Memoirs of an octogenarian Vision Research Scientist.

Ian Overington B.Sc. M.Inst.P., A.M.I.E.T.

I was born in the small town of Ilkley in the then West Riding of Yorkshire on March 3rd 1931. After basic Secondary school education at Ilkley Grammar School I obtained good marks in Higher School Certificate and studied for a degree in Electrical Engineering at Leeds University.

Whilst on holiday from my studies in the summer of 1951 I met my future wife to be, whom I married in 1953 – a marriage which lasted nearly 60 years until my wife's death at the end of 2011, the marriage producing 3 children.

After graduating in 1952 I joined the then Bristol Aeroplane Company (later to become part of British Aerospace Ltd.). I initially worked in fields associated with testing and installation of various forms of instrumentation associated with the Research phase of development of the Bloodhound Missile Defence System. This work particularly involved use of temperature sensors in what at the time was a new hostile environment (dealing with temperatures changing both spatially and temporally at high rates). Amongst other things this involved use of 16mm cine cameras as spatial & temporal detectors of energy with a high degree of accuracy, plus the testing, installation & recovery of cine cameras for in-flight recording of behaviour of such things as the ram jet engines in the test vehicles. How I wish that ***digital*** imagery and modern easy transmission of imagery had been available in those days - life would have been very much easier!!

Together the foregoing provided a good grounding in the ***real world*** (as opposed to what I have come to realise is a very artificial academic world in the field of research) over a substantial range of disciplines as well as electrical engineering. As a result of the studies on highly controlled use of the cine photographic process as a realistic remote sensor I also found myself involved in other forms of controlled spatio-temporal brightness measurement. From this I was able to publish a series of 6 papers in the Journal of Photographic Science in the late 1960's under the general heading of 'Photographic Detectors', as a result of which I was able to obtain full membership of the Royal Photographic Society. I was also able to present a paper at an international conference on 'Light and Heat Sensing' in Paris in 1963 (Ref. 1).

Then in around 1963 a small research team headed by me was offered a Contract from RAE (Royal Aircraft Establishment), Farnborough to carry out a study under the general heading of 'Visual Studies'. These vision research studies were initially loosely interactive with the foregoing experiences, in particular with reference to the controlled use of the cine photographic process in a highly controlled manner for accurate reproduction of incoming light.

The basic general brief for this Visual Studies Contract was to endeavour to improve on available abilities to predict detection ranges for military targets from low flying aircraft based on available experimental data which was becoming widely available at that time.

On first sight that didn't seem an unreasonable requirement. However, it rather quickly became apparent to my team & me that there were ***major*** differences between the widescale ***laboratory*** experimental conditions for which data were becoming increasingly available and the task in hand. At the time there were little (if any) available data to help us to put any sort of figure on the real world thresholds of detection. With hindsight, after some 50 years, it seems appalling how naive and restricted the open knowledge related to ***practical*** visual performance which was available at the time was. I am also horrified how naive our ***own*** thinking and knowledge was at that time. A very considerable effort was made to publicise our own progressive findings over an extended period (see Refs. 2 & 3 for summaries etc., also Ref. 4) and considerable relevant associated facts relating to imagery were published by others through the 1970's & 1980's in particular. However, despite all this I am even more horrified to find that not very much ***open*** progress seems to have been made in all these years - with the exception of mainly our own publications, which seem to be largely unknown despite our efforts at publicity.

Some basic observations need to be spelled out at this point.

