Scottish Vision Group Meeting
29 April – 1 May 2003
Abstracts

What motion transparency has taught us about visual information processing.

Frans Verstraten & Ryota Kanai, Utrecht University, The Netherlands

Motion processing is one of the most important faculties of the human brain. Currently, a lot is known about motion processing at the front-end level, often considered as the local motion stage. We are also gaining insight in higher stages of visual motion processing, often indicated by the term 'global motion processing'. One phenomenon that falls into this category is known as motion transparency or transparent motion (e.g. Snowden & Verstraten, TICS, 1999). I will review the current status and especially focus on recent findings that show that motion transparency can help us to investigate two surface grouping strategies. One based on feature integration and another on temporal synchronicity. Our stimuli consist of dots which are changing their direction periodically over a short distance at a constant velocity. This stimulus allows us to vary the relative strength of the two grouping strategies. The results show that when the timing of the direction changes is random (zero synchronicity between the dots), the configuration results in a compelling percept of streaming motion transparency that cannot be distinguished from continuous motion transparency. As the synchronicity of the dots' direction changes increases, the percept of streaming motion transparency changes drastically towards a percept of globally oscillating surfaces, that is, a pair of global surfaces is moving back and forth. The experiments also show that local signals, inconsistent with a surface interpretation, are prevented from further visual processing required for conscious perception of the dots.

The spatial extent of centre-surround suppression for complex stimuli

J. Scott McDonald, Wayne S. Smith, Yoav Tadmor, Newcastle University,
j.s.mcdonald@ncl.ac.uk

Natural images have a characteristic second order statistics since their average amplitude spectra fall with spatial frequency (f) as 1/f *alpha*, where *alpha* is about 1.0.
Previously, we have shown that the perceived contrast of a small central texture patch (1x1 deg) is suppressed when it is surrounded by natural images (4x4 deg) of varying *alpha* values and that this suppression is maximal around the *alpha* value of 1.0. Here we explored the spatial extent of this suppression.
Stimulus images consisted of a central texture patch (1x1 deg) centred on a background image. Images were either photographs or were derived from random luminance patterns (256 grey-levels), filtered to different *alpha* values. Stimuli were presented on a 21î monitor (120HZ) using a VSG graphics board. Contrast suppression was measured using a nulling paradigm.
First, we determined the *alpha* at which maximal suppression occurred using 4 different surround sizes (1.5x1.5. to 8x8 deg.). Suppression of the perceived central patch contrast increased with surround size but maximal suppression always occurred with surrounds of *alpha* = 1.0. Thus, the tuning to an *alpha* = 1.0 appears to be scale invariant at those surround sizes.
Next, we measured the spatial configuration of this suppression by surrounds with an *alpha* of 1.0 using surround stimuli of different area and of equal area but with different separation from the centre. We found that, for surrounds with the characteristic statistic of natural scenes (i.e. *alpha* = 1.0), the suppressive region is predominantly localised to 0.4 deg. from the edge of the central patch.

Ambiguous Stereograms and Prior Constraints on Correspondence Matching

Ross Goutcher & Pascal Mamassian, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: ross@psy.gla.ac.uk

Motivation: A priori rules or constraints on stereo correspondence matching have been central to the success of models of stereopsis. Several such constraints have been popular in the literature. However, little evidence exists as to how these constraints interact. Here we present a novel ambiguous stereogram and discuss its use as a tool for investigating the interaction of matching priors. **Methods:** Our new ambiguous stereogram offers alternate percepts of a single, opaque surface or two-surface stereo transparency. By alternating the contrast of horizontally adjacent image dots we are able to selectively bias matching toward either of these percepts. This allows us to quantify the relative strength of the constraints underlying each percept. We hypothesise that a single surface will be perceived more often than two surfaces, indicating that matching to the nearest disparity is a more powerful constraint than nearest neighbour matching. **Results:** Points of equal probability between one and two surface matches were obtained for each observer. As predicted, one surface (nearest disparity) matching exceeds two surface (nearest neighbour) matching for the majority of observers. Our results further suggest that same-polarity dot contrast biases – rather than limits – stereo matching when alternative matches are available. **Discussion:** We discuss the implications of these results for models of stereo matching. We also discuss further applications of our new stimulus and present pilot results from such experiments.
This project is funded by EPSRC Grant No. GR/R57157/01.

