Applied Vision Association, Annual General Meeting 31st March 2009. University of Birmingham, UK.

1. Transfer of perceptual learning in mid-level visual processing

David P McGovern, Ben S Webb & Jonathan W Peirce (Nottingham Visual Neuroscience, School of Psychology, University of Nottingham, University Park, Nottingham, NG7 2RD, UK; e-mail: jonathan.peirce@nottingham.ac.uk)

We know that human visual coding is not static, but rather a flexible, plastic process. One of the ways in which this experience-dependent plasticity is manifest is through perceptual learning - the ability to improve sensory performance on a given task after training. Previous work (e.g. Fahle, 1997, Vis Res, 37, 1885-95) has found that trained improvements in one task (for instance, in orientation discrimination) do not transfer to other, related tasks (such as the detection of curvature). We wondered whether this lack of transfer was dependent on the nature of the tasks used. We trained subjects to discriminate between two arrays of Gabor elements that differed in either; a) mean orientation, b) mean curvature, or c) global form. Before and after training on just one of these, participants were tested on all three variants and an equivalent contrast discrimination task (used to estimate the contribution of procedural learning). All groups demonstrated substantial learning on the task variant in which they were trained. However, contrary to previous reports they also showed transfer of learning to certain, closely related, variants. For instance, those trained to discriminate orientation, also improved on the curvature task, and less so on the global form task. Those trained on global form, also improved on the curvature task, but less so in the orientation domain. Contrast discrimination did not improve significantly in any of the groups. The data clearly demonstrate transfer of learning between discriminations that require observers to pool local visual information across space.

2 Probing edge blur perception with reverse correlation

Keith A May[1,2], William H McIlhagga[1] ([1]Bradford School of Optometry and Vision Science, Bradford University, Bradford, UK. [2]Present address: Department of Computer Science, University College London, Gower Street, London, UK; e-mail: keith@keithmay.org).

We investigated blur perception in human vision using reverse correlation. On each trial, subjects saw two edges: the *target* was a blurred edge with a Gaussian integral profile, and the *nontarget* was a sharp step-edge. 1D noise was added to each edge. Subjects had to identify the target. We found the mean difference between target and nontarget noise profiles for the correct trials and for the incorrect trials. The difference between these two mean noise profiles is the *classification image* (CI), which can be interpreted as the receptive field that was used to perform the task. Consistent with the N_3^+ model (a multi-scale model of edge perception (Georgeson et al, 2007 Journal of Vision **7**(13):7 1-21), our CIs approximated a Gaussian third derivative. The model filters the image with Gaussian first and second derivative operators, with an intervening half-wave rectifier; the scale, σ, of each channel is determined by the scales of its two derivative operators. Each channel's output is multiplied by σ^α. Peaks across space and scale indicate the position and scale (i.e. blur) of each edge element. For a Gaussian edge with scale σ_e, the peak occurs in the channel with scale $\sigma = \sigma_e \sqrt{1/(3/\alpha - 1)}$. In the original N_3^+ model, $\alpha = 1.5$, giving a peak in the channel matched in scale to the edge. Our CIs were wider than predicted by this model, suggesting a higher value of α. The N_3^+ model with the best-fitting α-value predicted responses on a trial-by-trial basis, and gave simulated CIs that fitted remarkably well to the psychophysical ones.

3 Neural Basis of Motion Perception and Visual Navigation

Saqib I. Khan, David S. Young, Andrew O. Philippides (Department of Informatics, University of Sussex, UK; e-mail: S.I.Khan@sussex.ac.uk)

