The identification of mammalian species through the classification of hair patterns using image pattern recognition

Submitted in partial fulfilment
of the requirements of the degree
Bachelor of Science (Honours)
of Rhodes University

Thamsanqa Moyo

November 7, 2005

Abstract

The identification of mammals using hair is important in the fields of forensics and ecology. The application of automated pattern recognition techniques to this process provides a means of reducing the subjectivity found in the process, as manual techniques rely on the interpretation of a human expert rather than quantitative measures.

The first application of image pattern recognition techniques to the classification of African mammalian species using hair patterns is presented. We design a classification sub-process for each of the two hair characteristics used to classify hair; scale and cross section patterns.

The scale pattern sub-process utilises a 2D Gabor filter-bank to extract features from a scale pattern image and these features are classified using a minimum distance measure. The cross section sub-process utilises Hu's moments to extract features and a minimum distance measure to classify these features.

Implementation of the scale pattern sub-process produces a best result of 64% accuracy when using a filter-bank of size eight and a best result of 76% when using a filter-bank of size sixteen. The implementation of the cross-section sub-process produces a best result of 48%. These results indicate that hair scale patterns are more diagnostic than hair cross section patterns in determining the identity of a hair through the use of automated pattern recognition techniques.

Acknowledgements

I would like to thank my supervisors Professor Shaun Bangay and Dr Greg Foster for their patient guidance and positive criticism. Many thanks also go to Professor Ric Bernard and Dr Daniel Parker of the Rhodes University Zoology Department for the welcoming me into their department and for their time taken to assist me with this project.

Finally, thanks go to my fellow students Matthew Marsh, Michael Horne and Fred Otten for their support and assistance throughout the year.

Contents

List of Figures

Chapter 1

Introduction

This project's objective is to investigate and apply image pattern recognition techniques that can be used to identify animal species through the classification of their hair patterns. The background to the project is provided next.

1.1 Background

The matching of mammal hair patterns as a means of identification is valuable in the fields of forensics and ecology Perrin and Campbell [1980]. Most types of hair consist of three layers of keratinized cells; that is, the cuticle making up the outer layer, the cortex forming the middle layer and the medulla resulting in the inner layer [Keogh, 1983]. Hair scale patterns are formed by the cuticle and hair cross-sectional patterns are formed by the cortex and medulla (Figure 1.1).The main type of hair used in the manual hair identification of Southern African mammals is guard hair, which is the long, coarse outer hair found in the coats of mammals . This type of hair shows the greatest variation in scale patterns used to identify hair. In addition, cross-sectional patterns are used to identify hair.

As a result of such practices , manual photographic reference systems and keys exist to aid ecological researchers identify mammals using hair patterns. Keys such as those developed by Keogh [1983] and Perrin and Campbell [1980] provide illustrations of common hair pattern classes (Figure 1.2) and associate the patterns with different species.

Interactive tools such as Hair ID [Brunner et al., 2002] take this concept a step further by requesting a user to select patterns, displayed on a computer screen that best match the hair patterns under observation.The Hair ID study utilises hair scale patterns, hair cross-section patterns, hair profile patterns, medulla type and the geographical origin of the hair sample to interactively

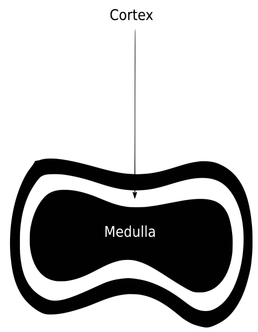

Figure 1.1: Hair cross section shape showing the medulla and cortex [Keogh, 1983]

aid a researcher in identifying an Australian mammal's hair. A classification is then provided by Hair ID on the basis of the patterns selected by the user.

This project follows the approach of Keogh [1983] and Perrin and Campbell [1980] by focusing on identification using hair scale patterns and hair cross-section patterns. This approach is favoured as it resembles that used by the Rhodes University Zoology Department, whom are to be the users of the system resulting from this study.

However, the identification of mammals through manual hair pattern matching techniques and interactive tools as with other manual pattern matching practices, relies on the subjective interpretation of a researcher as opposed to quantitative measures [Verma et al., 2002]. While independent observers may correlate their observations in order to provide some objectivity in their results, the initial observations carried out by each researcher still relies on their own interpretation. "The *science of pattern recognition*" [Verma et al., 2002] provides techniques which employ less subjective, quantitative measures based on the numeric and statistical analysis of patterns.

The first application of automated image pattern recognition techniques to the problem of classifying African mammalian species using hair patterns is presented in this project. This contribution removes the need for human interpretation from this process and replaces it with a less subjective interpretation based on the statistical and numerical analysis of hair patterns.

The following research questions are arise in the application of these techniques:

1. What image pre-processing of pattern images is required? Since automated techniques are used, methods of providing appropriate input to these techniques are explored.

2. What accurate representation of patterns can be employed in an automated system? In

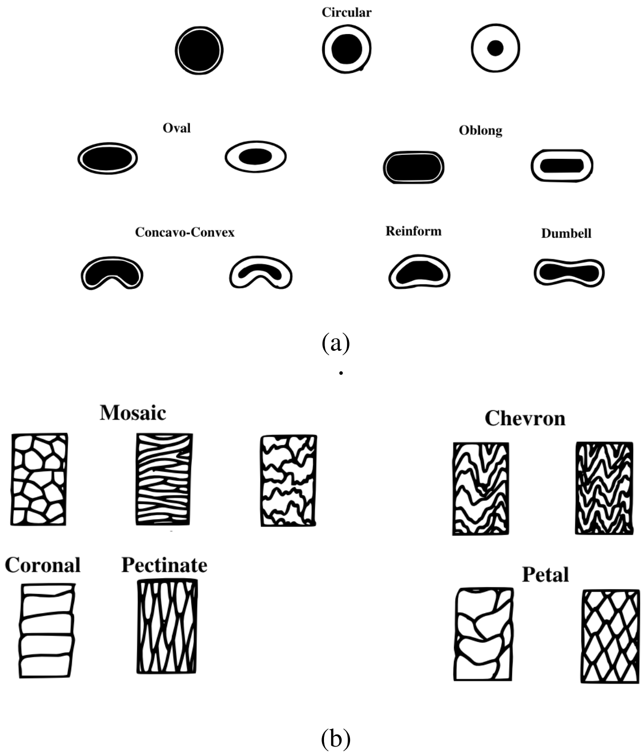

Figure 1.2: Common (a) Cross section and (b) scale patterns found in Southern African mammal hair [Keogh, 1983]

order to carry out a statistical and numerical analysis of a pattern, it is represented in quantifiable terms also known as features. The conversion of patterns contained in images into feature representations is investigated in this project.

3. How can patterns be accurately distinguished within an automated system? This question is addressed as classification depends on it being answered successfully . Therefore, techniques that may be employed to classify patterns are investigated in this project.

The work presented in this document addresses the above questions and is structured as detailed next.

1.2 Document Organisation

This document is organised as follows: Chapter 2 discusses the techniques found in the work related to this study. Chapter 3 details the design of the hair pattern recognition process developed in this study. Implementation details of this design are provided in Chapter 4 and an evaluation of this implementation are provided in Chapter 5. Finally Chapter 6 summarises the achievements and contributions of this study.

Chapter 2

Related Work

A discussion on current hair pattern recognition work found in literature is provided in this Chapter. Literature in terms of work focused on automated hair pattern recognition is limited in quantity. However , other applications using automated pattern recognition techniques that are easily adaptable to automated hair pattern recognition are available in the literature and these are also reviewed in this section.

2.1 Automated Hair Pattern Recognition

2.1.1 A Generic Automated Pattern Recognition Process

Figure 2.1 shows five design steps carried out in developing a generic classification system as proposed by Theodoridis and Koutroumbas [2003]. The sensor phase is concerned with the input and pre-processing of raw pattern images. The feature extraction step deals with the generation of single measurements and the combination these measurements into a feature vector. Since more features than necessary are generated, the feature selection stage will eliminate features that present low quality or redundant information. Classifier design involves the deployment of mechanisms that place patterns in their correct classes and system evaluation determines the performance of the system.

This approach is taken in reviewing the literature in this Chapter and designing a hair pattern recognition process.

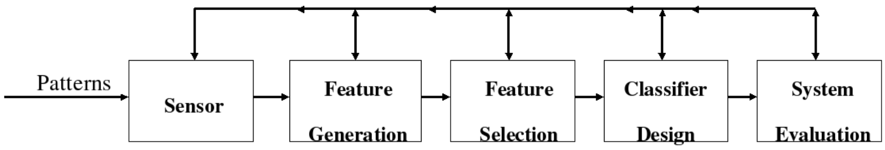

Figure 2.1: Phases carried out in designing a classification system [Theodoridis and Koutroumbas, 2003]

2.1.2 Hair MAP

The Hair Morphological Analysis (Hair-MAP) system is the closest study to this project [Verma et al., 2002]. Hair-MAP holds a database of hairs attributed to known individuals and determines whether an input human hair matches any of the hairs in the database. Therefore, the system uses hair to identify human individuals by taking advantage of the intra-species variation found in human hair. This variation is represented by measures based on the texture of the cortex, the medulla type, colour and shaft diameter of a hair.

The approach employed by the Hair-MAP study falls out of the scope of this project on two accounts: Firstly, this project aims to identify the species which an input hair belongs as opposed to the specific individual it belongs to. Secondly this project aims to utilise the inter-species variation of a hair which is best represented by measures of hair scale and cross-section patterns [Keogh, 1983] as opposed to the measures used by Hair-MAP. Therefore, techniques from studies which classify patterns similar to hair scale and hair cross-section patterns are utilised in this study. An overview of the techniques found in these applications as employed at each step of the generic classification system as mentioned in Section **??** is given next.

2.2 Scale Pattern Related Work

Iris recognition [Daugman, 2004, Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003] are applications which require the recognition of patterns similar to scale patterns. Two methods of carrying out comparisons are evident within these studies and these are described by Sanchez-Avila and Sanchez-Reillo [2005] as the biometric verification and biometric classification method .

Biometric verification determines an exact match between an input pattern and a training set. This is accomplished by rejecting all matches below a certain decision threshold value. Biometric classification assigns an input pattern to the class of the closest matching training set regardless of how dissimilar the input pattern and training set are. Therefore, the biometric verification method bases a correct match on the assumption that an exact copy of the input pattern exists in the training sets and the biometric classification method bases a correct match on the assumption that there exists a pattern in the training sets that is similar to but not necessarily a copy of the input pattern.

Biometric verification is the preferred method in studies such as iris recognition [Daugman, 2004, Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003], since the application of such biometrics involves the determining whether an input pattern can be assigned to a specific individual. However this method is inappropriate for the hair pattern recognition problem domain since it is concerned with determining the species of a hair pattern as opposed to the mammal individual it came from. In addition, given that the hair under analysis is gathered from ecological field work it is unlikely that an input hair pattern will belong to the same mammal individual which provided the hair patterns found in the training sets. Therefore, the biometric classification method is preferred for hair pattern recognition as the species of an input hair can be determined according which training set it most closely matches to.

Despite the preference for the biometric classification method, the techniques used in applications using the biometric verification method remain relevant to the problem addressed by this project, as this difference can be resolved through the elimination of a decision threshold. These techniques are reviewed in this section according to the steps provided in Figure 2.1.

2.2.1 Sensor techniques

Images containing fingerprint and iris patterns are converted to greyscale for processing in iris recognition [Daugman, 2004, Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003]. In addition, the hybrid fingerprint matching study by Ross et al. [2003] applies an image enhancement algorithm to the images resulting in clearer ridge definitions. The iris recognition study carried out by Sanchez-Avila and Sanchez-Reillo [2005] spreads the intensity values in the input image across the 255 gray levels available, to obtain a maximum representation of the variation within the image. Feature extraction techniques are applied to pre-processed images next.

2.2.2 Feature Extraction Techniques

The discovery that the receptive fields of the simple cells in a mammal's visual cortex can be modelled using Gabor functions has lead to the widespread application of Gabor filters in the field of computer vision [Lee, 1996]. Two such applications are feature extraction in iris recognition [Daugman, 2004, Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003].

A filter-bank of four filters is required for capturing global fingerprint texture information and a filter bank of eight filters is required to capture both local and global fingerprint texture information. This texture information is the frequency and orientation of a pattern and is captured by a Gabor filter as follows :

A Gabor filter in a filter-bank exaggerates the fingerprint ridges that are parallel to its orientation and evens out other ridges [Jain et al., 2000]. For example, a Gabor filter with an orientation of 0 degrees exaggerates all fingerprint ridges in an image that are parallel to the x-axis of the image and evens out the rest of the ridges. Therefore, each filter produces a resultant image of the original with the ridges parallel to its orientation exaggerated. This resultant image has a blurred appearance.

In order to compress this configuration information represented through the images created by the Gabor filter-bank, feature selection techniques are applied.

2.2.3 Feature Selection Techniques

Image tessellation is a technique used both in iris recognition [Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003] to compress the information held in a filtered image. The method used in the actual tessellation of an image varies within the literature, but the same objective of calculating a feature value from each square is realised in the iris recognition and ridge-based fingerprint matching studies.using the average absolute deviation from the mean [Jain et al., 2000].

These feature values are stored in a feature vector holding all the values calculated from a single filtered image. The collection of feature vectors representing the configuration of a pattern are referred to as the *IrisCode* [Daugman, 2004] and *FingerCode* [Jain et al., 2000] in iris recognition and fingerprint matching implementations respectively. The classification techniques discussed next utilise these feature vectors to place patterns into their appropriate classes.

