Automatic Species Identification of Live Moths

Michael Mayo[1] and Anna T. Watson[2]
[1]Dept. of Computer Science
The University of Waikato
Hamilton, New Zealand
[1]mmayo@cs.waikato.ac.nz
[2]kisalsera@yahoo.co.uk

Abstract. A collection consisting of the images of 774 live moth individuals, each moth belonging to one of 35 different UK species, was analysed to determine if data mining techniques could be used effectively for automatic species identification. Feature vectors were extracted from each of the moth images and the machine learning toolkit WEKA was used to classify the moths by species using the feature vectors. Whereas a previous analysis of this image dataset reported in the literature [1] required that each moth's least worn wing region be highlighted manually for each image, WEKA was able to achieve a greater level of accuracy (85%) using support vector machines without manual specification of a region of interest at all. This paper describes the features that were extracted from the images, and the various experiments using different classifiers and datasets that were performed. The results show that data mining can be usefully applied to the problem of automatic species identification of live specimens in the field.

1.   Introduction

The demand for computer-based systems that can automatically identify the species of live plants, insects or animals from digital images or recordings is only likely to increase in the future. Worldwide, there is a growing need for biodiversity monitoring, while at the same time the number of trained taxonomists declines [2]. Accurate species identification is also of critical importance in many practical areas, such as agriculture and border control, where pests and invaders must be identified swiftly and eradicated before becoming established as unwanted visitors in a country or agricultural region.

Can data mining techniques be applied to build useful, practical species recognition systems with the express purpose of assisting non-taxonomists in taxonomic tasks? Two significant challenges that data mining practitioners must face if they are to take up this challenge are firstly the need to handle a very high number of classes,

perhaps in the order of 1000-2000 classes for a typical, practical working system. And secondly, a species identification system must also be able to distinguish correctly between an individual from a species that the system has been trained on, and an individual belonging to a species that the system has not been trained on. In other words, the presence of novel species must be accurately recognized.

Traditional machine learning datasets, on the other hand, usually only have a small handful of classes. Furthermore, traditional classifiers do not usually need to detect novelty: they simply return the best matching class, even if the probability of a match is very low.

In this paper, the first of these challenges is addressed: to evaluate data mining techniques when applied to a species dataset with a high number of classes. Although 35 classes is not high relative to the ultimate goal (in which the number of classes can be counted in the thousands), it does represent a significant increase in the number of classes compared to traditional data mining domains and is therefore a good test-bed for study.

The results presented here show that data mining techniques can be effective for species identification tasks. In particular, the best classifier (a support vector machine) in a jackknife test achieved approximately 85% accuracy when no manual preprocessing of the images was performed. In contrast, the original analysis of these images used a neural network-based system to achieve approximately 83% accuracy in a jackknife test [1]. However, in that study, the best wing in each image was manually outlined and mirrored so that that best wings in all the images were in the same orientation, and could be cleanly segmented from the background. In the study presented here, we classified the raw images with no manual human intervention required.

In the following section, an overview of the automatic species identification field is briefly given. The live moth image dataset is then discussed in more detail in Section 3. Section 4 describes the features that were extracted from the moth images for this study, and Section 5 contains the results of the experiments. Section 6 is the conclusion.

2.   Recent Work on Automatic Species Identification

Researchers from a variety of different backgrounds have in recent times attempted to build accurate species identification systems. Most systems to date have a particular and focused domain, although the "grand challenge" in this area is to design a generic species identification system that can identify any organism on the planet from visual and any other type of available data – quite a lofty challenge.

One interesting aspect of this field is that data need not necessarily be limited to images. Animals may, for example, be recognized by the sounds they make, which is the basis for much research into bird species recognition (see, e.g., [3,4]). Furthermore, there is no reason why a human in the field who is trying to identify the species of an organism could not provide auxiliary information for the classification process in the form of answers to questions (e.g. "was this creature

found in a hole in the ground?"). However most research to date is currently focused on image analysis and classification.

