

Tina Memo No. 2006-007

Initial ideas presented at NCAF Meeting, Bristol, Sept. 2006.

Presented at CRV, Windsor, Ontario, Canada, 2008.

Quantitative Verification of Projected Views Using a Power Law Model of Feature Detection

Simon Coupe and Neil Thacker

Last updated
18 / 05 / 2008

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

- 2002-005 Tutorial: An Empirical Design Methodology for the Construction of Machine Vision Systems.
- 1996-001 Algorithmic Modelling for Performance Evaluation.
- 2006-002 A Statistical Framework for Detection of Connected Features.
- 2006-007 Quantitative Verification of Projected Views Using a Power Law Model of Feature Detection.
- 1997-003 Tutorial: Supervised Neural Networks in Machine Vision.
- 1995-003 Invariance Network Architecture.
- 1992-001 Combining the Opinions of Several Early Vision Modules using a Multi-Layer Perceptron.
- 2005-006 Curve Fitting and Image Potentials: A Unification within the Likelihood Framework.
- 1996-002 Tutorial: The Likelihood Interpretation of the Kalman Filter.
- 1994-003 Using a Switchable Model Kalman Filter.
- 2004-012 Tutorial: Computing 2D and 3D Optical Flow.
- 2005-011 Comparing the Performance of Least-Squares Estimators: when is GTLS Better than LS?
- 1994-001 Tutorial: Overview of Stereo Matching Research.
- 1992-002 Online Stereo Camera Calibration.
- 1995-002 Calibrating a 4 DOF Stereo Head.
- 2000-009 An Evaluation of the Performance of RANSAC Algorithms for Stereo Camera Calibration.
- 2001-011 The Evolution of the TINA Stereo Vision Sub-System.
- 2007-011 A Methodology for Constructing View-Dependent Wireframe Models.



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

Quantitative Verification of Projected Views Using a Power Law Model of Feature Detection

Simon Coupe and Neil Thacker

Abstract

We observe that conventional approaches to the construction of Likelihood models of visual appearance for image features are non-quantitative, precluding their use in tasks such as hypothesis testing for projected view validation. This document outlines a quantitative approach for verification of 3D objects' predicted edge features in images, which incorporates both the effects of image noise and local image structure. This approach supports the construction of a joint probability for the degree of conformity of image data to both edge orientation and location, without the need for arbitrary relative scale factors. The method has been validated on multiple views of man-made objects constructed from a variety of materials.

Introduction

Edge features are widely used throughout the computer vision community for representing and recognising objects in images. The projected edge-defined shapes of objects can be seen as compact and powerful shape descriptors, conveniently offering a high degree of invariance to environmental illumination conditions. Furthermore, edge representations are particularly useful at object boundaries, where the potential for arbitrary backgrounds in 3D scenes invalidates any approach based upon direct modeling of the grey-levels (i.e. across the boundary). In addition, quantitative spatial information is concentrated at edges, with surfaces offering little if any information to help determine and localise 3D scene structure. The ease with which humans can recognise objects from simple 2D line drawings provides anecdotal evidence for the potential significance of edge data for such tasks. Most machine vision researchers would be pleased to realise even a fraction of these recognition and interpretation capabilities on a computer.

Many techniques have been proposed with which to detect edge-based shapes in images and the prevalent attitude in computer vision is that this topic has largely ran its course. Consequently, the process of edge detection no longer represents a popular topic for conferences and journals. However, although some approaches (such as Canny) are popular, there is little consensus as to how best analyse image edge data in support of general scene interpretation tasks. In addition, there are significant failings with popular methods with regards to statistical and quantitative rigour. In particular, Likelihood models for detection and localisation are generally assumed, rather than derived from any process of image formation. In addition, we see no effort made to confirm the assumed distributions with real data [13]. Specific methods aside, the object detection process inherently requires a metric to evaluate the quality of any hypothesised model matches in order to perform the critical task of discounting any false interpretations of scene content. The lack of a quantitative statistical basis for edge detection makes this task unreliable, if not impossible, in the general case. The work presented here formulates a statistical framework with which to assess the quality of match between a model's predicted location and orientation and the underlying image evidence. The novelty of the work centres around a bootstrap approach to the definition of edge detection and localisation, together with quantitative testing of the assumed distributions. This approach supports combination of localisation and orientation information in validation tasks without the introduction of arbitrary weightings. The research was undertaken in support of development of the TINA model matching computer vision system [14].

