

Retinal Sampling, Feature Detection and Saccades; A Statistical Perspective.

Neil Thacker and Charles Leek.

Last updated
23 / 6 / 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.

Retinal Sampling, Feature Detection and Saccades; A Statistical Perspective.

N.A.Thacker and E.C.Leek. 23/6/2007

Abstract

This paper applies statistical design principles to a simple biological model of human vision so that we can more clearly interpret the apparent role of eye saccades. In doing so we show that many structural features of the biological system (such as the optical geometry of the retina) are not complications which need to be surmounted or accomodated by later processing, but very good strategies for minimising the resources required to construct a working image recognition system. Therefore, (and as expected) the structure we see can very easily be explained by evolutionary pressures. The work also suggests specific forms for intermediate computed quantities, chosen to match the representation of signals as pulse sequences in neuronal systems. The ideas presented have implications for the construction of artificial (computer) vision systems. The computational model is very closely related to (but not based upon) SIFT, but more strongly based on a consideration of vision as a process of measurement while also linking the idea of multi-scale analysis with biological structure.

Introduction

It is generally assumed that a primary function of eye movements (saccades) is to maximize information processing within the high resolution region of the fovea [20, 21]. The measurement of saccades has been widely used in psychophysical studies in domains such as visual search, reading, scene exploration and interpretation and 2-D pattern recognition [7, 3, 15], and in machine vision studies, for example, of knowledge-based scene analysis and human interactions with complex displays [4, 8, 10]. These studies have shown that eye movements can be influenced by a variety of factors, including low- level image statistics, prior knowledge and task requirements. Surprisingly, little is known about the relationship between eye movements and three-dimensional object recognition. The problem is of course complicated by the difficulty of understanding exactly what data is provided by the visual process for the task.

In this paper we examine three fundamental questions. First, can fixation data reveal preferences for specific types of image features? Second, are gaze preferences consistent between stimulus encoding and recognition? Third, are extracted features invariant across tasks and across changes in 3-D viewpoint? In order to interpret the results of this study we must first define the data we believe to be available from the retina. The first part of this paper therefore makes an argument for a simplified interpretation of biological function. The argument has been made that the only scientifically valid models of biological vision are those that identify the statistical principles which under-pin the analysis of image data. In particular the consideration of data as measurement and appropriate treatment of measurement uncertainly appear to be a fundamental requirement. We can also arrive at the same conclusion from the perspective of good engineering via application of a system design methodology. In this process individual computational modules are designed based upon statistical assumptions which match the properties of input data.

The process of human vision starts with detection of light in the retina and ends with semantic interpretations of scene content in the brain. To say that processes which operate in between are known to be very complicated would be an understatement. This document is necessarily light on details. Its purpose is to provide a computational framework for understanding the basic structure. This will be done so that the various statistical characteristics associated with measurement and representation are addressed, as they would be when applying a system design methodology. As we will see this approach places additional constraints on the possible biological model to those generally considered. In some cases this provides us with specific predictions of what needs to be computed and how.

We argue here that the requirement to construct invariant quantites for recognition, combined with the statistical nature of the data, in particular measurement stability and consideration of a quantitative understanding of the information present in the data, allows us to arrive at an hypothesis for a simplified equivalent form. Specifically a set of homogenous regional samples from an exponentially distributed set of scales.

There are several structural properties of the retina which are well known. In particular the retina is known to have spherical geometry with logarithmic sensitivity to intensity. It is also considered reasonable to assume (as we will do) that the input data has radially varying spatial resolution which is greatest in the center. The evidence for this comes primarily from observations of the structure of the striate cortex and secondary visual cortex [14] (cortical

magnification). This would appear to be in conflict with the known structure of the retina, where although there is such a distribution of sensors, they come in two specific types (the density of illumination sensors (rods) reduces to zero in the fovea, while the density of colour sensors (cones) are at their greatest). What we need to remember however, is that it is in principle possible to synthesise an intensity measurement from a combination of colour ones. Such a solution is clearly justifiable if the system has been evolved to minimise sensor resources and cones are somehow more resource intensive than rods. In other words, cones can perform the task of the missing rods, so we don't need to have both in the fovea. The implication is that evolution has found some value for having a high density of cones in the fovea, which may itself be set to provide a scale invariant system, with a different scale dependency. If this is the case then we may also expect this data to project to two different parts of the brain (for example the striate cortex and secondary visual cortex), as two scale invariant mappings imply two distinct computational tasks. The rod distribution is then explainable as providing the additional illumination measurements required to achieve scale invariant shape recognition.

This interpretation implies is that the structures found in the cortex are actually nothing more than a consequence of smoothly mapping the data required for scale invariant recognition onto a 2D surface, and that the spatial arrangement of these structures, in themselves, need play no important part in the recognition process. That is, that once adapted for use, the brain makes no computational use of the fact that two neurons are separated by a specific *physical* distance.