Perception of visual motion yields important information for controlling action. Neurophysiological studies of the visual motion pathway in primates have been well documented. A number of neural models (e.g. Simoncelli & Heeger, 1998, Vision Res., 38, 743-761; Perrone & Stone, 1994, Vision Res., 1994, 34, 2917-2938) illustrate how motion-relevant information is processed at different stages within the cortex. I have developed a hierarchal system for the computation and use of optic flow information for navigation. The first stage is based on Fourier decomposition of motion signals (Adelson & Bergen, 1985, J Opt Soc Am A, 2, 284 – 299; Watson & Ahumada, 1985, J Opt Soc Am A, 2, 322-341), the second stage estimates local velocity and the third pools local information for an estimate of global parameters of the optic flow field. The model is tested on the task of heading determination for navigation. Although the hierarchical approach is not novel itself, in contrast to previous studies I use only simple neurophysiologically plausible operators and mechanisms. The model demonstrates several properties described in other studies, namely: nonlinear spatiotemporal responses (Carandini et al., 1997, J. Neurosci., 17, 8621 – 8644); robust velocity tuning (Priebe et al., 2003, J. Neurosci., 23, 5650- 5661); and wide-field pattern tuning (Duffy and Wurtz, 1991, J. Neurophys., 65, 1346 – 1359). To conclude I discuss how such an approach can be useful in computer vision applications.

4 The shape and development of oculomotor inhibition revealed by saccade trajectory modulation

Eugene McSorley, Alice G Cruickshank, L A Inman (Department of Psychology, University of Reading, Harry Pitt Building, Earley Gate, Reading, RG6 6AL, UK; e-mail: e.mcsorley@reading.ac.uk)

Selecting a stimulus as the target for a goal-directed movement involves inhibiting other competing possible responses. The state of inhibition of the competing motor response to the distractor can be seen by examining oculomotor trajectories and landing positions. Individual saccades may initially deviate either toward or away from a distractor as a function of saccade latency with shorter-latency saccades deviating toward and longer-latency saccades progressively deviating away. Here we extend this by investigating the temporal development of oculomotor inhibition of distractors presented at different spatial locations. Targets were presented with distractors presented at various locations: 10, 20, 30, 60 and 120 angular degrees from the target. Fixation was removed from display at various times relative to stimuli onset. This manipulation is known to vary saccade latency independently of the influence of competing distractors. The deviation of saccade trajectories and their landing positions were measured. At close distracter distances (10, 20 and 30 deg) both the trajectory and landing position of the saccade deviated toward the distracter position reflecting an excitatory influence of the competing motor programme. At far distances (60 and 120 deg) trajectories deviated away from the distracters reflecting long range inhibition but their landing positions were largely accurate. We suggest that this can be explained by the metrics of the saccade reflecting a dynamic coarse pooling of the ongoing activity at the distracter location: saccade trajectory reflects activity at saccade initiation while landing position reveals activity at saccade end.

5 “To please and entertain the eye": Explorations in visual aesthetics

I. Christopher McManus, (Division of Psychology and Language Sciences and Division of Medical Education, University College London, Gower Street, London, WC1E 6BT, UK; e-mail: i.mcmanus@ucl.ac.uk)

The founder of psychophysics, Gustav Theodor Fechner, was also the founder of experimental aesthetics. While the former has flourished since its inception, the latter has mostly been marooned in the scientific doldrums, neither respected as empirical study, nor producing much in the way of substantive or useful results. That though does not mean the phenomena of visual aesthetics are lacking in interest. On a day-to-day basis, many more people work as designers and artists than study vision, and few visitors to art galleries are interested in basic visual mechanisms, whereas most will venture opinions on whether the works are visually attractive or not. In this lecture I will suggest that the mature discipline of visual science is now well-placed to address aesthetic questions. I will particularly look at preferences for simple geometric forms – one of Fechner's own interests – as well as the nature of visual composition, particularly as shown in photography. The challenge in both cases, as is so often the case elsewhere in psychology, is to understand how and why there are such large individual differences in aesthetic preference. It was Fechner himself who disputed that old maxim, *De gustibus non est disputandum*, and argued instead that, "It has to be possible to argue about taste" - - and hence, by implication, to study it scientifically.