i) Before the early 1960's the best that seems to have been possible in terms of predictive modelling of visual thresholds was largely restricted to a number of simple empirical laws such as Ricco's Law (for unresolved small discs); Weber's Law (for very large objects); Piper's Law (a simple empirical effort to bridge the gap between Ricco's Law & Weber's Law) and Bloch's Law (for images presented for short periods of time i.e. less than about 0.1 seconds. This meant that there were large portions of any ***continuum*** trend curves where the predictions were open to large errors.
ii) In real world situations one is very rarely looking for simple discs. Nearly always one would rather be looking for complex shapes, possibly with internal structure and likely against a structured background – see Ref. 2, Chapters 11 to 13, in particular).
iii) In the real world one is almost certain to be viewing for some relatively extended period of time (at least a few seconds, rather than less than 0.1 seconds).
iv) Real world tasks are most likely to be requiring a higher level of decision than just 'something there'.
v) In commonly experienced imagery associated with photography and television one is usually doing one's best to obtain the ***best possible*** image quality. But considerations of visual imagery require a ***total*** rethink. In the eye one has to acknowledge the fact that the sampling matrix is associated with sampling receptors which are very definitely of ***finite*** size (in modern thinking - since the advent of the computer age - synonymous with a rather ***coarse*** sampling matrix). But this means that, as opposed to attempts to obtain best possible image quality, the Shannon Sampling Theorem (which has been around since about 1950 – e.g. Ref. 3, Chapter 7.2) shows quite conclusively that, for optimum information transfer, there should be an optimum ***matching*** between the sampling frequency and the image blur. It so happens that the ideal match between ***foveal*** retinal receptor spacing (the optimum area for highest resolution image perception), and optical quality which yields 20/20 (6/6) vision on a Snellen chart, provides roughly this best match to satisfy the Shannon theorem! But on detailed consideration the retinal image itself providing this match is far from free from blur. It is this substantial blur, when coupled with the retinal sampling matrix, which provides the ideal opportunity to derive ***sequences*** of local fragmentary outputs actually ***across*** a blurred edge image. In turn this provides more than sufficient data to compute an edge position very simply to an accuracy of a small fraction of a receptor spacing - effectively what has now been defined as Hyperacuity (a word which did not exist in the 1970's & early 1980's when most of my research was being carried out!). It is stated in Wikipedia that the way which human vision accomplishes this is ***still*** awaiting a solution - several decades after we first sorted it out!!!
vi) By tradition it has always been accepted that, in ***mathematical*** image processing, the best that one can do is to plot brightness discontinuities to the nearest pixel. In fact, in 1986 a new operator for that purpose was proposed by John Canny. This has become the standard 'optimum' method of edge finding to this day (in particular in the MathCAD suite of image processing tools). But I met Canny around the time that he was proposing his new operator and illustrated to him that the visual system regularly achieves much better than effective single pixel location (and with much the same capability to separate edges as claimed by Canny). But he apparently chose to ignore what I was saying. An automatic result of the visual system's ability to achieve the now called Hyperacuity is that the visual system regularly achieves much better than single pixel accuracy – and my modelling of this capability automatically creates edge traces to better than 0.1 pixel local accuracy!
vii) An ability to sense exact movement of edges between frame samples to a fraction of the receptor spacing has also been found to be available for ***dynamic*** image sequences. Not only that, but a similar capability has been found to exist for ***stereo*** acuity in binocular vision – and this stereo acuity has been shown to be derived (all prior to the central cortex) by virtually the same (simple) circuitry used for motion sensing (e.g. Ref. 9; also Ref. 3, Chapters 5, 9 & 10)!
viii) We don't normally ***perceive*** 'noise' in normal viewing. However, there is ***always*** some level of noise present in ***any*** visual imagery - and it is this which is manifest as a reduction in

perceptive capability in low light situations. Hence in any attempted predictive modelling of visual performance it is imperative that some attempt is made to model this effective signal / noise ratio.
ix) No individual experimental result is in any sense absolute. There is always some degree of uncertainty (probability). This is usually assumed to be a Gaussian distribution function - and for most purposes this assumption is probably adequate. But this is on an assumption that a mean and standard deviation can be stated for the experimental results. However, what absolute values should this mean and standard deviation take? Over many years our own various studies have very positively shown that the probability function itself has ***several*** components - mainly free or forced choice and laboratory or field conditions. Together these variations in the mean for a probability curve can themselves be large - and in fact studies carried out by the Ronchi Foundation (Fondazione Giorgio Ronchi) in Firenze (Florence), Italy have shown conclusively that the whole probability curve for a given task in human vision for any one observer itself has a rather large fluctuation according to the time of day. So it is essential to make some attempt to allow for fluctuations in visual performance - ideally even for a ***particular*** individual (observer) according to time of day!
x) It has been found that the general probability curve obtained from a given experiment depends significantly on whether the observer is ***forced*** to make a decision or free to make his (her) own mind up (forced or free choice). Also there is some suspicion (not totally proven) that decision levels are also influenced by differences between laboratory and natural field conditions for the experiment, even after allowing for object & background situations.
xi) Yet again, it has been found over several years of experimentation that the whole concept of thresholds being controlled by the size (area) and contrast of an object against its surroundings ***is wrong***. The parameters to use for predicting threshold trends should rather be the maximum ***slope*** of the blurred edge of an object and the length of ***the portion*** of this profile which is relevant to the task in hand (be it simple detection or recognition). A further factor which has been confirmed to be massively more important - that of the prevailing ***sharpness*** of the image being presented on the retina - is then to be compounded with the prevailing object / background contrast. This is now somewhat obvious to me with hindsight, but has been rarely mentioned at all in reporting of threshold trends in years gone by - certainly not in the 1960's and 1970's. In conjunction with this, whilst not too important for ***low level*** decisions (namely detection thresholds for simple objects), as soon as the task becomes more difficult, involving recognition of object shapes, critical attention has to be paid to the ***variation*** of maximum composite ***slope*** around the perimeter of the target object. Although all this sounds very complicated, in actual fact, at least for relatively simple object shapes, it is not unduly difficult - and it was shown by us (and demonstrated in a paper which was refused publication in 1973, but which I have recently recovered and digitised) that good predictions were then possible.