Three main image-based systems have appeared in the literature. The first is DAISY (Digital Automated Identification System) [1,2], the second is SPIDA [5], and the third is ABIS [6]. The first two systems use a neural network for the core classification task. DAISY is generic and can be used for any type of image classification, and the system has in the past been used for classifying fish, pollen, and plants of different types, although the main focus appears to be moths. SPIDA, on the other hand, is specifically designed to distinguish between 121 species of Australian spider (it has also recently been tested on bumblebee wings with limited success). An interesting feature of SPIDA is that it is currently available on the Internet, and users can submit their own images of spiders for classification, although some expertise and equipment is required to obtain optimal images. Aside from providing a useful service to people, this is also a novel method of obtaining more training data. The third system, ABIS, is used to recognize bee species from the pattern on a bee's wings. It is based on support vector machines, and has achieved over 95% accuracy.

In general, all three of these systems have been trained on images taken from dead specimens. In a lab, dead specimens can be carefully positioned and photographed under consistent and ideal lighting conditions. In addition, details such as anatomical cross sections can also be provided as inputs to the system.

In the field, however, specimens are more often alive and moving than dead. Live specimens may not adopt the ideal pose required, they may move when the image is being captured, and the lighting conditions outside the lab may be poor and change unpredictably as a series of images of a group of images is taken. This tends to make the recognition task much more difficult. Furthermore, if the background is not guaranteed to be uniform (e.g. a moth is photographed when it lands on a tree branch), then the classification task becomes even harder because the system must distinguish the object of interest from the background. To our best knowledge, the Macrolepidoptera dataset used in this research is the only collection of images of live specimens taken in the field that is currently described in the literature.

Other systems have also been used for species identification. For example, Moyo et al. [7] discuss recognition of mammalian species in forensics from patterns occurring on a single hair when examined under a microscope, and Yuan et al. [8] report on the recognition of rat species from images of their tracks (which is paramount for determining the presence of invasive rat species that can cause the extinction of other species).

Researchers have also focused on the automatic identification of plant species, such as Samul et al. [9] who report on a method of recognizing the species of a tree given textural features, and Sogaard [10], who describes recent work on classifying weed plants from non-weed plants in order to minimize the wastage of herbicides in agriculture.

3.   The Macrolepidoptera Image Collection

The dataset used in this study is a library of live moth images created by the second author [1] over a period of nearly a year. A moth trap was set up in Treborth Botanical Garden, Gwynedd, UK, and cleared every morning. Captured live moths were photographed then released.

Figure 1 gives examples of two images in the dataset. Figure 1(a) is most typical of images in the dataset, in which a moth has been photographed against a nearly uniform background (although the lighting is not uniform from left to right). Figure 2(b) is an example of an image taken against a non-uniform background, in which the colour of the moth is actually quite similar to the colour of the background, making this image difficult to segment. Due to the practical nature of the task (photographing live moths) it is not surprising that all of the images do not have clean, uniform backgrounds. In practice, it would be desirable for a classifier to automatically learn the difference between a moth and the background, and for this reason images such as Figure 1(b) were retained in the dataset.

Figure 1. Images of (a) *Agrotis exclamationis* (uniform background) and (b) *Laothoe populi* (poorer background).

Each photographic image is 1024 x 960 pixels in resolution, and is a full 24-bit RGB image. In actuality, the complete image dataset contains examples of over 224 species of moth, but of those, only 70 species contain images of twenty or more individuals. Furthermore, because in the original analysis [1] the region in the image where the wing occurs was manually outlined in order to be separated cleanly from the background, and because this manual highlighting was extremely time consuming, only the best quality 35 classes, or a total of 774 images, were actually used. In this paper, we focus on only those 35 classes in order to compare our results with the results of the original analysis, and leave to future work a consideration of the expanded dataset with 70, or even all 224 classes.

4.   Extracting Features from the Images

The open source image processing toolkit ImageJ [11] was used to analyse the images in the Macrolepidoptera collection and extract from each one a feature vector containing 11,300 numeric features. Feature extraction was automated using ImageJ's macro language facility, which enables any of ImageJ's GUI features to

be invoked programmatically. The feature vectors were then written to a file in ARFF format, ready for further processing by WEKA [12].

To illustrate the feature extraction process, consider the image of an individual moth from the species *Polia nebulosa*, depicted in Figure 2. This image is interesting because the subject is off-centre to the right, which is typical for many of the images in the dataset. There is also some additional noise in the image, namely some distortions in the upper right-hand corner, and the moth has quite a distinct shadow due to the outside (uncontrolled) lighting conditions when the photograph was taken.