6 Do subjective equilibrium judgements correspond to physical reality?

Françoise Samuel, Dirk Kerzel, (Psychology Section, FPSE, University of Geneva, 40, boulevard du Pont-d'Arve, 1211, Geneve 4, Switzerland; e-mail: Francoise.Samuel@unige.ch)

Which visual attributes determine whether a 2D object is perceived as balanced or unbalanced? We presented irregular eight-sided polygons, termed “objects”, which had their lowest vertex resting on a support surface. Participants were asked to adjust the orientation of an object to unstable equilibrium. That is, its chances to fall to the left or to the right should be equal. In order to test the importance of the top and the base of the object, the stimuli were created with the following constraints. First, some objects were balanced when the top was exactly above the point of support; other objects were balanced when the top was eccentric with respect to the point of support. Second, some objects were balanced when the angles between the object base and the support surface were equal; other objects were balanced when they were unequal. The results showed that when the object top was above the support and the angles at the base where equal, participants were quite accurate. The unequal angles condition elicited a very slight tendency to rotate the object such that the angles between the object base and the support surface would be equal. In contrast, the eccentricity of the top induced a larger bias. Participants tended to position the top towards above the point of support, leading to a bias in the direction opposite to the eccentric top. Thus, the upper part of the object is taken into account more than physical laws require.

7 Social contexts alters perceptions of human movement: It matters where the action is going

Ava J. Senkfor (Department of Psychology, Wayne State University, 7th Floor, 5057 Woodward Ave, Detroit, MI 48202 USA; e-mail: Senkfor@wayne.edu)

Embodied perceptions encompass far more complicated components than analogous visual perceptions on the printed page. While investigations have evaluated perceptual biases and biological motion in real world environments, little is known about how social and non-social contexts alter these perceptions. Across multiple studies, participants viewed two brief animations of a real human body making a sequenced arm movement or control stimuli with non-biological motion - dot animations. Dot animations contained the same movement patterns, direction, and distances as the sequenced arm movements. Participants compared the distance traversed by the hands or dots. Experiment 1 showed superior accuracy in judging movement distance with biological motion over non-biological motion. Accuracies for both types of motion were differentially modulated by movement pattern, distance, and direction. The superior ability to detect distances from biological motion over non-biological motion may be achieved by using the body as an anchor point for distance evaluation. However, social contexts may change and/or facilitate the role of the body and alter our distance sensitivity. A second manipulation focused on the role emotion (good/bad news) has on our judgments. Participants were provided with a socio-emotional interpretation of movements in terms of good and bad news while maintaining the same movements, direction, and distances. This manipulation increased accuracy of distance judgments. Collectively, performance gains by body animations over the dot animations may be attributed to using the body as a metric for relational processing. Moreover, different contexts influence and shape our judgments and perceptions based on what information is being conveyed.

8 Blur adaptation explained by a norm-based model of contrast adaptation

Mark A Georgeson[1], Sarah L. Elliott[2], Michael A. Webster[3] ([1]School of Life & Health Sciences, Aston University, Birmingham, B4 7ET, UK; [2]Department of Psychology, University of California, Davis, CA, 95616, USA; [3]Department of Psychology, University of Nevada, Reno, NV, 89557, USA; e-mail: m.a.georgeson@aston.ac.uk)

Adapting to blurred images makes in-focus images look too sharp, and vice-versa (Webster *et al*, 2002, *Nature Neuroscience*, 5(9), 839-840). We asked how such blur adaptation is related to contrast adaptation. Georgeson (1985, *Spatial Vision*, 1, 103-112) found that grating contrast adaptation followed a subtractive rule: perceived (matched) contrast of a grating was fairly well predicted by subtracting some fraction k (~0.3) of the adapting contrast from the test contrast. Here we apply that rule to the responses of a set of spatial filters at different scales and orientations. Blur is encoded by the pattern of filter response magnitudes over scale. We tested two versions - the 'norm model' and 'fatigue model' - against blur-matching data obtained after adaptation to sharpened, in-focus or blurred images. In the fatigue model, filter responses are simply reduced by exposure to the adapter. In the norm model, (a) the visual system is pre-adapted to a focussed world and (b) discrepancy between observed and expected responses to the experimental adapter leads to additional reduction (or enhancement) of filter responses during experimental adaptation. The two models are closely related, but only the norm model gave a satisfactory account of results across the four experiments analyzed, with one free parameter k. This model implies that the visual system is pre-adapted to focussed images, that adapting to in-focus or blank images produces no change in adaptation, and that adapting to sharpened or blurred images changes the state of adaptation, leading to changes in perceived blur or sharpness.