2.2.4 Classification Techniques

Rapid matching is carried out using either the Hamming or Euclidean distance measures in iris recognition [Daugman, 2004, Sanchez-Avila and Sanchez-Reillo, 2005] and ridge-based fingerprint matching [Jain et al., 2000, Ross et al., 2003] respectively. An input image is placed in the class which has the least distance between the image's and the classes' feature vectors.

2.2.5 Evaluation of scale pattern related work

The false acceptance rate and the false rejection rate are evaluation methods used in biometrics where exact matching is the objective [Zhang, 2000].The false acceptance rate describes how often a positive match is reported when it is meant to be negative and the false rejection rate describes how often a negative match is reported when it is meant to be positive. The equal error rate describes the point where both false acceptance and false rejection rates are minimised and is calculated as the rate at which the false acceptance rate and the false rejection rate are equal.

Jain et al. [2000] from their ridge based fingerprint matching study, report a lowest false acceptance rate of 0.10% when the false reject rate in 19.32%, and a lowest false rejection rate of 2.83 when the false acceptance rate is 4.59%. The hybrid fingerprint matcher study carried out by Ross et al. [2003] which builds upon the work by Jain et al. [2000] reports an equal error rate of 4%.

The iris recognition study carried out by Sanchez-Avila and Sanchez-Reillo [2005] uses both the biometric classification and biometric verification method in classifying patterns. The biometric recognition method has genuine acceptance rates of 95.3% and 98.3% for feature vectors of size 256 bits and 992 bits respectively . The equal error rate is used in evaluating the biometric verification method. The lowest equal error rate of 3.3% is achieved from feature vectors with a bit length of 1860 bits and the highest equal error rate is lower than 10%.

The results provided in the studies reviewed in this section indicate that the techniques employed have potential to be successfully used in a hair scale pattern recognition process. Since these studies are concerned with evaluating the exact matching capability of their proposed solutions as opposed to the closest match, a different measure of evaluation is employed in this project. Therefore, the genuine acceptance rate is the preferred measure of evaluation as opposed to the false acceptance rate and false rejection rate.

Cross section pattern related work is reviewed next.

2.3 Cross section pattern related work

The recognition of patterns similar to hair cross-section patterns is found in studies such as tuberculosis bacteria identification [Forero et al., 2004] and rotifer identification [Yang and Chou, 2000]. These applications use biometric classification method as mentioned in Section 2.2. The techniques used in these applications are reviewed in this section.

2.3.1 Sensor

A morphological closing algorithm is employed in tuberculosis bacteria identification to complete the boundaries of a bacteria object in an input image [Forero et al., 2004]. This process is followed by the binarization of the image to enable the application of the feature extraction techniques mentioned next.

2.3.2 Feature Extraction

Moments are a mathematical concept that is applied in shape analysis to provide representations of a shape pattern based on the pattern's global configuration [Liao and Pawlak, 1996]. Geometric moments provide such an exact representation of the image they are calculated on, that the image may be re-built from these geometric moments .

Hu's seven moments fall under the category of geometric moments and are used for feature extraction in both tuberculosis bacteria identification [Forero et al., 2004] and rotifer identification [Yang and Chou, 2000]. They provide a rotation, scale, and translation invariant feature representation of shape patterns. These moments are calculated from the central moments, namely total mass, variance, skewness and the histogram sharpness [Theodoridis and Koutroumbas, 2003].

Compactness and eccentricity are additional moment based features that are used with Hu's seven moments to provide further classification [Forero et al., 2004]. Eccentricity measures the ratio between the maximum radius of an object and its minimum radius and compactness measures the degree to which an object's shape is similar to a circle [Theodoridis and Koutroumbas, 2003]. The use of these additional features is demonstrated in tuberculosis bacteria identification, where they allow for further classification when applied to the results achieved by classification from Hu's seven moments. However, before classification occurs, features containing the best discriminatory information must be selected using the techniques described next.

2.3.3 Feature Selection Techniques

The ratio of the inter-class variation to the intra-class variation of each moment is used in rotifer identification for feature selection [Yang and Chou, 2000]. Moments with inter-class variations greater than their intra-class variations are selected for use in classification. However the this analysis is dependant on prior knowledge of the types of shapes that are to be recognised. For example Yang and Chou [2000] based their analysis on 3 known types of rotifer shape classes. The techniques described next are applied to moments to obtain a classification of the patterns found in input images.

2.3.4 Classification Techniques

The tuberculosis bacteria identification study employs a classification tree that utilises information from the selected features [Forero et al., 2004] . The minimum Mahalanobis distance measure between the Hu's moment features of the input object and of each bacteria cluster, is recorded and compared in the top part of the tree to a threshold value . The lower parts of the tree compare the compactness and eccentricity features of the input object against known compactness and eccentricity values. The threshold value is calculated from running test classifications and the known compactness and eccentricity values are calculated from the observed shape characteristics of bacteria found in training samples. The leaf nodes of the classification tree, either correspond to a tuberculosis bacteria object or a non tuberculosis bacteria object to be rejected.

The first stage of rotifer classification involves eliminating the debris surrounding a rotifer [Yang and Chou, 2000]. This two class problem requiring objects to classed as either rotifers or debris can be carried out using either a similarity distance algorithm or probability histograms. The second stage of rotifer classification uses a back propagation neural network consisting of 7 input processing elements corresponding to the 6 selected moment features and the area of an object. In addition, a single hidden layer of 5 processing elements and an output layer of 3 processing elements corresponding to the known three types of rotifer, were deployed in the network. A total of 185 rotifer samples were used to train this neural network.

2.3.5 Evaluation of Cross-section pattern related work

Sensitivity is the rate of providing a positive match correctly and specificity is the rate of providing a negative match correctly [Nichol and Mendelman, 2004]. This measure is used in evaluating the tuberculosis bacteria identification study [Forero et al., 2004]. The results reported in this study show the best sensitivity and specificity values for bacteria samples are 91.43% +/-

2.75 and 100% respectively . However, these results must considered in the context that the classification is done using eccentricity and compactness values that have been observed from the training sets . This consideration raises uncertainty about the extensibility of the study's techniques outside the scope of the limited bacteria sample used and their usefulness in a hair pattern recognition process.

The best overall accuracy rate obtained in the elimination of debris at first classification stage of rotifer identification is 97.34% [Yang and Chou, 2000] . The reported performance of the artificial neural network employed at the second classification stage of rotifer identification is a genuine acceptance rate of 93.15% . These results indicate the high feasibility of utilising the neural networks in identifying shape feature based objects, such as hair cross-sectional patterns, from a noisy environment.

2.4 Summary

Table 2.1 shows a mixture of the techniques mentioned in this chapter, at the various stages of the generic pattern recognition process advocated by Theodoridis and Koutroumbas [2003]. The most important conclusion that can be drawn from this table is that only the means of evaluating each process is shared by both process. The stages where the actual processing takes place, that is from sensor to classifier design, use different techniques. This lack of overlap, provides a basis to conclude from the literature reviewed, that the different nature of scale and cross section patterns warrants separate processing for each.

The application of these techniques in the literature reviewed is specific to the problem domain of the studies mentioned in the review. Therefore, this project adapts and rejects some of these techniques during the design of a process suited for the hair pattern recognition problem domain. This design is given in the following chapter.

Table 2.1: Techniques found in literature and relevant to the hair pattern recognition study.

Stages	Hair Scale Pattern Processing	Hair Cross-Section Pattern Processing
Sensor	Image orientation standardisation, histogram stretching and image size standardisation.	Image binarization and a morphological closing algorithm.
Feature Extraction	Gabor filter-bank.	Hu's seven moments
Feature Selection	Filtered images are tesselated into squares and each feature is calculated from pixel values residing in a square.	Ratio of inter-class to intra-class variation of the features obtained.
Classifier Design	Euclidean distance measure.	Mahalanobis distance measure or artificial neural networks
System Evaluation	Genuine acceptance rate	Genuine acceptance rate

Chapter 3

Design

This chapter provides an implementation independent design of a hair pattern recognition process. This process is designed to allow flexibility in its implementation, therefore providing an opportunity to experiment with the various techniques. The design details of the overall design of the hair pattern recognition process and its two sub-processes in it are examined next.

3.1 Overall design

This section covers the issues affecting the entire design of the hair pattern recognition process. The collection of hair samples is discussed and this is followed by a discussion on premises established from this work. These premises influence the design goals provided in this section. Finally, the user interaction requirements of the user are provided in this section.

3.1.1 Hair sample collection

Samples of unknown hair are collected from field work such as scat analysis [Perrin and Campbell, 1980]. Since each sample may consist of hair from various species, hair in the sample is grouped using the immediately visible attributes of the hair such as hair colour. This grouping does not guarantee that hair from different species is separated, since some species have similar visible attributes. Once this grouping is done, the hair is passed to the sensor stage. The premises drawn from the collection of samples is provided next.

3.1.2 Premises drawn from hair sample collection

Four premises are established from the work carried out in collecting hair samples. These premises influence the overall design of the hair pattern recognition process and are mentioned here:

1. The first premise is drawn from the observation that a scale and cross-section pattern is in practice not obtained from the same hair. Rather, the scale patterns are obtained from one set of hairs in sample and cross section patterns are obtained from another set in the same sample.

2. The second premise drawn is that in this problem domain it cannot be safely assumed at the time of gathering a sample, that all the hair present in a sample are from the same individual. Since samples are gathered from scat or the surrounding flora, a determination on whether the hair in a sample is gathered from a single individual cannot be made. If this were possible, there would be no need for hair pattern recognition.

3. The third premise drawn is that it cannot be assumed that the all the hair in an input sample is from the same species. While it is possible to group hair according to its immediately visible attributes at the time of gathering the sample, various species may have hair with similar visible attributes. Therefore a sample may contain hair from multiple species with similar visible attributes.

4. The final premise drawn from observations during this manual preparation work is that cross section shapes offer relatively less variation amongst species as compared to scale patterns. While some species have unique cross sections which allows for discrimination using cross-sections, other species share cross-section patterns with little variation amongst species (Figure3.1) .

These conclusions impact the goals set in the design of the hair pattern recognition system described next.

3.1.3 Design Goals

A number of goals are set to guide the design of the hair pattern recognition system. Some of these goals are based on the work conducted in the previous chapter and are provided below:

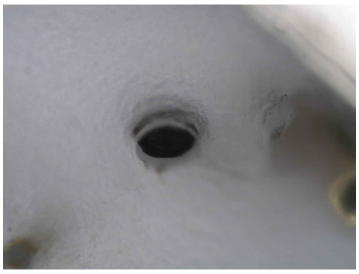

(a)

(b)

Figure 3.1: Similar cross sections from different species (a)Jackal Cross Section and (b) Zebra Cross Section

1. It is observed from the work conducted in sample collection that hair scale patterns vary in terms of texture and hair cross-section patterns vary in terms of shape. Given this difference in the two patterns, the first goal stated is the design of two distinct sub-processes. This split allows for a sub-process to be designed in a manner that is specific to the type of hair it handles.

2. Since the scale and cross pattern sub-processes are distinct, they produce their own distinct classification result. However, the second goal set is to avoid attempts at combining these two results to obtain a final classification. This goal is based on the first and third premises established in during hair sample collection which state that hair in a sample cannot be assumed to be from the same species. For example, it is possible that a hair from which a scale pattern is extracted is of a different species to one from which which a cross section pattern is extracted. In such a case, a combination of the correct result from each leads to a contradiction that needs to be resolved and this does not reflect the real situation since the results are mutually exclusive.

3. The hair pattern recognition process is designed to discriminate between species, but not between individuals. This goal is set as a result of the second premise established during hair sample collection that states it cannot be assumed that all the hair present in the sample is from the same individual. A classification process that takes intra-species variation into consideration produces mis-classifications as the hair used to build a training set is of a different individual to that of the input hair. Therefore, the hair pattern recognition process needs to be sensitive to the inter-species variation present in hair patterns but ignore any

intra-species variation.

4. The fourth goal of the hair pattern recognition process is to provide an extensible design. The design's extensibility is determined by whether the process is unconstrained in the species it can handle. The design provided in this chapter keeps the process it details flexible enough to handle any scale and cross section patterns as opposed to a limited set of patterns.

These goals influence the rest of the design discussed in this chapter. An important aspect of the overall design of the process are user interaction requirements.

3.1.4 User interaction requirements

The hair pattern recognition process designed in this study automates the manual pattern recognition process and this implies minimal user interaction. However, in order to meet the needs of a user, the researcher's requirements are taken into consideration. These output user requirements are elicited from discussions with researchers in the Rhodes University Zoology Department

A researcher making use of the process designed, needs sufficiently accurate information which can be used to corroborate findings from a manual hair identification process. This means that the results from each process should be given in a format that expresses the confidence that the hair recognition process has in determine a correct match. Such an expression of confidence is represented through the ranking of all possible matches.

The next section looks at the design of the first sub-process of the hair pattern recognition or scale pattern sub-process.

3.2 Scale Pattern processing

Since scale patterns are texture characteristic, the design of the scale pattern sub-process is based on techniques relevant to recognition of texture based patterns.This design begins with the sensor stage.

3.2.1 Sensor

The first stage, sensor, deals with the gathering and image pre-processing of hair images. Four cases taken into consideration are the rotation, scale, translation and reflection variations of captured images.

