Shape-from-shading and light source estimation in humans

By

Giacomo Mazzilli

A thesis submitted to

The University of Birmingham

For the degree of

DOCTOR OF PHILOSOPHY

School Of Psychology
The University of Birmingham
26/09/2013

Abstract

Light source estimation is very important for the interpretation of shape-from-shading by humans. We used a range of methods to characterise the way in which the type, and position of the light source can influence observers' performance in shape-from-shading tasks.

Firstly, we used classification images to discover people's priors for light source position using noise only stimuli. This cue-free approach uncovered the weakness of the light-from-above prior. We also examined the effect of varying the light source elevation on the perceived shape of isotropic and anisotropic surfaces, the impacts of lighting ambiguities on shape-from-shading and, finally, the interpretation of shadow regions.

We found that lighting priors are weighted by the visual system in a way that is inversely proportional to the strength of lighting cues in the stimuli, revealing that knowledge about the light source position is critical to perceiving shape-from-shading. Where ambiguous cues to lighting direction are present human vision seems to favour local cues over distal ones. We also showed that perceived surface shape varies with light elevation only in so far as elevation alters contrast. Finally we show that human vision does not treat shadows in the same way as objects.

Dedication

To my beloved parents and the way they supported me through this long trip.

Acknowledgement

I would like to give my heartfelt thanks to my supervisor, Andrew Schofield for the support and the generosity during my time as his graduate student.

I would also like to thank all the people of the School of Psychology at the University of Birmingham, my fellow students - Aidan, Kate, Katerina, Claudio, Max, Carlo, Karim, Matt, Satoshi, Peng - and my favourite (and only) labmate: Dicle. Conversations, meeting, lunches and trips with them have been an important part of my life.

I also would like to thanks all my good friends outside Birmingham that made me feel at home everyday even if miles away, and last but not least, Chiltern Railways for never letting me down on my weekly trips to London.

Most of all, I would like to thank Yannick just for being with me through it all, from beginning to the end.

GLOSSARY 8

1.INTRODUCTION 13
1.1 PHYSIOLOGY OF VISION 15
1.1.1 The Human Visual System *15*
1.1.2 Eye *16*
1.1.3 Retina *16*
1.1.4 Lateral Geniculate Nuclei *18*
1.1.5 Primary Visual Cortex *19*
1.1.6 Other Cortical Areas *21*
1.1.7 Shape from shading *22*
1.2 A 3D PERCEPTION MODEL 23
1.3 DEPTH PERCEPTION 25
1.4 AN INTRODUCTION TO SHAPE FROM SHADING. 28
1.4.1 Shading vs shadows *28*
1.4.2 Human processing of shape from shading *32*
1.4.3 The Perception of Shape-From-Shading *35*
1.4.4 The Role of Shading *37*
1.4.5 Boundaries and Contours *40*
1.4.6 Material Properties *43*
1.4.7 Type of Stimuli *46*
1.4.8 Role of Illumination *48*
1.4.9 Knowledge of the Light source *50*
1.4.10 The Type of Light Source *55*
1.4.11 Role of Shadows *57*
1.5 BAYESIAN AND PRIORS 59
1.6 RENDERING SCENES IN COMPUTER GRAPHICS AND EXPERIMENTATION. 64
1.6.1 Physics of illumination *64*
1.6.2 Ray Tracing *65*
1.6.3 Bi-directional reflectance distribution functions (BRDF)s *66*
1.6.4 Blinn-Phong model. *67*
1.6.5 Properties of the illuminants. *68*
1.6.6 Inter-reflections and ambient lighting. *71*
1.6.7 Summary *72*
1.7 GENERAL METHODS 73
1.7.1 Apparatus *73*
1.7.2 Classification Images *75*

2. A CUE-FREE METHOD TO PROBE HUMAN LIGHTING BIASES 89
2.1 INTRODUCTION 89
2.2 METHODS 94
2.2.1 Participants *94*
2.2.2 Apparatus *94*
2.2.3 Stimuli *95*
2.2.4 Procedure: *95*
2.2.5 Analysis *98*
2.3 RESULTS 102
2.4 CONCLUSION 109

3. THE EFFECTS OF LIGHTING DIRECTION AND ELEVATION ON JUDGEMENTS OF SHAPE-FROM-SHADING. 113
3.1 INTRODUCTION 113
3.3 METHODS 117
3.3.1 Participants *117*
3.3.2 Apparatus *117*
3.3.3 Stimuli *118*
3.3.4 Procedure *121*
3.3.5 Analysis *124*
3.4 RESULTS 126
3.5 MODELLING 146
3.5.1 Models *146*
3.5.2 Comparison of models *148*
3.5.3 Model results *148*
3.6 DISCUSSION 150
3.7 SUPPLEMENTARY INFO 153
4. IS THERE ONLY ONE SUN? 164
4.1 INTRODUCTION 165
4.3 METHODS 172
4.3.1 Participants *172*
4.3.2 Apparatus *172*
4.3.3 Stimuli *173*
4.3.4 Procedure *174*
4.4 RESULTS AND DISCUSSION 177
4.4.1 No Structure Condition *177*
4.4.2 Full Structure Condition *179*
4.4.3 No Local cues condition *184*
4.5 CONCLUSIONS 189
4.6 SUPPLEMENTARY FIGURES 192
5. THE ROLE OF SHADOWS ON VISUAL SEARCH 194
5.1 INTRODUCTION 195
5.2 MATERIALS AND METHODS 199
5.2.1 Participants *199*
5.2.2 Apparatus *199*
5.2.3 Stimuli *199*
5.2.4 Procedure *200*
5.3 RESULTS 202
5.4 DISCUSSION 205
6.DISCUSSION 208
6.1 THE PREVALENCE OF THE LIGHTING FROM ABOVE PRIOR. 209
6.2 THE EFFECTS OF LIGHT SOURCE ELEVATION 211
6.3 THE EFFECT OF MULTIPLE LIGHT SOURCES 215
6.4 THE INTERPRETATION OF SHADOWS 216
7. CONCLUSIONS 218
BIBLIOGRAPHY 222

Glossary

2.5D sketch: A state in visual processing proposed by David Marr in which the visual system is presumed to have computed surface orientation and depth but has yet to construct a full 3D, viewpoint independent, representation of the object.

Albedo: the diffuse reflectivity of a surface. The amount of light a surface reflects due to its diffuse – versus –specular properties. The diffuse reflectance of a surface.

Attached shadow: see also Shadow and shading: An attached shadow is a shadow on a surface caused by the surface itself. The difference between an attached shadow and shading is that in the former case the part of the surface in shadow received no light directly from the surface whereas a shades surface receives some direct illumination. Attached shadows may thus be regarded as special cases of shading. However, whereas shading often presents luminance gradients that are informative about surface orientation and curvature attached shadows do not as surface orientation can change within a shadowed region without any associated luminance change.

Beholder's share: a term coined by Jan Koenderink to describe the role of individual differences in shape perception.

Brightness: Perceived intensity of a light source or reflective surface

Cast shadow: see also Shadow: A cast shadow is a shadow on one surface caused by the presence of another object between the light source and the shadowed surface.

Collimated lighting: Parallel rays of light, light sources with no or minimal scatter distant light sources (eg the sun) are effectively collimated.

Crater illusion: an illusion (first described in print by Rittenhouse 1786) whereby a concave surface is seen as if convex due to being illuminated from an unexpected direction.

Diffuse illumination: Strictly diffuse illumination occurs when an object or scene received equal amounts of light from all directions. In practice this never occurs but there are cases, such as the light from the sky on an overcast day, where certain objects or surfaces, such as the ground, are effectively under diffuse illumination as all visible parts of the surface are illuminated equally from all directions.

Directional illumination: see also collimated lighting: Light that comes from one direction – unlike a punctuate source directional lights to do radiate light in all directions.

Frontal lighting: Light coming from directly in front of a surface. Light from the observer's own viewpoint; the two being the same if the surface normal is pointing directly at the observer.

Gloss: Glossy: colloquial for specular.

Illumination: The amount of light emitted by a light source/reaching a surface

Lambertian: Of a surface: A surface that reflects light equally in all directions. Although physically impossible to achieve many shape-from-shading algorithms assume that surfaces are perfectly Lambertian.

Light field: the pattern of light arriving from all sources (both radiant and reflective) at a given point is space.

Lighting from above prior (assumption): A bias towards perceived objects as if lit from above when there are insufficient cues in the sense to determine the actual direction of the light source.

Lighting prior: that direction of lighting assumed by an individual when there are insufficient cues to the actual lighting direction. Most people have a lighting prior that is above their own head but on average there may be a leftward bias to the lighting prior.

Lightness: Perceived or apparent reflectance of a surface, as in the albedo.

Lightness constancy: the ability of humans to maintain constant the perceived reflectance of a surface despite changes in illumination.

Luminance: Photometric intensity of a light source or surface.

Penumbra: see also shadow and umbra: A region of partial shadow caused by partial occlusion of a light source with finite area such that the area within the penumbra received light from some (but not all) parts of the light source.

Prior probability: Used formally in the Bayesian statistical estimation framework a prior probability expressed the uncertainty about an outcome before any evidence has been obtained. A 'flat' prior probability distribution assumes that all outcomes are equally probably (maximum uncertainty). In Bayesian models of perception a prior probability distribution that favours some outcomes over others may be used to express an inbuilt bias (or assumption) about the world.

Prior assumption: Informal expression of the notion of an inbuilt perceptual bias for certain outcomes.

Punctate light source: Strictly punctate illumination occurs when a light source emits light from a single point in space such that the light striking any given surface comes from one direction only. In practice this never occurs but is nonetheless a common assumption in rendering and image interpretation algorithms.

Reflectance: the proportion of light reflected by a surface.

Slant: the amount by which a surface or light source is oriented away from the frontal plane. That is a surface with zero slant will have its surface normal pointing at the observer. Slant is normally expressed in degrees of angle over the range 0-90° or 0-180° if some suppose that a surface might point away from the observer.

Shading: see also attached shadow: Shading describes the variations in the illumination received by a surface due to its orientation with respect to a directional light source. Shaded surface receive some light form the light source and are thus not in shadow but the intensity of the light received depends on the angle between the light source and the surface normal. Shading and shadows are however mathematically very similar especially when the light source has finitely small area (see umbra and penumbra) given that shadowed surface normally receive some indirect light. Both induce a multiplicative attenuation in the luminance of the object. Shading and penumbra both describe regions of partial illumination but differ as follows: in shading the partial illumination results from the orientation of the surface and can occur even when the light source is punctate (infinitely small) whereas penumbra result from the partial occlusion of a light with finite area such that the surface within the penumbra received light directly from some parts of the light source.

Shadow: Shadow describes a variation in luminance due to the occlusion of a light source. Typically when a surface is in shadow it received no light directly from the light source in question although it may receive light from other sources or indirectly via reflection off other surfaces.

Shape-from-shading: The process of inferring surface orientation from the pattern of luminance variations caused by shading.

Stereoscopic disparity: The difference in the retinal location of the image of a point on an object relative to the fovea. It is this difference that enables the visual system to estimate the relative position of each point in 3D space.

Specular: of a surface: a surface that reflects light in one principle direction relative to its own orientation. For a curves surface the observer will only see the specular reflections that happen to coincide with their own viewing direction.

Hence specular reflections appear as relatively small bright highlights on a surface. Mirrors a special case of specular surfaces.

Specular highlight: an apparent highlight on a surface caused by its specular reflectance component.

Specular reflection: The specular component of the light reflected from a surface.

Surface normal: Imagine a drawing pin attached to a surface with its flat head glued parallel to the surface. The pin will point in the direction of the surface normal.

Tilt: the direction, as if around a clock face, in which a surface or light source is oriented. Expressed in degrees of angle over the range 0-360° where 0° and 360° are identical directions. See also slant.

Tilt and slant: an angular coordinate system used to describe the orientation of a surface in 3D space.

Texture gradient: A change in the properties of a texture over space (Eg a change in texture element size or orientation). Often due to the orientation or changes in orientation of the surface onto which the texture is “painted”.

Umbra: see also shadow and penumbra: The portion of a shadow resulting from complete occlusion of a light source.

1.Introduction

The exceptional ability of the brain to reconstruct object shapes and surface reliefs has been matter of study for centuries. Some early studies in vision (Rittenhouse, 1786; Brewster, 1826) have shown compelling evidence that the visual system can misperceive the shape of surfaces in the presence of ambiguities. Perceiving shape-from-shading is a complex process that involves the reconstruction of a 3-dimensional object based on a 2-dimensional pattern of visual input information.

Studies of shape-from-shading have made extensive use of ambiguous stimuli in most experiments to unveil the strategies used by the brain in visual perception (Ramachandran, 1988; Sun and Perona, 1998; Mamassian and Goutcher, 2001; Adams, Graf and Ernst, 2004; Todd and Mingolla, 1983; Langer and Bulthoff, 2001). Even though shading is a reliable cue to 3-D perception, there are many other visual features that can improve or influence the perception of shapes: shadows, contours, edges, material properties, type of light and position of the light source. This thesis deals only with lighting and shadows and their contributions will be discussed below where we also make some reference to the influence of contours (the bounding edges of objects) and material properties (the properties of surfaces such as texture and colour).

The aim of this thesis is to understand the various way in which illumination can influence the perception of shape-from-shading. In fact, the type of illumination (directional Vs diffuse) and the illuminant position can greatly affect the shading pattern of a surface and lead the visual system to non-veridical shape interpretations. Nevertheless, the visual system can reliably perceive shapes in situations where the information is degraded or missing, thanks in part to prior knowledge.

This thesis will first attempt to clarify the role of lighting priors, then will focus on understanding how the position of the light source can affect depth perception and then investigate the brain's ability to cope with lighting ambiguities. Finally in chapter 5 the emphasis shifts towards understanding how the visual system treats shadows.

1.1 Physiology of vision

The human visual system is a complex apparatus that is able to produce a reliable and efficient representation of the natural world. In this chapter we aim to provide an anatomic and neurophysiologic overview of the system and a brief overview of visual processing of natural images.

1.1.1 The Human Visual System

Anatomical studies have helped out to map the keys areas involved in vision. These areas include indeed the eye, lateral geniculate nucleus (LGN), striate cortex (V1), and high-level visual areas such as V2,V3,V4, the medial temporal cortex (V5, MT), inferior temporal (IT) and posterior parietal cortex. We therefore provide a brief overview of each module/key areas.

The sketch in figure 1.1 depicts a simple example of the connections between the key areas.

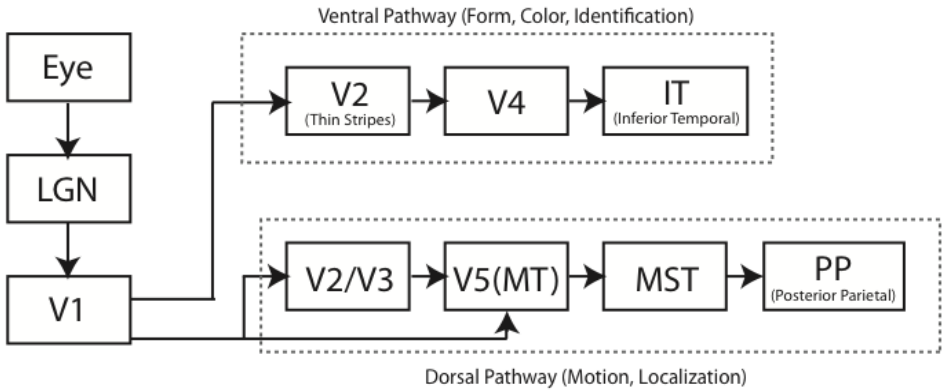

Figure 1.1 Sketch of key processing areas of the Human Visual System: the signal output from the eye is forwarded to the LGN to striate cortex (V1). Then the signal could go to the Ventral pathway or the Dorsal pathway. The ventral pathway is composed of V2, V4 and IT and it is believed to handle colour, form and recognition tasks. The Dorsal is composed of V2, V3, V5, MST and PP and it is believed to handle motion and localization.

1.1.2 Eye

The role of the human eye is to serve as a projector and converter for light into neural activity. Light enters the cornea, passes through the aqueous humor, then through the lens into the vitreous humor, and finally onto the photoreceptors (rods and cones) located at the back of the retina. At the centre of the retina, the fovea, is a region that contains the greatest density of cones and thus the highest acuity for spatial and color vision.

1.1.3 Retina

The retina is the place where the light signal is transformed to neural signal. It is composed of five layers: photoreceptors, horizontal cells, bipolar cells, amacrine cells, and ganglion cells. There are two types of photoreceptors: rods and cones and both are responsible for transducing light into neural activity. The rods are

responsible for vision in low-light (scotopic) conditions, whereas the cones function under normal (photopic) lighting and are responsible for color vision. The signal then passes to bipolar cells which receive input from the photoreceptors and provide output to the ganglion cells. Finally horizontal cells provide lateral connections between photoreceptors, and amacrine cells between bipolar cells and ganglion cells.