9 Loss of the visual object recognition does not disrupt tactile recognition: evidence for direct tactile activation of visual cortex.

Harriet A Allen, Glyn W Humphreys (Brain and Behavioural Sciences, School of Psychology, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK; e-mail: h.a.allen@bham.ac.uk)

Regions of the ventral occipito-temporal cortex are preferentially activated by intact (compared to scrambled) images of objects. Touching objects (as opposed to textures) also activates these regions (e.g. Amedi et al 2001 Nat Neuro 4 324-330). We used neuropsychological fMRI to probe whether dorsal regions of the lateral occipital cortex (LO) are directly activated in tactile object recognition (without visual object recognition). We tested a patient with visual agnosia due to large, bilateral lesions of the ventral occipito-temporal cortex (HJA) but spared dorsal LO. Despite his poor visual recognition, HJA's tactile object recognition was preserved. We measured brain activity in HJA and age-matched controls (N=7, aged 74 and over) whilst they viewed and touched objects and textures. We found overlapping activity in brain regions including LO and the cerebellum for tactile and visual stimuli for control participants. These included new regions not before considered multimodal. For HJA, there were overlapping regions in the intact dorsal region of the LO. Within a subset of the multimodal regions found in control participants, HJA only showed activity for recognised tactile objects, suggesting that these regions are specifically driven by successful recognition. Activation of dorsal LO by tactile input is not secondary to visual recognition and can operate without low level visual input.

10 The role of texture gradients in figure-ground segmentation

Bertamini Marco, and Andrew C. Shepherd (School of Psychology, University of Liverpool, Eleanor Rathbone Building, Bedford Street South, Liverpool L7 7AD, UK; e-mail: m.bertamini@liv.ac.uk)

Luminance and texture gradients at a boundary are informative about which side of the contour is the foreground. Palmer and Ghose (2008, *Psychological Science, 19,77-83*) have reported evidence that these gradients affect relative depth perception and figure-ground. Their explanation is based on how surfaces self-occlude. They concluded that gradients do not simply provide volume information about a surface because gradients orthogonal to the boundary provide volume information but do not lead to that region being selected as foreground, unlike parallel gradients (i.e. when boundary and gradient have the same orientation). We report data from two experiments that provide a replication using a simple checkerboard pattern. The procedure was closely matched to that in Palmer and Ghose. However, when a difference in orientation was introduced between the gradient and the occlusion boundary (a vertical gradient and a boundary tilted 6 deg from vertical), this difference did not significantly reduce the probability that the textured region was selected as foreground. This manipulation was introduced to test stimuli in which there is an accidental relationship between gradient and occlusion boundary (i.e. their orientation is neither the same nor orthogonal). Importantly, the 0 and 6 deg versions were easy to discriminate because when pitted against each other there was a preference for one of them (the 0 deg). These findings are problematic for the original proposal because when the gradient is not consistent with the occlusion boundary the textured region that is perceived as foreground should instead be perceived as background.

11 Stochastic Search on a Homogeneous Surface Texture

Alasdair D. F. Clarke[1], Patrick R. Green[2], Michael J. Chantler[1] ([1]The Texture Lab, Mathematics and Computer Science, Heriot-Watt University, Edinburgh, Scotland, EH14 4AS, UK, [2]School of Life Science, Heriot-Watt University, Edinburgh, Scotland, EH14 4AS, UK; e-mail: adfc1@hw.ac.uk)

Visual search is often thought of as being in some sense guided (Wolfe, 2007, Integrated Models of Cognitive Systems, Ed W. Gray, 99-119). Recent work by Najemnik and Geisler (2008, Journal of Vision, 8(3), 1-14) found that human performance was close to ideal when searching for a target hiding in noise. Using similar stimuli, we present a comparison between human search performance and that of a stochastic model. We use a search task involving naturalistic continuous surface textures (Clarke et al., 2008, Vision Research, 48(21), 2193-2203) and use the results from a Signal-Detection Experiment to model target detectability. Our stochastic model makes random saccades, weighted by empirically derived saccade amplitude and direction distributions coupled with the target detection model. We compare the model against human performance in terms of the number of saccades required by seven observers to find the target, and find a close match between the two. We also use the Voronoi method (Over et al., 2006, Behaviour Research Methods 38(2), 251-261) to analyse the uniformity of the spatial distribution of fixations and find that observers perform similarly to the stochastic model. This suggests that a stochastic process is sufficient to model human search strategies, and that inhibition of return, and other memory-dependent processes, do not have a large role to play in our search task.