Figure 2. Image of *Polia nebulosa.*

Figure 3 depicts the automatic processing and feature extraction pipeline, using Figure 2 as an example input. Features were extracted from the image at multiple points during processing, with 11,300 numeric features per image ultimately being generated. The features were a mixture of global image statistics (taken at various points in the processing of the image) and also local image statistics, generated by breaking the image up into "patches" and taking individual statistics such as the mean pixel values from each patch, in an attempt to account for different moth shapes and wing patterns.

The first step in feature extraction was to smooth the image six times. Each smoothing operating replaces each pixel in the image with the average intensity of its 3x3 neighbourhood of pixels, and this is repeated in the red, blue and green planes. Figure 3(a) depicts the results of this smoothing when applied to the image of *P. nebulosa* in Figure 2. The image was then transformed from RGB colour into an 8-bit greyscale image, and the edges were detected using the Sobel edge detector, as shown in Figure 3(b).

The next step was to apply an automatic binary thresholding scheme to each edge-detected image to produce a binary image such as the image depicted in Figure 3(c). The actual value of the threshold used was, for each image, set automatically using the iterative Isodata algorithm [13].

With the original RGB colour image now transformed into a binary image, a number of features were calculated. The features consisted of the number of white (foreground) pixels, the number of black (background) pixels, the ratio of the above two statistics (giving an overall indication of the moth's size), as well some

features related to the shape of the binarized moth such as the interior density, the standard deviation of pixel positions in the image, the skew, and the kurtosis.

To account for the fact that in many of the images, the moth is off-centre, the centroid of the binary image was also calculated. The centroid for the example image in Figure 3 (c) is (734, 465), and when the centroid is labelled on the image in Figure 3(d), it clearly gives a good estimate of where the centre of the moth is in relation to the image.
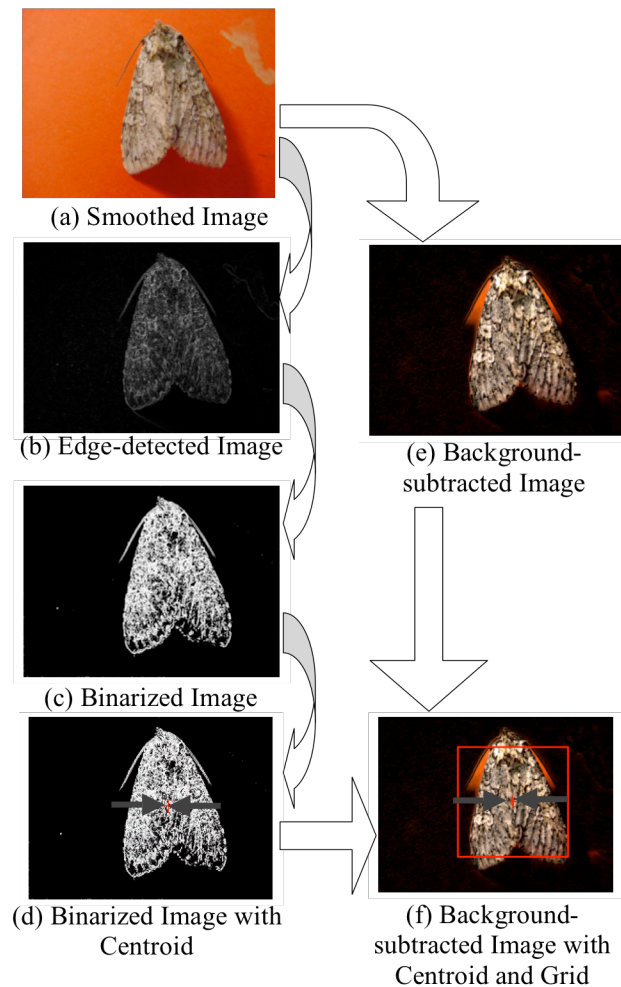

(a) Smoothed Image

(b) Edge-detected Image

(c) Binarized Image

(d) Binarized Image with Centroid

(e) Background-subtracted Image

(f) Background-subtracted Image with Centroid and Grid

Figure 3. The image processing and feature extraction process.

Once the centroid was calculated and the binary features extracted, the binarized version of the image was discarded (but the features calculated from it were retained) and all the further processing was performed on the original smoothed RGB colour version of the image in Fig. 3(a).