Background

Edge-based model matching has been of interest to computer vision researchers for many years [10], being popularly characterised by the 'hypothesise and verify' paradigm. It is commonly noted that the processes involved are susceptible to a high degree of uncertainty, due to factors such as sensor noise and low resolution pixelation. Accordingly, popular approaches to hypothesis verification are typically concerned with 'bounded error' models, where only loose constraints on the required degree of match between predicted and observed edge features are

utilised for verification. Since the problem of reliably and precisely locating arbitrary 3D objects itself remains unsolved (with detectable features rarely projecting exactly to the locations expected [7]), this has seemingly not been much of an issue to the computer vision community to date.

Many popular verification techniques are associated with the process of template matching, where some form of similarity measure between a specified template and a particular image edge region is minimised [8]. Optimisation of such a function can be utilised to best align the predicted and observed edge patterns, with the resultant difference in grey-levels indicating how closely the hypothesised model matches the image. However, detailed modeling of grey-level data has problems at object boundaries, where the edge defines the start of an arbitrary background which is by definition unknowable.

Also, verification has been based upon image plane distances. Thresholds are used to determine supposed valid instances of sought shapes. Chamfer and Hausdorff matching are perhaps the most prominent of these techniques, where the mean and largest individual distances between edge-point sets are respectively minimised [2, 6]. Otherwise, each predicted feature point can be validated by simply being on or near (to an estimated bound) a detected (binary) image edge point, where thresholds on the proportion of visible points can similarly be used to verify the validity of any hypothesised matches [1]. Although these methods provide a tuneable statistical system under restricted circumstances, they are not directly applicable to random images of objects, as the residuals from the predicted model values are not derived from a valid model of image formation.

The orientations of models' predicted edge points also convey a significant amount of information regarding model-match quality, as has been noted in a number of publications [8] [4]. Without such constraints, distance transform techniques, for example, are prone to registering significant Likelihood scores across their features in busy or highly textured images, where each predicted point would be close to an image edge. By ensuring that feature matches are only counted if at admissible orientations, any incorrect, coincidental feature match hypotheses can be more readily rejected. This is especially important when considering the typical flawed nature of objects' image-edge maps, where, as will be discussed, many of an object's features may not be detectable. Hausdorff methods have been adapted to account for edge orientation information, where a cost for disparate orientations is weighted into the distance transform function [8]. Similarly, truncation techniques have been developed where only a reduced set of the best point matches are considered [5]. It should however be noted that assessments of edge orientation are commonly dealt with without any account of propagated uncertainty [11, 9]. For weaker edges, such errors may be significant, thus requiring appropriate consideration.

Whilst the techniques described above do intuitively have some operational merit, i.e. accounting for whether each predicted edge point is supported to some degree by the image evidence, the inherent vagueness of each method can be seen to be a drawback in terms of precise and reliable quantitative verification. In summary, the similarity distributions observed for both appearance and geometric approaches will vary as an unknown function of the illumination, the object and its orientation.

The work in this paper is concerned with analysing how well each predicted edge feature point is supported by the underlying image edge evidence. By basing our statistical distributions directly upon the measured evidence we can attempt to construct quantitative tests which follow directly from our understanding of image formation and feature detection. The methods are developed to quantitatively embody any degrees of uncertainty, such as the reliability of detection and errors encountered in our analysis of edge orientation. This is proposed as being the only theoretically justifiable way to perform such a task.

There are several important benefits expected for this approach. First, by taking explicit account of the original image data it is expected to be more discriminatory than distance-based measures while being less dependent upon the details of image formation than appearance-based approaches. Secondly, the method is quantitatively testable. We are able to verify that our assumed probabilistic metrics are indeed truly reflective of the behaviour of sampled edge data for real-world examples of imaged objects. Thirdly, the theory gives an absolute calibration (i.e. relative weighting) for the orientation and localisation terms used in Likelihood construction. The difficulties in attempting this approach are however significant. A successful theory must deal with the effects of arbitrary illumination and changes of physical appearance under projection while defining calculations which are computable using minimal knowledge of expected image content, specifically; predicted edge location, orientation and local image data.