Here we attempt to describe the origins of radially varying density and retinal geometry as useful computational devices and compare this interpretation to the artificial sensors (eg: CCD arrays) generally used for computer vision. As the retina can provide only one set of image samples at any one time, building an internal representation of the world around us requires this sensor to be moved around the scene. This process is supported in the human vision system by the saccades. In this paper we provide evidence which excludes some of the less likely generators of saccadic eye movement.

Computing Invariant Quantities

In direct contrast to conventional electronic sensors (eg: a CCD array), the human retina is observed to have an irregular lattice of light sensors at spatially varying scales. This structure is often approximated as a log-polar detector such that the sensor resolution decreases as a function of radius. In order to understand why a detector such as the retina may have evolved it is necessary to consider the fundamental problems associated with visual recognition. This process is expected to terminate with pattern association, whereby a set of relationships constructed from the input visual data is matched to a pattern stored in visual memory. Much of the process of visual recognition is understood to be template matching. This is a simple idea which is difficult to make work effectively in practice. Such a process is potentially very memory intensive unless the spatio-temporal relationships are constructed in a way which eliminates unnecessary variation in the input image, such as illumination¹, rotation and scale. This is often referred to as construction of an **invariant representation**. Many researchers in computer vision would also like to build systems with invariance to perspective changes, induced by 3D sensor or object motion. Often this can be locally approximated for small rotations by affine invariance (invariance to linear skew).

Yet combination of measured values into invariant quantities solves only part of the problem. We must also consider the associated noise characteristics. Scale and illumination invariance provide specific examples of this problem. For a CCD array the simplest way to compute illumination invariant quantities q from measured intensities $h_i \propto I + N(\sigma)$ at location i , is to compute a ratio such as $q = h_i/h_j$. Distributions of these quantities can then be treated as patterns for template matching. However, using error propagation it is possible to show that the noise characteristics of q will then be $var(q) = \sigma^2(1/h_j^2 + h_i^2/h_j^4)$. This introduces spatially varying errors on the computed invariant quantities, which then have to be properly addressed during pattern matching, by for example committing extra memory resources. The simplest way to deal with this issue is to make the invariant quantities have homogenous error by modifying the measurement process, ie: $g_i \propto \log(I) + N(\sigma)$, as approximated in the human vision system. Then $q = g_i - g_j$ and $var(q) = 2\sigma^2$. One way to view this is to say, if we cannot deal with the variability introduced into invariant quantities by measurement noise then the extra information potentially gained by having a uniform sensor (h) provides no additional benefit to the alternative measurement system (g). Thus logarithmic sensitivity to light would appear to be a sensible strategy for an illumination invariant recognition system and one likely to be found in a system which has been optimised to minimise computational resources.

Another difference between an electronic sensor and the human eye is that the brain encodes data as pulses. As an additional consideration, we must therefore also match the measurement (noise) characteristics to this process. Assuming that the statistical character of a mean frequency f derived from a pulse sequence is Poisson ($var(f) \propto f$), we can represent a signal with uniform noise so that it has these characteristics by taking the square,

¹I use this term here to refer to a simple linear scaling of overall intensity, and not the more general illumination variation which can occur in real scenes.

ie;

$$f_g = g^2 \quad \text{and} \quad f_q = q^2$$

To do otherwise risks either losing information present in the original signal or expending resource representing the signal to an accuracy beyond that necessary to support the information. Thus consideration of the required invariance plus the noise characteristics tells us not only what needs to be computed but how it must be represented. Similar mechanisms are also described below for the matching of stored templates.

We will turn now to the issue of constructing scale invariant geometric representations of shape from data measured on the retina. As with the illumination example provided above, potentially a relationship constructed from features detected at one scale can have different statistical reliability (repeatability error) to the same feature constructed from data at a higher scale. In the limit we can scale an object by such a factor that an object projects entirely onto one sensor. Although this is an extreme example, it is an illustration that data from a fixed sized sensor has a finite limit of information. Though the mean of any invariant quantity might remain fixed, the variance around that mean will change as a function of image scale. This will prevent reliable and efficient scale invariant recognition, which we might describe as **variable scale sensitivity**. If we wish to perform recognition with a pattern matching approach, simply constructing invariant geometric quantities (or the kinds of approximate invariants often found in computer vision algorithms) is therefore not enough, as we need also to take account of the scale varying error on computed relationships. To ignore this problem can result only in a system which makes inefficient use of the available data.

