

Automatic Identification of Morphometric Landmarks in Digital Images.

Sasirekha Palaniswamy, Neil A Thacker and Christian Peter Klingenberg.

Last updated
22 / 4 / 2007

This document forms part of the **Recognition and Intelligence Series** available from www.tina-vision.net.

- 2007-001 Retinal Sampling, Feature Detection and Saccades: A Statistical Perspective.
- 2006-008 Statistical Principles for Selection of Computer Vision Algorithms as Modules for Visual Perception - Show Me the Errors.
- 1991-001 Designing a Layered Network for Context Sensitive Pattern Classification.
- 1997-002 Supervised Learning Extensions to the CLAM Network.
- 1996-003 Tutorial: Algorithms For 2-Dimensional Object Recognition.
- 1997-005 Speechreading Using Probabilistic Models.
- 2000-002 Solving Shape Based Object Recognition from a Computational Standpoint - Practical and Physiological Constraints.
- 1995-004 Assessing the Completeness Properties of Pairwise Geometric Histograms.
- 1996-004 Robust Recognition of Scaled Shapes Using Pairwise Geometric Histograms.
- 1996-005 Multiple Shape Recognition Using Pairwise Geometric Histogram Based Algorithms.
- 2007-007 Automatic Identification of Morphometric Landmarks in Digital Images.
- 1999-002 A Feature Representation for Map Building and Path Planning.
- 2001-015 Colour Image Segmentation by Non-Parametric Density Estimation in Colour Space.
- 2001-006 What is Intelligence?: Generalised Serial Problem Solving.
- 1994-002 A Correlation Chip for Stereo Vision.
- 1995-001 Specification and Design of a General Purpose Image Processing Chip.



Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

Automatic Identification of Morphometric Landmarks in Digital Images

Sasirekha Palaniswamy¹, Neil A Thacker² & Christian Peter Klingenberg¹

1: Imaging Science and Biomedical Engineering.

2: Faculty of Life Sciences.

Abstract

Our aim is to develop a completely automated and reliable system to identify morphological landmarks in digital images. The performance of the system is aimed to replicate manual digitization with equivalent accuracy and reliability, based upon a small number of training examples. The analysis system is constructed from four stages; a feature based detection of fly wing structure, correspondence matching based upon the pairwise geometric histogram (PGH) representation, global location of the wing using a Probabilistic Hough Transform (PHT), and finally local correlation based refinement of individual features. We evaluate this system and compare quantitative results to manually digitized data.

1 Background

Morphological landmarks are points that can be located precisely and establish an unambiguous one-to-one correspondence among all the specimens and are widely used in shape analysis [1]. Points like the tip of the nose or the outer corner of the left eye are possible landmark points of the human face. Analyses of shape investigate the arrangement of landmark points relative to each other. A substantial body of statistical methods is available for the analysis of configurations of landmark points [2].

This framework of shape analysis by landmarks is increasingly used in many biological and medical applications and widely applied in many other fields. The configuration of landmarks have helped identify the possible source of re-infesting specimens and encounter the epidemiologically challenging vectors of Chagas disease [3]. The potential of using geometrical morphometric techniques as an invaluable tool for recognizing taxonomic data is being explored [4]. Other scientific applications include investigating the study of size and shape to examine the effects of experimental treatments, genotype or other factors directly in the anatomical aspect. The use of landmarks has been adapted to specific biological contexts such as genetics [5, 6, 7]; geographic differentiation [8], and the study of morphological integration [9, 10].

The process of identifying the landmarks is an important and labour-intensive part of any such analysis. Presently, this is usually done manually. Plugins for the ImageJ software (for digitizing the standard sets of landmarks) increases speed and reliability over a completely unaided process [11]. However, there is still a requirement for an observer to manually identify each landmark point and therefore this process can be time-consuming, and quite often, the research questions are dependant on the duration of obtaining these data.