Scale pattern images are not rotation invariant hence each image is micrographed at a standard orientation that is employed consistently for all image capture. This standardisation assists in producing rotation invariant features in the feature extraction stage as only two orientations need are considered during classification, that is, an images' original orientation and its 180 degree rotation. At this stage the value of the original orientation chosen as the standard is irrelevant as long as it is consistently employed for all images. The 180 degree rotation conveniently produces the same image that would have been captured if the hair was micrographed from the opposite side.

The scale variance experienced in images is catered for through the resizing of images to standardised size. This allows the extraction features from the same number of pixels for all images.

The translation variation of the images is unimportant, as scale patterns are observed during the manual preparation work not to vary greatly along the length of a hair. Similarly the reflection variation of an input image is unimportant as it is unnatural for such a variation to take place. Such a variation would only exist if a mirror image of a hair is taken and since mirrored images are not captured, a reflection variation will not occur.

Finally after these cases are handled each image's histogram is stretched to represent the texture information contained in it using the entire intensity range. These images are passed to the feature extraction phase designed next.

3.2.2 Feature Extraction

Features are extracted from pre-processed images in this stage and represented in a numerical format. A 2D Gabor filter is a specialised Gabor filter that captures both local pattern frequency and orientation information [Ross et al., 2003]. The Gabor filter used in this study is an even symmetric 2D Gabor filter given by :

$$\begin{aligned} G_{\theta,f}(x,y) &= e^{\left\{\frac{-1}{2}\left[\frac{x'^2}{\delta_x^2}+\frac{y'^2}{\delta_y^2}\right]\right\}\cos\left(2\pi f x'\right)}, \\ x' &= x\sin\theta + y\cos\theta, \\ y' &= x\cos\theta - y\sin\theta, \end{aligned} \tag{3.1}$$

Ross et al. [2003].

The variables of the Gabor filter used by Ross et al. [2003] are specific to the extraction of features from fingerprints and these variables are adapted for use in this study as follows:

1. f is the frequency of each filter in the filter-bank and this is derived from the average distance between the ridges that form the scale patterns.

2. The values of δ_x and δ_y are set to the average thickness of the ridges that form the scale patterns .

3. The θ variable corresponds to the orientation of each filter and is the only variable that differentiates each filter in the filter-bank. The number of filters used in the filter-bank determines the values of θ.

A filter-bank of 2D Gabor filters is used to extract features from a pre-processed scale pattern image. A filter-bank of Gabor filters consists of a collection of Gabor filters which are only differentiated from each other by the value of their orientation variable. When an image filtered with the filter-bank, each filter will produce a resulting filtered image. For example filtering an image with a filter-bank of size 4 will result in 4 filtered images which correspond to the result of filtering the original image with each filter in the bank.

Given the standardised orientation of an input image done at the sensor stage , the image may only be at one of two orientations, that is, at its original orientation or its 180 degree rotated equivalent. In order to provide rotation invariant features, each filter is applied to the input image and its 180 degree rotated copy. Therefore filtering an input image with a bank of 2D Gabor filters in this study results in 2 x (the size of the filter-bank) filtered images being produced.

Filtered images present redundant information and in order to provide refined discriminatory information to the classifier stage, a feature selection stage is designed.

3.2.3 Feature Selection

Each image obtained from the Gabor filter-bank mentioned in the previous stage is passed to this stage. Feature selection is done by tessellating a filtered image into smaller squares as shown in Figure 4.9. A value for each square is calculated from the pixels in the square and considered a feature. The dimensions of an input pattern's feature vector depends on the size of the Gabor filter-bank used in feature extraction and the number of features obtained from tessellation. For example, a filter-bank of 8 filters whose resulting images are tessellated using 256 squares, results in an 8 dimensional feature vector of 8 vectors each of size 256.

This process of tessellation compresses the information contained in the filtered images and provides a better representation of the local and global variations in the scale patterns. This technique has been demonstrated to be effective in texture based fingerprint matching as carried out by Jain et al. [2000] and Ross et al. [2003]. Each feature vector is passed through to the classifier designed in the next stage.

3.2.4 Classifier Design

The classifier employed at this stage places patterns represented by the selected features into the class of the best matching training set. This matching is done by determining which training set is most similar to the feature vectors generated from an input image. The design of feature vectors and training sets is discussed here.

A rotation invariant feature vector of a scale pattern is produced at this stage through the combination of features extracted from the pattern's image at its original orientation and the corresponding features extracted from the image's 180 degree rotation. This combination is done by simple vector addition and results in a feature vector that takes into consideration both orientations at which an image may be input. Therefore, the vector calculated from a tessellated filtered image is added to the vector calculated from its 180 degree rotated equivalent which is filtered with the same filter in the filter-bank.

An example of combining the features obtained through filtering with a Gabor filter bank of size four and performing feature selection with four tessellation squares is shown in Figure 3.2. Feature Vector 1 represents the four dimensional feature vector of an input image at its original orientation and Feature Vector 2 represents the four dimensional feature vector of the input image rotated by 180 degrees. The resulting four dimensional feature vector shown is the feature vector which takes into consideration both orientations and is used to provide a rotation invariant represent of a scale pattern.

Finally, training sets are constructed by taking an average of the rotation invariant feature vectors from several individuals within a species. Constructing training sets in this manner evens out the intra-species scale pattern variation that exists within a species. However, this does not even out any inter-species variation that exists within the patterns as such variation is shared amongst the individuals within a species and hence maintained.

The design of the hair cross section process is described next.

3.3 Hair Cross-section Processing

Since cross section patterns are shape characteristic, the design of the cross section pattern sub-process is based on techniques relevant to the recognition of texture based patterns. This design as with the scale pattern sub-process design begins with the sensor stage.

Feature Vector 1

1	A	B	C	D
2	E	F	G	H
3	I	J	K	L
4	M	N	O	P

Feature Vector 2

1	A1	B1	C1	D1
2	E1	F1	G1	H1
3	I1	J1	K1	L1
4	M1	N1	O1	P1

Resulting Feature Vector

1	A+A1	B+B1	C+C1	D+D1
2	E+E1	F+F1	G+G1	H+H1
3	I+I1	J+J1	K+K1	L+L1
4	M+M1	N+N1	O+O1	P+P1

Figure 3.2: An example of feature vector combination to used to provide a rotation invariant feature. Feature Vector 1, Feature Vector 2 and Resulting Feature Vector are 4 dimensional vectors.

3.3.1 Sensor

The sensor stage is the point at which the shape of a cross section is extracted from an image. Therefore two issues which affect the performance of this stage are taken into the design considerations:

1. The fourth conclusion established in Section **??** states that some cross section patterns are similar amongst species and show minimal inter-species variation. Therefore, techniques employed at this stage need to be sensitive to small variations in cross section patterns in order to provide the feature extraction stage with a high quality representation of the patterns.

2. The second consideration taken is the noise present in cross section images. Noise in cross section images produces artifacts in the shapes extracted from the images and this results in a false representation of shape of a cross section pattern.Therefore, in order for successful shape extraction, techniques employed at this stage are robust against noise present in images.

The shape of a cross section pattern in an image is extracted from the background using image segmentation techniques. These techniques extract the complete outline of the shape from the background and place this outline in a new image. This new image is converted into a black and white where the inside of the shape outline is black and the rest of the image is white. This new image contains more of the shape filled in black than the white background, in order for processing to be accurately carried out in the next stage, feature extraction.

3.3.2 Feature Extraction

Hu's seven moments are applied at the feature extraction stage of the cross section sub process to provide a rotation, scale and translation invariant representation of cross section patterns . These moments are based on the central moments which are translation invariant [Theodoridis and Koutroumbas, 2003]. Central moments are calculated directly from the shape pattern represented in an image and are normalised in order to provide scale invariance in addition to the translation invariance property.

Hu's seven moments are calculated, as described by equation 3.1 in Section 2.3, using the normalised central moments. The calculation of these moments adds the rotation invariant property to the translation and scale invariant representation provided by normalised central moments.

Hu's seven moments have large value and in order to represent them with a minimal loss of accuracy they are represented as logarithms of their absolute values as suggested by Theodoridis and Koutroumbas [2003]. A discussion on the selection of appropriate moment features is given next.

3.3.3 Feature Selection

The literature reviewed in Section2.3suggests that the ratio of inter-class to intra-class variation provided by a moment feature may be used to select moment features. This comparison of moment features implies that a process employing this technique handles similar patterns which share some similar moment values. In such a situation, the removal of similar moments has no adverse effect since the information they present is of little use in discriminating between the patterns. However, while the cross section sub-process is designed to handle similar shapes, it also processes distinct cross pattern shapes. Therefore, this technique is rejected as moment features which are considered redundant during the matching of similar cross section patterns may become relevant when matching distinct cross section patterns.

Therefore, it is motivated that the selection of Hu's moments is undesirable for this problem domain and the prudent approach of employing classifiers that are robust to redundant information is taken. The design of these classifiers is discussed next.

3.3.4 Classifier design

Classification of the cross section patterns using moments is done by a distance measure. While the use of neural networks is described as a more favourable classifier in the Section 2.3.5, distance measures are preferred for the following reasons:

Firstly the number of known pattern classes needs to be known as this correlates to the number of output nodes of a neural network. Catering for a number of classes implies that is is known how many species the process will cater for and this is not the case as the design aims to be extensible to species not considered in the experiments carried out in this study.

Secondly the lack of feature selection at the previous stage means redundant data will be undetected and fed to the neural network. Such a scenario could lead to large performance penalties as irrelevant information is passed through the neural network.

Therefore, the use of distance measures is preferred as they are more robust to redundant data present in a small feature vector with a size of seven. In order to eliminate any intra-species variation that may be represented by the moments, the moment features from the images chosen

Table 3.1: Design used in the hair pattern recognition study based on the generic pattern recognition system stages.

Stages	Hair Scale Pattern Sub-Process	Hair Cross-Section Pattern Sub-Process
Sensor	Image orientation standardisation, histogram stretching and image size standardisation.	Image Segmentation and Thresholding
Feature Extraction	2D Gabor filter-bank.	Hu's seven moments
Feature Selection	Filtered images are tessellated into squares and each feature is calculated from pixel values residing in a square.	None required
Classifier Design	Minimum distance measure.	Minimum distance measure.

for building the training sets should be averaged together. A summary of this chapter is given next.

3.4 Summary

Table 3.1 shows the the final design of the hair pattern recognition sub-processes and the techniques applied at each stage. A comparison to Table 2.1 provided at the end of Chapter 2, indicates that two changes have been made during the design chapter, The first change is the elimination of the need for feature selection at the cross section sub-process, since feature selection is irrelevant to this problem domain. The second change is the generalisation of the distance measures at the feature selection stage. This generalisation is done in order to provide an implementation independent design.

In order for the design of the hair pattern recognition sub-processes to be tested, an implementation from which an evaluation of results is conducted is provided in the following chapter.

Chapter 4

Implementation

An implementation of the process designed in the previous chapter is detailed in this chapter. This implementation strictly adheres to the design, as the performance of the hair pattern recognition process designed is evaluated through it. The resources used to implement the entire process are described first and followed by a description the two hair patten recognition sub-processes.

4.1 Implementation resources

This section details the resources used in implementing the hair pattern recognition process in this project. The two types of resource discussed here are the development platform used and the hair sample selected.

4.1.1 Development platform

A test system is implemented in Java as a plug-in for a public domain graphics application, ImageJ [ImageJ, 2005]. This approach is chosen since ImageJ provides image processing functionality that allows development time and effort to be focused on the hair pattern recognition process. For example, the class abstractions of an image provided in ImageJ eliminate the need to develop functionality that accesses pixel data in an image.

A class diagram containing the plug-ins and classes developed for the system is shown in Figure 4.1. This diagram is not placed in the design chapter since it is not necessary to implement the hair pattern recognition process using an object oriented approach. However, the object oriented approach is used in this implementation since Java is the language required to develop plug-ins for ImageJ. The classes shown in Figure4.1 are explained below:

ImagePlus: This is an ImageJ class that provides an abstraction of an image.

CrossImageClassify: This class combines the functionality provided in the CrossFeatExtractor class and CrossClassifier class to classify a cross section input image.

CrossImageTrain: This class uses the functionality provided in the CrossClassifier class to add the feature vector of a known cross section input image to a training set.

CrossFeatExtractor: This class calculates Hu's moments and adds these moments to a training set.

CrossClassifier: This class calculates the Hamming and Euclidean distance between a training set and a cross section pattern image's feature vector.

ScaleImageClassify: This class combines the functionality provided in the ScaleFeatExtractor class and the ScaleClassifier class to classify a scale pattern input image.

ScaleImageTrain: This class uses the functionality provided in the ScaleFeatExtractor class to add the feature vector of a known scale pattern input image to a training set.

ScaleFeatExtractor: This class uses a set of Gabor filters, referred to in the Design chapter as a filter-bank, to extract features from a scale pattern input image.

ScaleClassifier: This class calculates the Euclidean distance between a training set and input image's feature vector.

GaborFilter: This class provides the implementation of an asymmetric 2D Gabor filter.

The nature of the hair sample used in this implementation is discussed next.

4.1.2 Hair sample used for implementation

The implementation provided in this project classifies input hair patterns into one of five classes. These classes are defined by the five initial species chosen for the study, namely: blue wildebeest, impala, jackal, springbok and zebra. These species are chosen for the following reasons:

1. Jackal and blue wildebeest have a scale pattern distinct from each other and the rest of the species in the sample. The inclusion of these two species in the sample allows the testing of the scale pattern sub-process's ability to distinguish between distinct scale patterns.