The physical structure of the eye might appear unnecessarily complicated in comparison to a simple colour CCD array. To begin with the retina is a curved (approximately hemi-spherical) surface whereas a CCD array is flat. In fact having a flat sensor causes some problems with image similarity. The optical model for a conventional electronic device is close to a pin-hole model. The geometric imaging process is described as perspective projection. Under such a model objects viewed at the centre of the field of view appear different to those at the edge. A spherical imaging surface on the other hand is rotation invariant and will produce an equivalent focused image of an object for any position in the field of view. This can be considered a simple form of perspective invariance. However, such a property pre-supposes a uniform sampling of the image, which is manifestly not the case for the human vision system. Another complicating factor is that our perception of the world must ultimately be mapped into a single consistent co-ordinate system. It has been suggested that there exist specific topological maps which represent locations of objects in both a “nose centered” and “body centered” framework, and are constructed with help of biological processes such as the vestibular ocular reflex. Little more will be said on this subject here, except to say that a logarithmic radial sensor potentially also complicates any process of mapping detected features into such maps. It would certainly be awkward trying to construct and maintain a single consistent “world” image from multiple views of a scene with locally varying resolution. Both this issue and that of spherical optical geometry will be further discussed in the sections below.




Figure 1: Computing scale invariant samples in the retina (a) and Scale invariant recognition (b).

One way to avoid the problem of variable scale sensitivity is by designing our image sensor so that it can deliver a scale invariant measurement. We don't need to believe that the biological solution explicitly computes these measurements at this stage, but if we cannot produce a scale invariant measurement system then we certainly cannot expect to produce any other scale invariant quantities². This is a statistical design constraint, we need to know that equivalent information is available at all scales so that we can compute the statistically equivalent quantities. As it is easier to consider the issue of scale invariance in cartesian geometry, we start on a planar sensor

²We are not talking here about the approximate invariants generally used in computer vision, but fully invariant quantities obtained by constructing a sensor so that such things are computable.

in what we will call the fovea where we assume the data has approximately homogenous spatial sampling³. We wish to synthesise a new sensor now at a lower resolution ($1/\alpha$). This process is illustrated in Figure 1a. The inner values of this new sensor can be computed from the fovea using a process of down-sampling. The most likely computational form for this process is a Gaussian convolution (or its irregular sampled equivalent) as this is the only scale invariant sub-sampling process, due to the central limit theorem. The outer values of the new region are unavailable in the fovea and therefore require additional sensors around the edge of the previous active region.

The new ring of sensors need only to have a spatial resolution which is a factor α larger than the ring of original samples on the edge of the fovea. To have better spatial resolution is potentially wasteful with resources. Down sampling once more, in order to construct the next equivalent sensor, we can again subsample the inner regions and we will need a new layer of measurements from areas now α^2 larger than the original. As this process continues we gradually construct a sensor surface, with ever decreasing spatial resolution. The mathematical process we have just described for positioning each ring of new sensor locations is

$$r_n = r_0(1 + k \sum_i^n \alpha^n)$$

Outside of the fovea, the radial function which characterises this process is obtained if the spacing of sensors is exponential outside of the fovea, (ie: a logarithmic sensor). This model differs from previous interpretations [14]⁴, in that it accomodates the uniform sample density in the fovea as an intrinsic part of the calculation of scale invariant shape representation. Making the adjustment from a planar sensor to another geometry, such as locations as angles ϕ on a spherical sensor, is mainly a case of applying the appropriate transform⁵.

In conclusion, the radial distribution of sensors is equivalent in terms of measurement to an homogenous sampling process at a fixed set of scales. As described below, this sensor down-sampling does not have to exist explicitly but can be combined with the process of feature detection without loss of the important computational properties. In particular, the spatial distribution of uniform sensors in the fovea can be used to compute the locations that additional sensors need to be placed when the same object is viewed at a larger scale in order to receive statistically equivalent input data. This now allows us to exploit the sampling properties of a spherical optical geometry. Such a scheme provides a partial invariance to projective deformation which is absent in a conventional electronic device. This interpretation may also explain how it is possible to represent and maintain locations for objects relative to the head and body, without having to worry about the spatially varying resolution of input image data. As effectively any detected feature is represented at one fixed scale.

An Overview of Statistically Based Feature Detection

We would like to take methods from computer vision as candidates for feature extraction in the brain. The above analysis makes no argument for what the layout of our scale equivalent sensor structure should actually be. Therefore we can just as well resample to a regular lattice without destroying the underlying computational properties. This allows us to look to the last 30 years of ideas developed by the computer vision community, in order to get a better idea of what might be required computationally to describe object shape. The topic of computer vision represents a large body of literature and in most cases the difference in behaviour and capabilities of the methods are heavily influenced by intended use. There are consequently as many feature detectors as there are applications. However, the general principles involved for much of this work can be characterised as template matching and interest operators. Both of these are physiologically viable methods for scene analysis. Template matching is analogous to the concept of receptive field patterns, while interest operators are more akin to the hypothesised mechanisms related to the possible role of micro saccades. The following section makes an argument for interpretation of these approaches as different aspects of the same underlying principle for information extraction.