With regard to spatial processing, bi-polar cells constitute the first stage of this complex process but, as they form the output of the eye, retinal ganglion cells are more important when considering the impact of retinal processing on the rest of the visual system. De Monasterio and Gouras (1975) used extracellular recording techniques to show that the receptive fields of ganglion cells consist of an antagonistic center-surround organization. This type of organization divides the ganglion cells in two types: on-center ganglion cell contains a central disc-shaped excitatory region, flanked by an annulus-shaped inhibitory region; while off-center ganglion cell exhibits a reciprocal spatial organization. The first step of spatial-frequency selectivity starts from the variety of spatial magnitude of these receptive fields. Given the antagonistic center-surround organization of ganglion cells, they respond to ratios of luminance, or contrast rather than to absolute luminance.

Ganglion cells have also been classified according to the layer in LGN to which their outputs are directed. As we will see in the next paragraph, there are M cells, with a characteristic high sensitivity to contrast, but reduced sensitivity to color which receive input from both rods and cones. On the other hand, P cells,

receive input from cones only and have a characteristic high sensitivity to color, but reduced sensitivity to contrast.

1.1.4 Lateral Geniculate Nuclei

The outputs of retinal ganglion cells travel through the optic nerve, and then pass through the optic chiasm where the outputs received from the ganglion cells are here directed to opposite hemispheres. The neural outputs could then link to superior colliculus, a region believed to be responsible for eye movements; or lateral geniculate nuclei (LGN), a portion of the thalamus composed of six laminar sheets of neurons that are divided in two different classes depending on the input they receive from ganglion cells.

Some neurons receive input from the M ganglion cells, and constitute two layers of the LGN: the Magnocellular layers. Here neurons with large receptive fields are responsible for coding achromatic contrast. Neurons in the other four layers, are called Parvocellular layers. Here neurons have smaller receptive fields and a high degree of spatial resolution; these parvocellular neurons, which receive input from the P ganglion cells, are highly sensitive to color, but as opposite to Magnocellular layer, they have a reduced contrast sensitivity and lower temporal resolution.

1.1.5 Primary Visual Cortex

After the LGN, the visual signal reaches the visual areas of cortex (visual cortex), with the majority of the projections from the LGN linking directly to the primary visual cortex (V1; also called striate cortex). The LGN of monkey does not project to cortical areas other than V1 (see Hendrickson, Wilson, and Ogren, 1978). V1, which is located in the posterior region of the occipital lobe, contains approximately 200 million neurons, more than 100 times the amount found in LGN and it is, in fact, the largest area of visual cortex.

V1 is composed of six major layers, with its fourth layer divided into four sub-layers (4A, 4B, 4Cα, and 4Cβ). Most input from LGN is directed at layer 4C (magnocellular input to layer 4Cα, parvocellular input to layer 4Cβ), which then sends the processed signals on to layers 2, 3, and 4B. Layers 2 and 3 contain blob and interblob regions; the blobs, which contain color-selective neurons, receive both parvocellular and magnocellular input (via layer 4Cβ; pri- marily parvocellular input); whereas the interblobs, which contain neurons that are sensitive to orientation but largely insensitive to color, receive only parvocellular input (via layer 4Cβ). Layer 4B receives magnocellular input (via layer 4Cα) and contains neurons which demonstrate both orientation selectivity and selectivity for direction of motion.

Hubel & Weisel (1968), in what is considered a seminal work in neurobiology, investigated the receptive fields of neurons in primary visual cortex of cat and monkey via extracellular recordings. Their results shown a difference in some of

the V1 cells and they classified them in three different classes: simple, complex, or hypercomplex based on their neural tuning characteristics and their degrees of nonlinearity:

• Simple cells, which comprise approximately 25% of V1 neurons, have elongated receptive fields and they are primarily selective for bars or edges of specific widths and orientations and give rise to both spatial-frequency and orientation tuning (Schiller, Finlay and Volman, 1976; Gilbert, 1977; De Valois, Yund and Hepler, 1982). Hubel and Wiesel (1968) proposed that their receptive fields might be constructed based on the outputs of several LGN neurons. They therefore called them edge and bar detectors. This classification is based on the fact that the response of simple cells to complex patterns can be "simply" predicted based on their responses to the pattern's constituent spots of light.

• Complex cells, are in fact cells that exhibit highly non-linear responses. Similarly to simple cells, the majority of complex cells respond to oriented bars and edges, but mostly with a preferred direction of motion. However, they do not respond to the individual spots of light that comprise the bars and edges. Furthermore, complex cells demonstrate some invariance about position since their response is not affected by small perturbations in the location of a stimulus. The receptive fields of complex cells have been proposed to be constructed based on the outputs of several simple cells. As these cells have similar spatial-frequency and orientation tuning characteristics, they would produce position invariance and direction-of-motion selectivity.

• Hypercomplex cells are the most selective cells in V1 and they have a characteristic preference for stimuli of limited length. In fact, their response to a bar or edge beyond the neuron's preferred length decreases and this phenomenon is defined as "end stopping". This phenomenon has suggested that hypercomplex cells could also be considered as end-stopped simple or complex cells.

1.1.6 Other Cortical Areas

The signal that comes to V1 is then sent to the next visual area V2 and on to the other higher-level areas of visual cortex, this time not necessarily in a serial fashion (Livingstone and Hubel, 1988; Zeki and Shipp, 1988). V2 is composed of three types of stripes: pale which receive input from interblobs, thin which receive input from blobs and thick stripes which receive input from layer 4B in V1. This distinction is important as each of those type of stripes is believed to process a different type of information, respectively: form, color, and depth information. The different path that is taken from the visual signal from V2 to the higher visual areas can define separate processing. In fact, output from the thin stripes of V2 are sent to V4 the majority of whose own output is directed at inferior temporal cortex (IT); this path define what is often called the "what" system as it is believed to be responsible for form perception and recognition. Output from the thick stripes of V2 project to medial temporal cortex (V5, MT; responsible for processing stereo and motion), which in turn projects to medial superior temporal cortex (MST; responsible for visual tracking), and then to

posterior parietal cortex (PP); this path is known as the “where” system which is believed to handle localization.

The mechanisms of these extrastriate cortical areas are mainly defined by their task specificity. A study from Lee, Sing, Mumford, Romero and Lamme (1998) suggests that the higher areas perform complex tasks such as pattern analysis and object recognition working in conjunction with V1. Rao & Ballard (1999) have suggested that higher levels function as predictive coders whose feedback connections to V1 carry the prediction and whose feedforward connections from V1 convey the prediction’s error. However, on the whole, the functional mechanisms of extrastriate visual cortex (and even much of V1) remain largely unknown.

1.1.7 Shape from shading

Within the network of cortical areas the caudal ITG areas in particular seem to play a particular role in shape from shading (Svetlana, Todd, Peeters, and Orban (2008). These areas are close to area MT/V5 suggesting that shape from shading may be a dorsal ‘where’ stream property. However, if shape from shading is processed in ITG there must be some feedback to other areas as we know that shape from shading and shape from stereoscopic information are intergrated in area V3b/KO (Dovencioglu, Ban, Schofield and Welchman, 2013)

1.2 A 3D perception model

The human visual system's ability to reconstruct 3D shapes has been a matter of study for a long time. In fact, the pattern of light that is projected onto the retina is a complex function of the scene and its lighting could lead to an infinite number of interpretations. Yet, animals and humans can interpret such ambiguous 2D information in a coherent manner to produce valid 3D percepts most of the time.

Psychology, neuroscience, computer vision, physics and mathematics have all developed their own way to approach the complexity of visual perception and to try to understand the mechanisms that underlie the process of 3D perception in humans. One starting point to understand how these processes are organized is to introduce some general notions about the visual system.

Many scientists have tried to model the organization of the visual system, and David Marr's principles of modular design (1982) have been widely recognized as the reference for most computational theory. Marr identifies the visual system as a modular information processing system in which the image input is encoded very roughly at an early stage (the raw primal sketch), then information is separated out with regard to different stimulus dimensions and dispatched to different specialised areas (modules) where they are analysed and, finally, combined to form an object representation (the 2.5D sketch). Modularity is key to this proposed design, but some alternative theories (such as Lennie, 1998)

suggest that the various different attributes of an image may be not separated and processed by different areas but that the analysis of the signal is connected at all stages.

It is important to note that in Marr's model, the final 2.5D sketch incorporates information about depth and surface orientation from an object centered point of view. At a later stage, objects are transformed to a universal (in terms of view point) 3D representation of the object. A major assumption, in this interpretation, is that surface orientation and distance in space are fundamental parts of the ultimate 3-dimensional percept. This feature of Marr's model has been the subject of recent debate especially from Pizlo (2008). In his book, "3D Shape: Its unique place in visual perception", Pizlo evaluate the key role of figure-ground organisation in machine vision and claims that Marr 2.5 sketch was an attempt to avoid this problem. In fact, Pizlo's point of view is that humans do not take surface orientation into account in the perception of object shape. However, even though it is not certain that surface orientation and distance in space are the most important features for 3D perception, it would seem likely that they play an important role in 3-D shape interpretation and specifically in the judgement of surface undulations within an object as opposed to overall object shape as defined by bound contours.

1.3 Depth perception

We have seen before that light enters through the cornea and is finally projected onto the retina. This optical information that reaches the retina is firstly processed as a limited two-dimensional image. As depth information is absent in the image recorded by either eye, the human visual system has to employ strategies to recover the lost depth information when constructing a 3D percept. While stereoscopic disparity, the difference between the location of an object projection across the left and the right eyes' retina, provides an obvious cue to depth, there are other cues that allow strong 3D percepts to be formed even in the absence of disparity information. These cues are often revealed in artworks where certain techniques are used to produce a 3D percept from a 2D image, taking advantage of pictorial cues that play a crucial role in visual processing. This class of visual cues includes contours, texture gradients, perspective and shading. A good example (in Figure 1.2) to highlight how strongly the brain relies on pictorial depth cues is displayed in certain artworks such as Escher's paintings, Bridget Riley's art or impossible objects (Penrose and Penrose, 1958). Early studies have found how much human observers can rely on single visual cues such as texture gradients (Gibson 1950), to produce strong and stable depth perception. In this study, Gibson showed how texture gradients not only provide information about depth, but they can also inform observers about the orientation of a surface in depth and about its curvature.

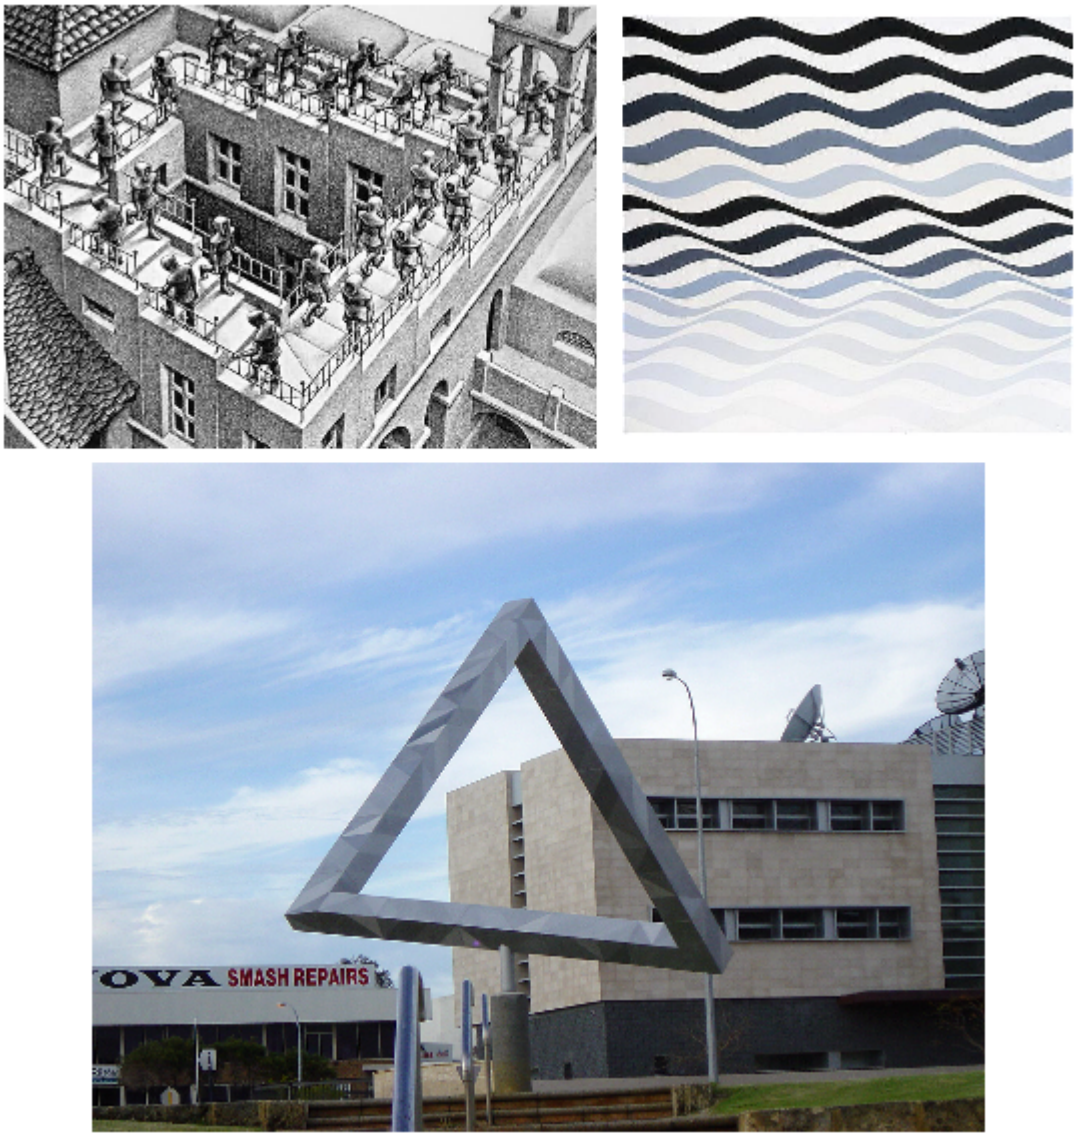

Figure 1.2. From top left, clockwise: M.C. Escher, Ascending and Descending (particular), (1960); B. Riley, Arrest 3, (1965); Roger Penrose, Impossible triangle sculpture, Eest Perth, Western Australia.

Marr's (1982) theory of visual perception suggests that the various depth cues are processed in different modules prior to some combination processes. It has been shown how human observers apply some recurrent methods (such as the ability to estimate local orientation away from the contour, the assumption that a surface patch is more likely to be elliptic or smoothness a priori constraint) to take advantage of all the possible combinations of depth cues that are present in the scene when perceiving 3D objects (Curran and Johnston, 1994; Mamassian and Kersten, 1996; Koenderink et al., 1996a). Vuong, Domini and Caudek (2006) have reported that the visual system can use shading information to drive the disparity module in the processing of depth estimation. This synergy across individual depth cues modules is critical in determining shape judgments since each of these visual cues are known to have their own domain. For example, contours can easily solve convexity/concavity ambiguities, shading provides more information about orientation and surface curvature while stereopsis discloses depth information more directly. In a more recent study, Lovell, Bloj and Harris (2012) demonstrated how shading cues and binocular disparity are optimally combined by the visual system in order to perceive shape. In particular, they found that, in opposition to more sophisticated and complex model, the resulting perceived shape could be well predicted by a model of maximum likelihood estimation of cue integration.

Evidence for this modular design has been reported by a recent imaging study that has localized a specific region (KO) to be responsible for the 3-dimensional

representation of objects (Tyler, Likova, Kontsevich, Wade, 2006). Other studies that associated psychophysical and imaging data have identified a subset of this area (V3B/KO) that plays an important role in the combination of single depth cues such as shading and disparity cues (Dovencioglu, Ban, Schofield and Welchman, 2013) and disparity and motion depth cues (Ban, Preston, Meeson and Welchman, 2012).