Developing an automated system for locating landmarks in digital images of *Drosophila* wings is largely significant, as it is an excellent model for the study of genetics of development and evolution of morphological form [12, 13]. They contain a wealth of interesting biological information and its simple, flat, two dimensional structure enables convenient handling. Therefore, automation has numerous advantages over a manual system, as it will not only diminish the labour needed for shape analysis, but it also will eliminate the source of error (mistakes made during digitizing and subtle differences between observers). Furthermore, automatic extraction of features from images can potentially change the way in which landmarks are chosen for morphometric studies. Whereas the traditional approach requires landmarks to be chosen a priori, based on outside knowledge of the study system, the approach using automated image analysis raises the possibility to identify and extract features from the total information contained in the images that are maximally informative in the context of a particular research project.

There have been previous attempts to automate the process of landmark location on the *Drosophila* wing. In [13], the operator initiates the process by marking two landmark positions and the system fits a series of spline curves to the margin and the veins of the wing, and defines the landmarks as the intersections of the splines. The drawback of this system is that the landmarks are not always at the exact location of the intersection of veins and the splines may not exactly match the veins (because of the "stiffness" of the spline interpolation) and the system has problems identifying wings of species with highly melanized spots at the intersections.

Another such system to locate the landmarks on digital images of bee wings is being developed at the Paris National Museum of Natural History [14]. This system applies the techniques of mathematical morphology and skeletonization to obtain the landmarks. However, using these techniques are not guaranteed to be robust. The method also requires human intervention in terms of loading the data and identifying the initial set of points to start the process and the pre-processing step includes certain parameters that have to be set by the operator.

Other similar automated systems include the Digital Automated Identification SYstem (DAISY) which was developed whilst attempting a novel approach to identify insect specimens from the images. Principal component auto-associative memories and trainable classifiers are exploited to identify closely related parasitic wasps based on their wing venation and pigmentation patterns [15]. This system has been designed to identify several organismal groups in real time and it successfully classifies data into morphologically similar classes and proves to be a very useful and practical tool for taxonomic identification of various species. The internal algorithms are based upon the use of a pairwise geometric histograms (PGH) representation, which is used to establish shape correspondence.

2 Introduction

We present an automated system for the analysis of edge based structure for use in morphometric studies. The current work takes a grey level image of *Drosophila* wing as input and extracts the coordinates of 15 landmarks (Figure 1). A typical shape analysis requires several hundred images and identifying these landmarks is a laborious process. An automated method to extract these features can potentially improve the methodology with which the landmarks are identified via improved standardization and accuracy.




Figure 1: 15 Landmark locations on *Drosophila* wing (Image courtesy: [11]).

The proposed method extracts the ridges (linear features such as wing veins) using the knowledge of their known grey level profile and the noise characteristics of the image. This approach has been shown to be statistically valid [16]. The ridges obtained are approximated by line segments and the geometric relationships between them are encoded in PGH, an approximation to the probabilistic density function for the geometric co-occurrences in the data [17]. Shape correspondence is determined by comparing and matching the pairwise histograms of the scene and model data. A probabilistic Hough transform (robust Likelihood) is used to determine the hypothesized landmark location (Figure 2). Sub-pixel estimation of the landmark location is performed by template matching, i.e., correlating a small region around the Hough estimated landmark location.

We show that a single training image with its landmark coordinates is enough to independently estimate the landmarks of any individual within a particular dataset. However, the reliability and accuracy of the method can be further enhanced by using multiple training images. Multiple estimates also offer the possibility of accuracy assessment, an important aspect of any scientific study. The precision, repeatability and robustness of the algorithm have been evaluated here as a pilot study. Although some predictions regarding reliability can be made with a small sample, a further study will be carried out on a larger sample to test the reliability of the system on scientific studies.

3 Methods

3.1 Data Acquisition

Data acquisition is carried out by mounting the fly wings in rows on a microscopic slide and flattening with a coverslip. The digital images are obtained using an appropriate digital camera mounted on the microscope and attached to a computer. A calibration image is generally obtained along with each of the dataset to standardize the difference in magnifications between different dataset. The anatomical landmarks can be easily collected in two dimensions from digital images and this approach is quite useful in evolutionary research as the landmarks

can be collected from non-model organisms or even fossils. The X & Y co-ordinates of these landmarks are usually obtained by manually digitizing the location of these co-ordinates appropriately based on their anatomical context. Specialized algorithms and plugins can be used to semi-automate the process and to enhance the speed of the digitization.