Processing of the smooth colour version of the image proceeded as following: first of all, the background was removed from the image using Sternberg's rolling ball algorithm [14]. The result is shown in Fig. 3(e). Next, global colour features were measured. Because Fig. 3(e) is a 24-bit colour image, separate measurements were taken for each of the red, green and blue planes of the image.

The measurements for each plane included the same measurements as taken from the binary version of the image (such as the skew and kurtosis), and also various global statistics: the mean, minimum, and maximum pixel intensities, and their standard deviation. Two colour frequency histograms were also encoded as features for each plane, one consisting of all 256 bins, and another condensed version consisting of only 16 bins.

Finally, the image was transformed from the RGB colour space to the HSB (Hue, Saturation, Brightness) colour space and all of the measurements described above were taken again, but this time using the HSB version of the image. Thus, the resulting set of global features included statistics and measurements taken from binary, RGB, and HSB versions of the image.

The next step in processing was to extract local features from the images. This is depicted in Figure 3(f). A grid of total size 600 x 600 pixels was centred over the centroid of the moth. The 600 x 600 square was then subdivided into 400 square patches, each patch approximately 30 x 30 pixels in size. For each patch, the mean, minimum, maximum, and standard deviation of the pixel values was calculated and added to the feature vector. This was performed firstly in RGB colour space, and then, as has been described above, the pixels in the region were transformed into HSB and the measurements were taken again. Thus an additional 9,600 local features were generated in this way and added to the global features generated previously from the binary and colour images.

5.   Results

Once the feature vectors had been generated using ImageJ and saved to a file, the machine learning toolkit WEKA [12] was used to perform further analysis. WEKA contains a number of different classifiers including naïve Bayes, instance-based learning, decision trees, random forests and support vector machines. Each classifiers was applied to the dataset with default parameters in a stratified 10-fold cross validation experiment. Table 1 gives results by classifier.

Classifier	Average Accuracy
Naïve Bayes	65.9%
J48	58.3%
IB1	71.6%
IB5	65.36%
Random Forests (n=200)	83.2%
SMO	85.0%

Table 1. Average prediction accuracy by classifier for Moth species classification.

Table 1 clearly shows that SMO, an implementation of support vector machines that builds a linear separating hyperplane, gives the best overall accuracy. Random forests with 200 randomly generated trees comes a close second, while Naïve Bayes, the decision tree learner J48, and instance-based learning with two different neighbourhood sizes (1 and 5 respectively) both perform poorly in comparison.

The next step was to perform a jackknife test (also known as leave-one-out-cross-validation) on the data using SMO. Jackknife tests give better estimates of how good a classifier is on a dataset, but it is only feasible to perform a jackknife test if the classifier runs relatively efficiently. A jackknife test with SMO in this case yielded a prediction accuracy of 84.8%, quite similar to the 10-fold cross validation result.

Overall, the SMO and random forest results compare favourably to the original analysis of this dataset in [1] in which DAISY achieved a classification accuracy of 83%, albeit with manual pre-processing of each image.

The next issue to consider was the number of attributes in the dataset, which is quite high at 11,300 features. Can the same performance be achieved but with a reduced number of features, thus incurring a smaller cost?

During informal testing before the experiments described here were carried out, it was observed that accuracy tended to increase as the number of cells from which local features were measured increased. That is, the present dataset generated by dividing the 600 x 600 grid over the centroid into 400 patches (that being 20 x 20 patches) gave better accuracy than when the same grid was divided into only 225 patches (being 15 x 15 patches). This is perhaps unsurprising, as a finer grid better captures the textures and patterns on a moth's wings, but at the cost of adding more features.

Experiments in which the local features were deleted completely leaving only global features such as the colour histograms, on the other hand, resulted in about 75% accuracy using SMO: a reasonably high contribution to performance from a relatively small proportion of the total number of attributes.

Another reason that the total number of attributes is quite high is that the same measurements were taken from both RGB and HSB versions of the moth images. It is not at all clear whether both colour spaces are needed for effective moth species classification, and if they are both not needed, it is not clear which colour space is the best. RGB and HSB were both used because transformations between them are straightforward.

With this question in mind, two new datasets that were subsets of the original dataset were generated. The first dataset contained all of the features derived from the binarized and RGB images only, and the second dataset contained all of the features derived from the binarized and HSB images only. A 10-fold cross validation experiment was performed using SMO, the best classifier as indicated in Table 1.