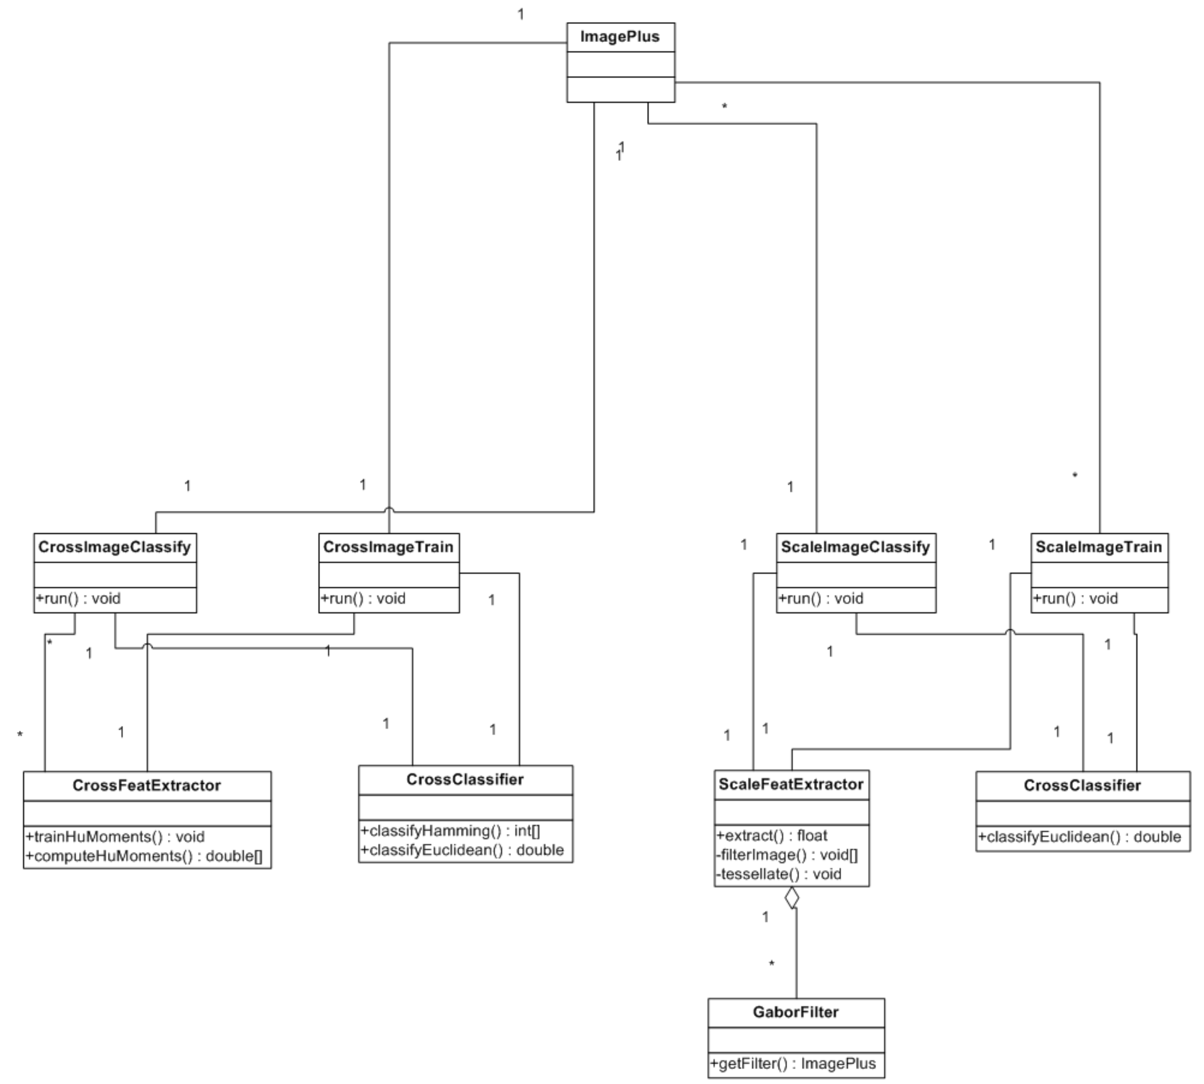

Figure 4.1: High level class diagram of the test model developed in the project

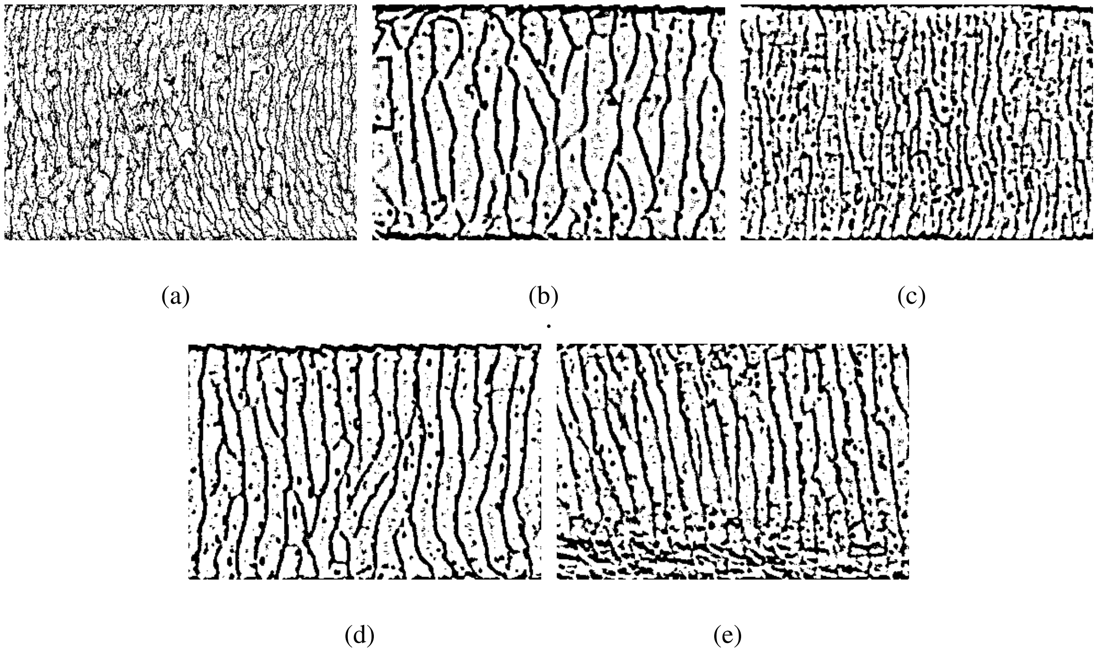

Figure 4.2: Variation in scale pattern. Images have been manipulated to emphasize the pattern. (a) Blue Wildebeest. (b) Impala. (c) Jackal. (d) Springbok. (e) Zebra

2. This sample also allows the scale pattern sub-process's ability to discriminate between species with similar hair patterns to be ascertained. Impala, springbok and zebra are the species included this sample that have similar scale patterns. Figure 4.2 illustrates the scale pattern variations of selected species.

3. The hair pattern recognition process's suitability for distinguishing between similar cross-section patterns is tested through the inclusion of zebra and jackal which have similar cross section patterns. Blue wildebeest, impala and zebra have cross section patterns distinct from each other and from zebra and jackal.

4. The inclusion of these species allows the ability of the cross section sub-process to discriminate between species with different cross section patterns to be determined. Figure 4.2 illustrates the cross section variations of selected species.

The hairs of these species are obtained from a collection of mammal hair in the Rhodes University Zoology Department. The hair pattern images from these images are captured in the

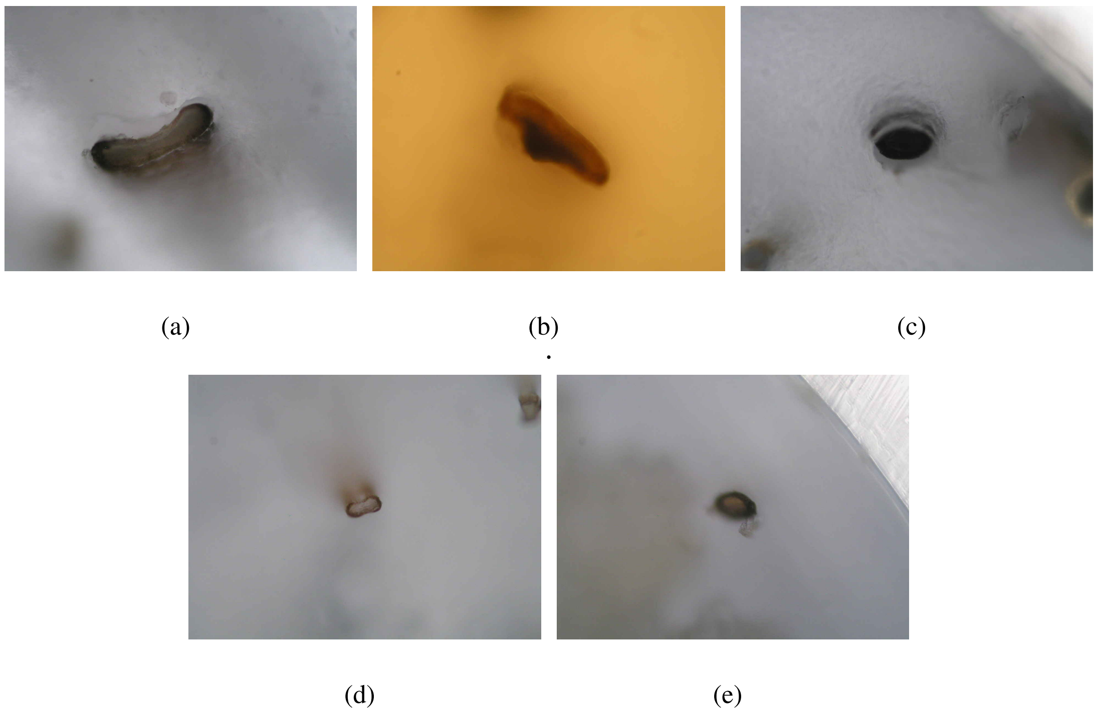

Figure 4.3: Variation in cross section pattern. (a) Blue Wildebeest. (b) Impala. (c) Jackal. (d) Springbok. (e) Zebra

using the equipment mentioned next.

4.1.3 Image Capture Equipment

Two types of image capture are explored in this project and these are through the use of a light microscope and a scanning electron microscope (SEM).

The SEM used in this project is a JEOL JSM 840. The light microscope is an Olympus BX series photomicroscope mounted with an Olympus Camedia 4040 Zoom digital camera. Hair patterns are found and focused on with the light microscope. Once in focus, the image is captured with the digital camera.

The difference between these two pieces of equipment, that is relevant to this study, is found in their zoom facilities. The scanning electron microscope has a digital zoom facility whereas the light microscope achieves has an analogue zoom facility. This difference is particularly relevant in the image capture work carried out in the implementation of the hair scale sub-process mentioned next.

4.2 Hair scale processing

The hair scale sub process designed in Chapter 3 begins with the sensor stage which deals with the pre-processing of raw input images. Before mention is made of this stage, the methods explored in capturing the raw input images are detailed.

4.2.1 Image Capture

Hair is prepared for scale pattern image capture using either the light microscope or SEM. This preparation and the decision on which image capture equipment is preferred for scale pattern image capture is provided in this section.

Preparation of a hair for scale pattern image capture with a light microscope begins with spreading a thin film of gelatin across a glass slide (Figure 4.5) . Single straightened hairs are laid in the gelatin and allowed to settle. Once the gelatin has settled, each hair is peeled off the slide, leaving a microscopic imprint of the scale pattern on the slide. Work carried out in this project suggests that the peeling off process should be carried out in a delicate manner, as it can easily result in artifacts in the pattern imprints left behind. The best method of peeling off a hair involves leaving an piece of each hair to hang off the edge of the slide. Whilst this technique at

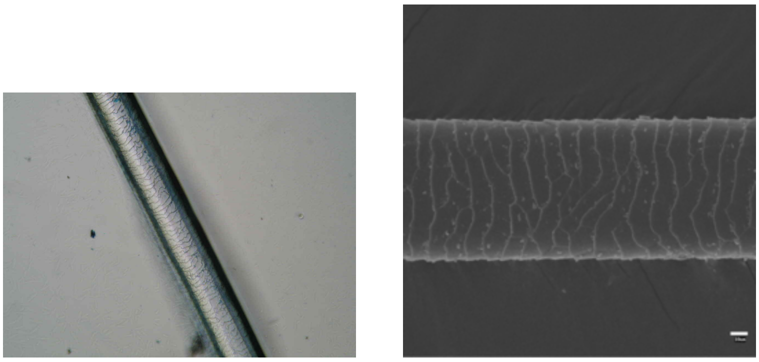

(a) (b)

Figure 4.4: Results of image capture using (a) light microscope (b) Scanning election microscope

times produces the negative effect of a hair breakage in the middle of the peeling process, the imprinted pattern is left less unaltered than if an instrument is used to lift the hair off the slide.

Preparing hair for scale pattern image capture with the SEM begins with placing straightened hair flat onto a carbon sheet of paper. This carbon sheet of paper is mounted onto a metal block and the entire mount is gold plated as shown in Figure 4.6. Since the diameters of the gold blocks are less than that of the length of the hair, a section from each hair is cut from the middle of the length of the hair and placed on the carbon paper. The gold plated hair is left in this state for analysis under an electron microscope.