12 Object recognition is slowed by shadows of inconsistent orientation

Gillian Porter, Louise Stimpson, Andrea Tales, P. George Lovell (Department of Experimental Psychology, University of Bristol, 12a Priory Rd, Bristol, Bs8 1TU, UK; e-mail: Gillian.Porter@bristol.ac.uk)

Demonstrations that shadows cast by light-from-above are not readily perceived (e.g. Rensink & Cavanagh, 2004, Perception 33, 1338-1358) typically use tasks where the viewer looks for the shadow. Here we examined whether the presence of light-from-above cast shadows affects search for a pre-determined target object among distractors. Individual white mugs, differing in shape, were photographed with real shadows cast at varying orientations to the camera, and without shadows using a ring flash. Combinations of these images were used to prepare search arrays in which we manipulated shadow presence, contrast, and consistency of orientation between items. Each display involved the target mug and 2, 5 or 11 heterogeneous distractors presented on a uniform background, with all mugs presented with their own shadows, if any. The absence or presence of directionally-consistent shadows, whether of low or high contrast, did not affect search slopes, although search intercepts were lower without shadows. The consistency of high-contrast shadow directions affected search speed, however: slopes were steeper when shadow orientations were a mix of 0°, 90°, 180° and 270° to the camera than either a mix of 0°, 15°, 30° and 45° or consistent shadow orientations. We conclude that object recognition is largely unaffected by shadow presence unless shadow discrepancies are large enough to be detected by coarsely-scaled shadow

processing mechanisms

13 Social inhibition of return.

Geoff G. Cole[1], Paul A. Skarratt[2], & Alan Kingstone[3]. ([1]Department of Psychology, Durham University Science Laboratories, South Road, Durham, DH1 3LE, UK. [2]Department of Psychology, University of Hull, Cottingham Road, Hull, HU6 7RX, UK. [3]Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, B.C. V6T 1Z4, Canada; e-mail: g.g.cole@durham.ac.uk)

Inhibition of return is said to occur in order to facilitate visual search. As social beings, might inhibition in humans' be influenced by the actions of others? Indeed, recent evidence shows that when two people alternate reaching responses to one of two spatial locations, person B may be slower to respond to a location to which person A has recently responded. This finding is thought to reflect a socially-based inhibition of return (SIOR) mechanism in which one inhibits to locations on the basis of another person's behaviour. We describe a series of experiments examining: 1) whether SIOR occurs if the observer infers rather than sees where the other person has reached to; 2) how long the inhibition lasts; and 3) whether the inhibition is applied to spatial locations or to objects. SIOR was observed both when the reaching response was visible and when the response location was only inferred, suggesting that SIOR can be induced without visual representation of another person's behaviour. We also found that the effect can last for at least 2400 ms after a response is observed. Finally, we found that SIOR can be based on an object as well as a spatial representation. These findings will be discussed in relation to existing views of inhibition of return.

14 Single click image segmentation using mean shift

Dragos Bozdog[1], Ionut Florescu[1], Rustam Stolkin[2] (Department of Mathematical Sciences, Stevens Institute of Technology, 1 Castle Point on Hudson, Hoboken, NJ 07030, USA; [2]School of Computer Science, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK; e-mail: ifloresc@stevens.edu)

We present a novel method for segmenting objects in images, for which a user must first designate the object by clicking on it once using a mouse. Our motivation is to use this method for initializing automatic target tracking algorithms in surveillance applications. Following a single click segmentation of a target of interest, the features of the target image region can then be learned automatically and passed to a tracking algorithm. Alternatively, single click segmentation has potential applications to editing and manipulation of photographs and video. We note that the human visual system can rapidly decompose scenes into individual objects and can accurately determine the extent of such objects. No doubt these processes partly rely on high level reasoning with semantic knowledge, however we suggest that a substantial part of this task can be achieved with simple low level image processing operations.