Finally there is one other significant matter which needs to be addressed if detailed ***physical*** modelling of ***perception*** of visual imagery is to be attempted. This is the fact that the entire visual image handling (from initial sensing of individual cone signals onwards) is actually carried out on an effectively ***hexagonal*** matrix, rather than a rectangular matrix as used in conventional computing. I have for many years had an approximate (but more than adequate) method of handling this on a conventional computer - and this is discussed in some detail in Ref. 3, Chapter 3. But this is advanced stuff and there is no point in even trying to outline this here. Should a reader wish to see visual demonstrations of the progressive visual processing using a hexagonal matrix, some animated gif images are provided in the section of my website (Ref. 4) which is headed 'hierarchical processing' (readily accessible from the Home Page).

During the late 1970's further work was carried out to refine much of our basic threshold modelling, including allowance for effects of such things as the quality of any external optics. A much extended mathematical model discussing this was published in Optical Engineering in early 1983 (Ref. 5). Meanwhile a series of open papers were published in various scientific journals, several of these also being presented at various International Conferences. Some of the more important of these are listed as Refs. 6 to 13.

Of these I feel that I need to single out one particular one for further immediate comment. That is Ref. 12, which concerns the simulation of random dot stereograms and kinematograms, which had become very popular in the early 1980's. These were, as the title of the paper suggests, pairs of mathematically random dot patterns which contained a hidden object, which either stood apart from the main pattern in the case of stereograms or moved ***as a block*** in the case of kinematograms (provided that the observer had adequate control of their eyes). I was never able to observe the supposed resultants myself due to a weakness in my stereo perception, but nevertheless I could envisage what I ***should*** be able to perceive. I was then able to illustrate the resultant perceived images via my developing vision model, which I felt was a very strong indication that the modelling was correct.

Around 1980 I was asked by the top management of the B.Ae. Research Laboratory to move sideward from the mainstream ***threshold*** performance studies and concentrate my efforts on trying to unravel the (by then growing) emerging data concerning the ***inner*** workings of the early visual system. This was (somewhat naively at the time) with the eventual intention of creating a physical model of the inner structure of the early visual system (prior to the optic nerve) - an aim which, although ***originally*** thought to be naive, by the end of the 1980's had become a realistic possibility (see later)

Around this time an increasing amount of literature was emerging which, rather than concentrating on laboratory thresholds of visual performance, was exploring the ***neural*** layers of the actual eyes themselves. This proved to be very relevant for the new tasks which B.Ae. had set me. Some of the most important general findings which came out of the change of direction in published vision material are as follows.

One of the most important findings (at least to me) was first published (to my knowledge) by Bergen & Wilson in 1979 (Ref. 14), this general finding being of up to ***four*** layers of early neural activity (still within the retina) which were of varying scales (resolution) and ***concentric*** with the individual retinal receptors. This was refined a few years later in a very lucid paper by Granlund (1983) (Ref. 15), who took the matter further into the possibilities of useful ***interaction*** between these various scales of receptive field.

Progressively this concept was further refined - and it was also found that these concentric layers were multiplexed with the tricolour cone receptors to create new retinal maps of varying resolution and with changing colour sensitivity. In the end this provided the resultant of a conversion of the incoming trichromatic (R, G, B) sensing into, instead, a set of ***opponent dipoles*** (red versus green, yellow versus blue and black versus white). It has subsequently been shown that this transform of colour sensing has very powerful repercussions in all sorts of ways.