The dependence of slant perception on texture orientation statistics

Paul Warren, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: paulw@psy.gla.ac.uk

PURPOSE: Many studies have addressed the problem of extracting surface orientation from texture. Few, however, have looked specifically at the information available from the statistics of the orientation distribution of linear texels. There is evidence that observers are able to extract this information (Dakin & Watt, 1997) and use it to recover slant and tilt under orthographic projection (Witkin, 1981). Here we assess the extent to which human slant perception is biased by texture orientation statistics. METHODS: Observers briefly saw patches of texture through a circular aperture subtending 20 degrees visual angle. The texture comprised of small line segments on a planar surface. Texel orientation on the surface followed either a uniform distribution (the reference) or a unimodal circular distribution with variable bias and spread parameters (the comparison). The plane was rotated about a horizontal axis and seen under perspective projection. In each trial observers saw a reference and comparison surface and were asked to indicate which had larger slant (top end away). Predictions were obtained by computing the closed form likelihood function describing the probability that a surface slant produces a particular image under perspective projection. RESULTS: Texture orientation statistics had a significant effect on perceived slant. Observers overestimated slant relative to the isotropic case when texel orientations were biased towards the horizontal, and inversely underestimated slant when they were biased vertically. CONCLUSIONS: Human observers rely on texture orientation statistics to recover surface orientation. Our data are consistent with a Bayesian model which uses an assumption of orientation isotropy.

Estimating dot motion behind an occluder

Erich Graf, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: erich@psy.gla.ac.uk

The human visual system often encounters situations in which moving objects disappear behind an occluding surface. In such situations it might be beneficial to predict where the object will emerge on the other side of the occluder. Here we investigate the information in the scene, and in the motion of the object, that is important to observers when performing such a task.

Observers were asked to track a small dot on the screen as it moved towards a vertical occluder, which was clearly visible. The point initially moved purely horizontally. At subsequent time steps the speed was held constant while the direction of the velocity was sampled from a unimodal probability distribution which was symmetric about the current direction. As soon as the dot reached the occluder it disappeared. Observers were asked to estimate where the dot would emerge on the other side of the occluder, as well as indicate a vertical interval around that point that they were certain would capture the dot as it emerged. Between trials the variance of the direction distribution was changed, increasing or decreasing the ?randomness? of the path. Occluder width and dot speed were also manipulated between trials.

We found that observers increased the size of the capture region as the variance of the direction distribution increased. The capture region also increased in size with increasing occluder width, however, we found no effect of speed.

Human observers extrapolate motion by integrating what is known about the characteristics of the path with information about the scene. We discuss and model this data in terms of a probabilistic account of motion perception.

Attentional effects and motion-induced masking

Bettina Friedrich, Maria Hadjigeorgieva, Pascal Mamassian, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: bettina@psy.gla.ac.uk

The position of a moving object can be perceived shifted in the direction of motion when its edge is equiluminious (Ramachandran & Antis, 1990; De Valois & De Valois, 1991; Whitney, 2002). Mamassian and Adams (VSS 2002) tested whether this would lead to a masking effect, i.e. if a target positioned directly in front of a moving object would be more difficult to detect as it is masked by the motion. This could not be shown consistently. We examine here whether this previous result can be explained by competing attentional effects to the front side of the moving object by using attentional cues in order to control the allocation of attention. Four grating motion fields modulated by stationary Gaussian envelopes (Gabors) were presented around a fixation cross. A target was presented next to one moving Gabor and observers had to report its polarity (dark vs. bright). Before the onset of the trial, the location of the target was hinted by a cue which was valid 75% of the time. As expected, the valid cues led to better discrimination of the stimulus than invalid cues did. After invalid cueing, targets were consistently better discriminated when they appeared behind rather than in front of the moving Gabor. Our results suggest that failures to find a masking effect can be due to more attention being allocated to the front of the moving object. When attention was drawn away from the moving object, a motion-induced masking effect appeared consistently.