3.2 Analysis

The analysis system is constructed from four stages; a feature-based detection of fly wing structure, correspondence matching based upon the PGH representation, global location of the wing using a Probabilistic Hough transform, and finally correlation based refinement of individual features. We evaluate this system and compare quantitative results to manually digitized data below.

3.2.1 Ridge Detection- An extension to Canny framework

The wing veins are extracted as ridge features, using a method which is a modification of more conventional edge detectors. We locate these features using a matched filter approach, approximating the vein profile as a Difference of Gaussians. Local maxima in response are then passed into a more conventional hysteresis threshold and linking system, based upon the popular Canny [18] system, in order to extract connected structures. This extracted edge map can be used to determine the precise location of landmarks. The uniformity of noise in the feature enhancement stage guarantees that this process is inherently stable. The ridge detector has been optimized for the task of locating landmarks by analyzing the specific characteristics of noise and scale stability. The whole process can be interpreted as a statistical null-hypothesis test for the presence of the defined feature [19].

3.2.2 Pairwise Geometric Histograms

The extracted edge-map is approximated by line segments and the geometric relationships between each pair of line segments are encoded in the pairwise geometric histograms. This is an approximation to the probabilistic density function,

$$H_i(\theta, d) = P(\theta_i - \theta_j, d_{ij} | e_i) \quad (1)$$

for the geometric co-occurrences of an edgel e_j given e_i as a function of relative angle $\theta_i - \theta_j$ and perpendicular separation d_{ij} . This is a well established method of shape representation based on recording the distribution of pairwise geometric relationships between local shape features which can support recognition and there is considerable robustness to the loss of data due to fragmentation noise and occlusion [17]. The method is also known to be complete, in that the original structure of the object can be reconstructed from the set of histograms describing a shape. This representation is invariant for portions of the same linear feature so that it can be constructed by considering a linearization of the edge map. The importance of a pair of line segments defining the representative shape can be encoded by entering the product of their lengths at the value of the entry. The entry is blurred along each axis to encode the uncertainty regarding the true position and orientation of each line segment. The scale of binning and extent of blurring defines the extent of allowable differences when matching similar shapes.

Shape recognition is done by identifying the correspondences between image and object features. Shape representation comprises many geometric histograms, each representing a single model feature. The degree to which a linear edge feature in the test image matches a particular model feature can be determined by comparing their histograms. The degree of match between them is given by the Bhattacharyya measure B_{ij} , which takes the form of a dot product correlation of the histograms of lines H_i and H_j .

$$B_{ij} = \sum_{\theta} \sum_d^{d_{max}} \sqrt{H_i(\theta, d) H_j(\theta, d)} \quad (2)$$

This can be related, via the χ^2 variable, to a maximum likelihood similarity metric and can be derived as an approximation to Fisher's Exact test as a method for comparing two distributions. The hypothesized matches can then be used as input to pose an estimation algorithm such as the generalized Hough transforms. Scale independent recognition can be achieved by representing an object at a range of scales [20]. However, this property of the PGH representation was not required in our study and therefore was not utilized.

3.2.3 Hough Transform

A probabilistic Hough Transform (robust maximum Likelihood), is used to make an estimate of the global position and orientation of each wing. Entries in the 2D location histogram are made according to the localization covariance, propagated from the errors on the constraint lines. This takes proper account of errors, resulting in improved robustness and more accurate determination of model position, orientation and scale in comparison to the more conventional form of this algorithm. The entries in the Hough arrays are constructed from pair of lines (n, m) , i.e., a tuple transform. The equivalent probabilistic form for the Hough transform $H(x, y)$ used to find the position of a model in a scene is given by the expression,


$$H(x, y) = \sum_n^N \sum_m^N \log(p(x, y|n)p(x, y|m)) = \sum_n^N \log(p(x, y|n)) \sum_m^N \log(p(x, y|m)) \quad (3)$$

so that the Hough entry can be considered as the square of the robust log Likelihood $L(x, y)$ for the localization of the object,