1.4 An introduction to shape from shading.

1.4.1 Shading vs shadows

Before we consider shape from shading we must consider the difference between shadows and shading. Shadows and shading are almost indistinguishable mathematically and their functional differences in real world scene is also rather nuanced. Here we adopt a somewhat extreme definition of shading because it fits well with our discussion of and experiments on shape-from-shading. A shadow is an area of relatively low illumination due to some objecting obscuring the path of light from a light source to a surface. The surface in question can be, planar, undulating or part of another surface. If the object is distinct from the surface the shadow is generally called a cast shadow. When an object casts a shadow onto one of its own surfaces this is called an attached shadow. The different geometries of cast and attached shadows may cause variations in the size and shape of penumbra regions but the

mathematics of such shadows is essentially the same. For matte surfaces the intensity of light reflected by a surface to the eye is the product of the surface reflectance and the intensity of illumination hitting the surface. Within their umbra regions, shadows produce a uniform reduction in illumination. In the extreme illumination within the umbra is reduced to zero but more commonly it is reduced to the ambient illumination level. Note if there are multiple light sources it is possible that a surface will be shadowed from one source but captured by the other. In this case the illumination in the shadowed region may not be uniform. However, this is a product of multiple light sources not shadows per-say. Illumination also varies in the penumbra region of a shadow (the region at the shadow's edge). Such regions are infinitely narrow for point light sources (which are infinitely small and not practically realizable) and larger when the light source covers a wide area. Such area lights can be modelled as multiple point sources in which case penumbra can be seen as a result of partial illumination by a sub-set of those sources. The widths of penumbrae also vary with the geometrical relationship between the light source, the object casting the shadow and the surface receiving it. However the change in illumination across a penumbra has relatively little to do with the orientation of the surface with respect to the light source – any relationship being somewhat indirect and mediated by other relationships in the scene.

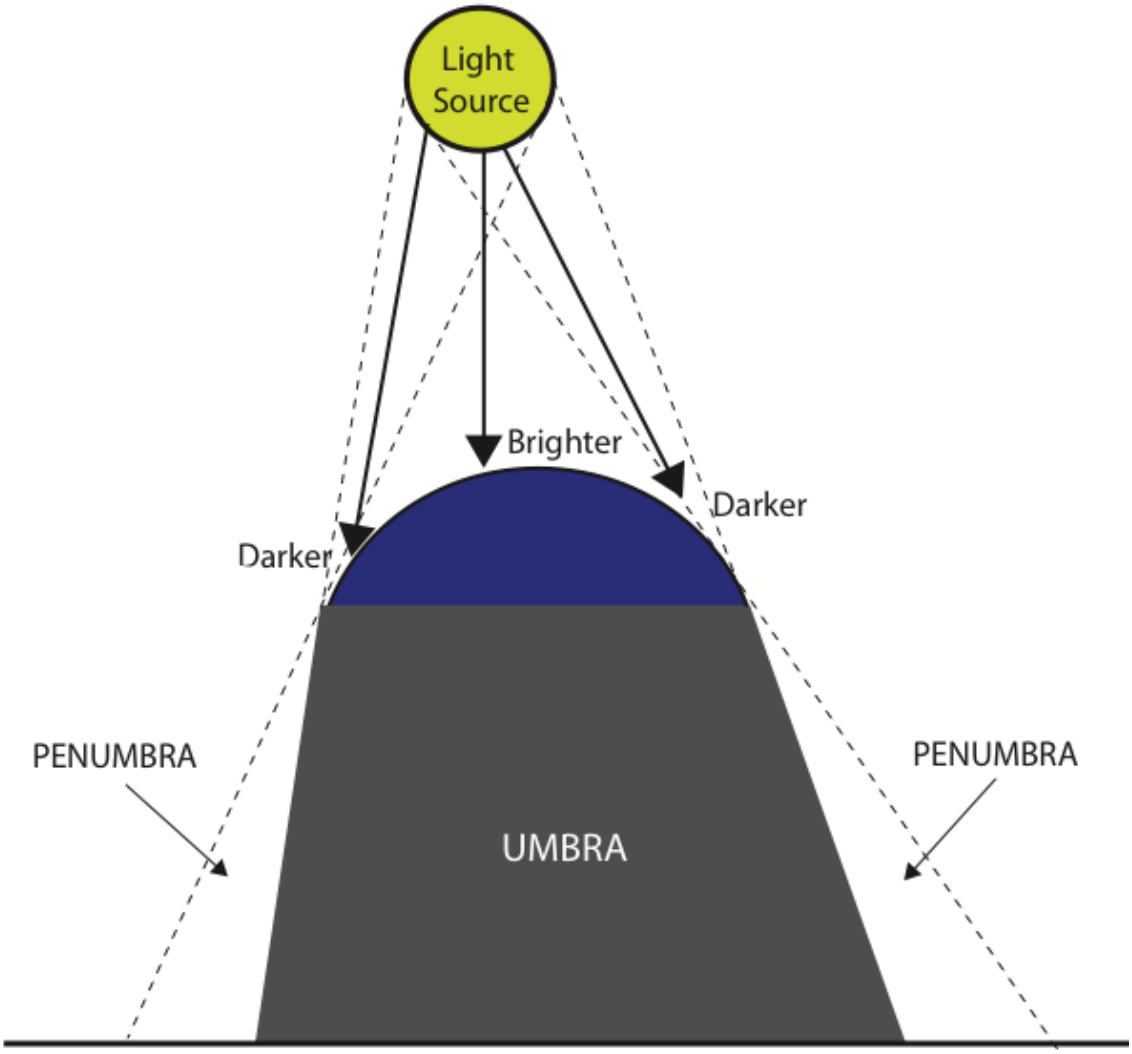

Figure 1.3. Surface brightness is dependent on surface orientation. Here surface's patch facing the light source would be brighter as it will receive more irradiance. Umbra and Penumbra areas are also highlighted

We adopt the following definition of shading here. Shading describes the variation in illumination due to the orientation of the surface relative to the light source when surface is not in shadow. The intensity of light received by a

surface is proportional to the cosine of the angle between the surface normal and the direction of the light source. That is $L=R.I.\cos(\phi)$ where R is the surface reflectance, I the strength of the illumination and ϕ the angle between the surface normal and lighting direction. This has similarities with shadows in that when the surface normal is orthogonal to the direction of the light source or turned away from it illumination will be zero. However for smaller angles there will be some illumination – the surface has less than maximal illumination even though it is not strictly in shadow. This will even in the case for a theoretical, infinitely small light source. A key aspect of shading is that the variation in illumination is mathematically linked to the orientation of the surface. Thus so long as surface reflectance does not also change, if one knows the direction of the light source and the degree of illumination / shading one can calculate the orientation of the surface. Critically such a calculation is not possible in areas of shadow umbra or penumbra. This relationship between surface orientation and shading has long held the promise that it might be possible to judge surface shape from shading. Indeed it is well demonstrated that humans can derive shape from shading to some extent. However, the conditions in which the relationship is reliable (eg uniform Lambertian surface reflectance, single light source illumination of a known direction) are extremely uncommon and even close approximations to the ideal very rare. Nonetheless humans do seem to derive shape from shading possibly using very simple algorithms aided by a number of prior assumptions.

1.4.2 Human processing of shape from shading

The human visual system uses many visual cues to interpret the shape of objects and shading plays a key role in perceiving surface orientation and curvature. In physical terms, shading refers to the luminance variations defined by the amount of reflected light of a surface in relation to its orientation with respect to the illuminant position. This definition highlights the most important features of shading: reflection, illumination, orientation and position of the light source. In fact, the perception of shape-from-shading (SFS) is highly influenced by small changes on the physical features of surface such as its reflectance and its orientation along with the type of illumination, the number of light sources and their position. Nevertheless, in real life, humans can perceive shape-from-shading reliably even though this process requires that a number of sub-tasks be solved in order to provide useful solutions.

Two common methods by which humans seem to derive shape from shading are the linear shading model and the dark is deep model. These seem to be the default processes by which shape from shading is calculated in the absence of other shape cues (Sun & Schofield, 2011).

Linear shading is approximated by setting surface slant proportional to luminance (Pentland 1989). If the possibility that luminance varies due to other reflective components such as specular reflection is ignored then this model will produce surface distortions in regions which actually contain specular

highlights. However, there is some evidence that human vision is indeed fooled by such variations (Nefs, Koenderink and Kapper, 2006). The model will also be fooled by changes in surface reflectance. Where these are obvious human vision is not so confused suggesting that it has a means to extract the shading component out of such a scene (Kindgom 2008). However, since simply 'painting' an illumination gradient onto a sheet of paper can produce an impression of shape from shading it is clear that human vision is confused by simple reflectance changes.

The second principle mode of shape-from-shading is the dark-is-deep model (Langer and Bulthoff, 2000). Here the height of a surface relative to its own average plane is proportional to luminance. This model for shape from shading is derived from the case when the light source is non-directional or diffuse. There is evidence to suggest that humans use both methods for extracting shape from shading and that they can switch between the two – or even use a combination of the two – depending on image cues to lighting type (Langer & Bulthoff, 2000; Schofield, Rock & Georgeson, 2011; Sun and Schofield, 2011).

There are a number of processes that need to be performed in parallel to execute SFS: the visual system has to decompose the luminance signal (that is, divide reflectance changes from illumination changes), discriminate shading (which reveals surface orientation and curvature) from shadows (which are optically similar to shading but which do not directly reveal surface orientation or

curvature), develop an estimation of the light field (the directions of key illuminants) and finally to make use of all the bounding contours that can help the segregation of surfaces or objects one from another.

The perception of SFS often also relies on other sources of information such as stored, prior assumptions that can be used when in presence of degraded or ambiguous visual stimuli. The effect of such assumptions is such that they sometimes lead to misperceptions (i.e. Visual illusions).

Shape from shading is one of the oldest topics in human vision. Rittenhouse (1786) and Brewster (1826), described a visual effect on depth perception that was caused by rotating stimuli through 180 degrees. Specifically, they discovered the unstable nature of convexity with respect to changes in the position of the light source relative to our prior for lighting from above (the crater illusion). These early studies are a clear example of the complexity of SFS, in which a given pattern of shading can be generated by an infinite combination of surface properties and light fields. Such ambiguities render both the task of interpreting shading (as conducted by the visual system) and the task of studying SFS (as conducted by vision scientists) impossible without the application of certain constraints and assumptions (Freeman, 1994; Johnson and Adelson, 2011). The assumptions employed by human vision and those employed by vision scientists are to some degree aligned. This is the reason why most of the studies in SFS have assumed surfaces to be Lambertian

(surface brightness does not change with the observer's view point) and light sources to be point-like and distant.

In recent times, the study of SFS has diverged into two branches: computer vision and human vision. Computer vision has tried to develop new methods to solve practical problems such as reflectance models or the representation of surface orientation (Horn, 1975; Pentland, 1989). On the other hand, human vision has studied the processes that can lead observers to perceive 3-dimensional images starting from 2-dimensional patterns of shading and how the visual system weighs the image features in order to have a stable and veridical percept of surfaces. The remainder of this introduction will focus on the perception of SFS in humans.

1.4.3 The Perception of Shape-From-Shading

Human observers are able to reliably perceive the 3-dimensional orientation of surfaces based on luminance gradients only. As shown in figure 1.2 luminance variations alone can produce a clear percept of a convex bump (left) or a concave object immersed in a scene composed of convex objects (right). As demonstrated by Sun and Perona (1998), SFS is a process that takes place at early stages of visual processing since reaction times are not affected by the number of objects in a scene.

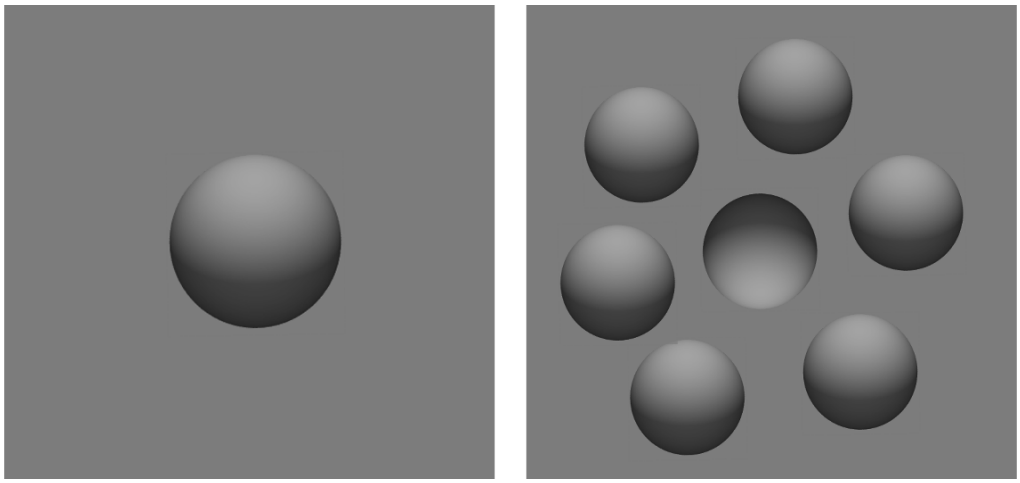

Figure 1.4. A scene showing SFS in humans

Interestingly, In figure 1.4 we can see at least two main aspects of SFS that have been the matter of study and can emphasize the role of prior knowledge: the assumption that there is only one light source and that light is coming from above. In fact, the figure on the left would be perceived as a convex sphere coming out of the surface since most observers would assume it as lit from above; on the right, the central figure would be perceived as a dent surrounded by bumps given the fact that observers will assume there is only one light source that is illuminating the whole scene. There are a series of constraints that need to be applied in order to solve SFS - which otherwise is a mathematically ill-posed problem - and one of the most important is the origin and the number of light sources. In fact, if we turn the page through 180 degrees, we realise that the shading patterns would lead to a “flip” in the depth

perception: objects perceived as convex are now perceived as concave and vice-versa.

1.4.4 The Role of Shading

One of the first experiments to investigate the crucial role of shading in perceiving 3-dimensional shape was run by Todd and Mingolla (1983). Assisted by improved technology for the manipulation of images, they studied how patterns of shading could provide enough information for the representation of 3-dimensional forms. In a series of three experiments they wanted to assess how the visual system weighs shading information. The first experiment was run to test the ability of participants to judge the curvature of cylindrical surfaces in which direction of the illumination and specular highlights varied by some degrees. Results revealed that observers did not completely rely on shading as their performance fell halfway between complete insensitivity to shading and perfect accuracy. In the other two experiments their results highlighted the pivotal role of illumination direction on shaded images and indicated in conditions where the surfaces' curvature was defined by textures in conjunction to shading or by texture only: the influence of shading was not very strong.

The conclusions of this study were important for the realization that models of computer vision such as Horn (1975) or Pentland (1989) were not completely

valid for human vision as they would not mimic human data. In particular they claimed that since computer vision models have imposed quite severe constraints (such as Lambertian surfaces) they do not provide a useful model to human behavior. Still, although many studies and in particular the experiments of this thesis have used Lambertian surfaces as images/stimuli, we are not sure that humans make full use of it so that their performance is strongly affected by the use of such surfaces.

A subsequent study (Mingolla and Todd, 1986), tested observers on the perceived curvature of ellipsoids. Their results showed that while observers were able to correctly report the 3-dimensional structure of the shading pattern, settings varied with the direction of the illumination. Similar results have been obtained by Mamassian and Kersten (1996). In their experiment, they used a croissant-like object – a computer rendered 3-dimensional stimuli - that included all local solid shapes such as elliptic, parabolic and hyperbolic, to find that the elliptic patch was less biased in terms of slant perception compared to the other two. They interpreted it as an indication that the visual system may assume a surface to be elliptic *a priori*.

The variation found by Mingolla and Todd's (1986) study demonstrates that the visual system tends to perceive surfaces reliably but it does not necessarily match the original shape of surfaces depicted in the visual stimulus. For instance, Koenderink, Van Doorn and Kappers (2001) showed that the information enclosed in an image can only disambiguate shape estimation up to

a certain level, after that, the observer has to apply some internal estimate of the missing information (their “beholder’s share”) in order to interpret the shape of surfaces.

In a previous study Koenderink, van Doorn and Kappers (1992) tested the observers’ variability for tilt and slant settings on shape judgements and found that participants were internally consistent for the given task but when comparing across observers, this consistency was less evident. Specifically, they obtained quite homogeneous data for tilt settings across observers but their variability was fairly high for slant settings. In both studies they found a significant fluctuation within observers for the same shape judgements and this outcome might suggest that shape-from-shading is highly unreliable. However, as Koenderink et al. (2001) pointed out, the shapes estimated by observers were, most of the time, affine transformation of one another (both across task and observers) and of the same basic shape. Therefore, we can associate these transformations with both “beholder’s share” and the use of prior assumptions.

These results suggest that prior assumptions are adopted by the visual system in a way that is determined by the stimulus and the measurement task. Koenderink et al. (2001) concluded that, setting aside affine transformations, the estimation of shape-from-shading is very reliable.

In conclusion, human observers can reliably interpret shape-from-shading up to an affine transformation of the perceived shape. Given the infinite numbers of interpretations of a shading pattern they need to make use of, some prior assumptions that may vary across tasks and result in big variations across participants and between tasks within participants. This notion should be taken into account and carefully analysed when choosing the task to be tested in any psychophysical study.

1.4.5 Boundaries and Contours

Even if shading represents a useful and reliable cue for human observers when perceiving surfaces, the estimate of shape-from-shading can be highly influenced by other sources of information such as contours. For instance, object recognition studies have shown that the perceptual organization at the basis of object representation is mainly due to the way boundaries are interpreted rather than by the physical luminance contrast of the boundaries themselves (Anstis, 1990;Tadin, Lappin, Blake, Grossman 2002; Albrecth,List, Robertson, 2008).