Meanwhile, in 1982 there was a flurry of excitement in vision research circles in the UK with the publication of a new book simply called 'Vision' by one David Marr of M I T (Ref. 16). This book gave a very lucid demonstration of how important the teachings of Shannon etc. were in a real practical visual environment. He clearly demonstrated how one might expect an extreme positional sensitivity from images containing 'sharp' edges in a visual environment if one had an adequate match between the sampling (by inference, by cone receptors) and the optical blur. But, quite some time earlier, I had already demonstrated that the retinal image is itself surprisingly blurred as imaged on the retina. Hence it turned out that there was a good natural match between the sampling intervals and the optical blur spread function, all this fitting neatly with Marr's concepts.

But the whole concept proposed by Marr depended on a ***second*** difference function of a blurred edge crossing zero at (ideally) the position of the true original edge. Sadly, despite the flurry of excitement and the wonderful promised potential, it became easy to demonstrate that things were not as rosy as might at first seem when one considered the implications of the sampling concept in ***2D space*** (as opposed to ***any*** form of ***1D*** situation, be it spatial or temporal). In 2D space, under some conditions, although the basic match of sampling and blur was without question, one could find that the whole ***differential trend*** across the blurred edge fell to one side of zero, thereby not showing any ***'zero crossing'*** at all! See Ref. 3, Chapter 4.1 for demonstrations of this.

However, if one looks carefully at the method of transmitting data from the retina to the central cortex, this, by definition, has to be in two parts - all ***positive*** fragmentary signals in one channel and all ***negative*** fragmentary signals in a second channel (because one cannot transmit negative energy)l. Thus one has available both a positive and a negative ***portion*** of all second difference signals. I have shown that, given these two ***partial*** signals (as must be readily available at the central cortex), one can easily recombine to yield the true ***second*** difference signal and ***also*** a conceptual ***first*** difference signal (Ref. 3, Chapter 4.2.2). Now it is found that a first difference edge image is free from the danger of distortion, so by measuring a first difference signal instead of a second difference signal (i.e. sensing the position of a peak rather than of a zero crossing) the problem is averted. Furthermore, all this information is potentially available ***before*** data leaves the retina! These alternative processing options were openly published by me in the 1980's (e.g. Ref. 9) and a whole chapter of Ref. 3 (Chapter 4) was devoted to just this topic. Also it is interesting to note that this sensing of first differences fits admirably with the ***early*** ideas of visual performance being controlled by 'contrast' (or brightness difference) of a simple object against it's surroundings, despite the erroneous idea of contrast & area!

Meanwhile it had originally been assumed without question that the input to the brain from the eye was simple trichromatic - with roughly equal numbers of red, green & blue (R, G & B) cones. But in more recent times there had been a growing questioning of whether this was strictly true - or whether, instead, one should be considering a system which was essentially consisting of opponent dipoles of complimentary colours. This grew in the 1970's and then, around 1987, a first paper was published showing that the B cones, rather than being in roughly equal numbers with the R & G cones, were instead very sparse (as few as only a small handful throughout the fovea). So the question of trichromacy or opponency was suddenly weighted very strongly towards opponency. Fortunately for me, although I was then in the process of compiling my second book, I was able to take due account of at least the ***main*** facts concerning this in my treatment of colour in the said book – and I believe the other implications as discussed in the said book (Chapter 14) still hold to this day (together with additional discoveries to be covered later)..

During the 1980's, as part of my new directive from B.Ae., and with assistance from one of my colleagues who was what one might call an early computer programming 'whiz kid', I was able to construct a ***computer model*** of what I believed to be the essential component parts of the early visual system. This took account of all the then known components of the early visual tract up to Area 17 of the Striate Cortex. This model was capable of simulating monocular still vision (in full colour), stereo vision & optical flow (greyscale) and providing interactive outputs of both fragmentary edges and associated local closed, segmented regions. From this I then found it possible to provide ***very simple*** supplementary software to recreate edge maps, region maps, optical flow maps and stereo maps as necessary & appropriate, in addition to reconstructed 'cartoon' representations of original input scenes. In addition I also found it readily possible to generate maps of local profile curvature, including location of corners

As a result of the publication of my second book in 1992, I received an invitation from Massey University, North Island, New Zealand to go there (all expenses paid) to present a paper at a forthcoming International Conference. From this I obtained a limited Contract with Massey University to provide a limited version of my vision simulation software for the New Zealand Forestry Organisation.