$$H(x, y) = L(x, y)^2 = \left(\sum_n^N \log(p(x, y|n)) \right)^2 \quad (4)$$

During array construction $H_{nm} = \log(p(x, y|n)p(x, y|m))$ is estimated from a 2D Gaussian distribution centered at the the position of the model hypothesized by the m, n th pair of scene line labels with variance propagated from the individual line location errors. This tuple-based construction helps to remove background noise from the Hough array and has some computational advantages. The variability of the line segmentation process and the uniform error on the scale estimates are independent and are adjusted to give a quantitative estimate of the hypothesized location of a pre-defined reference point from pair of scene lines. Training from example data involves recording the perpendicular distance, 'd', from each model line to the reference point. Consequently, for each pair of scene lines, extended lines at the appropriate perpendicular distance will intersect at the hypothesized position of the reference. Error at the point of intersection can again be estimated by standard error propagation.

Models can be located based on the positions, orientations and scales hypothesized by scene line labels. However, the orientations and scales of the models are not determined explicitly. This can be determined separately using 1-parameter Hough transforms. For each model position determined, a 1-parameter orientation Hough transform and 1-parameter scale Hough transform can be constructed from entries selected on the basis of consistency between the scene lines and model position. The orientation is determined from the difference in orientation between the scene line and model line to which it matched. Comparing the perpendicular distance from the scene line to the model position to this same distance in the model itself would yield the scale. Peaks in these Hough transforms would give the orientation and scale of the model at that position in the scene.



(a) Flywing image with model overlaid. (b) Peak in the Hough transform.

Figure 2: Hough transform located 15th landmark.

3.2.4 Template Matching

The above Hough scheme computes an estimate of landmark position based upon global wing shape. As we need to determine variations in shape for the morphometric study this estimate needs to be refined based upon local image evidence. To obtain this estimate, template matching is performed on the Difference of Gaussian image of the scene $D(I)$ and model (example mark-up) $D(M)$ data, over a small region around the Hough estimate for the

feature in the scene data. To save processing time during alignment, the scene data is rotated to match the model data using the Hough estimate, which is assumed to be sufficiently accurate for final location of the landmark. The use of the Difference of Gaussian images eliminates any image illumination offset and the matching is performed as a dot-product correlation in order to eliminate the effects of illumination scaling,

$$L_{h_x h_y} = \sum_x^R \sum_y^R D(M(x, y))D(I(x + h_x, y + h_y)) / \sqrt{\sum_x^R \sum_y^R D(I(x + h_x, y + h_y))^2} \quad (5)$$

where, R is the region size. This is directly equivalent to performing a least squares comparison of the image regions with one free grey level scale parameter. In this study, the denominator is presumed to be constant to save time on the computationally expensive calculations. The best possible match is identified and that location is transformed back onto the scene image. The least-squares difference between the two scale image regions is stored so that the best matching examples can be selected for final estimation of landmark position (see below). This quality control feature not only allows a check on the adequacy of the example mark-ups but also eliminates residual problems in alignment estimation, such as poor rotation estimates.

4 Results

4.1 Precision of manual digitization

The precision of manual digitization by an expert is determined by determining the deviation (difference between the value of each attempt and the mean) of 10 repeats of digitization of a single image. The outcome shows that it is within a range of +/- 1 pixel (Figure: 3).




Figure 3: Reproducibility of landmark location over 10 repeats-manual digitization.

4.2 Accuracy of the template matching relative to manual digitization

To test the efficiency of the template matching stage, the feature (landmark 12) that was significantly variant in comparison to other landmarks was taken. Figure 4(a) shows the positions of the landmarks located by the Hough transform relative to the landmark locations digitized manually. The Hough transform locates most of the landmarks within a range of +/-20 pixels. This provides an estimate of the range that the correlation search must operate over, and is used to define a 'window size' parameter. During initial testing, by ensuring that the window size is large enough, we can be certain that the landmarks can be located reliably. The refinement by the template matching strategy is shown in Figure 4(b) indicating the improvement in accuracy to be within a range of +/-6 pixels on the X-axis and +/-3 pixels on the Y-axis.