The use of the SEM for scale pattern image capture is preferred to the use of a light microscope. This preference is built on the output of both image capture techniques shown in Figure 4.4. The SEM produces images with a better representation of the scale pattern through its digital focus facility than as opposed to the analogue facility of the light microscope. As shown in Figure 4.4, the scale pattern is micrographed at a higher zoom level than that possible with the light microscope. Therefore,the images produced from the SEM are passed to the sensor stage described next.

4.2.2 Sensor

An initial sample of 100 scale pattern images, comprising 20 images of each selected species, is gathered using the SEM. The raw images are taken at magnifications varying from 600Xto 400X and are output at a size of 1024x1024 pixels. Images with the least noise in the sample are

Figure 4.5: Glass slide with gelatin film spread on top

Figure 4.6: Gold plated blocks and hair on carbon paper

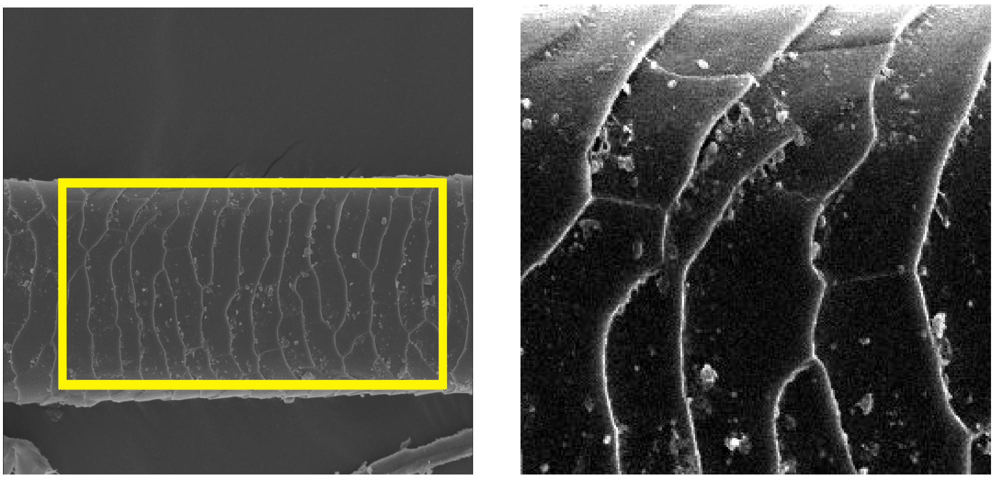

(a) (b)

Figure 4.7: (a) Raw input image with ROI emphasised and (b) pre-processed image converted to grey-scale, re-sized and histogram stretched.

selected for training and testing in this study. This results in a final selection of 40 images (15 images for training and 25 images for testing) from the initial sample which are pre-processed as follows:

In order to handle rotation variations, scale patterns are micrographed at a standard orientation. The orientation chosen in this project is shown in Figure 4.7(a), where hair is micrographed horizontally with its length parallel to the x-axis of the screen.

The second type of variations handled are the scaling variations that may lead to one pattern having more pixels analysed than another. A square region of interest (ROI) is manually selected by the user to segment the scale pattern from the background. Various sizes of ROI may be selected by a user and in order to standardise the size of the image analysed, the image resulting from the ROI is standardised to a size of 256x256 pixels. The new 256x256 image is converted to greyscale and its histogram is stretched. Figure 4.7 shows the process of converting a raw image to a pre-processed version. The pre-processed image is then passed to the feature extraction stage described next.

4.2.3 Feature Extraction

A copy of an image passed from the sensor stage is produced and rotated by 180 degrees. The original image and its 180 degree rotated copy are filtered with the Gabor filter-bank, resulting in images similar to those shown in Figure 4.8. This image shows the results of an image and its 180 degree copy that has been filtered by a filter-bank of four filters. Filtering of an image using a Gabor filter-bank is performed in this project through convolving the input image with each of the Gabor filters in the filter-bank. This convolution is done with a kernel set to a size of 15x15 pixels. The Gabor filter variables as explained in Section 3.2.2 are set to the following values:

1. The average distance between the ridges that form the scale patterns is calculated from 25 pre-processed scale pattern images (five images from each species) and amounts to 40 pixels. Therefore the frequency f is set to $\frac{1}{40}$.

2. The average thickness of the ridges that form the scale patterns is calculated from the same images used in calculating f . The average ridge thickness amounts to 4 pixels and therefore the values of $\delta_x$and δ_y are set to 4.

3. The values of θ are set to values ranging from 0 to 180 degrees. The difference between the values assigned to each of the filters is $180/n$, where n is the number of filters in the filter-bank. For example, the values of 0,45,90,135 degrees are set when four filters are used and the additional values of $22.5, 67.5, 112.5, 157.5$ degrees are included when eight filters are used.

Since eight filters perform better than four filters in fingerprint matching and this project investigates whether this applies to the hair scale sub-process, by testing both sizes of filter-bank. The project also implements a filter-bank of 16 filters to determine whether increasing the size of a Gabor filter-bank increases its performance.

The images that result from the filtering are passed to the feature selection stage .

4.2.4 Feature Selection

Each filtered image from the previous stage is tessellated into squares of size 16x16 pixels. Since fingerprint matching studies use either the average absolute deviation from the mean (AAD) or variance to calculate the value for each square, this project experiments with both. The size of the feature vector obtained from tessellation is 256, as there are 256 squares from which feature values are obtained. The feature vector of each image is passed to the classification stage.

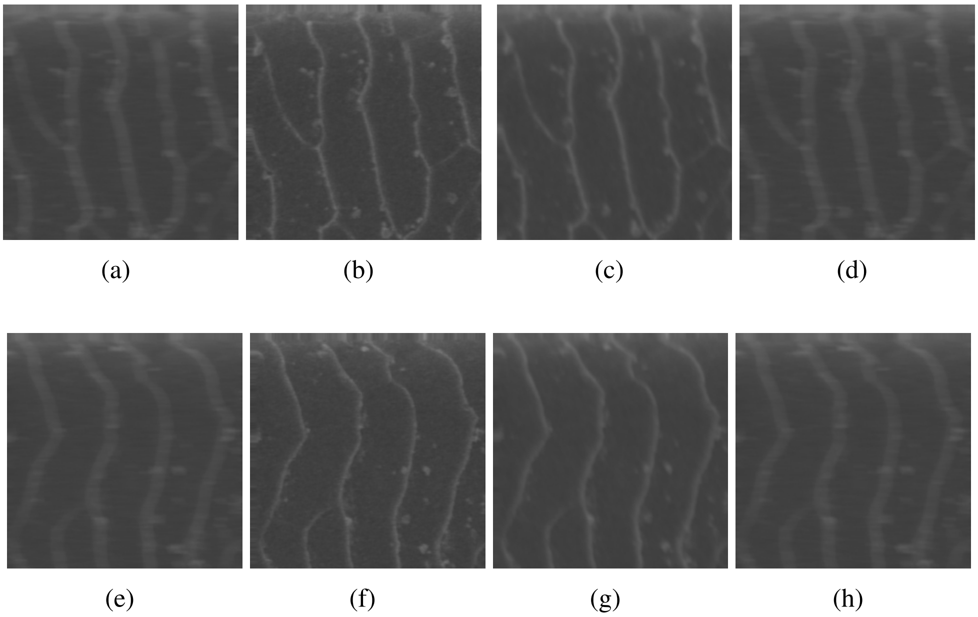

Figure 4.8: Filtered images of springbok scale pattern image. (a)-(d) correspond to the original image filtered with filters at orientations set at 0,45,90 and 135 degrees. (e)-(h) corresponds to the original image rotated by 180 degrees filtered with filters at orientations set at 0,45,90 and 135 degrees.

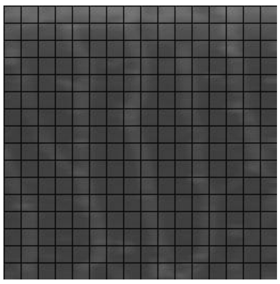

Figure 4.9: Tessellated image

4.2.5 Classification

The feature vectors obtained from each tessellated image are summed together with their 180 degree rotated equivalent, as prescribed in the design. For example, the feature vectors obtained from the images in Figure 4.8(a) and 4.8(e) are summed together as they were filtered by the same filter. This ensures that rotation invariant feature vectors are obtained as an image may be captured at one of two orientations.

The values assigned to a training set are obtained by taking the average of the feature vectors of the three images set aside for building the training set. This averaging smooths out the intra-species variation that exists in a single scale pattern. The training sets are serialised as a comma separated values in a file of their own.

An input hair pattern is classified by summing the Euclidean distances between its feature vectors and the training set vectors of each known hair class. For example, if an 8 dimensional feature vector is input, then the 8 Euclidean distances between the input 8 dimensional feature vector and the 8 dimensional training set vector are summed together. The training set with the least summed distance from the input feature vector is considered to be the best matching class of the hair under analysis. In order to show the confidence that the sub-process has in determining the best match, the results are ranked from the training set with the least summed distance to the one with the greatest summed distance from the input feature vector.

The hair cross section sub-process is discussed next

4.3 Hair Cross-section Processing

This section deals with the implementation of the hair cross section sub-process designed in Chapter 3. The methods used in preparing a hair for image capture are provided first and then followed by the rest of the stages designed according the stages of a generic pattern recognition system[Theodoridis and Koutroumbas, 2003].

4.3.1 Image Capture

Hair is prepared for cross section pattern image capture using either the light microscope or SEM. This preparation differs from that required for scale pattern image capture and is mentioned in this section. In addition, a decision on which image capture equipment is preferred for cross section image capture is made.

Preparing hair for cross section pattern image capture under a light microscope involves

(a) (b)

Figure 4.10: (a)Plastic pipette filled with wax and (b) sliced sections of the pipette

Figure 4.11: Metal clamp used to hold hair for gold plating

standing hairs in an upright position inside a plastic pipette (Figure 4.10a). The pipette containing the hair is filled with liquefied wax. Once the wax solidifies, thin sections of the pipette are cut resulting in thin discs of wax holding sliced sections of hair (Figure 4.10b). Therefore, the wax maintains the sliced hairs in an upright position from which cross sections may be viewed.

The preparation of hair for cross section analysis under the SEM requires gold plating of the hair. However, the challenge posed by this approach is the method used to place the hair in an upright position. This cannot be accomplished through the use of wax, therefore a metal clamp (Figure 4.11) is employed to hold the hair in an upright position. Once fastened in the clamp, the hair is gold plated.

Since the cross section sub-process is concerned with the shape of a cross section instead of its texture, the shape representations produced by the SEM and light microscope are considered to determine which capture equipment performs best.

Both the light microscope and SEM provide images with clear definitions of a cross section

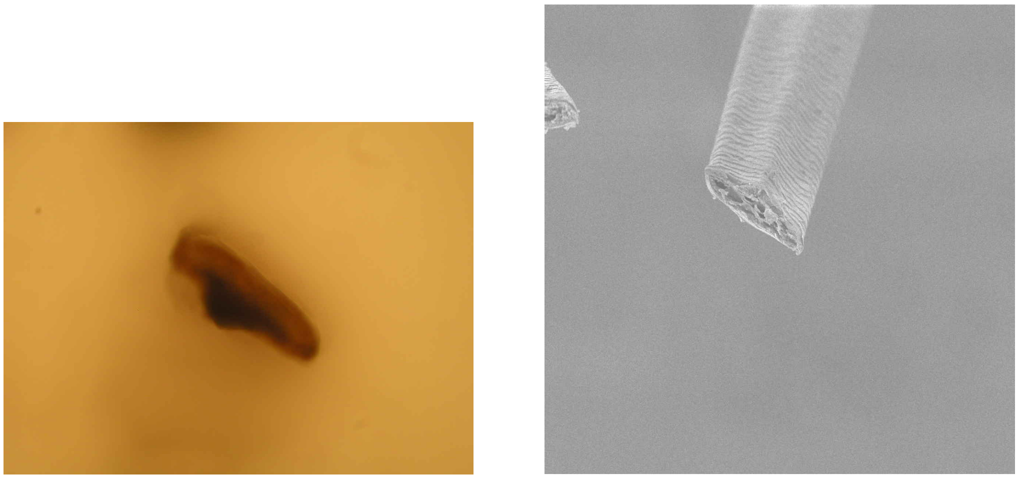

(a)

Figure 4.12: Impala cross sections taken with (a) light microscope and (b) SEM

shape, as shown in Figure 4.4. However, the light microscope provides a two dimensional representation of the shape as opposed to the three dimensional representation provided by the SEM. Since the moments employed at the feature extraction phase work on two dimensional shapes, the sensor stage segments a two dimensional shape from the three dimensional image produced by the SEM. This proves to be more challenging than segmenting a two dimensional shape from the two dimensional image provided by the light microscope. Therefore , the use of the light microscope for cross section pattern image capture is preferred to the use of the SEM. The images produced from the light microscope are passed to the sensor stage described next.

4.3.2 Sensor

An initial sample of 100 cross section pattern images, comprising 20 images of each selected species, is acquired using the light microscope. Of the 100 images captured, 50 (10 per species) are dedicated to training and the remaining images are used in testing the process. Images captured by the light microscope contain noise which may lead to distortions of the cross section shape if poor segmentation techniques are used. An image with such noise is shown in Figure 4.13. Therefore, robust segmentation techniques are investigated to provide the accurate segmentation required by the design of the sensor stage.