Soon after the millennium I became conscious that many people were starting to set up personal websites and I concluded that such a course must be highly desirable for me. Hence I set about creating a website to include my business history, a listing of my open publications and, on recommendations primarily of my contact in New Zealand, a step by step demonstration of the sequences of processing which I believe to take place in the early visual tract. By 2004 this was

more or less complete and was 'published' as Ref. 4. It was found that a website provided valuable additional ways of demonstrating some aspects of imagery which I believe have proved much more useful than the printed word with static diagrams.
Meanwhile the software ***originally*** created in the late 1980's was provided with a user friendly 'front end' by my son and was packaged as a saleable item – under the title of 'Simulated Human Vision' or SHV. Sadly there have never been any 'takers'. Recently I have had reason to run a check of this software on a modern computer and find that a ***total*** processing (fragmentary edges & associated closed regions) for a single still image of size 128 x 128 pixels is completed in less than 1 second (how much less I do not know because my recording facilities do not measure less than 1 second for total process completion). It should also be noted that this software, originally written prior to even the ***existence*** of Microsoft Windows, still runs on the latest version of Windows without any need for conversion!

As far as I am concerned the ***final*** piece of a complex jigsaw of interconnections was put in place with the publicising (by a firm called 'ColorCode3D') of a new form of ***anaglyph*** viewing in the autumn of 2009. In this the spectral data were split into effectively yellow (amber) and deep blue instead of red & green (or cyan) as had been adopted for many years (since Edwin Land's early anaglyph demonstrations in, I think, the 1960's).

To set the scene for this, so to say, it is first necessary to reiterate a fact concerning the optical image which has been known of for very many years but which seems to have been largely down-played. This is the fact that, because the optics of the eye are not in any way colour corrected (i.e. consisting of a simple, single element lens), the image formed on the retina is subject to an image blur strongly dependent on the part of the spectrum being imaged - technically, a large amount of chromatic aberration. The high street optician, when one has one's eyes tested, will always optimise one's sharp vision to be balanced between red & green, thus ignoring the blue end of the spectrum completely. When the image quality is explored as a function of ***spectral*** input it is thereby found that, for the extreme blue end of the spectrum, this is up to 2 dioptres out of focus! This fact had concerned me for several decades - and I have never got a satisfactory response from the opticians (to this day).

Of course, I presume that most people had assumed through the years that what we perceive in terms of colour (including blue) was there throughout the visual processes (me included). However, at least as far as ***edge*** detail is concerned (which had been shown to be the predominant factor in creating the visual image) this could no longer be assumed to include the blue end of the spectrum (despite the strong sensations of blue in perceived scenes). How then could we have such strong perceptions of blue?

The new form of anaglyph put the problem into what, in colloquial terms, may be said to be 'sharp focus'. That is, if one looked through the amber & deep blue filters individually two facts were inescapable. When viewed through the amber filter alone the whole scene took on a distinct yellow cast (fairly obviously) but was very sharp, whereas, the scene viewed through the deep blue filter was like looking at a scene on a very dark night! Yet, if the scene was viewed through the amber & blue filters together, the yellow cast totally disappeared and the scene to all intents & purposes looked absolutely normal (as well as sharp in all respects).

To me the foregoing could not be ignored and could only mean one thing - despite the view through the deep blue filter alone looking like a dark night it was still providing all the necessary input to reconstruct the true coloured scene (edges included). So this scene reconstruction must be being provided from one or more of the ***secondary*** receptive fields. If one then assumes that the B cones are sensing a very out-of-focus blue image (which is known to be the case), then this very blurred blue image can be considered to be a component of a secondary blue / yellow opponent channel. I have done considerable personal experimentation with original stereo colour image pairs and this concept works extremely well. These studies were greatly assisted by the fact that, also in 2009, a true stereo camera was marketed by Fuji (the Fuji REAL 3D stereo camera). With this (compact) camera, which is capable of taking either high quality still stereo pairs of images ***or*** standard TV resolution video clips, I have been able to amass a substantial library of pairs of both stereo stills and stereo video clips.

As a result of my studies with the new anaglyph presentations I felt that there were a number of inevitable facts.

Firstly the blue cone input could not play a significant part in high resolution colour perception. Hence high resolution scene brightness must be controlled by ***only*** R & G cone inputs - which lead straight to an explanation of why dayglow yellow works so well. The brightness of dayglow yellow has essentially the same perceptual brightness as white.

Secondly, because the B cone outputs do not contribute to the scene brightness, for high resolution the B cone outputs cannot have ***any*** brightness.