We acknowledge the support of EC grant RTN2-2001-00107.

Attention and Inference in Vision

Katherin Wilkin, School of Biology, University of Newcastle, NE2 4HH, Newcastle upon Tyne
Peter Andras, School of Computing Science, University of Newcastle, NE1 7RU, Newcastle upon Tyne, peter.andras@ncl.ac.uk
Alexander Thiele, School of Biology, University of Newcastle, NE2 4HH, Newcastle upon Tyne

We investigated the role of spatio-temporal priors (i.e., expected distributions of stimuli) on visual processing. We performed experiments with human subjects, in which subject's expectations about stimuli (an orientation discrimination task) and attentional state (full attention vs. reduced attention) were manipulated. Even in the full attention condition our results are in good agreement with theoretical predictions based on a Bayesian model of inferential information processing, where perception is systematically biased by expectations about spatio-temporal priors. Under dual task conditions, where attention was withdrawn from the orientation discrimination, we hypothesized that the effect of the priors will be increased, i.e. perception will be even more biased by priors. In other words lack of attention peaks the prior distributions, while the presence of attention flattens prior distributions. Our experimental results are in good agreement with this hypothesis.

Modifying visual priors with haptic information

Wendy Adams, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: wendy@psy.gla.ac.uk

We asked people to make shape from shading judgments. Observers appeared to be using a 'light from above' prior in distinguishing concave from convex objects.

During a training session, haptic (touch) information was presented concurrently with the visual stimulus. The haptic information was consistent with a light source direction that was shifted by 30¢ª. In a post-training session, observers' responses to purely visual stimuli indicated that they had changed their prior on light source direction towards that indicated by the haptic training.

Perception and Action in Drawing Circles

Maria Hadjigeorgieva, Bettina Friedrich, Frank E. Pollick, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: mariat@psy.gla.ac.uk

A circle viewed from the side will appear more circular than the ellipse projected onto the retina. This phenomenal regression to the real object is a basic result in shape constancy. We examined what happens to shape constancy when people are asked to draw circular shapes from different viewing directions. We addressed two questions. The first was whether the width-to-height ratio of the drawing motion would change as the viewing direction changed from frontal to oblique. The second was what form the velocity profile of the drawing would take. Of interest here is the fact that the 1/3 power law of planar drawing predicts a relationship between local velocity and shape such that the hand slows down to draw more curved sections. Moreover, that this relation provides motion at constant affined velocity so that the apparent velocity-curvature relation is invariant upon viewing direction. Participants were asked to draw with their fingertip 5 circles of radius 3 inches, centred at eye level, on a board oriented in the vertical plane while position of the fingertip was recorded at a rate of 200 Hz. Viewing was binocular and oblique angles obtained by rotating about the centre of the circle in the horizontal plane. In one condition participants received visual feedback of the drawing motion, in the other there was no visual feedback. When feedback was available participants could see the global shape revealed by a visible trace on the board. Results of global shape indicated that the width-to-height ratio

increased for the more oblique viewing directions. However this result was obtained for both feedback and no-feedback conditions, precluding a purely visual interpretation of the effect. Results of the velocity profile analysis showed that velocity was roughly constant with small variation. These results will be discussed within the context of models of visuo-motor processing which couple visual perception and visually-guided action.
We acknowledge EC grant RTN2-2001-00107

The role of spatial and temporal stimulus features in the perception of emotion
H M Paterson, F E Pollick, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: helena@psy.gla.ac.uk
Although conventional form cues are not present in biological motion displays, there are still spatial as well as temporal features in such stimuli. In this presentation I will show how these temporal and spatial features might act independently on perception of emotion. Previously (Pollick et al, 2001 Cognition 82 B51 - B61) we showed that emotions such as angry and sad are discriminated primarily according to temporal stimulus features, in comparison angry and happy are discriminated primarily according to spatial features. Here we extend this research by morphing movements separately in time and space to show that observers can tell apart sad and angry exclusively according to movement duration while using spatial features to discriminate angry and happy.