In SFS Erens, Kappers and Koenderick (1993a) demonstrated that observers found it difficult to discriminate between convex, concave, quadratic and hyperbolic structure when stimuli were constituted of patches of surfaces deprived of any boundary or lighting direction information. This inability to

recognize convex and concave surfaces could be circumvented by adding some lighting information in the form of cast shadows, while the problem persisted for quadratic and hyperbolic surfaces.

Knill (1992) provided a good example of the ambiguity of shape perception by creating stimuli in which stripes were painted on top of shading patterns. Specifically, the same shading pattern could be seen as two bumps illuminated from behind or four bumps illuminated from above only by changing the spatial frequency of the curved stripes.

The power of contours in shape perception has also been demonstrated by Mamassian and Kersten (1996). In their study they found that the shape of objects composed of shading cues only were perceived as only as well as silhouettes of the same object, suggesting that occluding contours can override shading cues in observers' judgements. An alternative interpretation that the authors did not mention is that shading did not provide enough additional information.

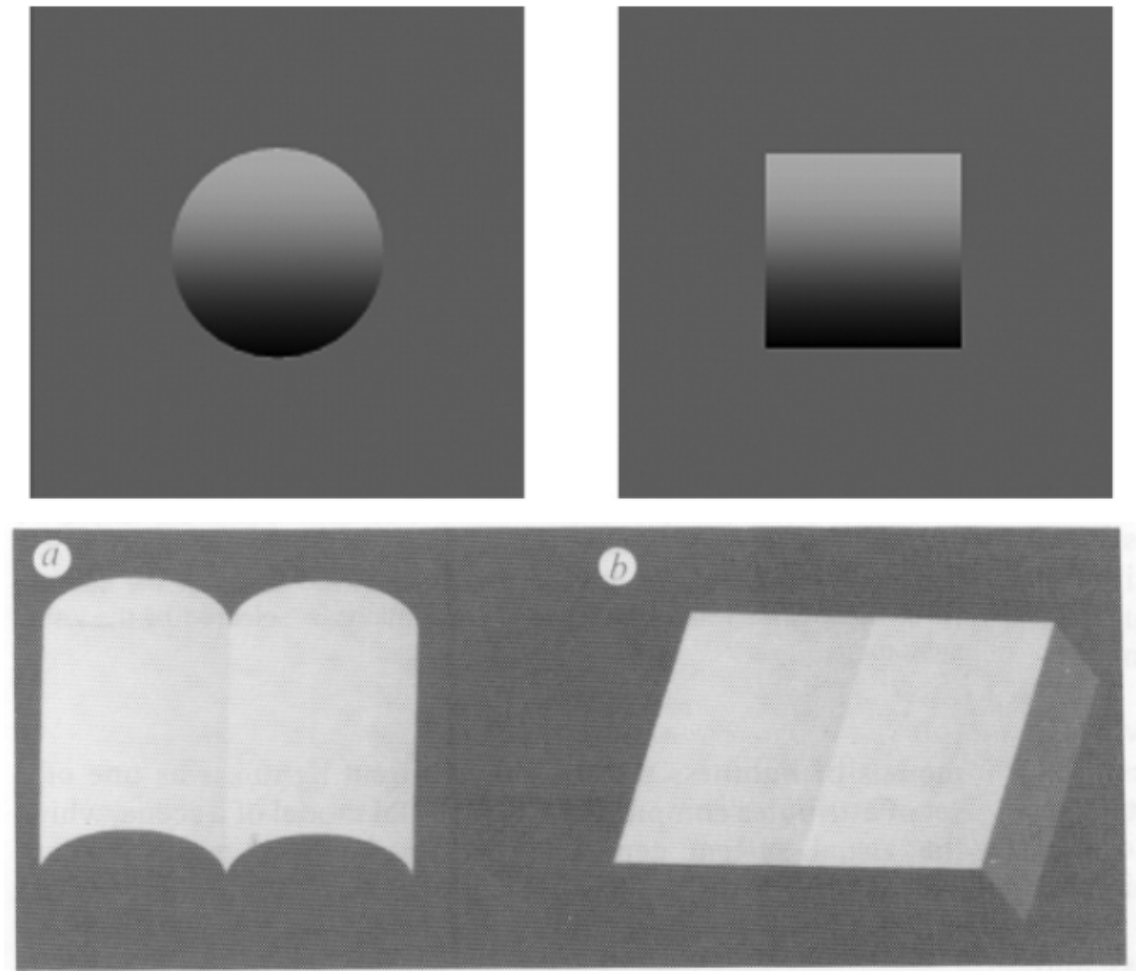

Figure 1.5. (Top) The same luminance profile could be perceived as either a disk or a cylinder depending on the edges outline (after Sun and Schofield, 2012). (Bottom) Apparent surface curvature does affect the perception of lightness (after Knill and Kersten, 1991)

A more recent study on boundary constraints (Sun and Schofield, 2012) found that the shape estimates often conformed to linear shading when simple periodic stimuli are presented to observers. Still, the perceived shapes were influenced by luminance edges in the image even in cases where truncating the image let only one cycle of luminance variation visible to the observers. They conclude that even when contours do not define occluding boundaries and are

not interpreted as surface marking, they can still influence shape perception (figure 1.5).

1.4.6 Material Properties

The luminance profile of a surface is generated by a combination of illumination conditions and the reflectance of the surface itself. In visual perception, we can directly observe luminance but we need certain perceptual processes to extract information about reflectance and illumination. Humans are able to estimate the reflectance of objects regardless of big variations in illumination: lightness constancy. A distinction must be made between photometric quantities such as illumination (the strength of the light source), reflectance (the amount of light reflected by a surface) and luminance (the amount of light reaching the observer) and their perceptual equivalents such as brightness (perceived intensity of a light source or surface) and lightness (the perceived reflectance of a surface). Illumination, reflectance and luminance are physical properties that are ot affected by context whereas brightness and lightness depend very much of the context surrounding a give object. A clear example is to compare white and black paper under different lighting conditions to notice the difference in reflectance between the two types of paper - we never perceive white as black and vice-versa so long as there is a range of surface reflectance present similarly, a black tile illuminated by a strong light source will appear white against a non-illuminated background (Radonjic, Todorovic and Gilchrist, 2010).

Gilchrist (1977, 1988) has shown that humans do not judge the surface lightness by the simple estimation of luminance but instead that lightness perception is often contextualized by global visual information and spatial arrangement. This contextualized ability of the visual system to use illumination information to approximate the reflectance properties of a specific surface, suggests that the representation of brightness and lightness may be segregated in the brain. Some studies agree that the image could be decomposed at a certain stage by the visual system and represented at different levels in accordance to their origin, namely illumination, reflectance and optical medium (a process known as layer segmentation Kingdom, 2008, Gilchrist, 2006, Andreson & Winaver, 2005). Kingdom (2008) has demonstrated how humans can conduct layer segregation with the help of many cues including, color, and texture. An alternative point of view comes from Zaidi (1998) that assume that the visual system could simply match objects across illuminants and derive illuminant shifts based on a pair of chromaticity sets. This interpretation means that observers are able to discriminate when the same objects are being viewed under different illuminants despite the fact that a discernible shift in the colour of the objects occurred.

Most studies of SFS have applied some reflectance constraints such as Lambertian (matte) surfaces (Mingolla and Todd, 1986; Curran and Johnston, 1996; Mamassian and Kersten, 1996; Khang Koenderink and Kappers, 2007). This particular reflectance type has been widely used because here the

incoming light is scattered equally in all directions so that the viewing point has no bearing on luminance. In real life though, the reflectance of surfaces has a specular (glossy) component too such that luminance does depend on the viewing angle.

Nefs, Koenderink and Kappers (2006) have tested observers using matte objects with added specular highlights. They found no difference in perceived shape between matte and glossy objects although they found some influence of illumination direction on the perceived surface shape. They claimed that as previously demonstrated by Erens, Kappers and Koenderink (1993b), shading cues can lose most of their strength if they are the only available cues. In certain circumstances, in fact, the visual system can by-pass and ignore the specular highlights.

To sum up, surface material can affect the perception of shapes but its effect is not completely clear. Some studies have reported that changes in the material can influence shape perception (Todd and Mingolla, 1983; Curran and Johnston, 1996) whereas others found no obvious difference between matte and glossy surfaces (Nefs et al , 2006; Nefs, 2008). One main difference between those studies is the presence of contours and edges in the stimuli tested. As reported by some studies (Mamassian and Kersten, 1996; Sun and Schofield, 2012) contours can override shading information in the perception of

SFS and this effect might have affected some experimental results (Nefs et al , 2006; Nefs, 2008).

1.4.7 Type of Stimuli

So far we have discussed how some fundamental features of the image can influence SFS. We have illustrated how material properties, boundaries, and shading need to be experimentally controlled in order to produce stimuli that will provide consistent results. For example, some studies have illustrated how shading may be ineffective when stimuli provide stronger visual cues such as boundaries or edges (Knill, 1992; Mamassian and Kersten, 1996).

In particular, studies that used very simple stimuli found that observers were not exhibiting shape constancy (Todd and Mingolla, 1983; Curran and Johnston, 1996; Khang et al., 2007) while opposite results were found for more complex stimuli (Nefs et al., 2006; Nefs, 2008). Koenderink et al. (1996) have used complex stimuli such as images of shaded sculptures representing human bodies to test observers' shape judgments. These type of elaborate stimuli not only provide shading information, but are rich in visual cues that could activate some higher level object recognition perception.

Prior to Koendrink's (1996) study, a typical experimental procedure would be to test very simple stimuli based on single cues to test shape perception and, along those lines, some studies have reported that shading cues can be easily overridden by other depth cues such as stereoscopic disparity (Bulthoff and

Mallot, 1988) or surface contours and outlines (Ramachandran, 1988; Knill, 1992). Koenderink et al. (1996) proposed that SFS might be a cue that needs some support in order to completely function, and, thence, an object provided with appropriate visual information, would certainly let observers make full use of shading component.

The problem of SFS is not only limited to the type of stimuli used, but also the inter-observers variances reported in many studies should be taken into account (Koenderink et al. 1992; Koenderink, et al., 2001; Mingolla and Todd, 1986). As a consequence, Koenderink et al. (1992) have tried to use easier tasks and to increase the number of tested locations but their results showed that such changes did not have an effect on inconsistencies. Intriguingly they found that those inter-observers' variations were not unsystematic but they rather seem to follow some structured design such as a scaling effect. For example, it has been shown that when resolving the ambiguous 2D shading of some stimuli, observers might use their own "beholder's share" (Koenderink et al. 2001). With the notion of Beholder's share, Koenderink aims to explain and account all those individual differences that may occur in observers' results in shape perception tasks.

Unfortunately so far, it has been impossible to keep apart SFS processing from the "beholder's share" and therefore, experimenters should take extra care when choosing the type of stimuli or the design to be used in order to minimize

its effect. As we will see in Chapter 2, the use of Classification Images technique may be a valid and useful way to minimize the effect of image features reliability.

1.4.8 Role of Illumination

A pattern of shading depends on physical features of the surface but also on the type of illumination and its position relative to it. Contours and edges can suggest the profile of surfaces or the conformation of objects, nevertheless the knowledge of the illuminant position can disambiguate the process of SFS.

To infer the role of illumination, Johnston and Passmore (1994) tested observers in a curvature discrimination task under different lighting directions. They found that the perceived curvature of a sphere was not influenced by tilt variations of the light source position (namely the changes in the declination of the light source along the vertical axis of the image plane), but that curvature thresholds increased when the position of the light source was closer to the viewing direction. It should be pointed out that under frontal lighting a Lambertian surface produces a luminance map with a much lower contrast than the one illuminated obliquely. There is a chance that the results found in the

experiment might be due to a decrease in contrast and consequently to a rather poor interpretation of the luminance profile. As we will see in Chapter 3, the perceived slant of a surface could be affected by the image contrast range. In this case the display gamut and the display range can have an effect on the perceived curvature of the surface.

Curran and Johnston (1996) conducted a series of experiments on the effect of illuminant position on perceived curvature. They tested whether the apparent curvature of an unambiguously convex surface was influenced by the position of the illuminant, revealing that perceived curvature decreased with frontal lighting producing an underestimation of the actual surface curvature. This misperception did not occur when the light source was above the viewing axis since observers were most accurate in this condition. This effect was not limited to matte surfaces as it was replicated on glossy surfaces (with weaker effect intensity).

Using more complex scenes Koenderink, vanDoorn, Christou and Lappin (1996a; 1996b) found that observers produced a veridical pictorial relief of sculptures when they were viewing degraded images (without shading) in presence of contours; on the other hand, their perception of images containing shading cues only were unstable as the perception of shape was varying along with lighting directions. Specifically, photographs of the same picture under

different lighting directions produced consistently dissimilar reliefs for each stimulus.

Christou and Koenderink (1997) demonstrated that observers' shape estimates for elliptical solids were influenced by the position of the light source so that the brightest regions were perceived as closer to the observer's position and darker regions were perceived as further away. This effect of illuminant direction underlines instability in shape constancy. In fact, knowledge of the light source position should secure a more stable perception of surface shape.

These findings suggest that shape constancy is largely influenced by the direction and type of illumination and therefore cannot be held in the processing of SFS under different illuminations.

Perceived shape is systematically influenced by changes on lighting direction. It could be that the visual system associates changes in the illumination to changes in the surface profile, therefore shape constancy should not be expected under variations of the lighting directions. This could be due to the fact that changes in lighting slant produce changes in image contrast which can be misattributed to changes in the surface profile.

1.4.9 Knowledge of the Light source

The position of the light source plays a key role in SFS processing since no shading pattern can be veridical without knowledge of the illuminant direction. As a matter of fact, most SFS algorithms in computer vision need the illuminant

direction to be known, while for human vision, the knowledge of the light source does not preclude SFS, although it may influence the outcome.

Todd and Mingolla (1983) found some evidence that shape estimation and judgments of lighting direction could be processed independently. Specifically, they asked observers to report the curvature of matte and glossy surfaces under different lighting directions and they found that while curvature estimates varied across the two surfaces types, estimation of the light direction did not. Specifically, specular highlights did not affect the perceived direction of the light source. In a subsequent study (Mingolla and Todd, 1986), they confirmed that observers' accuracy on shape judgments was poorly correlated with that of judgments on lighting direction since they found that the illumination direction influenced the judged shape of all ellipsoid surfaces, but the presence or absence of specular highlights did not have a significant effect. Also Mamassian and Kersten (1996) found that even in cases where observers could easily estimate the origin of the light source, they still reported the tilt of the light source with large variation. They concluded that in SFS tasks observers were not making use of the illumination direction.

Nevertheless, other studies have demonstrated that observers can deduce the direction of the light source from other visual features such as cast shadows and specular reflections (Mingolla and Todd, 1986; Liu and Todd, 2004). In a recent study, Gerardin, Kourtzi and Mamassian (2010) have shown that lighting

direction processing and shape processing may take place in different brain areas. For instance, early retinotopic areas showed an higher activation for light processing while higher cortical areas appeared to be involved in shape processing. Furthermore, they compared their results with previous light and shape models to find no correlation. These results made them suggesting that no area was involved in both light and shape processing together and made more clear that there is degree of independence between the two processes. Their conclusions strongly support the idea that SFS depends on light source estimation – as light source estimation occurs in earlier brain areas.

It is essential now to define the different ways in which the position of the light source can affect SFS and to focus in those studies that have described how the position of the light source direction could be detected by human observers. A direct way is to extract information about the light source straight from the image: the observer decides on surface shape and then decodes the illuminant position accordingly (Mamassian and Kersten, 1996); otherwise, the illuminant position could be directly derived from the image's lighting cues and could be computed by testing the light source estimation via gauge objects that could be setup from observers in terms of slant, tilt and light intensity in order to report the perceived light field illumination (Koenderink, Pont, vanDoorn, Kappers & Todd, 2007; Koenderink, vanDoorn & Pont, 2004). More recent imaging studies (Gerardin, Kourtzi and Mamassian, 2010) have found that the estimation of the

light source activates early visual areas and it is processed earlier than the representation of the 3D shape in the visual system.

An alternative, more indirect way, is to estimate the light source position by assuming the location where it should most likely be: the lighting prior. There are two known lighting priors: the first one is a directional prior, which assumes that the light is coming from above (some studies above-left) the observers point of view (Adams, Graf and Ernst, 2004; Brewster, 1826; Gerardin, Kourtzi and Mamassian, 2010; Mamassian and Goutcher, 2001; Ramachandran, 1988; Rittenhouse, 1786; Sun and Perona, 1998; Todd, 2004); the second prior is for a diffuse light source (Langer & Bulthoff, 2000; Tyler, 1997; Schofield, Rock and Georgeson, 2011).