Thirdly, because the 3D perceptual depth is powerfully shown by combination of a dark blue scene image and a yellow scene image, the 3D depth has to be coded into a ***low*** resolution scene representation.

Putting these three ideas together I came up with the idea that possibly 3D depth perception is controlled by comparison of the left & right eye very low resolution receptive fields. This idea is strongly supported by a simple experiment which anyone can try.

i) Find a view where there is both sharp foreground detail and sharp background detail;
ii) Accommodate to the foreground whilst trying to think of the background;
iii) Cover eyes alternately.

One should readily observe that the ***spatial relationship*** of the background and foreground objects is massively different for the left eye as compared to the right eye (possibly up to about 2 degrees). But this is a massive mismatch compared with the resolution of the eyes. Therefore the ***differential*** mismatch between the foreground and the background parts of the images must be being converted into a sensation of perceptual 3D depth. But in turn this must be operating because of the interaction of left and right eye ***low resolutions*** being combined.

With this in mind I attempted to create a sort of overlay of local ***differences*** of ***low*** resolution representations of the whole scene by left & right eye - and I discovered that, by overlaying this low resolution difference map on the high resolution image from one eye, I produced a ***high*** resolution image which did exhibit 3D depth (without any sign of ghosting). As I only completed this study about a couple of years ago I felt that I could no longer face all the difficulties of open publication of what must inevitably be a topic which will lead to many arguments. Hence, at least temporarily, I have presently opted to publish this instead as a series of short papers on my website. They are then available for open access for anyone who cares to be bothered to look into what must presently be a very contentious subject.

With all the variations and developments discussed above, what was originally thought to be a relatively trivial task back in the 1960's & early 1970's turns out to be a minefield! Nevertheless, looking back now, I feel that my research studies have been very valuable, if sufficient people will take the time to study the outcomes and to understand the necessary wide background behind them.

Most of the foregoing has been subject of specific studies by my team and me, since the 1970's and many open papers have been published as well as the major 1992 book. But there still appears to be virtually no acknowledgement of any necessity to make due allowance for several of the essential aspects of a real world scenario and realistic human vision in the majority of published laboratory experimental results by others, even in recent years!

References.

1. Overington I. (1963), 'The study of time and space variable surface temperatures of self-luminous bodies and other brightness parameters by colour densitometry', in 'Light and Heat Sensing' (AGARDograph 71), Chapter 12, Pergammon Press.

2. Overington I. (1976), 'Vision and Acquisition', Pentech Press, London.

3. Overington I. (1992), Computer Vision - a unified biologically-inspired approach, Elsevier North Holland.

4. 'www.simulatedbvision.co.uk'

5. Overington I. (1982), 'Towards a complete model of photopic visual threshold performance', Opt. Eng., 21, 002.

6. Overington I. (1983), 'Computer simulation of preperceptual processing of form', Proc. of the SPIE, Vol. 397, page 73.

7. Overington I. (1985), 'A simple, direct practical method of sensing local motion and analysing local optical flow', Proc. of the SPIE, Vol. 596, 145.

8. Overington I. (1985), 'A paradox - high fidelity from poor image quality', Proceedings of Machine Intelligence '85, IFS Publications, Bedford, U.K..

9. Overington I. and Greenway P. (1986), 'Accurate local form and motion extraction by a composite 1st and 2nd difference processor based on an interpretation of human vision', Proc. of the SPIE, Vol. 726, 188.

10. Overington I. and Greenway P. (1986), 'Theoretical and practical super-resolution performance in simple processing of local form, motion and stereo disparity', Proc. of the SPIE. Vol. 728, 10.

11. Overington I. (1987), 'Combined modelling of colour response and visibility', Proc. of the 21st Session of the CIE, Venice, June 1987, p. 86.

12. Overington I. (1987), 'Simple analysis of random-dot stereograms and kinematograms', Proc. of the SPIE, Vol.804, 144.

13. Overington I. and Greenway P. (1987), 'Practical first difference edge detection with sub-pixel accuracy', Image and Vis. Comp., 5, 217.

14. Bergen J.R. and Wilson H.R. (1979), 'A four channel model for threshold spatial vision', Vision Research, 19, 19.

15. Granlund G.H. (1983), 'Hierarchical image processing', Proc. of the SPIE, Vol. 397, 362

16. Marr D. (1982), 'Vision', W.H. Freeman and Company, San Francisco.

17. Overington I. and Voke J. (1977), 'Colour vision modelling - Interim report', B.Ae.D. Report ST18829.