Learning Biological Motion
Frank E. Pollick, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: frank@psy.gla.ac.uk

Invited Lecture by Julie Harris
How is binocular vision used for seeing 3-D motion?
Julie Harris, University of Newcastle

Several theories and models have suggested how the tiny differences between right and left eye views of the world (binocular disparity) can be used to help us see objects moving in depth. But can the human visual system use this information? In this talk I will outline what binocular information, in principle, could be used to estimate the direction of 3-D motion. Then I will describe experiments where we have measured both the accuracy and precision of human performance, including some using real objects moving in depth. Large and consistent errors in accuracy were found for almost all observers. This suggest that binocular visual information available is not used optimally in human vision, certainly for some 3-D motion tasks.

Dichoptic motion on different depth planes?
Martin Lages and Alexander Dolia, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: martin@psy.gla.ac.uk

Recent computational and neurophysiological findings suggest the existence of binocular complex cells that jointly encode motion and stereoscopic depth. Dichoptically presented sinusoidal gratings flickering in spatio-temporal quadrature can induce directional motion perception but the existence of an early binocular motion system is a debated issue. Here we investigated perception of dichoptic motion for stimuli defined on different depth planes. Discrimination of motion direction was best when spatial phase offset of the flickering gratings and binocular disparity resulted in spatial quadrature. No reliable discrimination performance was obtained for stimuli with static pedestal, flicker frequencies beyond 4 Hz, and disparities outside Panum's fusional area. Therefore the present result does not support the notion of an early binocular motion system that can integrate dichoptic motion on different depth planes.

Computational motion-stereo model and dichoptic motion stimuli
Alexander N. Dolia, Martin Lages, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: a.dolia@psy.gla.ac.uk

We used dichoptic motion stimuli to investigate a computational motion-stereo model (Qian & Andersen, 1997 Vision Research, 37, 1683-1698). The performance of the model is analysed for a set of model receptive fields tuned to different interocular phase differences and velocities. Despite well-established empirical evidence that sinusoidal flicker presented in spatio-temporal quadrature in a Wheatstone configuration is perceived as either left or right motion our simulations demonstrate that the integrated motion-in-depth model fails to produce the same result. For dichoptic motion the motion-stereo model predicts two opposite motions in different depth planes. Therefore, if the model is indeed correct, there ought to be a higher level system that makes a decision about motion and depth based on the output of the motion-stereo model.

Motion Transparency and Depth Perception
Pascal Mamassian, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8PG, Scotland, UK; email: pascal@psy.gla.ac.uk

Local and global interactions in motion aftereffects
Nicholas J. Wade, University of Dundee

The motion aftereffect (MAE) can be observed following protracted observation of a pattern that is translating, rotating, or expanding/contracting: a static pattern appears to move in the opposite direction (reviewed by Wade and Verstraten, 1998). The phenomenon has long been known, and it continues to present novel properties. One of the novel features of MAEs is that they can provide an ideal visual assay for distinguishing local from global processes. Motion during adaptation can be induced in a static central grating by moving surround gratings; the MAE is observed in the static central grating but not in static surrounds. The adaptation phase is local and the test phase is global. That is, localized adaptation can be expressed in different ways depending on the structure of the test display. These aspects of MAEs can be exploited to determine a variety of local/global interactions. The following questions are examined in a series of experiments. How local is the adaptation effect? Is the storage of MAEs influenced by local or global factors? What influence does phase alternation have on the interocular transfer of MAEs? Are attentional effects localized to the adapted or inducing region? Do similar interactions apply to MAEs from translation and rotation?