Scaled features can be computed at a range of spatial scales by processing regular equivalent sensors with a fixed set of feature detectors, or (more likely) by combining sub-sampling with the process of feature detection. The equivalent sample images need never be explicitly computed, but the resulting output of the system is equivalent to if they had been. Scale invariance can then be achieved by applying a selection process to choose the one scale most suitable for representation of the local image patch (Figure 1b). The selected representation at a single scale is then compared to stored visual memory. If the viewed object changes scale, by for example moving towards the viewer, then the scale at which selected features are detected will change so that the same, scale invariant,

³As discussed above, in the biological system this process is complicated by the presence of two distinct types of light detector, rods and cones.

⁴Where an adjustment to the logarithmic form $\log(r) \rightarrow \log(r+a)$ is made which breaks the pure scale invariance property exactly where we might expect it is most needed, in the fovea.

⁵Actually, this argument to apply exactly we need the initial definition of the sample on the fovea to be equivalent to an homogenous sampling on a spherical surface not a planar one. The difference here is negligible for a fovea of small angular extent.

spatio-temporal relationships are selected for pattern matching. This process is effectively what is advocated in the Scale Invariant Feature Transform (SIFT) [11].

The best known of the template based approaches include the Canny edge detector [2] and the Difference of Gaussian operator [13]. These methods apply combinations of image convolution operations in order to enhance selected features. Connected or isolated feature points can then be identified as maxima on these enhancement images, in this case step edges and ridge locations respectively. The main difference between these two approaches is that although a step edge detector will respond to all but the finest of ridge features, Difference of Gaussian processing will not enhance step edges ⁶ (the most common manifestation of an object boundary) and has an output which is more strongly dependant upon feature scale and illumination.

In order to take advantage of the inherent properties of illumination invariance, feature detection processes need to be computed as combinations of differences between measured values g . For many features defined in computer vision, such as conventional first derivative based edge detectors, this is clearly the case.

$$e = G(\sigma) \otimes \sqrt{(g_{i+1} - g_{i-1})^2 + (g_{j+1} - g_{j-1})^2}$$

where $G \otimes$ represents a Gaussian convolution of width σ . Interestingly e^2 is a smoothed local estimate of the Fisher Information associated with local image plane orientation.

$$\theta = \text{atan}(g_{i+1} - g_{i-1}) / (g_{j+1} - g_{j-1})$$

Therefore selecting local maxima of e also selects samples which maximise the accuracy or orientation.

Although using a multitude of template feature detectors, all matched to distinct feature types, is a possible algorithmic solution for the extraction of object structure, this approach raises an important question; What is a valid mechanism to combine the responses from the different detectors? This must be done in a way which provides generalisation for recognition of shape, not only for changes in illumination and object scale, but also over possible responses to changes in scene content (for example arbitrary possible backgrounds at object boundaries). Interest operators provide an alternative which is capable of detecting many characteristic feature types with one simple computation. The simplest interest operator would be a local variance estimate of image signal, which when applied at a range of spatial scales, is also useful as a descriptor of texture [5]. For example;

$$v = G(\sigma) \otimes (g - \langle g \rangle)^2 \quad \text{where} \quad \langle g \rangle = G(\sigma) \otimes g$$

which can be considered as either the estimate of signal to noise associated with local image variation, or the inverse of Fisher Information associated with a mean value. This simplifies to;

$$v = G(\sigma) \otimes g^2 - (G(\sigma) \otimes g)^2$$

Notice that these calculations embed directly Gaussian convolutions, which we have already stated are required for scale invariant resampling. Also, this has been done in a way which, unlike conventional Difference of Gaussian detectors, is capable of detecting step edges.

Another popular feature detector in computer vision is the corner detector and one approach uses the concept of an interest operator which is often based on the idea of auto-correlation. Corner detection can also be performed with templates, but is significantly more difficult than edge detection due to additional variation in orientation and corner shape [7]. The Harris corner detector [6] defines corner locations using the second order spatial variation of an auto-correlation around a point. However, auto-correlation can be interpreted as a log-likelihood for the degree of match between the original local image patch and a shifted version. In addition the matrix of second order behaviour is the second term in a Taylor expansion, so it can also be interpreted as the second derivative with respect to image location. As the Cramer-Rao bound is the second derivative of a log likelihood, this is the Fisher Information for spatial localisation.