Methods

The proposed approach is based upon the premise that a faithfully representative wireframe model has been accurately projected onto the corresponding image data. This is done so that we can concentrate on the statistical distributions associated with detectable edges. A process of lateral feature shifting has been adopted in this work to ensure that any edge features are optimally aligned with any underlying image edge evidence. This process accounts

for the systematic shifting of object boundaries due to illumination which often occurs in real images of man-made objects [7]. This work aims to formulate probabilistic terms which directly account for the quality of match between oriented edge feature points and their hypothesised image positions. Particular emphasis is placed on validating our assumed statistical models by getting good agreement between the expected and observed Likelihood distributions. Given such agreement, we have reliable, well-defined predictions for the expected behaviour of our assumed Likelihood terms for well-aligned edge-based object models. A process of statistical hypothesis testing can then be implemented to indicate the probability that a specified model point projects onto a valid oriented edge point. This evidence can then be accumulated across object models, at appropriate intervals, to support specific curve and whole model match verification.

It is acknowledged that for real-world object detection scenarios, many of the features predicted by a wireframe model may not be visible even for un-occluded objects. This can be due to, for example, specular highlights or lack of contrast across similarly illuminated fore- and backgrounds. This places extra emphasis on the integrity of any feature detection and verification processes and dictates that, at best, we will only ever be able to detect a restricted set of any predicted edge feature points. This process of course assumes that our wireframe models also faithfully represent any occluded features across changes in viewpoint. In this work, wireframe models are accompanied by feature visibility files which detail any viewpoint dependencies. We take this approach rather than using CAD like surface models as we believe it is more fundamentally suited to the learning of visible structure from arbitrary sequences of images. In particular, although it is possible to reliably infer the presence of a surface boundary using edges, it is not possible to reliably infer the presence of a surface from an area of smoothly varying grey-levels. Arbitrary illumination also precludes the possibility of making reliable estimates of 3D surface shape, thereby inhibiting the kinds of update process which would be necessary to learn a surface from multiple views. We envisage therefore that the first stages of object learning will involve modeling edge structure, with surfaces only being included in the representation once the view-based feature representation is already sufficient to recognise and localise an object.

Edge Orientation Analysis

The true distribution for estimation of the local gradient direction from separate x and y image derivatives has been described as a Von Mises distribution [15]. This is expected to be more detailed a model of the distribution than is needed for most practical purposes (i.e. away from low contrast regions). A first order (Gaussian) approximation can be derived using error propagation and tested empirically on real data. The local image orientation can be defined as

$$\psi = \arctan\left(\frac{dI/dx}{dI/dy}\right) = \arctan\left(\frac{u}{v}\right)$$

Given equal errors δ on the derivatives u and v

$$\delta_\psi^2 = \left(\frac{v\delta}{v^2 + u^2}\right)^2 + \left(\frac{u\delta}{v^2 + u^2}\right)^2$$

Since the variance is the square of the error, and taking 'r' to be the edge magnitude

$$\text{var}(\psi) \approx \frac{\delta^2}{r^2}$$

i.e. the error on the local edge orientation is inversely proportional to the edge strength r and proportional to the image noise δ . The edge strength can thus be considered as the Fisher information for calculation of image plane orientation [12] (i.e. the best locations to apply an orientation test).

The orientation term can be selected in order to match the error distribution on orientation measurement ($\text{var}(\psi)$)

$$p(\psi|x, y) = \frac{1}{\sqrt{2\pi \text{var}(\psi)}} \exp\left(-\frac{(\phi(x, y) - \psi(x, y))^2}{2 \text{var}(\psi)}\right)$$

with $\phi(x, y)$ representing the edge orientation of a feature pixel. This is our Likelihood for feature orientation (L_{angle}).

Given such a prediction of the distribution (Gaussian) of edge orientation residuals for well aligned edge features, we are able to perform quantitative hypothesis tests (H_{angle}) to indicate the probability that a specified feature point represents a valid instance of an edge match using the error function (erf()).

Edge Location Analysis

The Likelihood for edge orientation is fully quantitative, as we have a model for the way in which it is affected by image noise. The generation of a Likelihood term which corresponds to feature localisation is more problematic. It is impossible to predict the distributions of grey level data associated with an object viewed under arbitrary illumination. We therefore define a quantitative probability relating to the localisation of edge structure, based upon our understanding of the edge detection process combined with a bootstrap resampling of local image data.