The results were 83.3% and 84.2% respectively: hardly a difference at all, and very close to the 85% accuracy achieved when all of the attributes were included. This

result seems to indicate that duplicating measurements in more than one colour space is only likely to improve accuracy by a very small amount, and that the HSB measurements give slightly better quality predictions than the RGB measurements.

The next and more important question to investigate was which of the 35 classes in the dataset were the main contributors to the classification error rate. Statistics per class were recorded during the SMO jackknife test and they are presented in Table 2.

Class	TP Rate	FP Rate
Agrotis exclamationis	0.909	0
Alcis repandata	0.958	0
Anaplectoides prasina	0.864	0.001
Apamea monoglypha	0.917	0.005
Biston betularia	1	0
Cabera pusaria	0.875	0.012
Campaea margaritata	1	0.008
Charanyca trigrammica	0.952	0.001
Chloroclysta truncata	0.769	0.011
Cosmia trapezina	0.857	0.007
Crocallis elinguaria	0.923	0
Diarsia brunnea	0.818	0.003
Herminia tarsipennalis	0.7	0.009
Hydriomena furcata	0.792	0.005
Hylaea fasciaria	0.773	0.008
Hypena grisealis	0.571	0.016
Hypena proboscidalis	0.714	0.009
Idaea aversata	0.87	0.007
Lacanobia oleracea	0.905	0.003
Laothe populi	1	0
Lomospilis marginata	0.81	0.001
Noctua comes	0.85	0.007
Noctua janthe	0.81	0.007
Ochropleura plecta	0.9	0.001
Opistograptis luteolata	0.857	0.003
Orthosia gothica	1	0
Orthosia munda	1	0
Ourapteryx sambucaria	0.826	0.001
Polia nebulosa	0.857	0.001
Rivula sericealis	0.5	0.009
Sphrageidus similis	0.8	0.005
Xanthorhoe designata	0.85	0.005
Xanthorhoe ferrugata	0.682	0.001
Xanthorhoe montanata	0.762	0.009
Xestia triangulum	0.952	0

Table 2. True Positive (TP) and False Positive (FP) rates per class.

Looking at the true positive rates in Table 2, it is clear that some species such as *Biston betularia* and *Orthosia munda*, both with true positive rates of 1.0, are very easily recognised. Other species such as *Hypena grisealis* and *Rivula sericealis* however have unacceptably low true positive rates – 0.57 and 0.5 respectively in these cases.

There are two possible explanations for this: either the poorly recognized species themselves are physically very similar to one or more other species, in which case, individual misclassifications would be most frequently to those very similar other species, or the examples of the poorly recognized species in the dataset are of generally poor quality, in which case the misclassifications would be expected to be more or less uniformly distributed over the other species.

The confusion matrix produced by WEKA during the jackknife experiment with SMO was examined. The confusion matrix for this experiment is a 35 x 35 matrix showing how often instances of one class are either classified correctly, or if they are misclassified, it gives the frequency with which they are misclassified as another class. Although it is too large to reproduce the entire matrix here, some interesting observations can be made.

Consider first of all the two most poorly recognized species from Table 2, *H. grisealis* and *R. sericealis*. When *H. grisealis* is misclassified, the misclassifications (nine in total) are spread nearly uniformly over seven other species. Two examples of *H. grisealis* are classified as *R. sericealis*.

For *R. sericealis*, however, the misclassifications (ten) are highly skewed, with four of the erroneous misclassifications of *R. sericealis* images being to *H. grisealis*. It is not at all obvious from an examination of the misclassified images why they were misclassified. Figure 4 shows four images from the *H. grisealis* and *R. sericealis* classes. There are two images from each class, and one of the images in each case is classified correctly by WEKA while the other is misclassified.

Why were these two classes frequently confused, and why is the source of the confusion not apparent to the human eye? It appears that both *H. grisealis* and *R. sericealis* have quite distinct wing markings, which should easily distinguish them.

There are two possible explanations for this. Firstly, both species have very similar wing colourings, even if the pattern is different. It could be that the global colour histograms of both classes are very similar as a result, which would certainly hinder accurate classification.