Figure 4.14 shows the segmentation techniques experimented with during this project. Fig-

Figure 4.13: Cross section image with noise

ures 4.14 (a) to 4.14(d) all show failed attempts at segmenting the cross section pattern shown in Figure 4.13. The results from these techniques are either not robust to noise or cannot reproduce the complete outline of the cross section pattern. Therefore, the GrabCut image segmentation technique implemented by Marsh [2005] is utilised to perform image segmentation. Figure 4.14(e) shows the results of image segmentation using the GrabCut technique. This technique is utilised in this project as follows:

Since the GrabCut technique is implemented in the Gimp image processing application, images are first opened in the Gimp application and segmented using the GrabCut technique. This technique requires the user to specify the background and foreground sections of the image as shown in Figure 4.15. Once a cross section pattern is extracted, it is placed in a new image which is then imported into ImageJ for further processing. The user selects the area where the cross section resides using a rectangular ROI (Figure 4.16) and the image in this ROI is converted to a black and white image using the auto-thresholding functionality in ImageJ. The white in the new image defined by the ROI corresponds to the background and the area in black corresponds to the cross section pattern (Figure 4.14(e)). Since the calculation of moment features in the feature extraction stage requires that the amount of background pixels of an image are minimised, the smallest possible ROI that contains the shape is selected. This minimises the amount of background pixels in the resulting image passed to the feature extraction stage detailed next.

(a) (b) (c)

(d) (e)

Figure 4.14: Image segmentation techniques investigated (a) automatic thresholding (b) maximum entropy thresholding (c) colour thresholding (d) edge detection using the Sobel operator combined with automatic thresholding (e) GrabCut segmentation

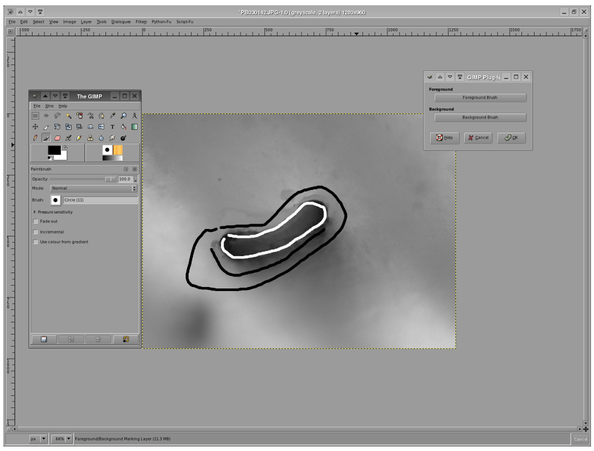

Figure 4.15: GrabCut technique

Figure 4.16: ROI selection cross section pattern

4.3.3 Feature Extraction

The central moment values of total mass, variance, skewness and the histogram sharpness are calculated on an image passed from the sensor stage using the Moment Calculator plug-in by Richard [2001]. The class also normalises these central moments as required by the design of this stage.

Since Hu's seven moments are derived from normalised central moments, this class is modified to calculate Hu's seven moments using the normalised central moment values calculated on an input image. There is no need for any extra processing of these features to handle the variations present in input cross section images as the seven Hu moments are rotation, scale and translation invariant.

It is observed from the work carried out in this project that the double primitive in Java can easily store the moment values without a loss in precision, since the largest moment values observed in the project are approximately equivalent to $5\text{x}10^{51}$.

The feature vectors created at this stage are passed to directly to the classification stage as dictated by the design, since no feature selection is needed. This classification stage is mentioned next.

4.3.4 Classification

Two distance measures are experimented with in this project and these are the Euclidean and Hamming distance measure. The Euclidean distance measure is shown to be effective in the scale pattern process and is used here with an expectation of similar success. As with the scale pattern process, the least Euclidean distance corresponds to the best match between an input vector and a training set.

The Hamming distance is chosen since it classifies vectors by counting how many elements most closely match between two vectors. This allows each moment value to be given a equal weighting during classification. A training set that has the greatest number of closest matching elements is chosen as the best match.

In order to show the confidence that the sub-process has in determining the best match, the results are ranked from the best matching training set to the worst matching.

Finally, in order to eliminate intra-species variations, the training set of each species is formed by averaging the feature vectors of the 10 images set aside at the sensor stage for the creation of the training set.

A summary of this implementation detailed in this chapter is given next.

4.4 Summary.

The evaluation of this implementation is given in the next chapter. Despite keeping a strict adherence to the design provided in Chapter 3, there are a number of questions that arise from the implementation of the design and these are:

1. What is ideal size of the Gabor filter-bank implemented at the feature extraction stage of the scale pattern sub-process? The work detailed in this chapter implemented filter-banks of size four , eight and sixteen.

2. Which of the two feature selection methods (AAD and variance) implemented in the scale pattern sub-process leads to better classification results?

3. Which of the two distance measures (Euclidean and Hamming distance) provides better classification of moment based features in the cross section sub-process?

These questions and the overall research questions posed in this project are addressed through the results provided in the next chapter.

Chapter 5

Evaluation

This chapter evaluates the hair pattern recognition process designed in this project by analysing the results achieved through its implementation. The first measure used to evaluate this implementation is the genuine acceptance rate which was determined in Chapter 2 as the most appropriate measure of evaluation of the hair pattern recognition process. The second measure used is a test of consistency in the results. This measure is used since, researchers need to trust the results produced by the sun-processes. Therefore, the consistency of each sub-process is determined through analysing the amount of times it changes the classification result of the same pattern. The results of each sub-process are evaluated and explained using these measures.

The scale pattern sub-process is evaluated first followed by the cross section pattern sub-process.

5.1 Hair Scale Pattern Sub-Process Evaluation

The hair scale pattern sub-process is tested using Gabor filter-banks of size 4, 8 and 4. In addition, the variance and AAD are used to perform feature selection on the information passed from all the filter-banks. A total of 25 images (five from each species) are repeatedly classified three times in order to ascertain the consistency of the classification results achieved. Since this testing produces a large amount of data, the evaluation of this sub-process will entail evaluating the feature selection methods and then drilling down into the results of the best performing method.

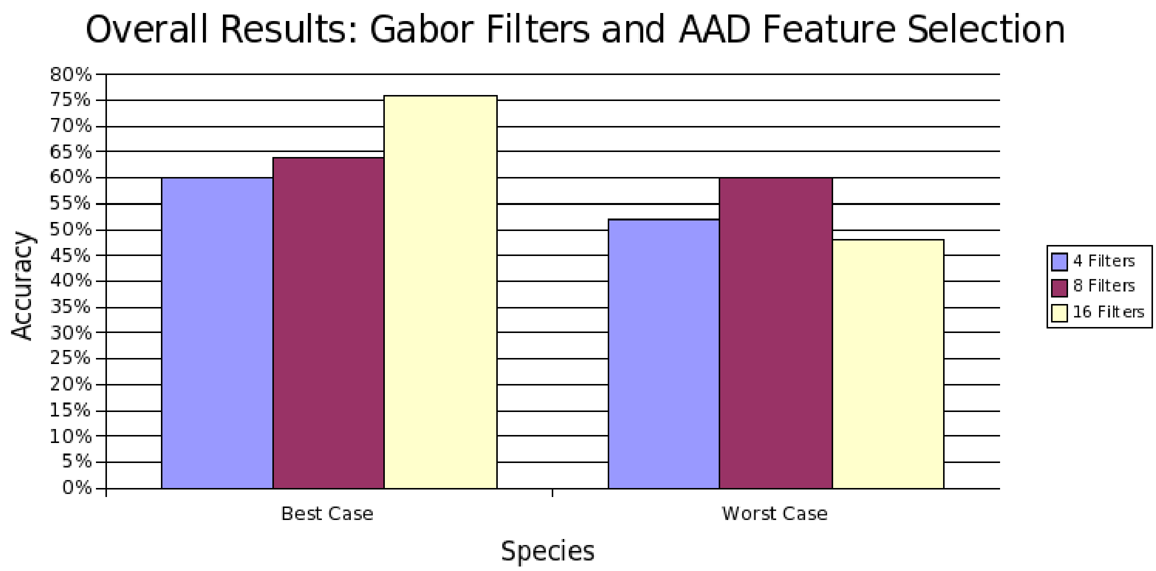

Figure 5.1: Chart showing best and worst results achieved with the AAD method for feature selection

5.1.1 Feature Selection Evaluation

Figures 5.1 and 5.2 show the overall classification results for all sizes of filter when using the AAD and variance calculations at feature selection. The best case part of the charts refers to the best result achieved from the three times the sub-process was run over the 25 test images. The worst case refers to the worst classification result achieved from the three test runs. As can be seen from these charts, the AAD calculation produces better accuracy than the variance calculation. The worst case results produced by the use of the AAD measure compare to the best results produced by the variance. In addition, Figure 5.3 shows that the variance measure leads to more changes in classification result. Therefore, it is concluded that the AAD calculation is a better feature selector than the variance calculation. A drill down is performed on the results produced by the use of the AAD feature selector in the evaluation of filter-bank size provided next.

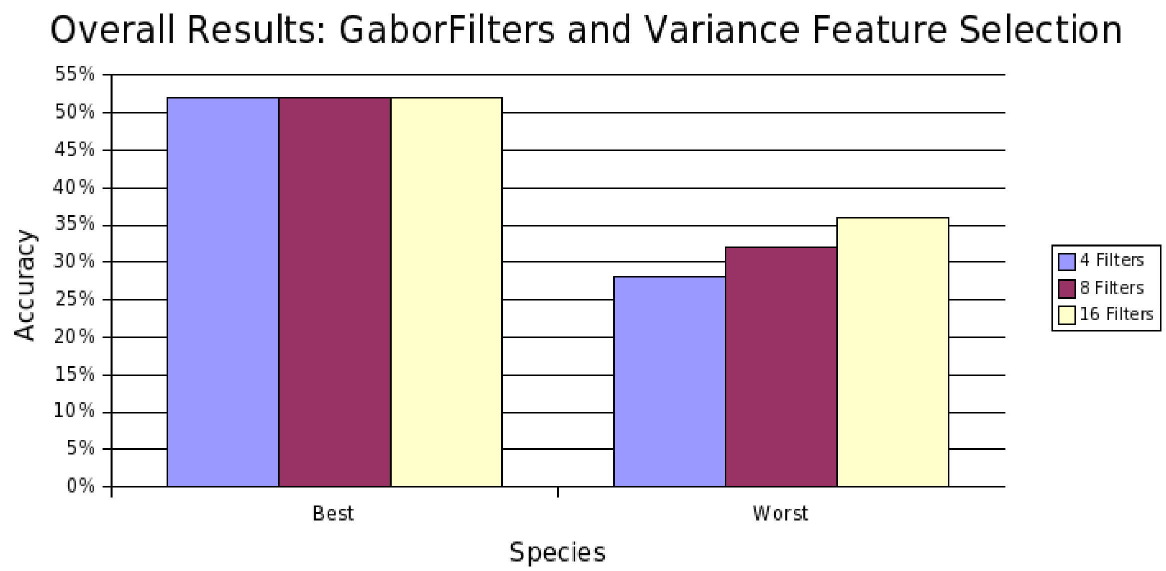

Figure 5.2: Chart showing best and worst results achieved with the variance method for feature selection

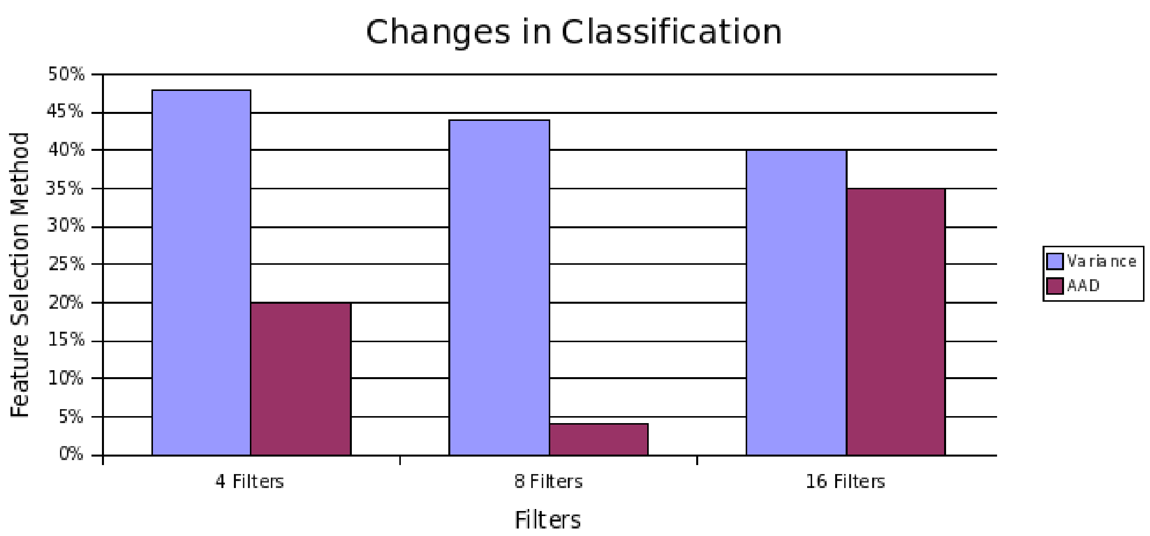

Figure 5.3: Chart showing changes in classification produced by the different filter sizes.

