparison algorithm we use. In addition, we describe a nearest-neighbor method for metric spaces that significantly speeds up the comparisons needed for this approach. In Section 6, we describe the hardware and UIs that we have constructed to allow the user to navigate the search results. We also describe our ongoing work on experimental augmented reality (AR) UIs for the system. We present a challenge problem for the computer vision community in Section 7 and describe our plans for a future system in Section 8.

2 Related Work

2.1 Massive Digitization Efforts

The amount of digital information available on-line has recently increased dramatically. For example, our group has digitally photographed (at high resolution) each of the 90,000 type specimens of vascular plants in the US National Herbarium at the Smithsonian, where the images are now available at <http://botany.si.edu/types/>. Complementary efforts include those of the New York Botanical Garden (120,000 high resolution images), the Royal Botanical Gardens, Kew (50,000 images, including 35,000 images of type specimens), and the Missouri Botanical Garden (35,000 images of plants). Recently, a consortium of museums and research institutions announced the creation of the Encyclopedia of Life (<http://www.eol.org>) to someday house a webpage for each species of organism on Earth.

2.2 New Means to Access Data

Traditionally, biologists use field guides and dichotomous keys to assist in species identification. Field guides contain pictures and textual descriptions of known species. Dichotomous keys provide a decision tree based on features of the organism, with species at the leaves of the tree. Although valuable, neither solves the problem of identification, as field guides are difficult to search, and dichotomous keys contain questions that are daunting to the non-expert and difficult even for experts to answer with certainty. Electronic versions of these tools have been available for a long time (Pankhurst [16]; Edwards and Morse [6]; Stevenson et al. [21]). Electronic keys have been created through character databases (e.g., Delta: <http://delta-intkey.com>, Lucid: <http://www.lucidcentral.org>). Some of these guides are available on-line or for downloading onto PDAs (e.g., Heidorn [11]), while active websites are being developed that can continually be revised and updated (e.g., <http://botany.si.edu/>

pacificislandbiodiversity/hawaiianflora/index.htm). While valuable, these electronic systems do not solve the fundamental problems faced by traditional tools.

2.3 Visual Search

Automatic recognition systems promise to greatly enhance access by using images as search keys—this, we believe, is the real key to making any such electronic field guide truly groundbreaking. There has been a good deal of work on identifying plants, primarily using leaf shape (see Nilsback and Zisserman [15], though, for recent work using flowers). Abbasi et al. [1] and Mokhtarian and Abbasi [14] present a method for classifying images of chrysanthemum leaves. Saitoh and Kaneko [18] use a neural network to classify wild flowers based on shape and color. Wang et al. [23] use what they call the centroid-contour distance, combined with more standard, global descriptions of shape. Ling and Jacobs [13] introduce shape descriptions based on the *Inner Distance*, which they combine with shape contexts (Belongie et al. [5]), and show that the resulting IDSC outperforms many other approaches on two large leaf datasets. More recently, Felzenszwalb and Schwartz [8] have presented a hierarchical shape matching algorithm that performs even better on a publicly available leaf dataset (Söderkvist [20]). However, since this method is significantly slower, a fast version of the IDSC seems to be the best approach currently available for a large-scale, real-time identification system. We present experiments with this algorithm using data sets that are ten times the size of those used in Ling and Jacobs [13].

This paper is the first complete description of our system. A preliminary version of our system was described in the botanical journal *Taxon* [2] to introduce these ideas to biologists. Work on UIs for automated species identification has been described in [24], [25]. Many components of the current system have not appeared in any previous publication, including our segmentation algorithm and our use of nearest neighbor algorithms. Finally, our datasets and experiments are described here for the first time.

3 Datasets

An important objective of our project is the development of standard, comprehensive datasets of images of individual leaves. Currently, the only large leaf image dataset available to vision researchers is a collection of 15 species with 75 leaf images per species (Söderkvist [20]). This dataset is useful, but insufficient for testing large-scale recognition algorithms needed for species identification. The datasets that we have collected have an order of magnitude more species and are well suited for testing the scalability of recognition algorithms. They also provide complete coverage of species in a geographical area. We have made them available for research use at <http://herbarium.cs.columbia.edu/data.php>.

Leaves were collected by field botanists covering all plant species native to a particular region, and entered in the collections of the US National Herbarium.