In summary we now have three definitions of feature detector based upon quantitative measurements which define positions of maximum information; local variance for spatially varying intensity, edge strength for orientation, and interest operators for spatial location. In fact, if we consider feature detection as a template based approach supported in the biology by receptive fields, then grey level scale, orientation and location are the only measurable quantities possible. It makes sense to suggest that if restrictions on processing capacity (for example finite connectivity) results in the need for the brain to identify a subset of features in order to solve quantitative tasks, then those which make the largest quantitative contribution (ie. those which maximise some aspect of Fisher

⁶The response to a step edge from a DoG filter (as advocated in SIFT), is exactly zero at the position of the edge. Although it takes large positive and negative values on either side, the locations of the maxima are systematically shifted as a function of the Gaussian kernel size and therefore cannot be considered as spatially consistent.

information) are the ones it should be using and the ones we should be basing any model of scene interpretation upon. As with the case for intensity, we can also determine how the brain would need to encode these values as neuronal pulses, so that the noise characteristics of the signal are matched to the way the data is passed. For the case of e , which error propagation tells us has spatially uniform noise, this gives us an equivalent frequency value of $f_e = e^2$. While for v the noise characteristics already match Poisson behaviour so, $f_v = v$.

Finally, illumination invariance of these measures and also for colour is entirely reliant upon logarithmic sensitivity to light. As with spherical optical geometry, this is a property which is lacking in a conventional electronic sensor. It is becoming increasingly obvious that when it comes to getting a simple solution to visual analysis tasks, the biological sensor has characteristics which make a lot of sense. Indeed, the analysis of data from a conventional colour CCD array will be difficult by comparison.

Shape Recognition

Illumination invariant colour representations can be constructed by looking at ratios of colour signals, in direct analogy to the mechanism outlined above. Interestingly, the clear functional separation of rods and cones, together with their different spatial arrangements in the retina, precludes the possibility of a unified colour and shape recognition system in which colour is used from the earliest stages to enhance the shape detection process. This would seem to suggest that colour information is used mainly for point classification processes, while intensity (mainly obtained by the rods but supplemented by cones in the fovea) is intended for the analysis of structure. This observation makes predictions regarding the computational processes associated with recognition of shapes described by pure colour (no intensity difference) boundaries, ie: they have to be analysed differently to the more conventional case⁷. This also precludes the early integration of colour and motion cues⁸. This is in direct contrast to recent computer vision work on recognition from colour images (for example colour based edge detection), and suggests that more realistic physiological models of shape and motion analysis will be generated in computer vision if we avoid colour cameras.

Assuming that the process of shape recognition is based upon conjunctions of detected features, then the above description of a multi-scale feature detection process eliminates the need for scale invariance. If we also eliminate the possibility of full 3D rotation invariance (on grounds of in-homogenous error characteristics), the required invariances for a shape representation are therefore translation and rotation within the sensor “plane”. An ideal representation of shape would be one which supported the reconstruction of the shape up to an unknown position and orientation. Such a representation has been previously described as “complete”. Although simple regional histograms of local image orientation (as used in SIFT) are not complete [17], the property has been established for the representation scheme referred to as “pairwise geometric histogrammes” [16]. This approach provides an encoding of local shape as a 2D frequency distribution of relative angle against perpendicular distance. These histogrammes also have the statistical characteristics necessary to perform a meaningful comparison to templates stored in a frequency based computational system, such as supported by biological neurons [18]. Matching of the best template w_j is achieved using the following Bhattacharyya similarity measure;

$$B_j = \sum_i \sqrt{f_i \cdot w_{ij}}$$

which can be directly related to the log probability of the similarity between sets of frequency coded data via the Matusita measure [19]. As with the case of the equivalent retinas, geometric histogrammes do not necessarily need to exist as physical topographic maps, as equivalent calculations can be performed without requiring an explicit physical structure. However, benefits might be gained from having regular structure with regard to physical construction (such as compactness) and adaptation (such as training efficiency). The publications which provide extensions to this representation and support matching of extended line features (key-points) across multiple scales [1] pre-date more recent non-statistical approaches [11] by almost a decade.

Methods: Investigating Patterns of Eye Movement

The simplest hypothesis for the role of saccades is that we move our eyes in order to build up a high resolution measurement of the scene. This hypothesis can be immediately excluded by observing real eye movements, which

⁷For example; if you look at a yellow circle on a blue background (in good light), and the luminance is adjusted to be equal, the boundary between the blue and yellow disappears (The Liebmman effect 1927).

⁸If two or more small spots, separated by up to 4° are briefly lit, in rapid succession, a single spot wil appear to move back and forth, even when the dots are illuminated in different colours (the Phi Phenomenon 1912, Kolars and von Grunau, 1976)

do not uniformly scan the potential view field, but seem instead to be drawn to particular visual features, movement or objects. A more sophisticated hypothesis for the role of saccades in visual exploration is that we saccade to areas which are expected to have useful spatial information for the interpretation of shape or structure [20]. We would therefore expect the eye to saccade to those features which are most useful for this task. As we have only low resolution data available in the periphery of the retina, we must assume that this is somehow used to predict the most useful places for fixation. Although we may not know precisely what the human vision system does, we suggest here that we can take the standard feature detection processes as characterised by interest operators and template matching approaches as indicative of those features which would be useful for the purpose of extracting image structure. We can then see to what extent the saccadic process targets locations in images which contain structural features in order to examine our initial hypothesis.