5.1.2 Gabor Filter-bank Evaluation

Figure 5.1 suggests that increasing the size of the filter-bank leads to better classification performance. The best performance of of a filter-bank of size 4 is 60 % compared to 64 % for a filter-bank of size 8 and 76 % for a filter-bank of size 16.

However, a deeper analysis is carried out in an investigation to determine how each size of filter-bank classifies each species. The data obtained from this drill down is shown in Figures 5.4 and 5.5. It is expected that the highest genuine acceptance rates are found in classifying blue wildebeest and jackal since these species have different scale patterns. A comparison of the best and worst results shown in these charts respectively indicates the classification information produced by Gabor filters allows for the comfortable discrimination between species with distinct scale patterns. However, the same analysis indicates that changes from correct classification to incorrect classification and vice-versa occur when dealing with species with similar scale pattern. The overall changes produced by each of the filters is shown by the blue bars in Figure 5.3. From these charts it is concluded that a filter-bank of size eight produces the most stable classification results when faced with classifying similar scale patterns.

From this data it is concluded that both the global and local pattern information of a scale pattern extracted by a filter-bank of size 8, lead to better classification accuracy and consistency than when global pattern information only is extracted by a filter-bank of size 4. In addition, while the extra information extracted by a filter-bank of size 16 may lead to higher classifications, it it less consistent than both filter-banks of size 4 and 8. Therefore, a Gabor filter-bank of size 8 is confirmed as the best size for feature extraction.

The evaluation of the cross section pattern sub-process is considered next.

5.2 Hair Cross section Sub-Process Evaluation

The hair cross section pattern sub-process is tested using both the Euclidean and Hamming distance measures. A total of 50 images (10 from each species) are repeatedly classified three times in order to ascertain the consistency of the classification results achieved. An evaluation of both distance measures and their classification results of each species is carried out first.

5.2.1 Evaluation of the classification of each species.

Figures 5.6 and 5.7 show the best and worst classification accuracy of the cross section subprocess when applied to each species. The genuine acceptance rates for each species are, lower

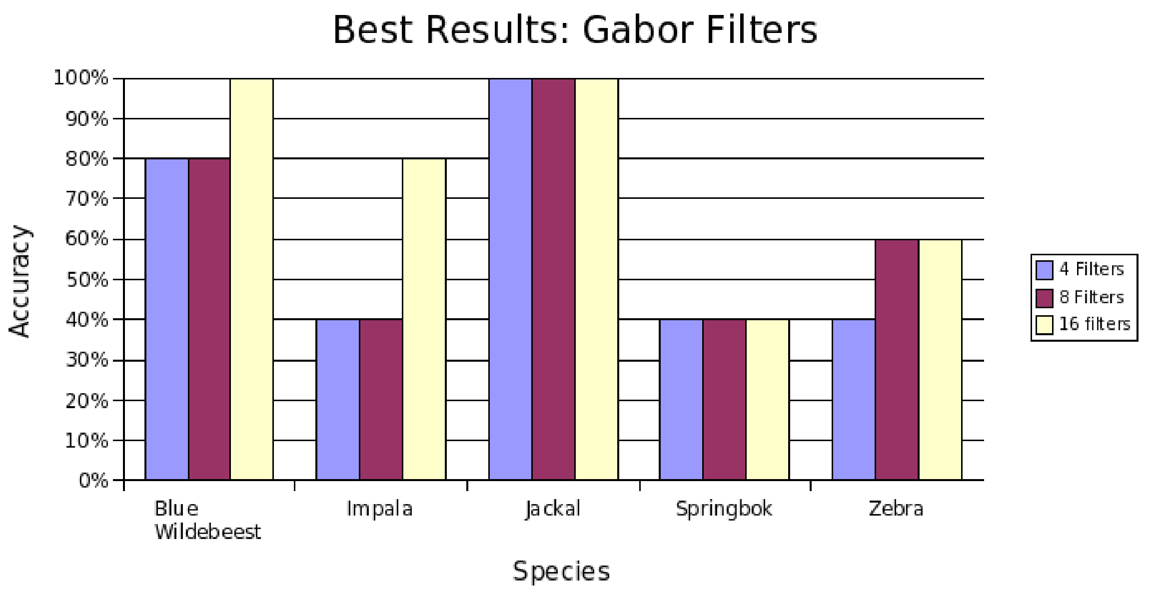

Figure 5.4: Chart showing best classification results for each species using the AAD for feature selection

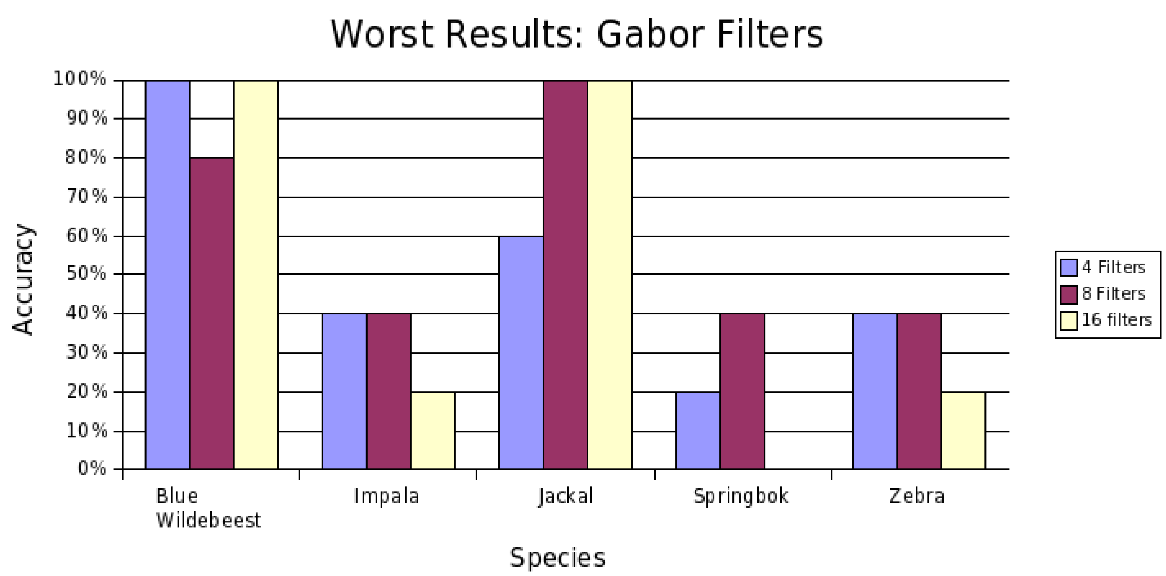

Figure 5.5: Chart showing worst classification results for each species using the AAD for feature selection

than those achieved by the scale pattern sub-process since cross section shapes show less inter-species variation than scale patterns.

However, the low matches found in blue wildebeest and impala are unexpected since these species have distinct cross section shapes. This anomaly is explained by the representation of the shapes obtained after image segmentation. Despite, having the appearance of a distinct shape to the human eye, blue wildebeest, impala and springbok cross section shapes are elongated and rounded at both ends. The rest of the variations such as the bump in the middle of an impala cross section, or the slightly curved nature of the blue wildebeest cross section are in some images too small to be represented through Hu's moments.

However, the relatively high matching scores of zebra and jackal when using the Hamming distance measure are unexpected, since these species have similar cross section patterns. A closer inspection of the results reveals this similarity as all matches ranked second after a correct jackal classification are reported as zebra. The best reason provided in attempt to explain this result, is that Hu's moments are better at distinguishing between rounded objects than elongated objects. However, this reason is flawed because if the Hu's moments cannot represent small variations in elongated objects that the human eye can see, then they cannot be expected to pick up variations in rounded objects that the human eye cannot. Therefore, the answer to this result remains an open issue. An evaluation of the distance measures used is provided next.

5.2.2 Evaluation of Euclidean and Hamming distance classifiers

Figure 5.8 and 5.9 indicate that the Hamming distance measure is more accurate and consistent than the Euclidean distance measure. The overall best genuine acceptance rates are 34% and 48% for the Euclidean and Hamming distance classifiers respectively. The Euclidean distance produces changes in classification of 40% as compared to 28% by the Hamming distance during the three test runs.

This is explained by the fact that the Hamming distance measure assigns an equal weighting to each moment feature during classification. Therefore, the moment feature values in a feature vector are considered to be off the same classification value in the matching process. However, the Euclidean distance combines the distances between the elements of two feature vectors in determine a matching score. Therefore, it is concluded that assigning an equal weighting to Hu's moment features, leads to a more accurate comparison of moment feature vectors. As a result, the Hamming distance is the preferred classifier employed by the cross section sub-process.

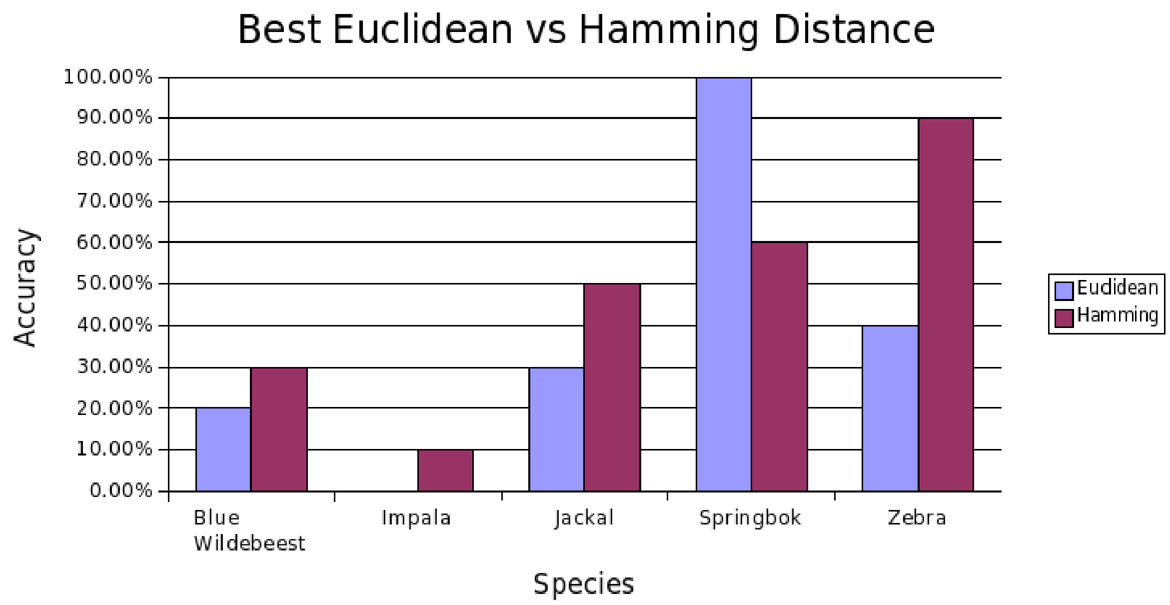

Figure 5.6: Chart showing best results of the cross section sub-process.

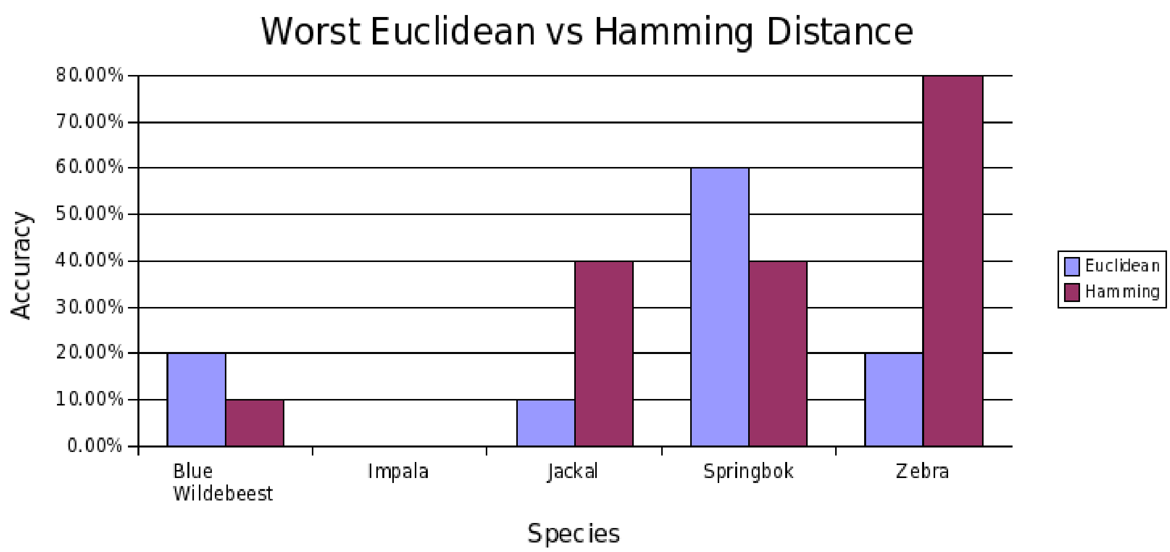

Figure 5.7: Chart showing the worst results of the cross section sub-process.