The number of leaves per species varied with availability, but averaged about 30. After collection, each leaf was flattened by pressing and photographed with a ruler and a color chart for calibration. Each side of each leaf was photographed with top and bottom lighting. The leaf images were then automatically resized to a maximum side dimension of 512 pixels. Because manual processing of multiple, large datasets is impractical, we developed systems to automatically crop images to remove the ruler, color chart and empty space, and then to segment the images to separate the leaf from the background, as described in the next section. The results were inspected by hand, and a small number of erroneously processed images were removed from the dataset. The datasets consist of the cropped isolated leaf images, as well as the corresponding segmented binary images. To date, we have collected the following three single leaf datasets, each representing different regional flora with about 30 leaves per species:

Flora of Plummers Island: 5,013 leaves of 157 species. Provides complete coverage of all vascular plant species of Plummers Island, MD, an island in the Potomac River near Washington, DC, which has long been studied by botanists.

Woody Plants of Baltimore-Washington, DC: 7,481 leaves of 245 species. Provides complete coverage of all native woody plants (trees and shrubs) of the Baltimore-Washington, DC area.

Trees of Central Park: 4,320 leaves of 144 species. Provides complete coverage of the trees of Central Park in New York City.

Finally, it is often critical for botanists to access more complete type specimens when identifying species. When a new species is discovered, a cutting of branches, leaves, and possibly flowers and fruit is collected. This specimen becomes the *type specimen* that is then used as the definitive representative of the species. Type specimens are stored in herbaria around the world. As part of this work, we have helped to complete the digitization of the complete Type Specimen collection of vascular plants at the US National Herbarium:

US National Herbarium Type Specimen Collection: 90,000 images, covering more than one quarter of all known plant species. Each specimen has been digitally photographed under controlled lighting to produce an 18 megapixel image. These are online in lower resolution formats at <http://botany.si.edu/types/>.

4 Segmentation

In our automatic identification system, a user photographs a leaf so that its shape may be matched to known species. To extract leaf shape, we must begin by segmenting the leaf from its background. While segmentation is a well-studied and difficult problem, we can simplify it in our system by requiring the user to photograph an isolated leaf on a plain white background. However, while we can require users to avoid complex backgrounds and extreme lighting conditions, a

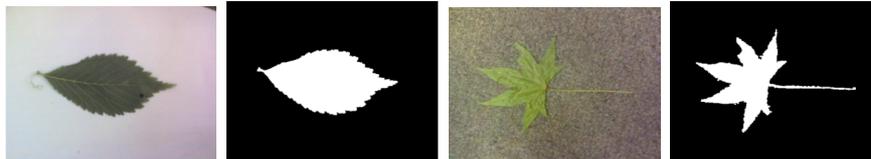


Fig. 3. The first and third images show input to the system, to the right of each are segmentation results. We first show a typical, clean image, and then show that segmentation also works with more complex backgrounds.

useful segmentation algorithm must still be robust to some lighting variations across the image and to some shadows cast by leaves.

Unfortunately, there is no single segmentation algorithm that is universally robust and effective for off-the-shelf use. We have experimented with a number of approaches and achieved good performance using a color-based EM algorithm (see, e.g., Forsyth and Ponce [9]). To begin, we map each pixel to HSV color space. Interestingly, we find that it is best to discard the hue, and represent each pixel with saturation and value only. This is because in field tests in the forest, we find that the light has a greenish hue that dominates the hue of an otherwise white background. We experimented with other representations, and colored paper backgrounds of different hues, but found that they presented some problems in separating leaves from small shadows they cast.

Once we map each pixel to a 2D saturation-value space, we use EM to separate pixels into two groups. First, during clustering we discard all pixels near the boundary of the image, which can be noisy. We initialize EM using K-means clustering with $k = 2$. We initialize K-means by setting the background cluster to the median of pixels near the boundary, and setting the foreground cluster to the mean of the central pixels. Then, in order to make the segmentation real-time, we perform EM using 5% of the image pixels. Finally, we classify all pixels using the two resulting Gaussian distributions. The leaf was identified as the largest connected component of the foreground pixels, excluding components that significantly overlap all sides of the image (sometimes, due to lighting effects, the foreground pixels consist of the leaf and a separate connected component that forms a band around the image). In sum, to get effective results with an EM-based approach has required careful feature selection, initialization, sampling, and segment classification. Figure 3 shows sample results.