In brief, participants ($N = 24$; Mean age = 22.67) first viewed sets of six novel 3-D objects each containing one principal component and three sub-components or volumetric parts (see Fig 3a). 12 objects (6 targets and 6 distractors) were presented from three different viewpoints (0, 120, 240 degrees) each for 10 seconds while eye movement patterns were recorded. Following the Learning Phase, participants performed a recognition memory task in which they had to discriminate learned from unfamiliar objects, presented either at practiced (0, 120, 240 degrees) or novel orientations (60, 180, 300 degrees) in depth. Behavioural responses (accuracy and Reaction Times (RT)) were recorded. Eye movement data were recorded on a Tobii ET17 remote eye tracking system running at a data acquisition rate of 50 Hz. Experimental stimuli were viewed from a distance of 60 cm at a screen resolution of 1280 x 1024.




Figure 2: Novel objects (a) and reaction times for recognition (b).

Results

Accuracy of target detection in the Test Phase was very high (range 80- 94.17 %). As expected, targets were detected more accurately at the practiced viewpoints, $F(1, 23) = 26.01, p < .001$. Mean RTs for correct trials are shown in Fig 3(b). A 2 (Familiar vs Unfamiliar) x 3 (Viewpoint) repeated measures ANOVA showed that RTs were faster for practiced over unfamiliar viewpoints, $F(1, 23) = 13.73, p < .001$. The main effect of Viewpoint was not significant. There was no interaction.

Analyses of eye movements were conducted by initially pre-processing raw gaze data by applying spatial and temporal filters to remove micro-saccades and drift. Fixations were defined as eye movements that remain within the same circular region (diameter 60 px) for a minimum of 100 msec. Filtered data were used to compute fixation frequency across participants for each stimulus. Figure 4 shows a time series fixation frequency plot for all participants overlaid onto the original stimulus image. The data are grouped into 10 epochs corresponding to the first 500 msec post stimulus onset, and then for each 1000 msec thereafter. The data show that participants rapidly fixate on image regions which appear to correspond to salient 3-D image segment points around object sub-components. Figure 5(a) shows an analysis of the consistency of fixations across changes in the 3-D viewpoint for two of the test stimuli. This shows that participants consistently search for and fixate the same 3-D image segmentation points across viewpoint, despite changes in the low-level image properties of these locations (e.g., vertex types). Figure 5(b) shows an analysis of the consistency of fixations between the Learning and Test Phases for two items. As in the previous analysis, participants fixate the same image regions between phases.

Conclusions

This paper has sought to explain saccadic eye movement within a framework which includes some of the more obvious structural features of the human vision system. It seems to be possible to account for many observed properties, including logarithmic intensity sensitivity, and spherical optical geometry, in terms of construction of invariant representations which take account of measurement error. The kinds of algorithms generally developed in the area of computer vision, with regular input lattices and at fixed scales, may seem a world away from irregular




Figure 3: Time development of fixations.

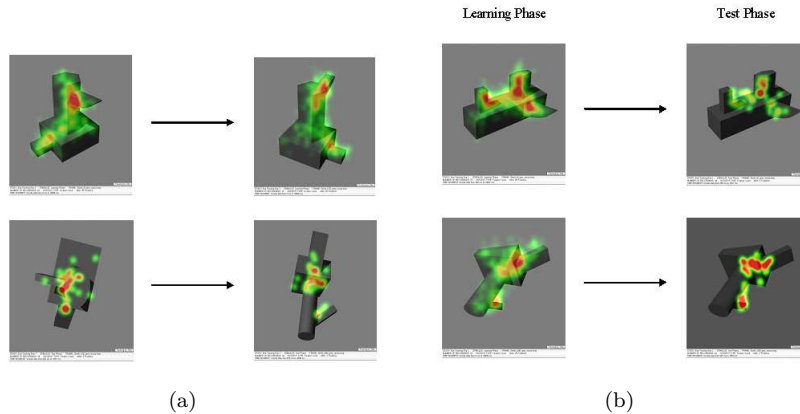


Figure 4: Changes across viewpoint (a) and consistency between learning and test phases (b).

sampling of the retina and saccades. However, it seems possible to replicate the process of scale analysis by simply processing at multiple scales and selecting one result. This opens the possibility of applying insights from image analysis to interpretation of visual biological. The process of sub-sampling combined with feature detection will produce a range of feature detectors whose main characteristic is that a feature detected at one resolution will be equivalently detected at another. This appears to be a fundamental requirement in order for a multi-scale system to operate consistently. However, if we consider these features as receptive fields the details of their construction may look quite complicated and will have a range of sizes on the retina. We therefore need to be able to appreciate the process of scale invariant feature detection before we can interpret the range of receptive fields observed in biological systems in terms of function.