Our single click segmentation algorithm is based on the mean shift procedure (Comaniciu et al, PAMI, vol. 24-5, pp. 603-619, 2002). All pixels in the image are mean shifted until convergence, according to user specified bandwidth parameters (see Bozdog et al, these proceedings). Next, the pixels from a small region surrounding the user clicked pixel, are identified as a cluster of points in the mean shifted feature space. The clicked object is defined by an ellipsoid in this mean shifted feature space, with ellipsoid axes given by the eigenvectors of the covariance matrix, and size governed by the quartile of a relevant chi-squared distribution. Additional pixels, neighboring those in the original small area are examined to see if they lie inside this ellipsoid. If so, they are added to the cluster and the ellipsoid size and shape are updated. The object region is progressively expanded by repeating this process with pixels neighboring the identified object region, until no more pixels can be added anymore.

15 Optimal parameter selection for mean shift type segmentation algorithms

Dragos Bozdog[1], Ionut Florescu[1], Rustam Stolkin[2] (Department of Mathematical Sciences, Stevens Institute of Technology, 1 Castle Point on Hudson, Hoboken, NJ 07030, USA; [2]School of Computer Science, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK; e-mail: ifloresc@stevens.edu)

All computer vision algorithms are dependent on key algorithmic parameters. The success or failure of the algorithm is often dependent on subtle variations in the choice of these parameter values. However, in the field of computational vision, the practice of researchers is predominantly to modify these values by hand until good results are achieved on a particular set of test images, and then present these as "successful" results.

In contrast, we note that the human visual system continuously and rapidly adapts itself to a wide range of different scenes, viewed objects, lighting and other imaging conditions. This suggests, firstly, that such a system possesses mechanisms for automatic online parameter selection and, secondly, that sufficient information is contained within image sequences themselves, to enable such parameter selection. This work develops a formal statistical methodology for objectively deriving such parameters in a structured manner, and is directed towards automatic parameter selection by the vision algorithm itself. We present a

case study where we examine optimal parameter selection for a single click target segmentation algorithm (see Bozdog et al, these proceedings), based on the mean shift segmentation procedure (Comaniciu et al, PAMI, vol. 24-5, pp. 603-619, 2002). We show that the same set of parameter values cannot be universally optimal for all objects in all image sequences. However, we show that within the same image sequence (situation) it is possible to have one set of parameter values behaving in an optimal way for all objects in the sequence. Furthermore, we show that it may be possible to group the image sequences by situations from the perspective of having the same optimal parameter values. We conclude that it is possible to gather preliminary information about the image sequence itself (e.g. degree of clutter, range of the color space, dispersion of colors, scene dynamics, etc.) which will indicate the situationally optimal parameter set.

16 Sex Differences in Facial Preferences

Daniel Gill (Department of Statistics, The Hebrew University of Jerusalem, Jerusalem 91905, Israel; e-mail: gill@mta.ac.il)

Previous studies have suggested that the appraisal of facial attractiveness is determined by traits such as: averageness, symmetry, sexual dimorphism (Rhodes, Annu. Rev. Psychol. 2006. 57:199–226). The current study inspects discriminative facial attributes between female and male observers in attractiveness preference task. Female and male participants were asked to give their preferences for given pairs of identical sex stimuli. The procedure was carried out in six different sessions where each session included one type of stimuli: image fragments of eyes and eyebrows, fragments of noses, fragments of lips, fragments of hairdo, low-pass filtered greyscale faces with hairdo cut-off (the blurred faces do not contain features' details but retain their spatial relations and face outline), and full images in their original format. This procedure was carried out with male-male and female-female pairs. Correlations between the covariates (eyes, nose, lips, hairdo, and low-frequency) and the whole face response reveal significant differences between sexes. Females' whole-face preferences had higher correlation with eyes, lips, and hairdo than males' preferences both in male and female faces. Male observers' whole-face preferences had higher correlation with low-frequency preferences both in male and female faces. In addition, Female observers demonstrated higher correlation of whole-face preferences with eyes, lips, and hairdo when making preference decision between male faces than when making preference decision between female faces. These results suggest that males and females use different strategies either in same-sex and opposite-sex attractiveness appraisal.