Although we did not rigorously evaluate competing segmentation algorithms, we would like to informally mention that we did encounter problems when attempting to apply graph-based segmentation algorithms to these images (e.g., Shi and Malik [19], Galun et al. [10]). One reason for this is that these algorithms have a strong bias to produce compact image segments. While this is beneficial in many situations, it can create problems with leaves, in which the stems and small leaflets or branches are often highly non-compact. The segmentation algorithm that we use goes to the other extreme, and classifies every

pixel independently, with no shape prior, followed by the extraction of a single connected component. It is an interesting question for future research to devise segmentation algorithms that have shape models appropriate for objects such as leaves that combine compact and thin, wiry structures with a great diversity of shape.

5 Shape Matching

Our system produces an ordered list of species that are most likely to match the shape of a query leaf. It must be able to produce comparisons quickly for a dataset containing about 8,000 leaves from approximately 250 species. It is useful if we can show the user some initial results within a few seconds, and the top ten matches within a few seconds more. It is also important that we produce the correct species within the top ten matches as often as possible, since we are limited by screen size in displaying matches.

To perform matching, we make use of the *Inner Distance Shape Context* (IDSC, Ling and Jacobs [13]), which has produced close to the best published results for leaf recognition, and the best results among those methods quick enough to support real-time performance. IDSC samples points along the boundary of a shape, and builds a 2D histogram descriptor at each point. This histogram represents the distance and angle from each point to all other points, along a path restricted to lie entirely inside the leaf shape. Given n sample points, this produces n 2D descriptors, which can be computed in $O(n^3)$ time, using an all pairs shortest path algorithm. Note that this can be done off-line for all leaves in the dataset, and must be done on-line only for the query. Consequently, this run-time is not significant.

To compare two leaves, each sample point in each shape is compared to all points in the other shape, and matched to the most similar sample point. A shape distance is obtained by summing the χ^2 distance of this match over all sample points in both shapes, which requires $O(n^2)$ time.

Since IDSC comparison is quadratic in the number of sample points, we would like to use as few sample points as possible. However, IDSC performance decreases due to aliasing if the shape is under-sampled. We can reduce aliasing effects and boost performance by smoothing the IDSC histograms. To do this, we compute m histograms by beginning sampling at m different, uniformly spaced locations, and average the results. This increases the computation of IDSC for a single shape by a factor of m . However, it does not increase the size of the final IDSC, and so does not affect the time required to compare two shapes, which is our dominant cost.

We use a nearest neighbor classifier in which the species containing the most similar leaf is ranked first. Because the shape comparison algorithm does not embed each shape into a vector space, we use a nearest neighbor algorithm designed for non-Euclidean metric spaces. Our distance does not actually obey the triangle inequality because it allows many-to-one matching, and so it is not really a metric (eg., all of shape A might match part of C , while B matches a different

part of C , so A and B are both similar to C , but completely different from each other). However, in a set of 1161 leaves, we find that the triangle inequality is violated in only .025% of leaf triples, and these violations cause no errors in the nearest neighbor algorithm we use, the AESA algorithm (Ruiz [17]; Vidal [22]). In this method, we pre-compute and store the distance between all pairs of leaves in the dataset. This requires $O(N^2)$ space and time, for a dataset of N leaves, which is manageable for our datasets. At run time, a query is compared to one leaf, called a pivot. Based on the distance to the pivot, we can use the triangle inequality to place upper and lower bounds on the distance to all leaves and all species in the dataset. We select each pivot by choosing the leaf with the lowest current upper bound. When one species has an upper bound distance that is less than the lower bound to any other species, we can select this as the best match and show it to the user. Continuing this process provides an ordered list of matching species. In comparison to a brute force search, which takes nine

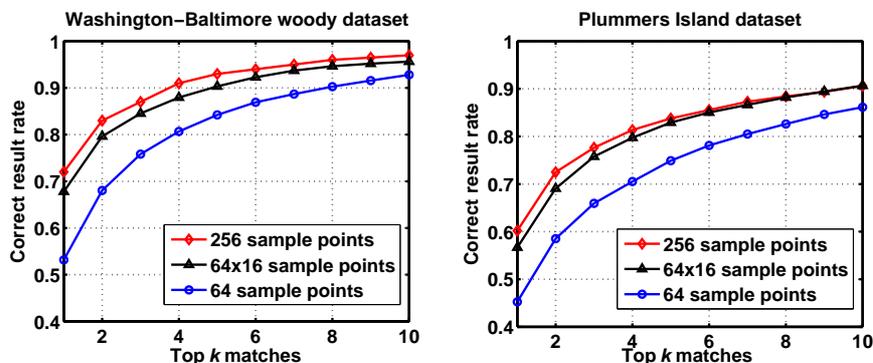


Fig. 4. Experimental results for two datasets.

seconds with a dataset of 2004 leaves from 139 species, this nearest-neighbor algorithm reduces the time required to find the ten best matching species by a factor of 3, and reduces the time required to find the top three species by a factor of 4.4.