The suggested model is very close in essence to SIFT and we may therefore be tempted to consider testing SIFT as a theory of vision for comparison with psychometric tests, such as eye movement. However, SIFT is too general an idea to be able to perform a meaningful test. It has many free parameters and no theoretical basis from which to specify them, (simply executing the distributed version is not a solution). Pinning them down requires a statistical interpretation of what it is supposed to be doing and a quantitative relationship to biology. Our theory explains how to model the quantitative information available from the retina, including illumination sensitivity and scale dependency. The analysis also blows big holes in both the ideas of using DOG filters and orientation histograms, on the basis that they don't make good use of the available information. Without consideration of these issues

any SIFT algorithm is just one selection from an infinite number of candidates. Under application of the scientific method this precludes sensible interpretation of the results from any comparison with perceptual data.

Ultimately the brain must use detected features for the construction of shape representations. The brain will need to analyse the incoming spatio-temporal data to extract compact descriptions for the purpose of accurate prediction and categorisation. Analysis of the statistical nature of the data tells us that it is not possible to construct a representation which is invariant to every form of variation produced during image formation. However, invariances are key to the construction of efficient vision systems, as the more we can correctly generalise from data we have already learned and understood, the easier it is to interact with our environment. The development of invariant recognition processes could be invoked as an implicit target during the process of human evolution, thereby explaining the kind of structure we see on the retina today.

Our study supports the following conclusions: (1) Human visual biology is consistent with a simple structural hypothesis (based on multi-scale samples) for the construction of invariant recognition systems which opens the way for interpretation of retinal data using conventional machine vision approaches, (2) Fixational eye movement patterns during 3-D object recognition are not random, but rather structured and highly consistent among observers. (3) There is remarkable consistency in the patterns of fixational eye movements shown for 3-D novel objects across both changes in viewpoint and between learning and test phases. (4) These locations show evidence of ‘top down’ selection and are **not** the low level features generally constructed for machine vision. The specific locations which people repeatedly look at are difficult to characterise in terms of simple image features. For example, confounding factors, such as shadows are immediately ignored. If anything the eye tends to be drawn to 3D volumetric primitives or surfaces. Observers tend to consistently track the same image segmentation points despite variation, across viewpoints, in the low-level properties (e.g., edge density and vertices) at those regions.

These conclusions are not dependent upon the particularly simple nature of our stimuli, and tell us that tracked features are a long way from the input visual data, in terms of processing. They imply a high level representation of 3D structure which is already available prior to eye movement. The most striking observation is that fixated locations are often on surfaces, which in our data at least contained no low level information. At first sight this may seem to be at odds with a feature based analysis of shape. However, these conclusions can be reconciled with a feature based analysis if we take a view based approach to recognition, whereby sets of features within a focal region are used for shape representation such as a learned set of geometric co-occurrences (such as the PGH). We can have every reason to believe that fundamental understanding of the problems involved in extracting shape information from conventional images is potentially of direct relevance to understanding high level processing in human vision. This being the case, locations of saccadic fixation should contain valuable information which can help identify these high level processes. This is an avenue we intend to explore further.

References

- [1] A.P.Ashbrook, N.A.Thacker, P.I.Rockett and C.I.Brown, ‘Robust Recognition of Scaled Shapes Using Pairwise Geometric Histograms.’, *proc.*, BMVC 95 Birmingham, 503-512, July 1995.
- [2] J. Canny., A computational approach to edge detection., *IEEE Transactions on Pattern analysis and Machine Intelligence*, 8(6),679-698, 1986.
- [3] A.T. Duchowski, *Eye Tracking Methodology: Theory and Practice*. Springer. London, 2003.
- [4] J.M. Findlay and I.D. Gilchrist, *Eye guidance and visual search*. In G. Underwood (Ed.), *Eye Guidance in Reading and Scene Perception*. Oxford. Elsevier, 1988.
- [5] R.M. Haralick, *Statistical Image Texture Analysis*, *Handbook of Pattern Recognition and Image Processing*, Ed. T.Y Young and K.S. Fu, Academic Press, Orlando, 247-279, 1986.
- [6] C.Harris and M.Stephens., “A Combined Corner and Edge Detector” *Proceedings of the Fourth Alvey Vision Conference*. 147-151, August 1988.
- [7] J.M. Henderson, C.C. Williams, M.S. Castelhana, R.J. Falk, *Eye movements and picture processing during recognition*. *Perception and Psychophysics*, 65, 725-734, 2003.
- [8] R.S. Johansson, G. Westling, A. Backstrom, J.R. Flanagan, *Eye-hand coordination in object manipulation*. *Journal of Neuroscience*, 21, 6917-6932, 2001.
- [9] A.J.Lacey, N.A.Thacker and N.L.Seed. ‘Smart Feature Detection Using an Invariance Network Architecture’, *proc.*, BMVC 95 Birmingham, 327-336, July 1995.
- [10] M.F. Land, M. Hayhoe, *In what ways do eye movements contribute to everyday activities*. *Vision Research*, 41, 3559-3565, 2001.
- [11] D.G.Lowe, *Distinctive Image features from Scale-Invariant Key-points*, *Int. Jou. Comp. Vis*, 2004.
- [12] S. Martinez-Conde, S.L. Macknik and D. H. Hubel., *The Role of Fixational Eye Movements in Visual Perception*, *Nature Reviews Neuroscience*, 5(3), 229-238, March 2004.