Secondly, and perhaps more importantly, the moths were not all photographed at the same uniform distance. Rather, the "best shot" was taken of each moth, so some moths may appear larger relative to others than they actually are. Figures 1(a), 1(b), and 2 are all examples of images in which the moth size is fairly representative of the size of most of the moths in the collection's images. Clearly, however, if a moth photograph is taken from a further distance, then the moth will occupy less of the fixed size 600 x 600 grid of local features, which would hinder classification. Therefore it is much more likely that the spatial distribution of the moth's distinctive wing markings would become harder to recognize or "blurred"

to the system in comparison to representations of the wing patterns in moths photographed at a closer distance.

H. grisealis classified as:

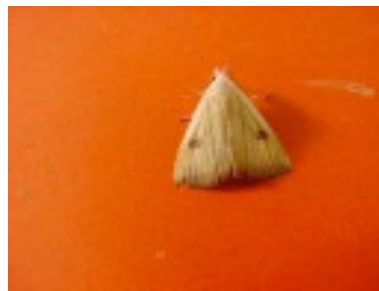|	
(a) *H. grisealis*	(b) *R. sericealis*

R. sericealis classified as:

(a) *H. grisealis*	(b) *R. sericealis*

Figure 4. Correctly and incorrectly classified examples of *H. grisealis* and *R. sericealis*.

With this in mind, a refinement of the feature extraction process to consider in the future is size normalization. The size of the moth (actually a count of the number of foreground pixels in the binarized version of the image) was already recorded for each image during feature extraction; this value could be used to regulate the size of the grid from which the local features are extracted, so that the grid effectively becomes smaller for smaller species of moth, and larger for larger moths.

Examples of other frequently misclassifications are given in Figure 5. Out of seven misclassifications of *Xanthorhoe ferrugata*, four images were incorrectly classified as *Xanthorhoe montanata*. Again, this could be explained by the fact that the *X. ferrugata* appears slightly smaller in the images than the other moths, and therefore there is less information in the non-global features extracted for this class. Interestingly, none of the *X. montanata* images, which is of the usual size in the images, were misclassified as *X. ferrugata*.
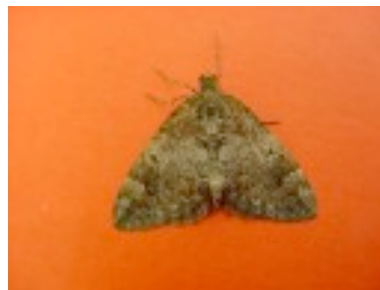
(a) *X. ferrugata*

(b) *X. montanata*
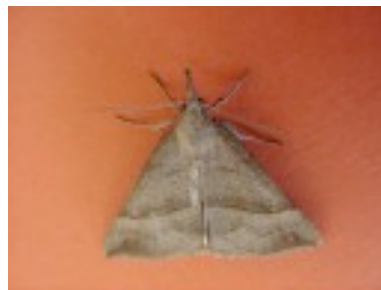
(c) *H. proboscidalis*

(d) *C. truncata*

Figure 5. Further examples of incorrect classifications of (a) *X. ferrugata* as (b) *X. montanata* and (c) *H. proboscidalis* as (d) *C. truncata*.

Finally, another common confusion was between *Hypena proboscidalis* and *Chloroclysta truncata*, with four misclassifications of former as the latter and two misclassifications in the opposite direction. There is no clear explanation as to why this happened.

6.   Conclusion

In summary, this paper has reported on an investigation of a Macrolepidoptera image dataset using the ImageJ image processing toolkit [11] for feature extraction and the WEKA machine learning workbench [12] for classification. The purpose of the investigation was to determine whether data mining techniques could be effective for species classification, where the number of classes is higher than usual (in this case, there were 35 moth species).

This dataset is interesting because the images are of live moths. Whereas dead specimens can be carefully laid out, pinned to a board and imaged in a uniform way under lab conditions, there is much more variability and "messiness" in images taken from live samples. For example, live moths can easily move one of their wings as a photograph is being taken, resulting in a blurred image. Necessity may also dictate that a perfectly uniform background cannot always be used when

imaging a live specimen. And finally, live specimens may not always be centred in an image.

Despite this, the results do indicate that data mining techniques are effective in this domain. A jackknife test achieved accuracy of 85% on the image collection using support vector machines, with random forests coming a close second in terms on classification accuracy. This performance exceeds the classification accuracy achieved in a previous investigation of the same dataset [1].