We have tested our algorithm using both the Plummers Island and Baltimore-Washington Woody Plants datasets. We perform a leave-one-out test, in which each leaf is removed from the dataset and used as a query. Figure 4 shows performance curves that indicate how often the correct species for a query is placed among the top k matches, as k varies. In this experiment, we achieve best performance using $n = 256$ sample points for IDSC. We reach nearly the same performance by computing the histograms using $n = 64$ sample points averaged over $m = 16$ starting points. The figure also shows that using $n = 64$ points without smoothing significantly degrades performance. Using 64 sample points is approximately 16 times faster than using 256 sample points. The correct an-

swer appears in the top ten about 95%–97% of the time for woody plants of Baltimore-Washington and somewhat less (about 90% of the time) for the flora of Plummers Island. This is in part because shape matching is not very effective at discriminating between different species of grass (which are not woody plants). Overall, these results demonstrate effective performance. It seems that most errors occur for species in which the overall leaf shape is not sufficiently distinctive. We plan to address these issues by using additional cues, such as small scale features of the leaf margin (e.g., toothed or smooth) and the shape of the venation (vascular structure).

6 User Interfaces and Hardware

We have developed several prototype UIs to integrate the individual pieces of the matching system and investigate the performance of our interaction techniques and vision algorithms in real world situations. These prototypes are the result of collaboration with our botanist colleagues in an iterative process that has included ethnographic study of botanical species identification and collection in the field, user centered design, interaction technique development, and qualitative and quantitative feedback and user studies. We have pursued two primary research directions. The first focuses on existing mobile computing platforms for ongoing botanical field studies. The second develops mobile AR systems that are not appropriate for field use in their current form, but could provide significant advantages as hardware and software mature.

The conceptual model we use in our mobile computing platform is an extension of existing paper field guides. The system provides access to a library of knowledge about the physical world, and the physical leaf is the key to that information. In the AR prototype, virtual images representing matched species appear adjacent to the leaf in the physical world and can be manipulated directly through tangible interaction. In this case, the conceptual model is enhanced perception: the leaf anchors information embedded in the environment and accessed through augmented reality.

6.1 Mobile Computing

Our initial prototype, LeafView (Figure 1), provides four tabbed panes for interaction: browse, sample, search results, and history. The browse pane provides a zoomable UI (ZUI) (Bederson et al. [3]) with which the user can explore an entire flora dataset. When the user photographs a leaf with the system, the image is immediately displayed in the sample pane with contextual information including time, date, GPS location, and collector. The segmented image is displayed next to the captured leaf image to show the user what LeafView “sees” and provide feedback about image quality. As results are found, they are displayed with the original image in the search results pane. Each species result provides access to the matched leaf, type specimens, voucher images and information about the species in a ZUI to support detailed visual inspection and comparison, which



Fig. 5. AR user interface viewed through a video see-through display.

is necessary when matching is imperfect. Selecting a match button associates a given species with the newly collected specimen in the collection database. The history pane displays a visual history of each collected leaf, along with access to previous search results, also in a ZUI. This represents the collection trip, which can be exported for botanical research, and provides a reference for previously collected specimens. Making this data available improves the long term use of the system by aiding botanists in their research.

LeafView was built with C#, MatLab, and Piccolo (Bederson, et al. [4]). Our first versions of the hardware used a Tablet PC with a separate Wi-Fi or Bluetooth camera and a Bluetooth WAAS GPS. However, feedback from botanists during field trials made it clear that it would be necessary to trade off the greater display area/processing power of the Tablet PC for the smaller size/weight of an Ultra-Mobile PC (UMPC) to make possible regular use in the field. We currently use a Sony VAIO VGN-UX390N, a UMPC with an integrated camera and small touch-sensitive screen, and an external GPS.

6.2 Augmented Reality

AR can provide affordances for interaction and display that are not available in conventional graphical UIs. This is especially true of Tangible AR (Kato et al. [12]), in which the user manipulates physical objects that are overlaid with additional information. Tangible AR is well matched to the hands-on environmental interaction typical of botanical field research. While current head-worn displays and tracking cannot meet the demands of daily fieldwork, we are developing experimental Tangible AR UIs to explore what might be practical in the future.