The light-from-above prior was reported for the first time by Rittenhouse (1786). In his research the author describes a powerful effect in which a surface's reliefs as perceived from shading are inverted (physical convexities are here reported as concavities) by simply rotating their orientation by 180 degrees. Some studies (Ramachandran, 1988; Sun and Perona, 1998) have confirmed the effect and have located the direction of the assumed light source to the average location of the sun (above observers' head). More recent studies have made use of more subtle stimuli (Mamassian and Gouthcer, 2001; Gerardin, de Montalembert, Mamassian, 2007) in order to reveal a more precise location of this bias: above-left. However, there are large individual differences across

observers (Adams, 2007) and the light prior can be altered throughout training (Adams, Graf and Ernst, 2004).

A clear example of the light-from-above bias is the crater illusion (Figure 1.6), here a depression in the ground - crater - looks like a mound by flipping the image.

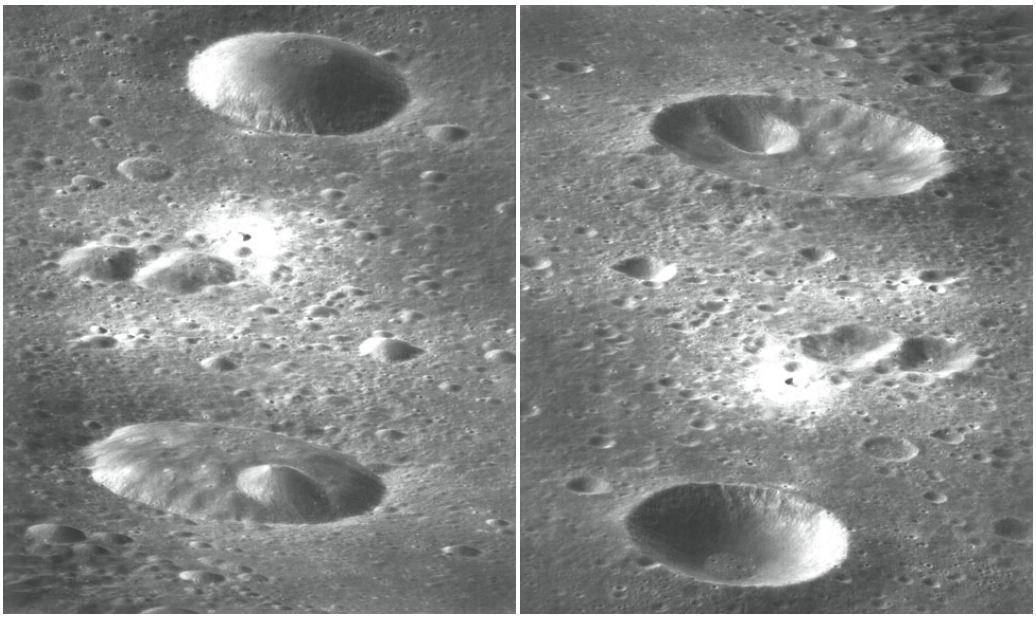

Figure 1.6 The crater illusion: perceived mounds (on the left) could be perceived as craters (on the right) by a 180 degrees flipping of the image.

However, the light-from-above assumption is not the only cause of the crater illusion. In fact, there are other assumptions that are equally plausible to justify this effect such as the convexity prior (Hill and Bruce, 1994; Langer & Bulthoff, 2001). Liu and Todd (2004) have tested naturalistic rendered images of concave and convex surfaces with ambiguous shading profiles and they found that the convexity prior was stronger than the light-from-above prior. Furthermore, their results indicated that shape perception was strongly

influenced by perceptual biases even in presence of cast shadows or specular highlights giving the impression that stimuli features were not highly influential on SFS. It seems necessary to record that the depicted scenes were strongly reduced in content - they contained one object only - and therefore could have led to a stronger reliability on prior assumptions.

Koenderink and vanDoorn (2004) have reported that the convexity prior can be interrupted by global shading of the image. This study followed the one by Erens, Kapper & Koenderink, (1993a) who demonstrated that shading in areas surrounding an ambiguous stimulus could help to indicate the direction of the illumination and thus to disambiguate the perception of shape. Recently, Morgenstern, Murray and Harris (2011) examined how the visual system combines priors with local and global visual lighting cues by using some more complex stimuli. They found evidence for a maximum likelihood combination of assumed and manifest illumination but they also observed that the influence of the lighting priors varied along with lighting cues strength: specifically, prior assumptions could be easily overridden by relatively weak local lighting cues.

These results suggest that the visual system does not strongly rely on lighting priors, but rather that shape perception is dominated by the observers' estimates of the actual lighting direction.

1.4.10 The Type of Light Source

Most SFS studies related to light-from-above priors assume that the default lighting assumption is directional and that there is only one light source

(Ramachandran, 1998; Sun and Perona, 1998; Mamassian and Goutcher, 2001; Adams, Graf and Ernst, 2004).

Tyler (1997) tested observers using radial sine wave stimuli under different lighting and reflectance conditions and concluded that the object interpretation for human vision assumes a default diffuse illumination. In an attempt to assess the visual system's ability to estimate lighting diffuseness and direction, Langer & Bulthoff (2000) used highly articulated isotropic surfaces to demonstrate that observers were able to recognise differences between diffuse and directional lighting and then process SFS accordingly.

Schofield, Rock & Georgeson (2011) used linear sine wave surfaces in an SFS paradigm to reveal that changes in the perceived location of surface peaks relative to luminance peaks were not in accordance with neither diffuse nor directional lighting assumptions, but rather with a combination of the two. People had a distinct from above component in the default light source but it also had a strong diffuse component like the sun in the sky.

To sum up, it seems like humans can estimate the direction and diffuseness of lighting quite well when in presence of isotropic surfaces, but what happens when surfaces are anisotropic? Koenderink, vanDoorn & Pont (2007) tested observers on isotropic rough surfaces in which directional lighting produced anisotropies in the shading pattern. Their results showed that anisotropies in the

textures dominated observers' perception by leading them to estimate the light source direction incorrectly, but consistently to be orthogonal to the predominant orientation of the shading pattern.

In conclusion, shape-from-shading is largely influenced by the isotropy of the depicted surfaces: in globally isotropic scenes humans are good at estimating the direction and the diffuseness of the light source, while in presence of anisotropy, the dominant orientation of the surface undulations strongly guides the perception of the lighting direction.

1.4.11 Role of Shadows

Shadows are regions of surfaces that are not reached by direct light. They can be generated by the obstruction of opaque objects (cast shadows) or by the surface itself which occludes another part of it from the light source (attached shadows).

Attached shadows are very useful to the human visual system in order to extract information about the shape of objects particularly in regards to local shape perception; cast shadows carry on information about the shape of objects but also information about the global scene, such as the global shape of surfaces or the lighting setup. In SFS cast shadows are very informative visual cues to disclose the spatial layout of a depicted scene.

In order to extract the information necessary to veridically perceive the spatial layout of a scene, shadows need to be segmented from the background and recognised as cast shadows. To do so, it is sufficient to detect an area that is darker than its surround (Cavanagh and Leclerc, 1989).

In their study, Cavanagh and Leclerc (1989), tested observers in a series of experiments to reveal which criteria were used to identify shadow regions. Their results suggest that the visual system uses only luminance information to process shape from shadows and does not take into account color, texture, motion or depth. They concluded that shadow interpretation can be useful in retrieving the 3-dimensional object shape especially to disambiguate concavity/convexity and in some cases to provide information about surface reliefs and object shapes.

Mamassian, Knill and Kersten (1998) have argued that shadows are not very informative as static cues to reveal object shape but they are perceptually more relevant when in motion (Kersten, Knill, Mamassian and Bulthoff, 1996). This is due to a “stationary-light-source-constraint” that encourages observers to perceive objects as moving in situations where other interpretations are still valid (i.e. motion of the light source or the background surface or the viewpoint). Nevertheless, cast shadows are very useful as cues to infer the position of the light source (Mingolla and Todd, 1986; Liu and Todd, 2004).

1.5 Bayesian and Priors

Our visual system is capable to provide a stable and unambiguous percept of the world even though the amount of information that is processed by the eye is objectively very ambiguous. Visual perception is specialised in examining the complex and ambiguous nature of visual processing and has adopted some methods to reduce and interpret visual stimuli in order to evaluate in a more controlled and limited contest the behaviour of the visual system. The idea of perception as unconscious inference has its origin in Helmoltz (1867). He realised that objective ambiguities arises, if several different objects produce the same image description of image features, and the visual system resolves this ambiguity through built in knowledge: this is the way the brain can infer properties of objects (Kersten, Mamassian, Yuille, 2004). Bayesian inference provides a framework for modelling artificial and biological vision. Bayes' formula implies that the perception is a trade-off between image features and the prior probability. The role of prior probability that is often called "prior knowledge" or more simply "priors" is very important in psychophysical experimentation and even though it is pivotal to visual perception, observers are in most cases not aware of it. Nevertheless, all the implicit assumptions made by the brain to solve ambiguities can be revealed through specific experimentation. Specifically, in some cases, the perception of some ambiguous stimuli may be highly prior driven, while otherwise could be more data driven. This could explain why in case of very ambiguous image features

the perception is more influenced by the prior (Mamassian and Goutcher, 2002). In the study of Shape-from-shading for example, priors can refer to assumptions about light directions (Adams, Graf and Ernst, 2004; Brewster, 1826; Gerardin, Kourtzi and Mamassian, 2011; Mamassian and Goutcher, 2001; Ramacahandran, 1988; Ramachandran and Kleffner, 1992; Rittenhouse, 1786; Sun and Perona, 1998), but also to the perception of solid shape. It has been demonstrated by study about object geometry that the perception of stimuli is consistently biased toward convexity rather than concavity (Liu and Todd, 2004).

To highlight the complexity of three-dimensional perception, it would be useful at this point to take a look at the Necker cube illusion in figure 1.7 (below).

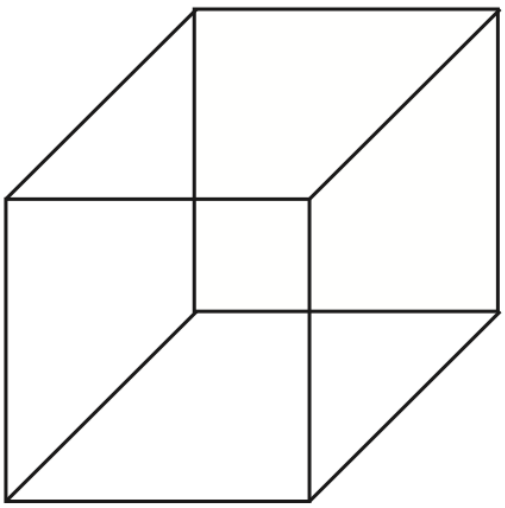

Figure 1.7. The Necker cube: a wire frame cube with no depth cues. This multi-stable figure can provide mutually exclusive interpretations.

The cube showed in the figure above is generally perceived as a cube with the bottom left square as the frontal face, or with the very same square as the base of a cube seen from above. Some questions arise when observers have to deal with such a figure: in first instance, why the visual system only takes into account two interpretations of the figure? And secondly, why it is not possible to combine the two interpretations simultaneously but only exclusively? We can reply to the first question by saying that the brain excludes a priori (priors) some interpretations over others, and, to the second question, by saying that the most likely interpretations of ambiguous stimuli are processed in the basis of perceptual decisions. Given the infinite number of interpretations that an ambiguous figure can produce, the Bayesian approach is widely regarded as an appropriate model to interpret the mechanisms that underlies visual perception and the use the brain makes of sensory information.

The Bayesian modelling is mainly based on two elements: likelihood function and priors distribution.

The likelihood function includes all of the sensory information we have. In the case of the 3-d perception, it is based on the possible set of objects that could be consistent with a given image data. By definition, the possible set of objects is an infinite set, therefore the likelihood function of an ideal observer is constrained to the most probable shape of the set of objects that could define the image data.

While the Likelihood function is based on the image features, the priors distribution represents the prior belief of the observer as the frequency of occurrence of various object shapes.

We can then compute the posterior distribution based on likelihood function and priors distribution:

Posterior$_{\varphi}$ (θ) = C x likelihood$_{\varphi}$ (θ) x prior$_{\varphi}$ (θ)

Where C is a constant introduced to make the posterior a probability distribution function. Therefore, the posterior probability of an event θ is the likelihood weighted by the prior probability of the happening of θ.

The formula is directly derived from Bayes' Theorem based on conditional probabilities:

$p(\theta|\varphi) = p(\varphi|\theta) \times p(\theta)/p(\varphi)$

In fact, C is equal to the inverse of $p(\varphi)$. Since it is needed to make the posterior distribution a real distribution. The result of this equation – the posterior distribution – is the estimate of the probability distribution of an unknown event θ after both prior distribution and the sensory data (likelihood) are taken into account. As displayed below in figure 1.8, the posterior distribution is located between the peaks of the likelihood function and the prior distribution. As previously said: Bayes' formula implies that the perception is a trade-off between image features and the prior probability.

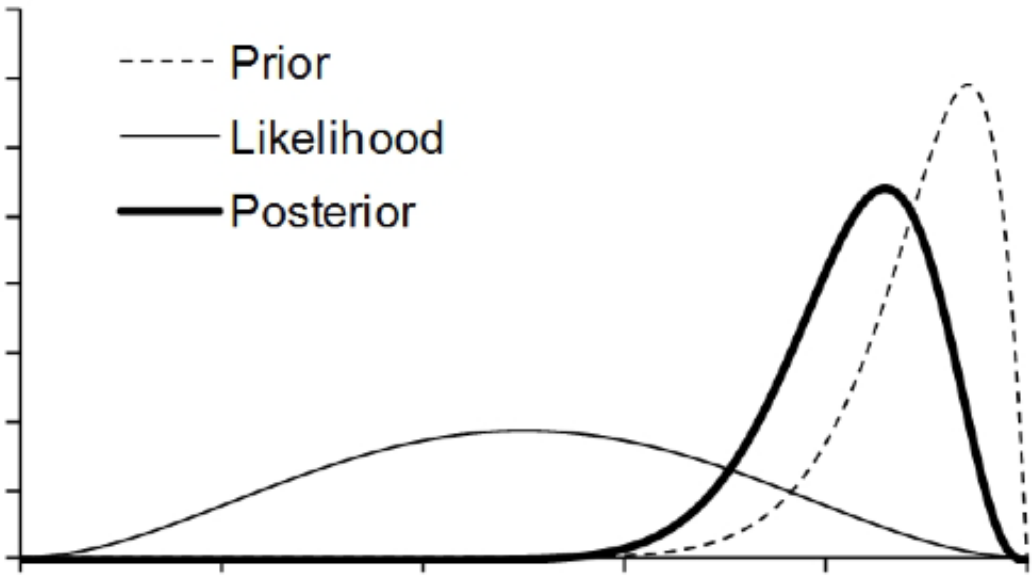

Figure 1.8. The Posterior distribution is a balance between the prior and the likelihood function. The peak of the posterior distribution lies between the peaks of prior and likelihood function.

The Bayesian approach could be applied to explain the behavioural data found in psychophysical testing. We know most Shape-from-shading studies agree that the human visual system has a bias for lighting from above. This assumption is responsible to the crater illusion, whereby a concave surface is seen as if convex due to being illuminated from an unexpected direction, and it indicates that the visual system make use of additional knowledge that goes beyond the stimulus information.

A study by Mamassian, Landy and Maloney (2004) reviewed the main principles that stay at the basis of Bayesian modelling in 3D perception. In their study,

they weighted the role of likelihood function, prior assumptions and decision making rules to conclude that Bayesian-inspired models provide an explicit example of the interaction between sensory information and prior assumptions.

Although their results are relevant and interesting they do not clarify if Bayesian approach take into account how priors constraints and likelihood function are influenced by the context, or if the priors are stable over time and conditions. As demonstrated by Adams, Ernst and Graf (2004) priors can change and their results suggest that even though the Bayesian approach can be usefully used as a reference model to explain the human visual perception, still this powerful tool can implies some strong limitations based on context.

1.6 Rendering Scenes in computer graphics and experimentation.

Before outlining our general method it is useful to deviate to describe some of the methods used to render shaded scenes in the literature and the possible deficiencies inherent is such methods.