The edges which we are most interested in for object detection are continuous step edges which define the contours of objects. By definition, a pixel can therefore be deemed to contain such an edge feature if its gradient value (edge strength) is both above the noise floor and greater than M of its 8 immediate neighbours. We can therefore assess the probability of any location being defined as an edge if we are prepared to sample the local distributions of gradients in image data.

The probability that a given pixel gradient value g is greater than a specific value x is defined by

$$P(g > x) = \int_{-\infty}^x p(g, y) dy$$

Where $p(g,y)$ is the probability density as a function of gradient value y , i.e. the noise distribution. Taking a bootstrap approach based upon the local region, the probability that a particular pixel would be greater than a threshold t and a sample $(x_0 \dots x_N)$ drawn from the distribution of N local pixels can therefore be defined as

$$\begin{aligned} R(g) &= \frac{1}{N} \sum_i^N P(g > x_i, g > t) \\ &= \frac{1}{N} \sum_i^N P(g > x_i) \int_t^{\infty} p(g, y) dy \end{aligned}$$

which can be regarded as the ‘soft rank’ of the central pixel in comparison to its neighbours¹. As a ranking process is a form of histogram equalisation, this must generate a uniform distribution in R where the local gradient calculation is above the specified threshold.

In order to define the probability of detecting an edge we need a definition for the sampling process which generates edge data. If we define this process as involving some number of a random selection of gradients being less than the central value then the probability that we will detect a feature at g will be a polynomial function of R (Appendix A). We can consider this, equivalently, as a random local spatial reorganisation of the data which was used to compute R . We approximate an edge detection process by defining the probability P_{max} that the central pixel would be larger than M neighbours², so that $P_{max} \approx R^M$. Further, as the distribution over R is uniform, application of Bayes theorem tells us that $P_{max}(R) = P(edge|R)$ is also proportional to the Likelihood of getting a particular value of R ($P(R|edge)$).

The hypothesis probability that the observed pixel is drawn from the sample of edges is obtained by integrating the Likelihood function so that

$$H_{edge} = P_{max}^{M+1/M} = R^{M+1}$$

Now, for a quantitatively valid hypothesis probability (i.e. in the situation where all of our assumptions hold and our definition of edge probability is correct), the histogram of H_{edge} must be uniform. Under these circumstances the distribution of $\log(H_{edge})$ will be exponential. Therefore, it follows that for data which is sampled according to the bootstrap model, the distribution of $\log(R)$ must be an exponential scaled by a factor of $M + 1$. We leave M as an empirical parameter to be determined from sample data. The consistency of this value across multiple data sets thereby provides information concerning the general validity and consistency of the assumed model. Determination of this factor from sample data, allows the calculation of a quantitative hypothesis test ($H_{edge} = R^{M+1}$) and the appropriate normalisation for the log Likelihood term ($-\log(L_{edge}) = -M \log(R)$) for use in object localisation. This is therefore essentially a **power law** model of statistical feature detection.

Hypothesis Testing of Edge Location and Orientation

Hypothesis testing is a fully quantitative method to indicate how well sample data conforms to a predefined model. In this case, our statistical models relate to edge location and orientation, which are expected to correspond to (log-

¹This is an approximation, see Tina memo 2009-008 for a derivation using ‘equal variance’ Likelihood.

²This simplification has a biological analogy, in that it is consistent with theories of neuronal function such as those developed by Grossberg. These utilise a ‘winner-take-all’ process, mediated by lateral inhibition.




Figure 1: Cumulative distributions of sampled data for edge location (a) and orientation (b), overlaid with the theoretically predicted distributions (solid curves).

) exponential and Gaussian distributions respectively. Probabilities resulting from hypothesis tests are expected to be uniformly distributed. In order to validly combine the terms for orientation and location, the following renormalisation formula can be derived [3]. Given n quantities, each having a uniform probability distribution $p_{i=1,n}$, the product $p = \prod_{i=1}^n p_i$ can be renormalised to have a uniform probability distribution $F_n(p)$ using

$$F_n(p) = p \sum_{i=0}^{n-1} \frac{(-\ln p)^i}{i!}$$

i.e.

$$H_{angle,edge} = (H_{edge}H_{angle})(1 - \ln(H_{edge}H_{angle}))$$

where H_{edge} and H_{angle} are the hypothesis probabilities for edge location and orientation respectively, estimated as shown above.