4.3 Accuracy of the automated system

The accuracy of the system is assessed using multiple reference images (multiple models). The results show that the landmarks can be located more precisely in cases where the model features are a good match to the scene data. However, it is quite unlikely that the model chosen would be suitable for all the features to be estimated. Therefore, it is important to choose a set of appropriate model features in the training data that can best match with the given test dataset. This can be achieved by computing the degree of least-squares match between the model and the test feature and taking an average of the best matches available. Since, the precise location of

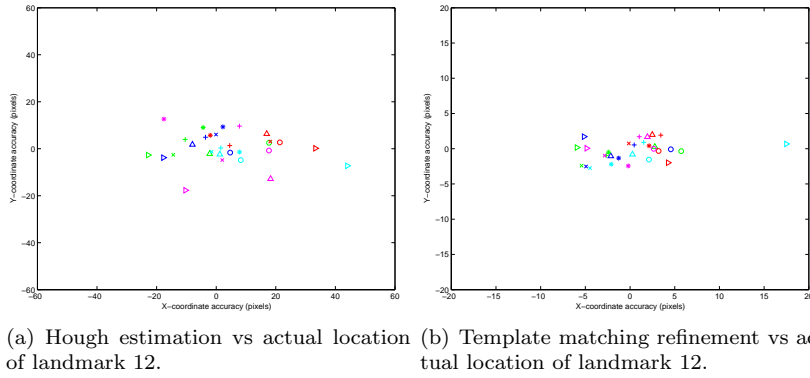


Figure 4: Hough transform & template matching performance.

certain features (eg., landmark 12 due to its structural complexity) can be quite challenging in contrast to most other features, the degree of least-squares match can also be applied as a quality control approach to determine the adequacy of the selected training examples.




Figure 5: Best Least Squares Match.

The Figures 5(a) & 5(b) shows the system performance with an average of the best 3 of 5 reference images and best 5 of 11 reference images using the least squares match. It can be seen that the accuracy of the system has improved with an increase in the number of reference images. Most of the landmarks in the sample dataset have been located within ± 3 pixels accuracy. It should be noted that the outliers are mainly contributed by one of the test image which clearly indicates that the reference images used were not necessarily a good match for feature location in that particular image. In such cases, increasing the number of reference images would considerably improve the accuracy.

4.4 Robustness of the system

The robustness of the system was tested by locating landmarks in an image with additive noise of $\tilde{10}$ times that of the original image. The system is robust in locating the landmarks within ± 4 pixels accuracy and we can therefore be confident that the system is quite stable to noise, well beyond the level normally present in this dataset. This is presumably because of the large degree of smoothing applied during the feature detection and correlation matching stages.

5 Discussion

This paper describes a system which can be trained from a few example images to automatically estimate the location of key features in ‘veined’ structures, such as insect wings. The performance of the automated system can be compared to human performance in terms of accuracy of landmark location. Although, the range of the system accuracy is nearly twice that of the manual digitization, the accuracy of the system can be considerably improved by using relevant models. The performance of the system is sufficiently accurate to allow it to replace the time consuming process of manual digitization, which is common to all morphometric studies.

The current system is capable of providing a set of 15 landmark locations based upon one reference image in about 3 minutes on a SUN Sparc Ultra 5 workstation. The time taken by a human to mark up one image using the standard mark-up tools is about 40 seconds though maintaining this speed across a large dataset might be regarded as unrealistic. We expect that the speed of the system can be optimized, however, the question of trade off between the speed and precision may arise (as more reference images may be needed to achieve the manual accuracy).

The pilot study can be scaled up with minor modifications and this automated method would be used in a scientific study with a large dataset comprising of 1600 images of different species of *Drosophila*. This analysis should enable us to test other performance aspects of the system, such as its reliability, and to evaluate any difficulties regarding the practical use of this dataset. The generic nature of object recognition and feature location incorporated in this automated system enables easy modification to locate features in a variety of other organisms. The method is intrinsically robust to changes in shape and based firmly on the statistical interpretation of data analysis. The system will be tested for its efficiency in locating the landmarks even in a scenario where the features to be located are quite complicated and beyond manual capabilities (eg. debris/bristles lying across one of the feature would be a major hurdle for manual digitization). Such an automated method will benefit major research groups in the morphometrics community and will easily be transferable to research groups in other relevant field of study. The automation of shape analysis has major potential advantages regarding standardization as the landmarks can be located without any manual intervention and will make large scale studies easily feasible [5, 8, 9]. The algorithms will be made available as an open source package via our website .