1.6.1 Physics of illumination.

The pattern of light that reaches the eye from a given natural scene depends on a large number of complex interactions. In theory a full physical mode of a

scene could produce a rendered image indistinguishable from the original and while certain ‘surfaces’ such as fluids and fibrous surfaces like hair and velvet are notoriously difficult to model in principle the physics are sufficiently well understood that they could be modeled to a close approximation. However full physical rendering takes time and so a number of short cuts are employed in common rendering packages. Such short cuts can be divided into two broad categories. Physical approximations involve leaving out some aspect of the physical rendering process and, possibly, replacing it by an approximation. Alternatively perceptual rendering uses existing knowledge of the visual system to avoid the use of a full physical model while achieving a perceptual metamer (or close approximation to a metamer) of the original scene / surface. Since perceptual rendering is a relatively new concept which, by definition, requires reasonably good knowledge of human visual processing we will not discuss it further. It is enough to say that in studies of human visual perception, where the aim is to determine how the visual system responds to some stimulus, there are few cases where perceptual rendering would be appropriate unless the aim was to test the efficacy of the rendering process itself.

1.6.2 Ray Tracing

Rendering engines of use to vision science (including the popular RADIANCE program and the PovRay package used in part of this thesis – Persistence of Vision Pty. Ltd., Williamstown, Victoria, Australia. http://www.povray.org/) employ reverse ray tracing methods as part of the rendering process. Here light

paths are traced from the observation point to services in the scene and then back to the light source(s). On interaction with objects it is necessary to estimate the reflectance of the surface so as to determine their likely contributions to any pixel value in the image. While methods for estimating natural reflectance functions exist most renderings employ approximate reflectance models for surfaces – we describe such models in the next few paragraphs. One must also consider inter-reflections. If the absence of a direct path back from a surface to a light source does not mean that it receives no light from that source. It is probable that light from any given source will be reflected via several surfaces (possibly more than once each) before arriving at the surface being rendered. The extent to which such inter-reflections are modeled varies widely. Clearly consideration of a infinite number of inter-reflections for each surface point rendered is impractical but equally allowing no inter-reflections at all fails to model natural lighting well. We discuss possible solutions to this dilemma later.

1.6.3 Bi-directional reflectance distribution functions (BRDF)s

With regard to physical rendering the amount of light reflected into the eye from a surface is a complex function of the physical properties of the surface, the amount of light striking it, the orientation of the surface, the composition of the light source and the position of the observer. One way to describe the reflectance properties of a surface is the bidirectional reflectance distribution function (BRDF, Nicodemus, 1965). This function (which can be measured for

natural surfaces) describes the ratio of the amount of light reflected in a certain direction to the intensity of light arriving at the surface from a certain other direction. In a rendered scene there is only one view point and that defines the direction of the outgoing ray at every point on the surface. For a single directional light source the direction of the incoming ray is also easily calculated and each surface point and entering these values into the BRDF will produce the require luminance value at each point on the surface. Multiple light sources, and less-directional light sources can be modeled by repeated application of the BRDF for each light source or partial light source. While it is possible to use measured BRDFs in rendering engines most use some approximation to typical surface BRDFs.

1.6.4 Blinn-Phong model.

The most common reflectance model is the Blinn-Phong model (Phong, 1975; Blinn, 1977). The Blinn-Phong model captures both diffuse (Lambertian) reflectance and specular reflectance. Diffuse surfaces reflect light equally in all directions making the direction of the viewer irrelevant. The diffuse reflectance component in the Blinn-Phong model depends on the orientation of the surface relative to the direction of the light source. Specular reflections are also known as mirror reflections. Mirror surfaces are the opposite of diffuse surfaces they reflect light only in one direction – the reflection of the light source direction. So if light is incident on a mirrored surface at an angle of 45 degrees to the surface

normal it will be reflected at -45 degrees. Thus the specular component depends only on the orientation of the viewer relative to the orientation of the reflected light ray. Perfectly mirrored surfaces reflect light only in one direction and that light will only be seen if the viewing direction is the same as the reflected direction. Scatter at slightly rough specular surfaces can be modeled with a dispersion term which effectively spreads the range of directions from which a specular reflection will be scene. In practice with little dispersion a point light source will appear as a discrete highlight on the surface only when the viewing angle is correct whereas with more dispersion the size of the highlight will grow. Between them the diffuse and specular components of the Blinn-Phong model together with suitable values for the albedo (paint colour) of a surface can produce reasonable approximations to many naturally occurring BRDFs, particularly matte and polished surfaces. However, the Blinn-Phong model does not deal with refractive surfaces, diffraction or filament services such as hair and velvet. Pov-Ray offers a partial solution to this deficiency in the form of selective forward ray-tracing – also known as Photon Mapping to pre-calculate a version of the light source prior to reverse ray trace rendering. This method is claimed to deal with refractive surfaces.

1.6.5 Properties of the illuminants.

As well as modeling the reflectance properties of objects in a scene is it necessary to model the light sources illuminating a scene also. Perhaps the

most comprehensive measure of lighting for a scene is the light field (Adelson & Bergen, 1991; Debevec, 1998; Gershun, 1939; Moon & Spencer, 1981) which is a function giving the illumination reaching each point in a given region of space from all directions. The light field might be regarded as the lighting equivalent of the BRDF. One can imagine constructing a light field but taking a photometer and placing it a each point in space in turn and measuring the light coming from each direction, repeating for all points. Alternatively a 360 degree panoramic camera or scanning camera might also be used for such purposes greatly speeding the measurement process. Even so measurement of a full natural light field is still somewhat impractical. A simplification of this concept, the light probe, is essentially a sample of the light field taken at just one point in space. Dror, Willsky, and Adelson (2004) report example light probes for a range of settings captured using an omni-directional camera.

It is relatively easy to convert a light probe into the model light source for a scene. One simply imagines the light probe image pasted onto the inside of a sphere containing the scene to be rendered. Backward ray tracing will indicate a point on the sphere as the appropriate light source for a given viewpoint-object interaction and the luminance and colour values at that point in the light probe image determine the intensity and colour of the light source. This technique combined with a good model of surface reflectance can produce very realistic renderings. However, light probes are measured / estimated at a point in space Morgenstern, Geisler and Murray (2014) note that objects do not exist at a point

but occupy a volume of space. To fully render an object using this method then one should measure the light field for all area of space occupied by the object and then use the appropriate light probe for each point on the object to be rendered. There are also issues with the limited dynamic range of the cameras used to capture light probes. Dynamic range can be overcome but leads to a high measurement overhead. Morgenstern, Geisler and Murray (2014) propose a rather coarser measurement tool in the form of a multi-directional photometer which has high dynamic range but low resolution. Nonetheless this device might be sufficient to measure incident light over a sphere in space sufficiently well to produce good rendering of a similarly sized object as if located at that point in space.

While light probes might be a good tool for rendering objects as if in a specific scene or possibly into some generic natural scene in most cases rendered objects are depicted in some imaginary scene. Most rendering programs have the ability to model multiple light sources of various kinds (eg point sources, spot lights, area lights). If the imagined scene is indoors such light sources might be sufficient especially if the rendering program takes inter-reflections into account and thus models the light reflected from walls. However many renderings exclude inter-reflections and use alternative rendering tricks to mimic ambient lighting as we outline below. For outdoor lighting, it is relatively easy to model the sun as a point or small area source but the sky represents a hemispherical diffuse source and these are not easy to model. Further,

reflections of the ground plane and buildings require the modeling of inter-reflections and this is often finessed in rendering programs. For good rendering of either indoor or outdoor scenes either a light probe light source or the modeling of inter-reflections is preferred however these steps are often missed in rendered images. We now discuss two alternative approaches which produce pleasing images but which may not capture the physics of lighting well.

1.6.6 Inter-reflections and ambient lighting.

The simplest approximation to diffuse light sources and for dealing with inter-reflections is to lump all diffuse or indirect lighting into a single ambient illumination term and apply this to every object in the scene. Effectively one multiplies the reflectance of the surface by the ambient illumination term and add this to the amount of light reflected from the surface to the observation point. This method avoids areas of completely black shadowing – which are unrealistic in most scenes but does not capture the subtleties of diffuse lighting at all well. For example Schofield et al 2011 (see also Langer and Bulthoff, 2000) have shown that a corrugates surface rendered under a well modelled 'sky' source produces a small luminance variation across the surface. Simply setting an ambient lighting term as is often the default for rendering programs would not produce such variations.

One technique to model indirect illumination is to calculate the light reflected from each surface due to direct illumination and then create new virtual light sources for each object corresponding to the light they reflect and re-render the scene with these additional light sources. This is roughly equivalent to considering one inter-reflection per light ray in a reverse ray trace method.

1.6.7 Summary

In conclusion rendering methods can be used to simulate natural scenes under natural lightning. The best way to achieve this is involves using measured BRDF's to represent surface reflectance properties and measured lighting probes to represent natural lighting. However these techniques are seldom used in vision science. Rendering with the Blinn-Phong reflectance model can produce acceptable results for a family of matte and polished surfaces. Modelling natural lighting with diffuse and point source lighting can be achieved but is often not implemented. Similarly inter-reflections should be considered but are often absent. The use of an ambient term to replace the effects of indirect and diffuse lighting is common but does not represent natural lighting especially well.

1.7 General Methods

1.7.1 Apparatus

Except were noted, visual stimuli were constructed and displayed using C++ and were presented on a Sony 520GDMF CRT monitor using a VSG2/5 graphics card (CRS Ltd, Rochester, UK). For experiments 3 and 4, stimuli were images rendered via PovRay (Persistence of Vision Pty. Ltd., Williamstown, Victoria, Australia. http://www.povray.org/). E-Prime was used to present the stimuli of Experiment 4 and record reactions times.

Experiments took place in a dark room in order to eliminate any cue to lighting direction. Participants viewed the screen while seated on a chair via a monitor positioned 114cm from their eyes (unless differently stated in each experiment method section). A chin rest was used all the time to minimize head movements.

The screen refresh rate was 120Hz and the resolution was set to 1024x768 pixels. Before each experiment the monitor's gamma non linearity was estimated using a CRS-ColourCal Photometer and corrected using look up tables in the VSG.

1.7.1.1 Standard Procedures

Participants were tested in separate sessions of at most 45 minutes each, the number of sessions varied for each experiment.

They all gave their consent to take part in experiment and were aware that could have left the experiment at anytime. They were paid £6 per hour of testing and their visual acuity and stereo-vision was tested before taking part to the experiment, as a minimum of 50 Arc sec of stereo-vision acuity was requested.

1.7.1.2 Analysis

All data were pre-processed by C++ and analysed via MATLAB, Excel or SPSS. In experiment 1 (chapter 2) data were fitted using MATLAB CurveFittingToolbox and statistics was run on SPSS. All graphs and plots were generated via MATLAB and then modified via Adobe Illustrator.

In experiment 2 (Chapter 3) photographs of real surfaces were processed via MATLAB to change format and size and to normalize the contrast. The noise textures and the gauge figures were added online via C++, data were analyzed using MATLAB polar plots and then processed via Adobe Illustrator.

1.7.2 Classification Images

1.7.2.1 Classical Approach

In classical psychophysical studies, a great variety of behavioral methods have been used to understand the mechanisms of visual processing. Classification images (CI) is a technique that makes use of reverse correlation to characterize the strategies / templates used by observers in visual tasks. It was firstly introduced using Vernier acuity tasks (Ahumada, 1996; Beard and Ahumada, 1998); a test where observers are shown two bars and their task is to detect whether they are aligned or not. Weak targets are presented embedded in relatively strong random noise samples.

The noise samples (excluding any target) are then accumulated according to the participant's responses. There are four possible stimulus-response categories for each trial. Target-present trials can lead to 'Hits' where the observer gives a positive response, or 'Misses' in case of negative responses. Target-absent trials can lead to 'False Alarms', if the observer gives a positive response, or 'Correct Rejections'.

It has been shown (Ahumada, 1996) that the correlation between the luminance at each pixel and the observer's response can be found by accumulating the

classification image across trials based on the four stimulus-response categories:

$$(1)\ c = (n^{12} + n^{22}) - (n^{11} + n^{21})$$

Where n^{SR} is the average of noise samples in a stimulus-response class of trials**:** for example n^{12} is the average of noise samples that could be identified as 'Miss' where the stimulus contained target 1 (i.e. target present)

and the observer identified it as target 2 (i.e. target absent); n^{11} is the average of noise samples that could be identified as "Correct Rejection" where the stimulus did not contain target 2(i.e. target absent) and the observers identified correctly as target 2 (i.e. target absent); n^{21} is the average of noise samples labeled as "Hit" where stimulus contained target 1 (i.e. target present) and the observers identified correctly as target 1(i.e. target present); finally, n^{21}labeled as "False Alarm" where the stimulus did not contain target2 (i.e. target absent) but the observers identified as target 1 (i.e. target present).

The CI method reveals the observer's template by making use of influence of the random noise sample on the perceptual decision. In a task performed at threshold where observers can either respond correctly or incorrectly, the noise can coincidentally be correlated with target features and lead the observer to

incorrectly identify a signal even when it is not present or vice versa. In particular n^{12} and n^{22} samples should highlight the critical features that the observer uses to recognize a stimulus as being target 2, as both noise signals led the observer to believe that the target 2 was present. In the same way n^{11} and n^{21} should highlight features that the observer took to be similar to target 1 and dissimilar to target 2. The difference between the two accumulated samples could be then seen as positive and negative pictures of the most pertinent features of the image for the tested observer. Accumulating over many noise samples reduces the effect of noise features that are irrelevant to the observer and thus more clearly exposes those that are relevant. The noise can be further reduced by adding one accumulated image to the negative of the other. This is the basis of equation 1.

The roots of classification images are to be attributed to Ahumada and Lovell‘s (1971) study, where they developed a similar method for auditory psychophysics. In their experiment, they investigated what stimulus features where used by observers to report the presence or absence of a 500Hz signal tone in a computer-generated Gaussian noise signal. In the case of a positive response, the observer also has to report his/her grade of confidence on a four-step rating scale. They conducted a multiple regression analysis to extract the weights that observers gave to different frequencies to judge the presence or the absence of the signal.

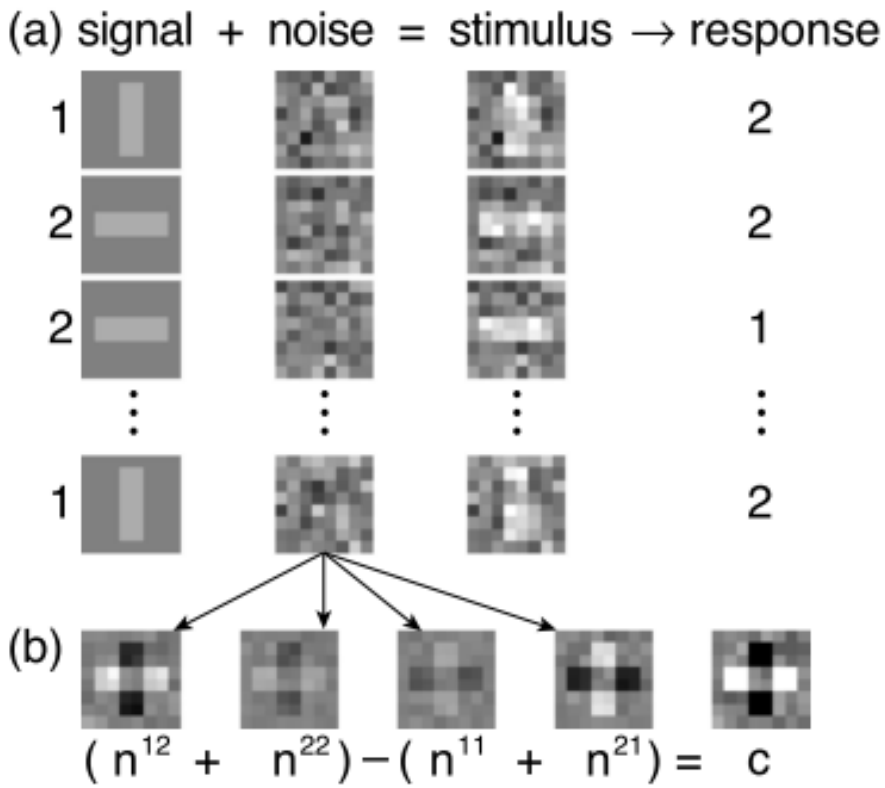

Figure 1.9 A model of the_observers' performance in a classical classification images task. On each trial (a) a signal is embedded in noise to produce a stimulus. The participants generate a response based on their personal biases. The analysis (b) is then based on Equation 1: noise patterns are added and averaged together in categories based on the observer's responses; the resulting classification image (c) is a template of the filter used by the observer to perform the task. (Murray, 2011)

Ahumada and Beard later introduced the standard classification image method used in visual psychophysics in an experiment where they tested humans in Vernier discrimination tasks (Ahumada, 1996; Ahumada and Beard, 1997, 1998). In this experiment observers needed to judge whether a segment was aligned or not to another line segment at a fixed position. Their results showed that humans rely heavily on the line segment that has the same position for both intervals even though he provides no direct information about the correct response. The influence of these studies has been crucial to the development of the CI image technique as a psychophysical method.