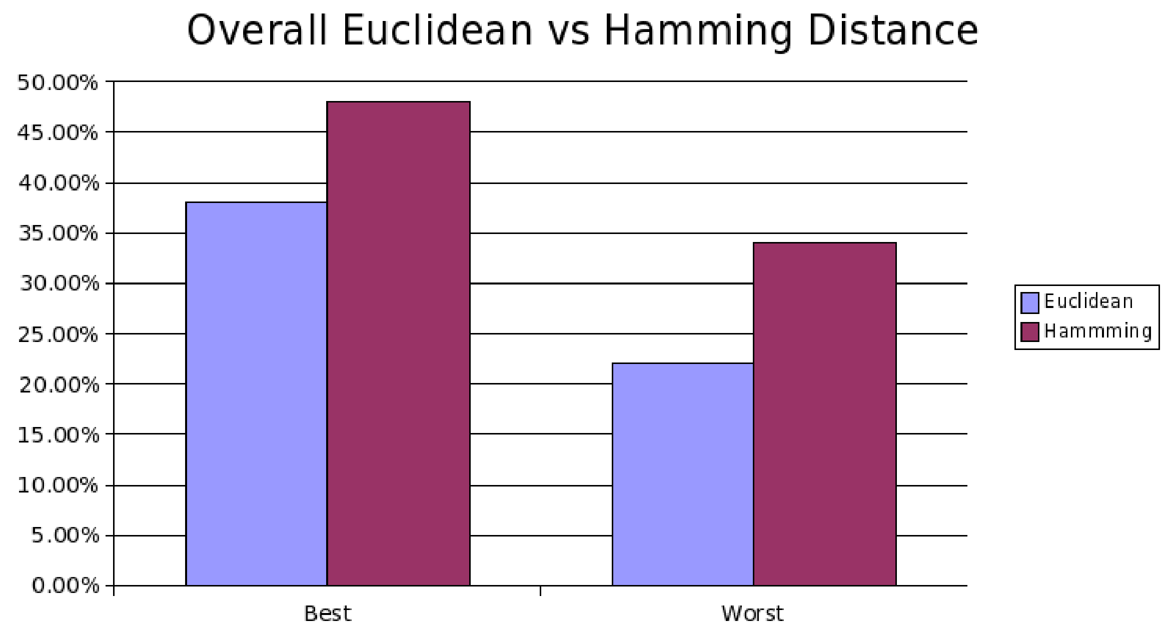

Figure 5.8: Chart showing overall best and worst results of the cross section sub-process

5.3 Summary

Table 5.1 shows the combination of techniques in both sub-processes that provide the best performance . The global and local scale pattern information extracted by a 2D Gabor filter-bank of size 8 and selected using the AAD, provides a best genuine acceptance rate of 64%. This combination of techniques produces only one change in classification during three test runs and this indicates that it is the most consistent scale pattern technique combination explored. Since the lowest scale classification rates are achieved in classifying hair from species with similar scale patterns, it is asserted that the mis-classifications are a result of the limited samples used in creating the training sets. The limited training samples limit the detail of the inter-species scale pattern variation represented by the training sets. Therefore, it is concluded that the scale pattern sub-process is limited by the quality of variation presented by the training sets as opposed to flaws in its design and implementation.

The cross section sub-process produces a best genuine acceptance of 48% when the combination of techniques shown in Table 5.1 is employed. This genuine acceptance rate combined with a change in classification of 28% during the test runs, shows that this sub-process performs

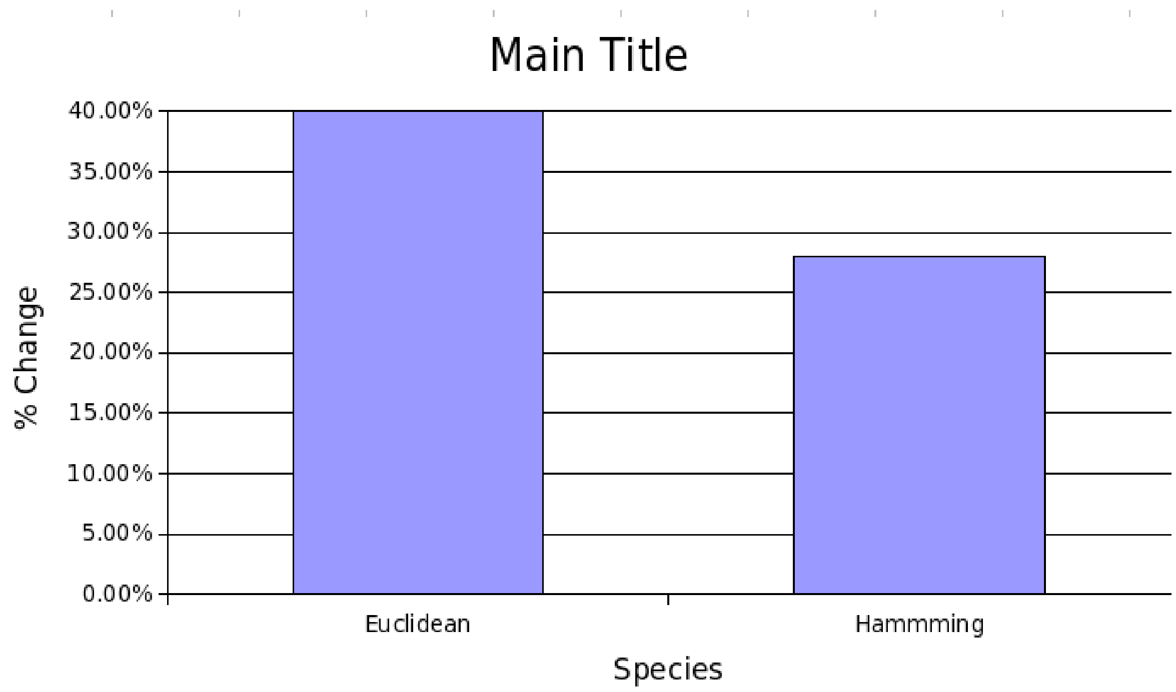

Figure 5.9: Chart showing the changes in classification experienced in the cross section sub-process

Table 5.1: Best performing combination of techniques

Stages	Hair Scale Pattern Sub-Process	Hair Cross-Section Pattern Sub-Process
Sensor	Image orientation, standardisation, histogram stretching and image size standardisation.	GrabCut image Segmentation and auto thresholding.
Feature Extraction	2D Gabor filter-bank of size eight.	Hu's seven moments.
Feature Selection	Filtered images are tessellated into squares and each feature is calculated using the AAD.	None required.
Classifier Design	Least Summed Euclidean distance measure.	Hamming distance measure.

worse than the scale pattern process. This is expected since human researchers also find cross section patterns of limited diagnostic use as compared to scale patterns. However, more careful segmentation using the GrabCut technique than that done during the implementation of the cross section sub-process may improve these classification results.

The conclusion of the entire project is provided next.

Chapter 6

Conclusion

The first application of automated image pattern recognition techniques to the problem of classifying African mammalian species using hair patterns is presented in this project. The hair characteristics used, that is scale and cross section patterns, are statistically analysed in their own sub-process to provide a classification of hair. This contribution is novel since previous studies relied on subjective human interpretation as opposed to objective statistical and numerical analysis. The conclusions drawn from this project are presented in this chapter. Finally, this document is concluded with a mention of future work.

6.1 Conclusions

This section presents the conclusions drawn from the work carried out in this project. The conclusions presented are structured according to the research questions answered during the project.

6.1.1 Sensor

The first research question answered by this project is concerned with determining the techniques applied in the image pre-processing of hair pattern images. The process designed in this project presents image size and orientation standardisation as the two techniques required to handle the variations that occur during the input of the scale pattern image. It is also found that the scanning electron microscope produces more appropriate images for the scale pattern sub-process than the light microscope.

The GrabCut image segmentation technique is shown to be the most effective segmentation technique explored in extracting hair cross sections from an input image. Careful segmentation with this technique results in the extraction of a complete cross section shape regardless of the noise encountered in an image. This noise is introduced through the methods employed in preparing a hair for input using light microscope which is preferred for cross section image capture.

The conclusions drawn from answering the second research question are provided next.

6.1.2 Feature Extraction and Selection

This project's second research question requires an investigation into feature representations of hair patterns. This project presents a 2D Gabor filter-bank of 8 filters as the most appropriate feature extractor examined for hair scale patterns. This classification is shown to improve when the feature selection technique of using the AAD to calculate values for each square in image tessellation is used. A filter-bank of size 16 produced both the best and worst classification results, and this indicates that its results cannot be trusted.

Hu's seven moments when calculated on a thresholded cross section image are shown to represent the shape description of the cross section. However, these moments are found to be insensitive to the small inter-species variations present in the cross sections of some species.

The conclusions drawn from the final research question are given next.

6.1.3 Classification

The final research question answered in this project requires the identification of classifiers that will place hair patterns into their correct classes. The use of the summed Euclidean distance is shown as an adequate classifier for scale pattern classification. The consistency presented in the classification of features obtained from a Gabor filter-bank of 8 filters, leads us to conclude that the best genuine acceptance rate of 64% can be increased if better training sets are used than those implemented in this project.

The Hamming distance measure is shown in this project to be a better classifier of Hu's seven moments than the Euclidean distance. It is established that the Hamming distance performs better since it weights each moment equally during classification. This classification is, as expected, lower than that of scale pattern classification since cross section patterns have less inter-species variation than scale patterns. However, it cannot be concretely established why the similar zebra and jackal cross section patterns have a higher classification rate than species with distinct cross section patterns.

6.2 Future Work

Since this project presents the first automated pattern recognition techniques to the domain of hair pattern recognition, there is scope for future work. The species used to test the hair pattern recognition may be increased to test the process's scalability. The number of patterns used to build the scale process training sets may be increased to test whether this results in more accurate scale pattern classifications. Finally, higher order moments may be derived and tested to ascertain whether they improve the classification of cross section patterns.

Bibliography

Hans Brunner, Barbara Triggs, and Ecobyte Pty Ltd. Hair id an interactive tool for identifying australian mammalian hair. CD-ROM, 2002. Publisher:Csiro Publishing . Last Accessed: 09 October 2005.

John Daugman. How iris recognition works. *IEEE Transactions on Circuits Systems and Video Technology*, 14(1):21–30, 2004.

Manuel G Forero, Filip Sroubek, and Gabriel Cristóbal. Identification of tuberculosis bacteria based on shape and color. *Real-Time Imaging*, 10(4):251–262, 2004. URL `http://www.sciencedirect.com/science/article/B6WPR-4D1YXKM-2/2/905d9e5eedbe2dd6eeb5c72efff4bc2e`.

ImageJ. ImageJ. Online, 2005. Available:http://rsb.info.nih.gov/ij/. Last Accessed: 26 September 2005.

Anil Jain, Salil Prabhakar, Lin Hong, and Sharat Pankanti. Filterbank-based fingerprint matching. *IEEE Transactions on Image Processing*, 9(5):846–859, 2000. URL `citeseer.csail.mit.edu/article/jain00filterbankbased.html`.

Hillary J Keogh. A photographic reference system of the microstructure of the hair of Southern African bovids. *South African Journal of Wildlife Research*, 13(4):89–132, 1983.

Tai Sing Lee. Image representation using 2D Gabor wavelets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(10):959–971, 1996. URL `citeseer.ist.psu.edu/article/lee96image.html`.

Simon X. Liao and Miroslaw Pawlak. On image analysis by moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(3):254–266, 1996.

Matthew Marsh. Implementing the grabcut segmentation technique as a plugin for the gimp, 2005.

Kristin L Nichol and Paul Mendelman. Influence of clinical case definitions with differing levels of sensitivity and specificity on estimates of the relative and absolute health benefits of influenza vaccination among healthy working adults and implications for economic analyses. *Virus Research*, 103(1-2):3–8, 2004. URL `http://www.sciencedirect.com/science/article/B6T32-4C8PDJP-2/2/7e1c9a544c6ac95f7548457b6f839973`.

M R Perrin and B S Campbell. A key to the mammals of the Andries Vosloo Kudu Reserve (Eastern Cape), based on their hair morphology, for use in predator scat analysis. *South African Journal of Wildlife Research*, 10(1):1–14, 1980.

Francois Richard. Moment calculator plugin. Online, 2001. Available:http://rsb.info.nih.gov/ij/plugins/moments.html Last Accessed: 5 November 2005.

Arun Ross, Anil Jain, and James Reisman. A hybrid fingerprint matcher. *Pattern Recognition*, 36(7):1661–1673, 2003. URL `http://www.sciencedirect.com/science/article/B6V14-47XWGYR-7/2/911498f4913fac9dbfeec6b6f25bf04d`.

C Sanchez-Avila and R Sanchez-Reillo. Two different approaches for iris recognition using Gabor filters and multiscale zero-crossing representation. *Pattern Recognition*, 38 (2):231–240, 2005. URL `http://www.sciencedirect.com/science/article/B6V14-4DKGY44-3/2/b46840422b76505175ff8c945bc12dd8`.

Sergios Theodoridis and Konstantinos Koutroumbas. *Pattern recognition*. Elsevier Academic Press, second edition, 2003.

Mool S Verma, Lorien Pratt, Chidamber Ganesh, and Christy Medina. Hair-MAP: a prototype automated system for forensic hair comparison and analysis. *Forensic Science International*, 129(3):168–186, 2002. URL `http://www.sciencedirect.com/science/article/B6T6W-46Y4PX6-4/2/9fbdfee40c26967951da61b8bea2b5e3`.

Chan-Yun Yang and Jui-Jen Chou. Classification of rotifers with machine vision by shape moment invariants. *Aquacultural Engineering*, 24(1):33–57, 2000. URL `http://www.sciencedirect.com/science/article/B6T4C-41P170Y-3/2/af0927ebcbdfe97aaffd2b7490527e21`.

David D. Zhang. *Automated Biometrics*. Kluwer Academic Publishers, 2000.