References

- [1] F.L. Bookstein. *Morphometric tools for landmark data: Geometry and Biology*. Cambridge University Press, UK, 1991.
- [2] I.L. Dryden and K.V. Mardia. *Statistical Shape Analysis*. Chichester, John Wiley and Sons, 1998.
- [3] J.P. Dujardin, C.B. Beard, and R. Ryckman. The relevance of wing geometry in entomological surveillance of triatominae, vectors of chagas disease. *Infection, Genetics and Evolution*, 7:161–167, 2007.
- [4] J.M. Becerra and A.G. Valdecasas. Landmark superimposition for taxonomic identification. *Biological Journal of the Linnean Society*, 81:267–274, 2004.
- [5] E. Zimmerman, A. Palsson, and G. Gibson. Quantitative trait loci affecting components of wing shape in drosophila melanogaster. *Genetics*, 155:671–683, 2000.
- [6] C.P. Klingenberg and L.J. Leamy. Quantitative genetics of geometric shape in the mouse mandible. *Evolution*, 55:2342–2352, 2001.
- [7] C.P. Klingenberg, L.J. Leamy, E.J. Routman, and J.M. Cheverud. Genetic architecture of mandible shape in mice: effects of quantitative trait loci analyzed by geometric morphometrics. *Genetics*, 157:785–802, 2001.
- [8] A.S. Gilchrist, R.B.R. Azevedo, L. Partridge, and P. O’Higgins. Adaptation and constraint in the evolution of drosophila melanogaster wing shape. *Evolution & Development*, 2:114–124, 2000.
- [9] C.P. Klingenberg, A.V. Badyaev, S.M. Sowry, and N.J. Beckwith. Inferring developmental modularity from morphological integration: analysis of individual variation and asymmetry in bumblebee wings. *American Naturalist*, 157:11–23, 2001.
- [10] C.P. Klingenberg, L.J. Leamy, and J.M. Cheverud. Integration and modularity of quantitative trait locus effects on geometric shape in the mouse mandible. *Genetics*, 166:1909–1921, 2004.
- [11] Klingenberg’s Lab. www.flywings.org.uk.
- [12] C.P. Klingenberg. Morphometrics and the role of phenotype in studies of the evolution of developmental mechanisms. *Gene*, 287:3–10, 2002.
- [13] D. Houle, J.G. Mezey, P. Galpern, and A. Carter. Automated measurement of drosophila wings. *BMC Evolutionary Biology*, 3(25):1471–2148, 2003.
- [14] D. Tharavy. Personal communication. *Paris National Museum of Natural History*.
- [15] P.J.D. Weeks. Species-identification of wasps using principal component associative memories. *Image and Vision Computing*, 17:861–866, 1999.
- [16] S. Palaniswamy, N.A. Thacker, and C.P. Klingenberg. A statistical approach to feature detection in digital images. *Leeds Annual Statistical Research Workshop*, pages 146–149, 2006.
- [17] A.C. Evans, N.A. Thacker, and J.E.W. Mayhew. The use of geometric histograms for model-based object recognition. *Proceedings of the 4th British Machine Vision Conference*, pages 429–438, 1993.
- [18] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern analysis and Machine Intelligence*, 8(6):679–698, 1986.
- [19] S. Palaniswamy, N.A. Thacker, and C.P. Klingenberg. A statistical framework for detection of connected features. www.tina-vision.net/docs.php, 2006.
- [20] A.P. Ashbrook, N.A. Thacker, P.I. Rockett, and C.I. Brown. Robust recognition of scaled shapes using pairwise geometric histograms. www.tina-vision.net/docs.php, 1996.