- [13] T. Peli and D. Malah., A study of edge detection algorithms., Computer Graphics and Image Processing, 20, 1-21, 1982.
- [14] E.L.Schwartz, Spatial Mapping in the Primate Sensory Projection: Analytic Structure and Relevanve to Perception, Biol. Cyber., 25, 181-194, 1977.
- [15] K. Schill, E. Umkehrer, S. Beinlich, G. Krieger, C. Zetzsche, Knowledge-based scene analysis with saccadic eye movements Human vision and electronic imaging IV; Proceedings of the Conference, San Jose, 520- 532. 1999.
- [16] N.A.Thacker, P.A.Riocreux, and R.B.Yates, ‘Assessing the Completeness Properties of Pairwise Geometric Histograms’, Image and Vision Computing, 13, 5, 423-429, 1995.
- [17] N.A.Thacker and J.E.W.Mayhew, ‘Designing a Network for Context Sensitive Pattern Classification.’ Neural Networks 3,3, 291-299, 1990.
- [18] N.A.Thacker, I.A.Abraham and P.Courtney, ‘Supervised Learning Extensions to the CLAM Network.’ Neural Networks Journal, 10, 2, pp.315-326, 1997.
- [19] N.A.Thacker, F.Ahearne and P.I.Rockett, ‘The Bhattacharyya Metric as an Absolute Similarity Measure for Frequency Coded Data.’ Kybernetika, 34, 4, 363-368, 1997.
- [20] L. Walker and J. Malik. Sequential information maximisation can explain eye movements in an object learning task. Journal of Vision, 4, 744a. 2004.
- [21] A.L. Yarbus, Eye Movements and Vision. New York. Plenum Press, 1967.

Questions

The following questions were asked following presentation of this work at BMVC 2007 and at an internal IBSE group meeting.

- You have focussed the criteria for system performance strongly on noise processes. When I design systems I see the dominant source of noise as illumination problems and missing data. Can you comment how this approach addresses these kind of issues?

Put simply it cannot. Designing a system which accounts for these larger scale degradations in data quality is something which needs to be addressed separately. See my previous publications regarding geometric histograms. However, if you ignore the noise charateristics you will not have a stable recognition system anyway. You need to address this issue from the beginning.

- Vision is a capability of many species and has evolved many different solutions. Can this statistical approach be generalised to justify the biology in other animals, such as insects?

You need to take the arguments on a case by case basis. In particular, unless you think that bees (for example) need to form internal represenations of 3D objects then there is no reason to suggest they need a scale invariant representation.

- I was always taught that the central foveal region was there to allow us to more closely examine objects of interest. You seem to be saying that this explanation is wrong.

Yes. You can make the observation that having higher spaitial resolution supports detailed inspection. You cannot use this to make a case for the specific form of spatial resolution dependency. Any radial dependency which has higher sensitivity in the fovea would permit the same level of “explanation”.

- When you say that something like a view based representation is necessary to explain patterns of saccades, are you claiming that responses to visual stimuli do not change gradually under small rotations of objects?

No. I’m claiming that the consistency of identification of 3D features cannot be accounted for by low level features alone, which change too much between views in these experiments.

- You say that geometric histograms would be a suitable way to construct shape representations, but aren’t there a lot of other models which could do the same, Biedermans geons for example?

I would say that geons have two specific problems. The first is that although they seem to do a good job of describing psychophysics data, they are not specific enough to allow us to understand how to build a working model. The second is, that these theories were not constructed using any concept of information content or statistical analysis. Of course, to some extent these two observations are interlinked. On the specific choice of geometric histograms: this representation is complete and based upon justifiable statistical matching processes, it also happens to be biologically plausible. I believe these characterisitcs are necessary for any sensible theory, and any other approach with these same characteristics would be indistinguishable, as it could only make the same predictions for any experiment. It is therefore valid for me to hypothesise this approach as a statistical model of the system.