In one of our Tangible AR prototypes (Figure 5), a leaf is placed on a clipboard with optical tracking markers and a hand-held marker is placed next to the leaf to initiate a search. The results of matching are displayed alongside the physical leaf as a set of individual leaf images representing *virtual vouchers*, multifaceted representations of a leaf species that can be changed through tangible gestures. As the user passes the hand-held marker over a leaf image, the

card visually transforms into that leaf’s virtual voucher. The visual representation can be changed, through gestures such as a circular ”reeling” motion, into images of the type specimen, entire tree, bark, or magnified view of the plant. Inspection and comparison is thus achieved through direct spatial manipulation of the virtual voucher—the virtual leaf in one hand and the physical leaf on the clipboard in the other hand. To accept a match, the virtual voucher is placed below the leaf and the system records the contextual data.

Different versions of our Tangible AR prototypes use a monoscopic Liteye-500 display, fixed to a baseball cap, and a stereoscopic Sony LDI-D100B display, mounted on a head-band, both of which support 800×600 resolution color imagery. The system runs on a UMPC, which fits with the display electronics into a fannypack. The markers are tracked in 6DOF using ARToolkit (Kato et al. [12]) and ARTag (Fiala [7]), with a Creative Labs Notebook USB 2.0 camera attached to the head-worn display.

6.3 System Evaluation

Our prototypes have been evaluated in several ways during the course of the project. These include user studies of the AR system, field tests on Plummers Island, and expert feedback, building on previous work (White et al. [24]). In May 2007, both LeafView and a Tangible AR prototype were demonstrated and used to identify plants during the National Geographic BioBlitz in Rock Creek Park, Washington, DC, a 24-hour species inventory. Hundreds of people, from professional botanists to amateur naturalists, school children to congressmen, have tried both systems. While we have focused on supporting professional botanists, people from a diversity of backgrounds and interests have provided valuable feedback for the design of future versions.

7 Challenge Problem for Computer Vision

One goal of our project is to provide datasets that can serve as a challenge problem for computer vision. While the immediate application of such datasets is the identification of plant species, the datasets also provide a rich source of data for a number of general 2D and silhouette recognition algorithms.

In particular, our website includes three image datasets covering more than 500 plant species, with more than 30 leaves per species on average. Algorithms for recognition can be tested in a controlled fashion via leave-one-out tests, where the algorithms can train on all but one of the leaf images for each species and test on the one that has been removed. The web site also contains separate training and test datasets in order to make fair comparisons. Our IDSC code can also be obtained there, and other researchers can submit code and performance curves, which we will post. We hope this will pose a challenge for the community, to find the best algorithms for recognition in this domain.

Note that our system architecture for the electronic field guide is modular, so that we can (and will, if given permission) directly use the best performing methods for identification, broadening the impact of that work.

8 Future Plans

To date, we have focused on three regional floras. Yet, our goal is to expand the coverage of our system in temperate climates to include all vascular plants of the continental U.S. Other than the efforts involved in collecting the single leaf datasets, there is nothing that would prevent us from building a system for the U.S. flora. The visual search component of the system scales well: search can always be limited to consider only those species likely to be found in the current location, as directed by GPS.

In addition, we have begun to expand into the neotropics. The Smithsonian Center for Tropical Forest Science has set up twenty 50-hectare plots in tropical ecosystems around the world to monitor the changing demography of tropical forests. We aim to develop versions of the system for three neotropical floras: Barro Colorado Island, Panama; Yasuni National Park, Ecuador; and the Amazon River Basin in Brazil. This domain demands algorithms that not only consider leaf shape, but also venation (i.e., the leaf's vascular structure). Initial results are quite promising, but we have not yet developed a working system.

Finally, we have developed a prototype web-based, mobile phone version of our system, allowing anyone with a mobile phone equipped with a camera and browser to photograph leaves and submit them to a server version of our system for identification. We hope to develop a touch-based version on an iPhone or Android-based device in the near future. We feel that it should soon be possible to create a mobile phone-based system that covers the entire U.S., usable by the general population.

Acknowledgements

This work was funded in part by National Science Foundation Grant IIS-03-25867, An Electronic Field Guide: Plant Exploration and Discovery in the 21st Century, and a gift from Microsoft Research.