A significant difference between this study and the previous one is that in the current study, the feature extraction and classification were entirely automatic, whereas previously a human was required to segment the best wing of each moth from the rest of the image before classification could take place. Although WEKA would still probably require moth/background segmentation to be performed if the background of the images are highly cluttered, what has been shown here is that this is not necessary if an effort is made to image the moths against a uniform background, even if lighting conditions are variable.

The most important factor in the success of any machine learning-based image classification system is the features that are extracted. Global features are useful because they are invariant to rotation. However, further accuracy gains can only be achieved by adding local features to the dataset, which is necessary to capture the patterns on the moth's wings. This was achieved by calculating each moth's centroid, centring a square over the centroid, and taking samples from 200 patches inside the main square. This proved mostly effective in capturing wing patterns, except that it may have disadvantaged smaller species that occupy less space in the image. Future work should address the problem of making local feature selection invariant to the size of a moth.

Another interesting avenue for future research is the colour space used to encode image features. In this research, it was found that extracting image features in both RGB and HSB colour spaces, as opposed to just one colour space, boosted accuracy by a small amount. Most images are encoded in RGB colour simply because it is the most commonly used and accessible colour space. However, other colour spaces may be much more suitable for the task of species identification. A more thorough investigation of this possibility needs to be carried out.

Finally, future research must also address the scalability of applying machine learning methods to species recognition. Can support vector machines, the best classifier identified in this study, scale to thousands of classes and still remain efficient and accurate? The main impediment is that support vector machines are two-class classifiers, and therefore many of them must be trained in order to solve multi-class problems. This scalability question needs to be explored.

In summary, we are excited about the application and success to date of data mining techniques in the area of automatic species identification. We plan to continue investigations in this area into the future, using both the Macrolepidoptera dataset described here as well as datasets containing images of other types of species.

Acknowledgements

References

1. Watson A., O'Neill M & Kitching I. 2004. Automated identification of live moths (Macrolepidoptera) using Digital Automated Identification System (DAISY). *Systematics and Biodiversity* 1(3): 287-300.
2. Gaston K. & Oneill M. 2004. Automated Species Identification: why not? *Phil. Trans. R. Soc. Lond.* **B** 259:655-667.
3. Chesmore D. 2004. Automated bioacoustic identification of species. *An. Acad. Bras. Cience.* 76(2): 435-440.
4. Somervu P. & Harma A. 2004. Bird song recognition based on syllable pair histograms. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'04)* Montreal, Canada.
5. Russell K., Do M., Huff J., & Platnick N. 2005. *Introducing SPIDA-web: wavelets, neural networks and Internet accessibility in an image-based automated identification system.* Draft yet to be published.
6. Arbuckle B., Schroeder S., Steinhage V. & Wittmann D. 2001. Biodiversity Informatics in Action: Identification and Monitoring of Bee Species using ABIS. In *Proc. 15th International Symposium for Environmental Protection*, Zurich, pp. 425-430.
7. Moyo T., Bangay S., & Foster G. The identification of mammalian species through the classification of hair patterns using image pattern recognition. In *Proc. AFRIGRAPH 2006*, Cape Town, South Africa, pp 177-181.
8. Yuan G., Russell J, Klette R., Rosenhahn B., & Stones-Havas S. Understanding Tracks of Different Species of Rats. In B. McCane (Ed.) *Proc. Int. Conf. Image and Vision Computing 2005* (IVCNZ), pp. 493-499.
9. Samul A., Brandle J. & Zhang D. 2006. Texture as the basis for individual tree identification. *Information Sciences* 176: 565-576.
10. Sogaard H. 2005. Weed Classification by Active Shape Models. *Biosystems Engineering* 91(3): 271-281.
11. ImageJ: Image Processing and Analysis in Java. http://rsb.info.nih.gov/ij/
12. Witten I. & Frank E. 2005. *Data Mining: Practical Machine Learning Tools and Techniques* (2nd Ed.). Morgan Kaufmann.
13. T.W. Ridler & S. Calvard. 1978. Picture thresholding using an iterative selection method. *IEEE Transactions on Systems, Man and Cybernetics*, 8:630–632.
14. Sternberg S. 1983. Biomedical Image Processing, *IEEE Computer* 16(1):22-33.