17 Looks like text messages take longer to comprehend than send.

Lisa Mayberry[1] & Kenneth C. Scott-Brown[2] ([1]School of Psychology, University of Dundee, Dundee, UK, DD1 4HN; [2]Centre for Psychology, School of Social and Health Sciences, University of Abertay Dundee, Dundee, DD1 1HG. UK; e-mail: L.Mayberry@dundee.ac.uk)

Keeble & Hazel (2004, *Perception*, **33**, 152-152) have assessed the extent to which punctuation can influence reading speed, but how does word construction itself influence the time taken to read prose? In particular does the style of abbreviated text routinely composed by 'texters' influence subsequent reading speed, and does lifetime exposure to texting affect time taken to read text passages? We recorded passage reading speeds (n=32) for two different social groups of people: students (mean age: 21 years; range: 19-24 years) and academic staff (mean age: 45 years; range: 33-64 years). Students acquired their first mobile phone under the age of 20 (mean acquisition age: 14 years; range: 12-16 years), whilst academic staff acquired their first mobile phone over the age of 30 (mean acquisition age: 37 years; range: 30-54 years). The prediction was that the latter group would be less able to adjust to the language of abbreviated text messaging, and consequently require longer to read a passage composed in this manner. Reading speed was assessed for each of four short (250 words) passages of text composed either in Standard English, contractions, phonological approximations or a mixture of approximation and contractions. Identical target words (n=45) were used in each passage but each passage comprised a different narrative thanks to variation of target word order. Reading for meaning was ensured by means of post-reading comprehension test. Across the experiment, each passage served in each condition in a counterbalanced within subjects design. Reading speeds were faster for unadulterated text than for all abbreviated passage types. Results demonstrate the deleterious effects of 'txt' abbreviations on reading speed, however the difficultly found in reading text does not appear to vary across the two populations sampled: no significant differences were found between academic staff and students.

18 Suppression and summation in contrast gain control for human vision

Tim S Meese, Mark A Georgeson, Daniel H Baker, David J Holmes, Kirsten L Challinor & Robert J Summers (School of Life and Health Sciences, Aston University, Birmingham, B4 7ET UK; e-mail: t.s.mesee@aston.ac.uk)

Over the last ten years our understanding of early spatial vision has improved enormously. The long-standing model of probability summation amongst multiple independent mechanisms with static output nonlinearities responsible for masking is obsolete. It has been replaced by a much more complex network of additive, suppressive and facilitatory interactions and nonlinearities across eyes, area, spatial frequency and orientation that extend well beyond the classical receptive field (CRF). A review of a substantial body of psychophysical work performed by ourselves (20 papers), and others, leads us to the following tentative account of the processing path for signal contrast. The first suppression stage is monocular, isotropic, non-adaptable, accelerates with RMS contrast, most potent for low spatial and high temporal frequencies, and extends slightly beyond the CRF. Second and third stages of suppression are difficult to disentangle but are possibly pre- and post-binocular summation and involve components that are: scale invariant, isotropic, anisotropic, chromatic, achromatic, adaptable, interocular, substantially larger than the CRF and saturated by contrast. The monocular excitatory pathways begin with half-wave rectification, followed by a preliminary stage of half-binocular summation, a square-law transducer, full binocular summation, pooling over phase, cross-mechanism facilitatory interactions, additive noise, linear summation over area and a slightly uncertain decision maker. The purpose of each of these interactions is far from clear, but the system benefits from area and binocular summation of weak contrast signals as well as area and ocularity invariances above threshold (a herd of zebras doesn't change its contrast when it increases in number or you close one eye). One of many remaining challenges is to determine the stage or stages of spatial tuning in the excitatory pathway.