1.7.2.2 The Linear Observer Model

To better understand the classification image method, it would be useful to introduce and explain the linear observer model. Many different papers have concluded that the linear observer model is a good starting point to comprehend the mechanisms involved in the classification image method (Abbey, Eckstein, Bochud, 1999; Ahumada, 2002; Murray, Bennett and Sekuler, 2002; Solomon, 2002).

In a yes-no experiment where two signals s^1 and s^2 are added to noise N we would have two stimuli:

(2) $$I^i = s^i + N$$

We will consider I^i, N and the following variables as $n*1$ column vectors even though the stimuli in the experiment are two-dimensional images.

According to the linear observer model, the observer's decision (d^1 or d^2) will be based on the dot product of the stimulus with two templates, t^1 and t^2 (which are internal representations of the signals). In this process, the observer may add internal noise (z^1 and z^2) to the dot product. The resulting decision variables will be:

(3) $$d^1 = t^{1T}(s^i + N) + z^1$$

(4) $$d^2 = t^{2T}(s^i + N) + z^2$$

Where T indicates the matrix transpose operator so $p^T q = \sum_i p(i)q(i)$ is the dot product of column vectors p and q.

In this model the observer perceives s^1 if d^1 plus some personal bias (B) is greater than d^2

$$r = \begin{cases} 1 & if \quad d^1 + B > d^2 \\ 2 & otherwise \end{cases} \tag{5}$$

And therefore $r = 2$ would be:

$$t^{1T}(s^i + N) + z^1 + B < t^{2T}(s^i + N) + z^2 \tag{6}$$

Which is equivalent to:

$$(t^2 - t^1)^T (s^i + N) + (z^1 + z^2) > B \tag{7}$$

The last equation means that the observer's decisions are based on template differences $w = t^1 - t^2$ and the relative internal noise z with a variance of δ_z^2 that is twice the variance of z^1 and z^2:

$$d = w^T(s^i + N) + z \tag{8}$$

$$r = \begin{cases} 1 & if \quad d < B \\ 2 & otherwise \end{cases} \tag{9}$$

Murray (2011) depicted a decision space map based on linear observer's strategy to show the categorization applied by the observer. This representation

has been used to explain the standard weighted method used in equation 1 to give an unbiased estimate of the template. The linear observer model could be conveniently used to simplify the aspects of human observers performance in a classification image task since conclusions are simply deduced from the fact that the final templates show robust correlation between the observers' responses and the stimulus area with the most informative region of interest. It is important to point out that we should be aware that is not a complete explanation of the method and the mechanisms involved in it. In particular, many aspects of human visual processing such as second order vision, spatial uncertainty and many others are certainly non-linear. In this concern, some studies have tested in more detail the linearity assumption in the classification image method as it turned out to be valid (Abbey and Eckstein, 2002; Murray, Bennett and Sekuler, 2005).

1.7.2.3 The Applications in Visual Psychophysics

In the last fifteen years many studies have made extensive use of classification image techniques in visual psychophysics. The peculiarity of this method is that the noise itself represents the experimental manipulation. Thus, it gives the opportunity to explore both the effect of noise and key features of the images in visual processing. This technique also has some disadvantages such as the need for several thousands of trials to create a robust and discernible

classification image. In this concern, some studies have tried to find a solution by simply reducing the number of pixels used, or by applying radial averaging to two-dimensional classification images (Abbey and Eckstein, 2002). Their analysis let them decrease the number of trials needed up to an average of 2000 trials per observer in order to have a precise estimation of the final CI. Another approach is to avoid the use of bi-dimensional Gaussian noise in favor of one-dimensional noise patterns to be used as both test stimuli and noise (Levi and Klein, 2002).

The type of noise used is also critical to the classification image method: Abbey and Eckstein (2000) showed how the use of non-white noise could affect the observed results (i.e. blurred noise pattern produces blurred CIs) in a discrimination task, and also that observers' templates (CIs) change as a function of noise amplitude (Abbey and Eckstein, 2009).

Many studies on the other hand have used classification images to reveal perceptual organization in visual processing. Gold, Murray, Bennett & Sekuler (2000) have studied the perception of illusory and occluded contours by using a shape discrimination task (figure 1.10). They took inspiration from Kanisza and Gerbino (1982) to induce observers to perceive a square with different characteristics (fat or thin) depending on small rotations of corner orientation. The corners could be either illusory contours or partially occluded but the

classification images nevertheless showed how observers relied on both of them as much as they relied on actual luminance-defined ones.

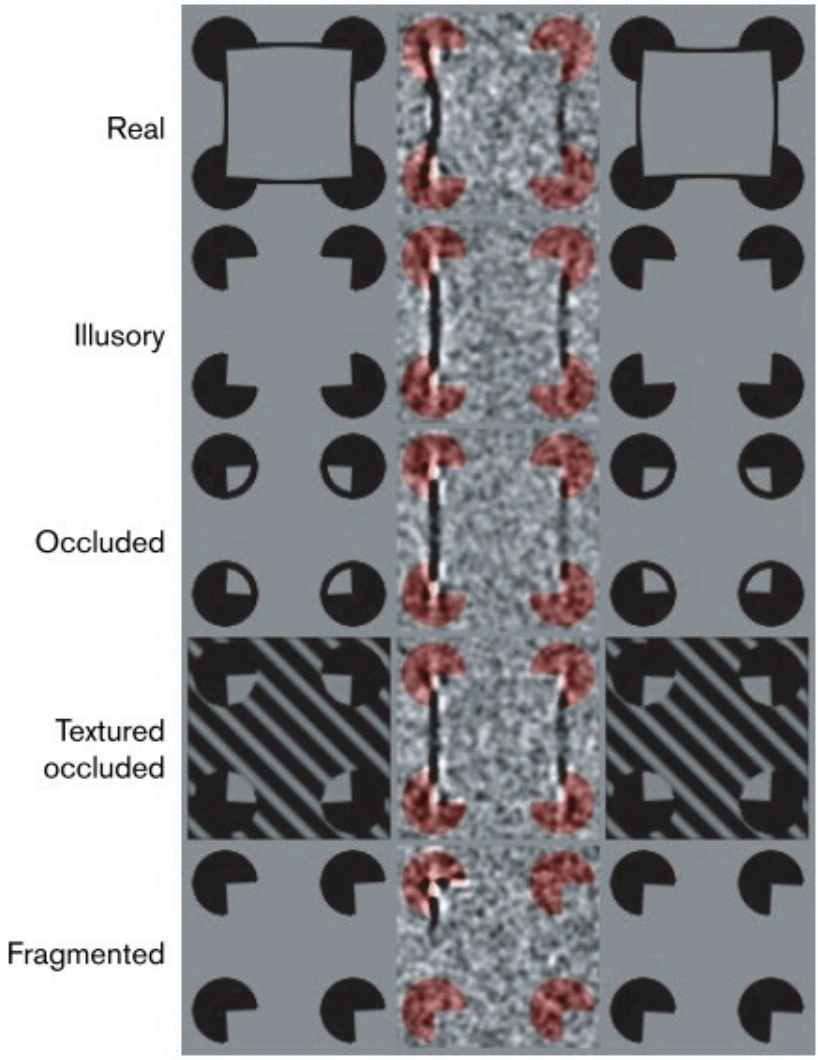

Figure 1.10. From left to right: thin stimuli, average classification images and fat stimuli. Each row corresponds to different conditions as explained by labels. The templates reported in the central column are averaged and smoothed classification images from three observers. Red Inducers have been superimposed on templates. (from Gold et al. 2000)

Keane, Lu and Kellman (2007) used classification images to reveal spatiotemporal contour interpolation. Contour interpolation is the process where

separated object fragments are perceptually unified or connected even if there is no physical connection between the two (Kanisza, 1976,1979). Their study showed that when observers have to interpolate illusory contours over space and time, both the time necessary for completion and the influence of noise along the contours are not strongly affected by spatiotemporal dimensions.

1.7.2.4 The Noise Only Method

As explained before, the role of noise is crucial to the classification image method. In this regard, the experimenter has to be very specific about the task that needs to be performed, especially if it involves the detection of a signal present only in some of the trials. An early study from Wiener (1958) showed that noise alone could be used to analyse the behavior of a 'black box' system (a system whose internal workings are unknown), and he also suggested that this approach could be extended to humans.

Taking inspiration from this study, Gosselin and Schyns (2003) tested observers in a task based on unstructured white noise only: no target stimulus was ever presented. They designed two experiments where observers had to report the presence or the absence of a given target embedded in noise. Observers were led to believe a signal would be present in some of the trials, even though no signal was ever added to noise. By using reverse correlation they wanted to extract the observers' internal representations in a letter (Experiment 1) and a smile (Experiment 2) detection task. Their results (figure 1.11) showed that not

only could observers apparently detect the target letter (the letter S), they also did so with a preferred, unique font and style. The three observers produced templates that best matched an uppercase Courier New bold S, a lowercase Verdana regular-style S and an uppercase Arial Bold S respectively.

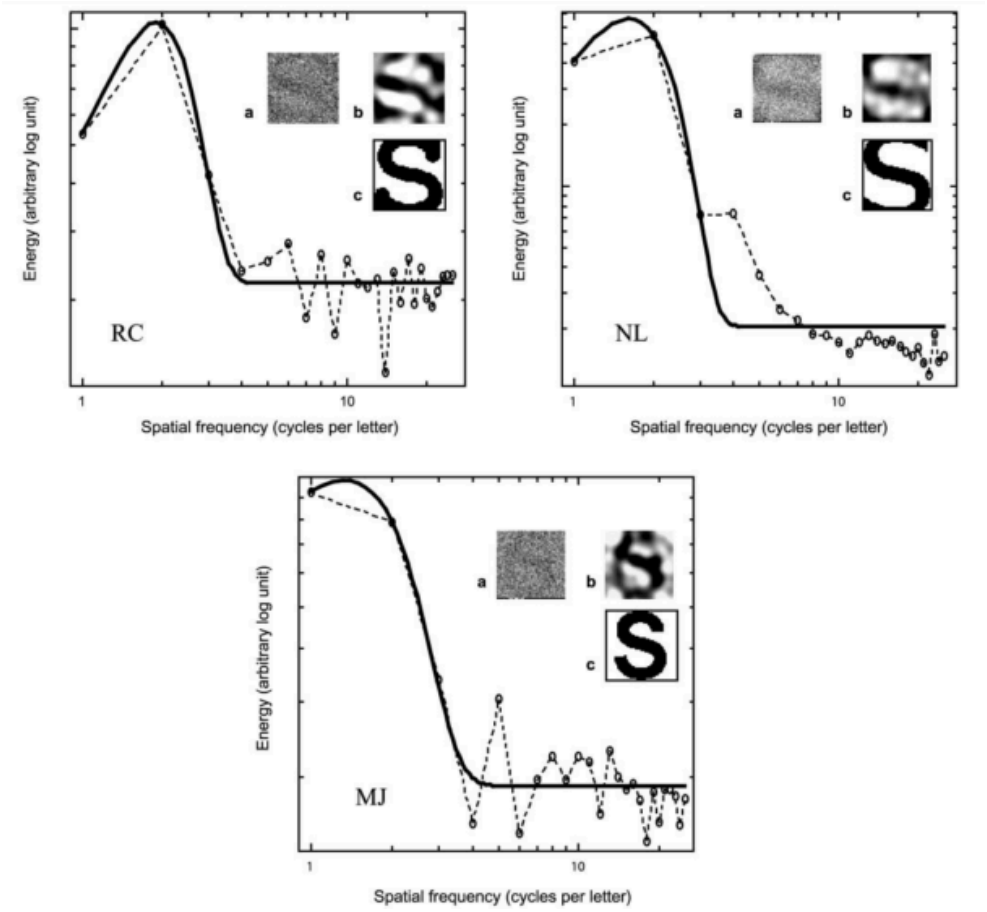

Figure 1.11 Data of three observers from Gosselin and Schyns (2003). The graphs show the distributions of the squared amplitude energy for different spatial frequencies of the final template. The solid lines are best Gaussian fits. a) raw classification images (CI), b) filtered CIs, c) font best matches.

In a similar task, Gosselin, Bacon and Mamassian (2004) showed that observers can perceive (as in, report the presence of) 3D structures in stereograms formed from disparity noise only. The linear observer model can be applied to the noise only method with some changes.

In this case, during the experiment, the observer has to match two vectors on each trial: a stimulus of dimensionality k and a template vector T, which corresponds to the internal representation of the signal (i.e. the letter S). Assuming that the observer's response is a linear function of this match, we can arrange the matrix X as $n*k$, where n is the stimulus vector of the experiment. A linear equation will describe the observer's behavior in the task:

(10) $$d = T * X + z$$

where d is an n-dimensional response vector and z is the internal noise with estimate $E(z) = 0$ and variance $V(z) = \delta^2 * I$.

The least square estimate of T is

(11) $$T = (X'X') - 1X'd$$

Since the stimulus vectors are uncorrelated. We then have

(12) $$(X'X')^{-1} = (kI)^{-1} = k^{-1}I$$

Consequently,

(13) $$T = k^{-1}X'd$$

If we take the constant k aside, and assume that the responses could only have values 1 and -1, the last equation (13) could be reduced to a sum of all the

stimulus vectors that led to positive responses (1) and subtracting from it the sum of all the stimulus vectors that led to negative responses (-1).

The resulting classification image will be:

$$c = (\sum_i n^{pos}) - (\sum_i n^{neg}) \quad (14)$$

The use of noise only stimuli is critical for final templates on a classification image task. Since white noise has the same energy across the whole image at all spatial frequencies, there is no bias that could favor any response in the experiment. The interpretation and therefore the “superstitious perception” is completely due to internal representation of the observers: as for Gosselyn and Schyns (2003), those features of the image that highly correlate with the structure of the letter S imagined by the observer. As mentioned before, since white noise could be considered a bias-free or empty stimulus, it could also generate an empty final template in cases where the observers do not apply a constant strategy or if they simply respond at random.

In our experiment (Chapter 2) we used the noise only classification images method to reveal observers’ lighting prior in a shape-from-shading task. The aim of the experiment was to access observers’ prior in a more direct/bias-free way. Even though white noise does not engage the visual mechanisms involved in shape-from-shading, the use of the noise only method can help to have access and directly measure people’s templates for interpreting shaded surfaces.

In conclusion, noise only classification image method can be usefully used in cases where experimenters want to access observers’ internal templates not

only for simple objects such as letters (Gosselyn and Schyns, 2003), but also in tasks where higher cognitive functions are involved such as faces (Gosselyn and Schyns, 2003) and lighting priors (Mazzilli and Schofield, 2013).

2. A cue-free method to probe human lighting biases

This chapter introduces the application of the classification image technique to shape-from-shading tasks. Previous studies have made use of ambiguous stimuli in order to extract observers' assumed light source. Here, we wanted to access people's lighting prior more directly by establishing the template they would employ to detect a shaded object in absence of any visual cue to object shape.

2.1 Introduction

People readily perceive patterns of shading as 3D shapes (Langer and Bulthoff, 2000; Ramachandran, 1988; Sun and Perona, 1998; Todd and Mingolla, 1983; Todd, 2004; Tyler, 1997). Due to the generalized bas-relief ambiguity when extracting shape-from-shading people must simultaneously estimate the shape of the surface and the nature of the light source. In many cases**,** cues in the image will be insufficient to resolve all of the ambiguities present and in such cases the human visual system may employ one of a number of prior assumptions based on ecology and experience. One such assumption is the lighting-from-above prior (Adams, Graf and Ernst, 2004; Brewster, 1826; Gerardin, Kourtzi and Mamassian, 2011; Mamassian and Goutcher, 2001; Ramachandran, 1988; Rittenhouse, 1786; Sun and Perona, 1998; Todd, 2004).

Here, in the absence of extrinsic cues to lighting direction, ambiguous shading patterns are interpreted as if lit by a light source that is above the observer's head.

When presented with ambiguous stimuli, the human visual system must decide on a percept by making assumptions about the scene which may not be correct. This is the case in Figure 2.1 where a concave dip lit from below is often perceived as a convex bump lit from above.