References

1. Abbasi, S., Mokhtarian, F., and Kittler, J. 1997. Reliable classification of chrysanthemum leaves through curvature scale space. *Proc. Scale-Space '97*. 284-295.
2. Agarwal, G., Belhumeur, P., Feiner, S., Jacobs, D., Kress, W.J., Ramamoorthi, R., Bourg, N., Dixit, N., Ling, H., Mahajan, D., Russell, R., Shirdhonkar, S., Sunkavalli, K., White, S. 2006. First steps towards an electronic field guide for plants. *Taxon*, 55:597-610.
3. Bederson, B. 2001. PhotoMesa: A zoomable image browser using quantum treemaps and bubblemaps. *Proc. ACM UIST 2001*, 71-80.
4. Bederson, B., Grosjean, J., and Meyer, J. 2004. Toolkit design for interactive structured graphics. *IEEE Trans. on Soft. Eng.*, 30(8):535-546.
5. Belongie, S., Malik, J. and Puzicha, J. 2002. Shape Matching and Object Recognition Using Shape Context, *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 24(4):509-522.

6. Edwards, M., and Morse, D.R. 1995. The potential for computer-aided identification in biodiversity research. *Trends in Ecology and Evolution*, 10:153-158.
7. Fiala, M., 2005. ARTag, a fiducial marker system using digital techniques. *Proc. CVPR 2005*, 590-596.
8. Felzenszwalb, P. and Schwartz, J. 2007. Hierarchical matching of deformable shapes, *Proc. CVPR*, 1-8.
9. Forsyth, D. and Ponce, J. 2003. *Computer vision: A modern approach*. Prentice Hall, Upper Saddle River, NJ.
10. Galun, M., Sharon, E., Basri, R., and Brandt, A. 2003. Texture segmentation by multiscale aggregation of filter responses and shape elements, *Proc. CVPR*, 716-723.
11. Heidorn, P. B. 2001. A tool for multipurpose use of online flora and fauna: The Biological Information Browsing Environment (BIBE). *First Monday*, 6(2). http://firstmonday.org/issues/issue6_2/heidorn/index.html.
12. Kato, H., Billinghamurst, M., Poupayrev, I., Imamoto, K., and Tachibana, K. 2000. Virtual object manipulation of a table-top AR environment, *Proc. IEEE and ACM ISAR*, 111-119.
13. Ling, H. and Jacobs, D. 2007. Shape Classification Using the Inner-Distance, *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 29(2):286-299.
14. Mokhtarian, F. and Abbasi, S. 2004. Matching shapes with self-intersections: Application to leaf classification, *IEEE Trans. on Image Proc.*, 13(5):653-661.
15. Nilsback, M. and Zisserman, A. 2006. A visual vocabulary for flower classification, *Proc. CVPR*, 1447-1454.
16. Pankhurst, R. J. 1991. *Practical taxonomic computing*. Cambridge University Press, Cambridge, UK.
17. Ruiz, E. 1986. An algorithm for finding nearest neighbours in (approximately) constant average time. *Patt. Rec. Lett.*, 4(3):145-157.
18. Saitoh, T. and Kaneko, T. 2000. Automatic recognition of wild flowers. *Proc. ICPR*, 2:2507-2510.
19. Shi, J. and Malik, J. 2000. Normalized Cuts and Image Segmentation, *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 22(8):888-905.
20. Söderkvist, O. 2001. *Computer vision classification of leaves from Swedish trees*. Master Thesis, Linköping Univ.
21. Stevenson, R. D., W. A. Haber, and R. A. Morris. 2003. Electronic field guides and user communities in the eco-informatics revolution. *Conservation Ecology* 7:3. <http://www.consecol.org/vol7/iss1/art3>
22. Vidal, E. 1994. New formulation and improvements of the nearest-neighbour approximating and eliminating search algorithm (AESA). *Patt. Rec. Lett.* 15(1):1-7.
23. Wang, Z., Chi, W., and Feng, D. 2003. Shape based leaf image retrieval. *IEE Proc. Vision, Image and Signal Processing*, 150(1):34-43.
24. White, S., Feiner, S., and Kopylec, J. 2006. Virtual vouchers: Prototyping a mobile augmented reality user interface for botanical species identification. *Proc. IEEE Symp. on 3DUI*, 119-126.
25. White, S., Marino, D., and Feiner, S., 2007. Designing a mobile user interface for automated species identification, *Proc. CHI 2007*, 291-294.