This metric can be evaluated for any sampled edge feature point relative to a model feature prediction to reflect the probability that that point was indeed representative of a valid model match. This evidence can then be accumulated across sets of representative curves to support verification of detected object models in images. For instance, thresholds can be set against the proportion of model features passing the hypothesis test at a specified confidence limit. Since many edge features may not be present due to unpredictable environmental illumination conditions, we define a feature to be present if the majority of its individual features are detected. This in turn implies that our predictions of appearance need only include curved structures when at least 50% of their length is expected to be visible. This allows for a fair degree of freedom in the initial construction of view-based models, which is important if we expect the system to refine model parameters (i.e. learn) upon repeated presentation of an object.

Results

The key assumptions that we are interested in testing here, relate to our assumed terms for edge localisation and orientation. In particular, theoretical and observed distributions should agree if the statistical formulations are valid.

A selection of 8 objects made from a variety of materials using a number of fabrication processes have been imaged in stereo from multiple viewpoints and 3D wireframe models with associated view dependency files were constructed. These were then approximately aligned using stereoscopic cues and final alignment performed (during what is effectively a re-calibration process) (Figure 2).

The cumulative sampled distributions for the data were plotted for each object and the overall scaling of each distribution (M) determined - to bring each into best quantitative agreement with the theoretical cumulative distribution curves (Figure 1, a and b). The scaling factors needed to achieve this level of agreement are shown in Table 1.

The results represent the average level of conformity between our statistical model and the data across all curves in each object. It confirms that we have reasonable agreement between the theoretical and assumed location and orientation distributions following object specific scaling. Although the scaling factors are significant they are