19 Using texture to recover shading and reflectance information from real images.

Peng Sun and Andrew J Schofield (School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK; e-mail: PXS315@Bham.ac.uk)

In computer vision, shape-from-shading is the process of recovering surface orientation from luminance changes in a scene. However, luminance changes are ambiguous in real images and may be due to shading or to reflectance changes. This ambiguity is a problem for those shape-from-shading algorithms that assume uniform reflectance. Fortunately, reflectance changes are often associated with changes in other surface properties such as hue and texture. Here we present an algorithm for separating the shading and reflectance components in greyscale images based on texture variations. Our algorithm exploits the same rule as appears to be used by humans to assist in shape-from-shading tasks: luminance changes that are coincident with changes in the contrast of a visual texture are more likely to be due to reflectance changes than those that are not (Schofield et al, 2006, Vision Research, 46, 3462-3482). This in turn arises from the multiplicative nature of shading: contrast does not change with shading. We first estimate luminance gradients in a low-pass filtered version of the image. These gradients are then classified as 'reflectance' or 'shading' depending on the presence of coincident contrast changes as found from a texture segmentation algorithm. Unfortunately the estimated positions of texture edges do not always match exactly with their associated luminance changes. We solved this problem by introducing an edge-width estimation mechanism that provides tolerance to such mismatches. The final shading map is obtained by reintegrating the 'shading' gradients, while the reflectance component is obtained by dividing the original image by the resultant shading map. The algorithm can separate shading and reflectance when a texture is present and the degree of shading is not so great as to reduce texture contrast below usable levels.

20 Motion in depth from interocular differences in relative direction of motion

Christa van Mierlo[1], Eli Brenner[2], Jeroen B.J. Smeets[2]. ([1]School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK. [2]Research Institute MOVE, Faculty of Human Movement Sciences, Vrije Universiteit, Van der Boechorststraat 9,1081 BT, Amsterdam, The Netherlands; e-mail: c.vanmierlo@bham.ac.uk)

When studying manual responses to changes in slant we found that a stereoblind subject responded to changes in binocular disparity. To find out how this was possible we asked three stereoblind subjects to judge the direction in which a horizontal transparent virtual cylinder was rotating. The stereoblind subjects performed better binocularly than monocularly. In contrast to the controls, the stereoblind subjects were also better when the dots defining the cylinder were different for the two eyes. Next, we examined the influence of various perturbations on how they judged the direction of rotation from motion differences between the eyes. Although all perturbations influenced the subjects' performance, performance only dropped to the monocular level when the cylinder expanded and contracted along its axis as it rotated. We conclude that stereoblind subjects use interocular differences in relative direction of motion to judge motion in depth, and do so without matching individual points in the two eyes.

21 Learning image regularities for contour detection: when do you need a teacher?

Jiaxiang Zhang[1], D. Samuel Schwarzkopf[2] and Zoe Kourtzi[1] ([1]School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK. [2]Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square
London WC1N 3BG. UK; e-mail: J.Zhang.1@bham.ac.uk)

It has been suggested that the visual system is optimized for the detection of frequently occurring regularities that typically define contours in natural scenes (e.g., collinearity). However, our previous work has shown that observers may learn to exploit other image regularities (i.e. orthogonal alignments at an angle to the contour path) for contour detection that typically signify discontinuities. Here, we ask whether this learning-dependent optimization requires task-specific training or may result simply from frequent exposure to statistical regularities. We tested detection performance (i.e. two-interval forced choice task) for two different contour types that comprised of elements either collinear with, or orthogonal to, the contour path. Training followed one of three different procedures; observers performed the detection task with (supervised) or without (unsupervised) auditory feedback, or were exposed to the stimuli while performing a contrast discrimination task (exposure). Our behavioural findings show that detection of collinear contours improved for all procedures, while detection of orthogonal contours improved only following supervised training. Analysis of fMRI responses using multivariate pattern classification methods showed that the discrimination of contour orientations in higher dorsal and temporal visual areas was enhanced following detection-based training or exposure for collinear contours, but only following supervised training on orthogonal contours. These findings suggest that the functional optimization of visual circuits depends on the behavioral relevance of image regularities; that is, learning of regularities typical of natural contours may occur during frequent exposure, while learning of discontinuities for contour integration requires task-dependent training with feedback.