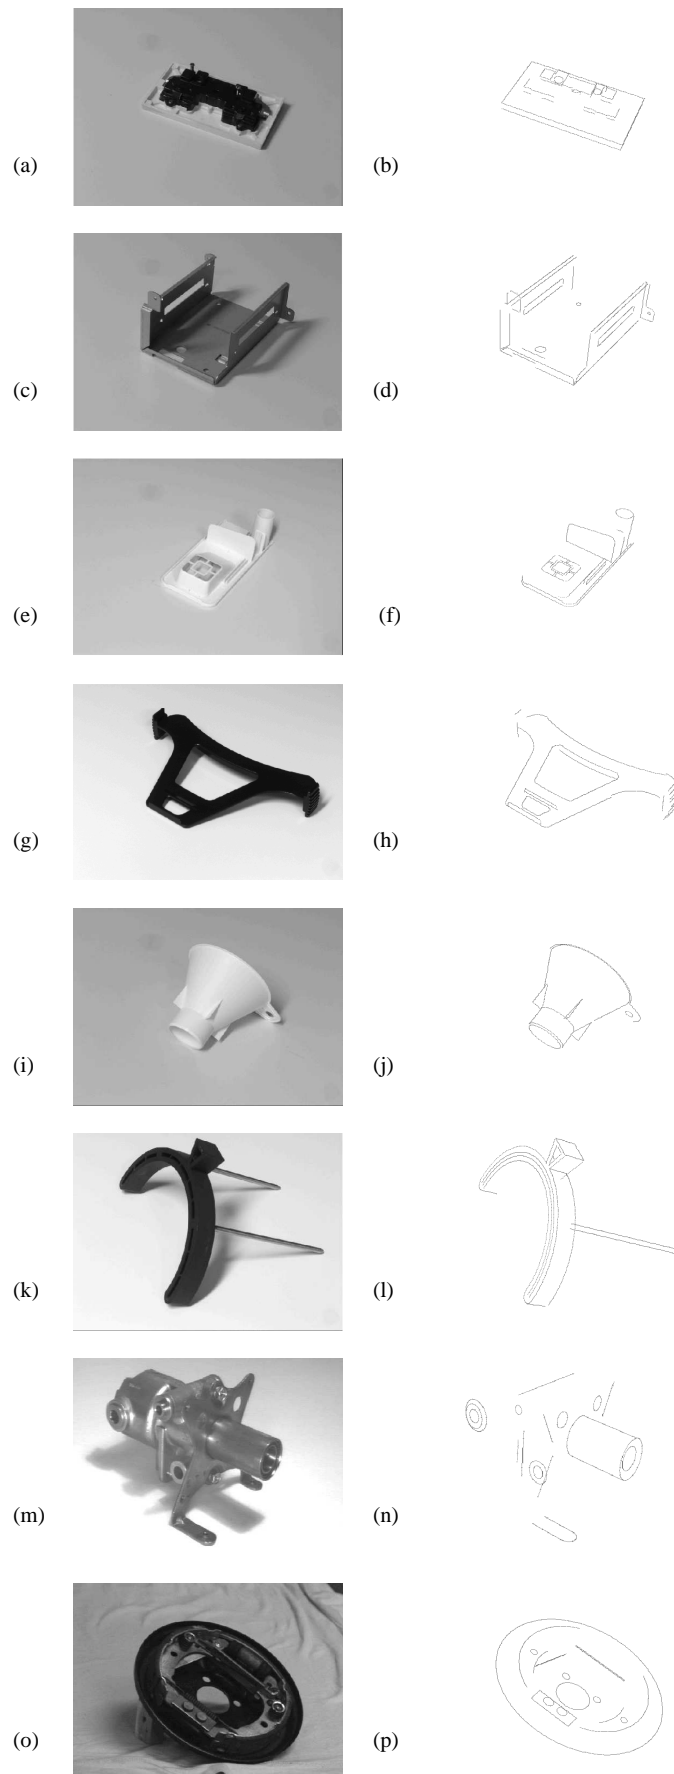


Figure 2: *Test objects and the features reprojected from the corresponding view-based 3D models.*

Object (as Fig. 2)	Edge Scale	Orient. Scale	Edge D-value	Orient. D-value
(a) (b)	5.6	1.8	0.15	0.05
(c) (d)	6.0	1.4	0.15	0.05
(e) (f)	3.2	1.4	0.09	0.07
(g) (h)	4.8	2.2	0.14	0.04
(i) (j)	3.6	1.4	0.09	0.11
(k) (l)	4.0	1.9	0.11	0.04
(n) (m)	3.6	1.3	0.09	0.07
(o) (p)	4.0	1.7	0.09	0.05

Table 1: *Relative scalings of the cost function terms required to bring the observed data distributions into best agreement with the assumed distributions, accompanying the maximum distribution-separation terms (D-values) required for the Kolmogorov-Smirnov test.*

reasonably consistent across the objects. The average scaling levels can thus be factored into our calculations. We believe that this scaling is predominantly due to contamination from non-edge pixels in the bootstrap estimates of the Likelihood terms. Knowledge of these values allows the Likelihood terms for edge location and orientation to be appropriately weighted during object localisation, so eliminating what would otherwise be an arbitrary scaling parameter in the similarity measure.

Given a known number of sample points, these plots can be used to predict the conformity of the model to the data via the Kolmogorov-Smirnov test. For example, for a (maximum vertical) difference between curves of 0.1, 50 sample points along a simple curve will produce a significance value of 0.13 which is insufficient to exclude this distribution as a description of the data. As the number of samples on an extended curve increases however, the strength of the statistical test will also increase until it eventually fails. Clearly, at this point, the approximations of the theory to the sample distributions will begin to have a detrimental effect on the performance of the alignment. Put another way, there is little point in sampling so much data along a curve that we go beyond the ability of the statistical method to make good use of it, but until that point, the simple model is a sufficient approximation of the feature detection process.

Since our assumed Likelihood models can be seen to closely reflect those observed, we can be confident that an associated hypothesis test for a sampled edge feature point will give a good indication of how likely it is that that point represents a valid match. Although object specific scaling may be required to get the most accurate interpretations of verification, factoring of an average weighting term across all objects still supports reliable verification interpretations. This can be demonstrated by projecting a full wireframe model (including occluded features) onto a corresponding image and by checking that each visible feature is accounted for and vice versa. This is demonstrated in Figure 3. This process is utilised by the TINA vision system to aid feature detection for the automated construction of view dependency files. The criteria utilised for this experiment are that a feature is displayed as being present if at least 45% of its related pixels pass a combined hypothesis test for location and matching orientation with a confidence level of 0.01%. The confidence level allows virtually any suspected edge feature to be passed, while the 45% threshold accounts for feature degradation due to factors such as partial occlusion or specular interference. The features used for modeling are straight line segments and circular features that are split along their major projected axes. As mentioned, prior to verification, the position of each such feature is locally optimised using a lateral feature shifting procedure, so that the distribution of feature Likelihoods can be sampled at the location determined by the image evidence. As can be seen, just about every expected feature is then correctly identified and any missing features can be explained away by illumination artifacts. None of the occluded features were flagged as being present with these thresholds. Similar results were achieved across the models featured in Figure 1.

The proposed verification metric has been successfully implemented in the TINA model matching computer vision system to reliably select the best of any competing match hypotheses in terms of the proportion of predicted edge feature points passing the hypothesis test at a specified confidence limit. The validity of any leading hypotheses can be verified by demanding that the majority of features are accounted for.

Conclusions

This document has outlined an approach for the statistical testing of the predicted locations of objects' defining edge features in images. Quantitative metrics have been formulated to account for both the matching locations and orientations of such projected edge features. These metrics have been validated by checking that data distributions sampled from real data correspond to those theoretically predicted. The mean value for M found in this paper (4.85)

is consistent with the value (4.6) which best approximates H_e (Appendix A). A combined statistical hypothesis test for corresponding edge location and orientation has been implemented to indicate the probability that a sampled data point represents a valid instance of an edge feature match. Accumulation of this edge point-based information across sets of representative curves has been shown to support reliable verification of the validity of any wireframe model match hypotheses.




Figure 3: Figure (b) represents the full 3D wireframe model used to represent the object depicted in (a). Figure (c) shows those features verified as being present in (a) according to the verification strategy outlined in the Results section.

References

- [1] M. Boshra, H. Zhang. Accomodating Uncertainty in Pixel-Based Verification of 3-D Object Hypotheses. *Pattern Recognition Letters*, 20(7):689-698, 1999.
- [2] G. Borgefors. Hierarchical Chamfer Matching: A Parametric Edge Matching Algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(6):849-865, 1988.
- [3] P. A. Bromiley, M. L. J. Scott, M. Pokrić, A. J. Lacey, N. A. Thacker, Bayesian and Non-Bayesian Probabilistic Models for Magnetic Resonance Image Analysis, *Image and Vision Computing, Special Edition; The use of Probabilistic Models in Computer Vision.*, 21:851-864, 2003.
- [4] T. F. Cootes, C. J. Taylor. On Representing Edge Structure for Model Matching. *Proceedings of IEEE Computer Vision and Pattern Recognition Conference*, 1:1114-1119, 2001.
- [5] D. M. Gavrilu. Multi-Feature Hierarchical Template Matching Using Distance Transforms. *Proceedings of the 14th International Conference on Pattern Recognition*, 1:439-444, 1998.
- [6] D. P. Huttenlocher, G. A. Klanderman, W. A. Rucklidge. Comparing Images Using the Hausdorff Distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850-863, 1993.
- [7] A. Lacey, N. A. Thacker, P. Courtney, S. Pollard. TINA 2001: The Closed Loop 3D Model Matcher. *Proc. BMVC 2001*, 203-212, 2001.
- [8] C. F. Olson, D. P. Huttenlocher. Automatic Target Recognition by Matching Oriented Edge Pixels. *IEEE Transactions on Image Processing*, 6(1):103-113, 1997.

- [9] C. F. Olson, D. P. Huttenlocher. Recognition by Matching Dense, Oriented Edge Pixels. *Proceedings of the International Conference on Computer Vision* , 91-96, 1995.
- [10] L. G. Roberts. Machine Perception of Three Dimensional Solids. *PhD Thesis*, MIT Department of Electrical Engineering, 1963.
- [11] C. A. Rothwell, A. Zisserman, J. L. Mundy, D. A. Forsyth. Efficient Model Library Access by Projectively Invariant Indexing Functions. *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*, 109-114, 1992.
- [12] N. A. Thacker, E. C. Leek. Retinal Sampling Feature Detection and Saccades; A Statistical Perspective. *Proc. BMVC 2007*, 990-1000, 2007.
- [13] N. A. Thacker, A. Clark, J. Barron, R. Beveridge, P. Courtney, W. Crum, V. Ramesh, C. Clark. Performance Characterisation in Computer Vision: A Guide to Best Practices. *IJCV*, 109:305-334, 2008.
- [14] TINA - Open Source Image Analysis Environment. www.TINA-Vision.net
- [15] V. Ramesh. Performance Characterisation of Image Understanding Algorithms. *PhD Thesis*, University of Washington, 1995.

Thanks to Dr Visvanathan Ramesh from Siemens Corporate Research (USA) and the EPSRC for project funding.