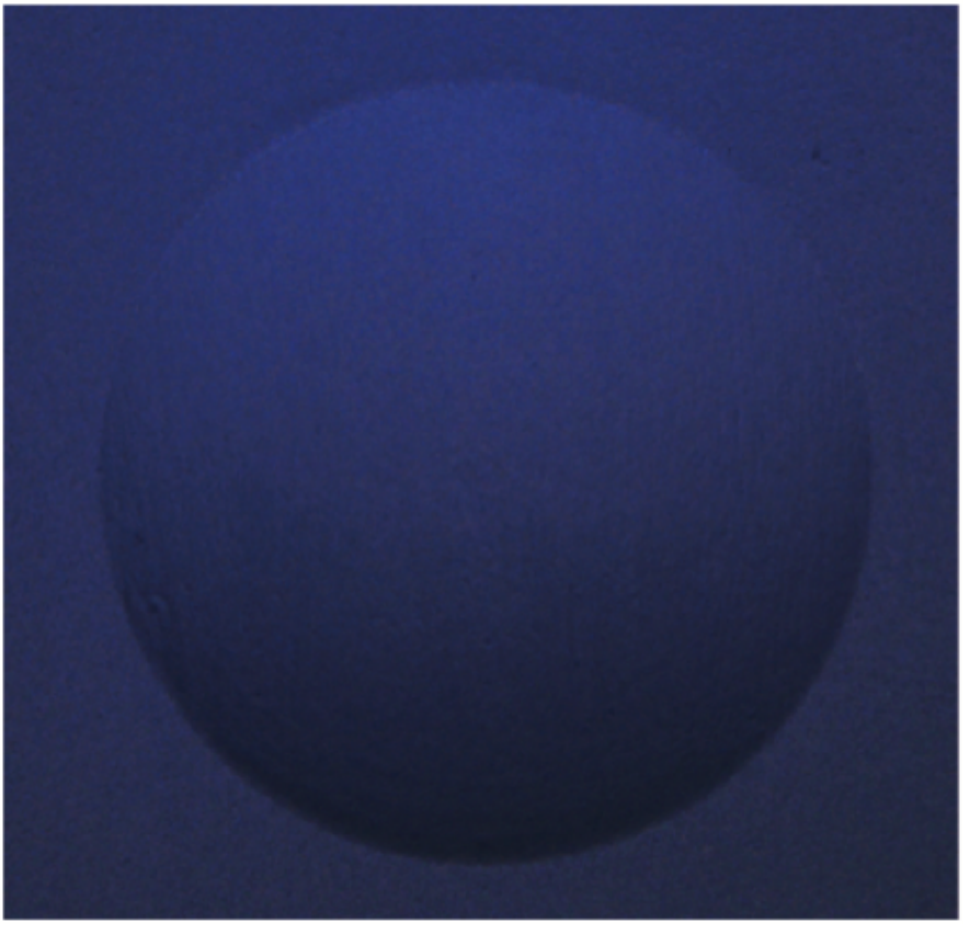

Fig 2.1. A photograph of a physical concavity lit from below which appears as a convexity lit from above.

In the absence of extrinsic cues to lighting direction, ambiguous shading patterns are often interpreted as if lit by a light source that is above the observer's head (Ramachandran, 1988; Sun & Perona, 1998; Todd, 2004; Rittenhouse, 1786; Brewster, 1826; Mamassian and Goutcher, 2001; Adams, Graf and Ernst, 2004; Gerardin, Kourtzi, Mamassian, 2010). However, the lighting-from-above prior is not the only lighting assumption that humans adopt (Tyler, 1997; Langer & Bulthoff, 2000; Schofield, Rock and Georgeson, 2011), and lighting assumptions can be overridden by assumptions about object shape (Liu & Todd, 2004), and by lighting cues in the image itself (Morgenstern, Murray and Harris, 2011). In many studies showing robust lighting-from-above priors the stimulus is clearly shaded by a strong directional light source but in an ambiguous manner; such that the observer is forced to choose between two interpretations, one of which happens to be consistent with lighting-from-above while the other is not. In such cases lighting-from-above 'wins' but it may not be so robust in other circumstances. Recent studies (Liu & Todd, 2004; Morgenstern, Murray and Harris, 2011) have demonstrated how easily the lighting-from-above prior can be overridden. We test the robustness of the lighting-from-above prior using stimuli that have no shading structure and no lighting cues. Thus we directly probe the observers' internal templates or assumptions for interpreting shaded objects in the absence of any interference or bias from stimulus features.

We address this question using classification images as a tool to probe the lighting prior. This technique (see Murray, 2011 for a review) has been used in psychophysics in order to elucidate the nature of templates or filters within human vision by correlating observers' decisions with apparently uninformative noise features over a large number of trials. It has been used to investigate problems such as Vernier acuity (Ahumada, 1996), illusory contours (Gold et al., 2000), and letter discrimination (Watson, 1998). The classical CI approach has been extended to include stimuli composed of noise only images (Gosselin & Schyns, 2003). In this study observers were instructed to detect the presence of a target in white noise; no signal was ever added to the white noise patch but observers were led to believe that the stimuli included a target signal. As with the classical method, the resulting classification image represents the template used by the participant to perform the task. If the observer does not apply a systematic approach, the resulting template should have the same statistical properties as the noise. A systematic approach would be indicated by structure in the template and this structure reveals both the observers' impression of what the target should look like given the instructions, and their strategy for detecting it. The use of noise only stimuli is crucial to our experiment since it is impossible for such stimuli to bias perception. Thus we directly measured people's templates for interpreting shaded surfaces.

Observers' templates were derived by accumulating the noise samples leading to positive responses and subtracting the accumulation of the unselected noise

samples. Local features in the noise samples that happen to support a certain percept, in this case the percept of an illuminate bump, will appear often in noise images that are deemed most bump like by the observer and will accumulate in the average of all the positive samples whereas uninformative noise features will tend to cancel each other out. Likewise features that consistently lead to noise images being unselected (non-bump features) will accumulate in the average of the unselected images. Thus the subtraction of the two accumulated images will produce a template of those features that observers use to 'discriminate' bumps form non-bumps: that is their bump template. Since shaped bumps tend to have a highlight, templates for such bumps should contain a bright feature indicating the part of the imagined surface that is directed towards the imagined light source. If observers apply the lighting-from-above prior we might expect such highlights to lie above the central row of their templates given that they were informed that objects would always be convex. A diffuse or frontal lighting prior might result in a central highlight, this being less well defined in the diffuse case and a light from below expectation would result in a highlight below the midline of the template. To test the dependence of the template on task demands, we tested three control conditions: large bumps (diameter 4 deg), cylinders (2.5deg x 4deg), and white disks (2.5deg). We expected to find larger offsets for highlights in the large bump condition; horizontally elongated templates for the cylinder condition; and flatter templates with less distinct highlights and no offset in the white disks condition (flat white disks do not imply any shading highlight so we might expect

participants to revert to the overall appearance of the disk – ie flat, central and whie – for their template).

Templates were typified by bright approximately Gaussian blobs, which varied systematically with the instructions given. To reveal the observer's lighting bias, we analysed the intensity levels of each template by fitting a Gaussian to the central column and by taking the position of its peak as the location of the illumination highlight that observers were aiming to detect in the stimulus. We judged the effective lighting prior by measuring the distance of the peak from the middle of the template.

2.2 Methods

2.2.1 Participants

Seven naive observers participated in the small bump condition; four of these also took part in the large bump condition; two observers participated in the white disk condition and one in the cylinder condition. All observers had normal or corrected-to-normal vision. All participants were students at University of Birmingham except TY and HB who were employees**;** all except the employees were paid £6/h. They all gave their consent to take part in the study

2.2.2 Apparatus

Noise samples were created and displayed in monochrome (grey) using Visual C++ 6.0 (Microsoft).. Stimuli were presented on a Sony 520GDMF monitor

using a VSG2/5 graphics card (CRS Ltd, Rochester, UK). The image size was 10x10 degrees of visual angle at the viewing distance of 114cm (resolution = 512x512 pixels). The monitor's gamma non-linearity was estimated using a CRS-ColourCal Photometer and corrected using look up tables in the VSG.

2.2.3 Stimuli

All stimuli were noise samples with an approximate 1/f amplitude spectrum in order to simulate the spectral content of natural scenes without introducing systematic structural features. Noise here is defined as a random signal (image) composed of sequences of uncorrelated intensity values with a mean grey level of. The use of 1/f noise produces some local correlations between pixel values because low spatial frequencies dominate such that adjacent pixels are likely to have similar values but there was no overall structure in any image and no correlation between pixels at the same location in pair of images. The use of 1/f noise results in a blurred final template (Abbey and Eckstein, 2000), however the use of this type of noise was deemed preferable (to say white noise) as people have a natural tendency to perceive shapes in 1/f noise but not so white noise. That is it is more plausible to ask observers to detect a bump hidden in 1/f noise than it is to ask the same question while presenting white noise.

2.2.4 Procedure:

We used a 2 Interval forced choice (2IFC) paradigm in which observers were told that a target would be added to the noise in one of two intervals chosen at

random and were asked to identify the interval containing the target. However, no targets were ever presented. Rather, the two presentation intervals contained noise only images and the participants indicated which one “looked most like” an imagined target. Each trial consisted of 4 images each of 150 ms duration: the initial fixation cross was followed by the first stimulus, another cross, then the second stimulus (figure 2.2). At the end of each trial the screen was set to mid grey pending the observer’s response.

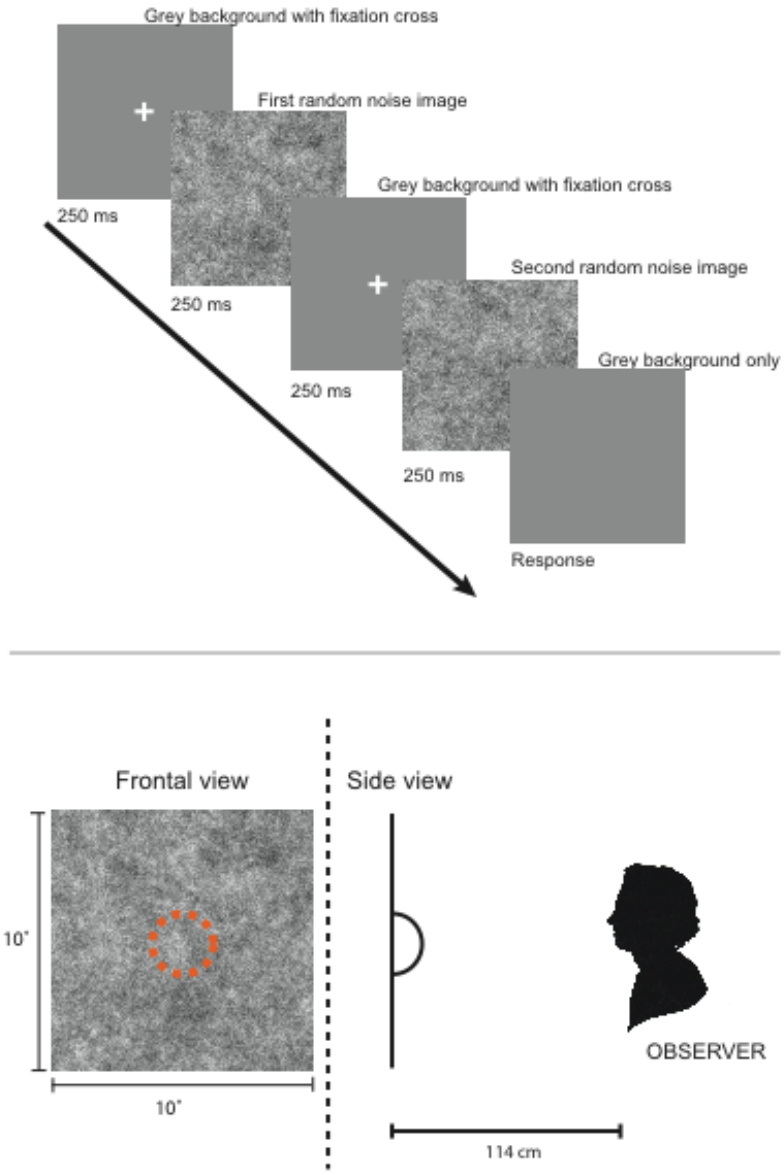

Figure 2.2. Top: experimental timeline 2 Interval forced choice. Observers looked at a central fixation cross embedded on grey background before each of the two noise samples were shown; a final grey background only was then shown up to observer's response. Bottom: frontal view showing relative size of small bump targets and side view sketch shown as part of instructions.

As no signal was ever presented we extended great caution when giving instructions because they represent the only information the participant has as to the nature of the target. At the start of the first session in any given condition

observers were shown a side-view cross-sectional sketch of the target without any lighting information. They were asked to imagine what this shape might look like in a frontal view. At no point did the observers see an example shaded target in frontal view. They were also instructed that the target, within each condition (session), had constant shape, size and location and that – with the exception of the flat disk - they were always convex objects coming out of the screen. They were asked to fixate the centre of the screen throughout the experiment.

We tested four target conditions: (i) small bumps, the imagined target was a raised bump with a diameter of 5 cm (2.5 deg); (ii) big bumps, the target was a raised bump with a diameter of 8 cm (4 deg); (iii) white disk, the target was a flat white disk with a diameter of 5 cm (2.5 deg); and (iv) cylinder, the target was a 5 cm high (2.5 deg) and 8 cm wide (4 deg) semi-cylinder

2.2.5 Analysis

The classic classification images (CI) technique typically uses weak target stimuli embedded in strong noise fields and asks observers to detect some target feature. The noise samples (excluding any target) are then accumulated according to the participant's response on each trial. There are four possible stimulus-response categories for each trial. Target present trials can lead to 'Hits' where the observer gives a positive response, or 'Misses' in case of negative responses. Target absent trials can lead to 'False Alarms', if the observer gives a positive response, or 'Correct Rejections'. It has been shown

(Ahumada and Beard, 1998) that the correlation between the luminance of each pixel and the observer's response can be found by accumulating the classification image across trials based on the four stimulus-response categories: CI=(Hit+False alarm)-(Miss+Correct rejection). One simply adds or subtracts the relevant images on a pixelwise basis. Since we had no targets in our experiments there were only two stimulus-response pairings and templates were computed by accumulating all the images that led to positive responses ('Yes images') into one pool and all the images that led to negative responses ('No images') into another pool. The classification image is then given by the equation (2.1):

(2.1) CI=(Yes images)-(No images)

As with the classic method, the resulting image represents the template of information used by the subject to perform the task. If the observer does not use a systematic approach, the resulting template should have the same properties as averaged noise. If not, the template indicates the presence of structures that underlie the perception of a target and more importantly, given the lack of a target in our stimuli, the observers' impression of what the target should look like given the instructions. We further adapted the CI method to use a two interval forced choice design such that on each trial the participant saw two similar noise samples (no target stimulus) and indicated which looked most like the imagined target. Thus each trial contributed an image to each of the positive and negative pools.

Templates comprised 512x512 images. We fitted a Gaussian curve to the central column and central row of each template (see Table 2.1). The peak of the best fit Gaussians gives the location of any highlight in the template while standard deviations indicate the size of such features. We estimated people's lighting prior as the difference between the centre of the screen and the peak of the best fit Gaussian.

PPS Initials	Variance	Mean (peak position)	Amplitude	R-Squared	Sessions
	Small Bump				
TY	30.81	215.9	58.15	0.6103	5
PS	43.97	244.9	87.8	0.9083	4
EL	42.29	256.2	156.3	0.9091	6
KG	43.69	266.4	80.78	0.8247	8
HB	40.31	254.2	232.1	0.9187	7
JH	41.88	258	99.41	0.7601	10
NK	46.27	242.9	52.32	0.6148	7
	Big Bumps				
TY	41.77	217.4	56.49	0.7594	5
PS	42.4	228.5	97.21	0.8356	4
EL	51.03	226.7	55.17	0.6456	6
KG	45.75	263	36.43	0.4056	6
	White Disk				
JH	44.59	243.9	159.1	0.738	4
HB	46.88	256.8	105.9	0.7438	4
	Cylinder				
HB	39.77	254.5	95.99	0.8262	4

Table 2.1. Gaussian Fit data for each participant and condition. Mean peak is an index of their lighting prior bias (as 256 is the centre of the image a smaller number would result in light-from-above preference while a bigger number would subtend a light-from-below bias); R-Squared is an index to the goodness of fit that highlights if the observer performed the task by using a consistent strategy across trials; Amplitude and variance are relative to the fitted curve and Sessions indicates the number of sets observers needed to produce a discernable template.

The Classification Images method, and in particular any biases in the noise pattern generation process, was tested in simulation using the same procedure used in testing human observers. We wanted to make sure that the use of 1/f amplitude spectrum did not influence the observers' performance (other than enabling them to perform the task set) and that any template resulting from equation (2.1) would be unbiased by the noise samples presented. The two noise stimuli were multiplied with a Gabor mask in each simulated trial and the pixel values in the post-mask stimuli summed to derive a score for each image. The image with the highest score was added to the positive image set, while the image with a smaller score was added to the negative set. That is, we used a simple linear model in which the model observer calculates the dot product of the stimulus with the two templates and makes a decision based on resultant differences. 'Observer' responses are simulated but with a known, reliable template. A simulated classification images template was then derived in the normal way. In order to complete this process we then repeated the same procedure using a circular template. The simulated classification images clearly showed templates similar to the ones used in their generation, indicating that the method works as intended provided observer templates are sufficiently stable. Finally when no mask is used in simulation, and the two images in each trial are assigned to the pools at random, the pooled images have the appearance of a noise sample verifying that there was no bias in the noise samples used.

2.3 Results

Observers produced clear target templates for each condition despite the fact that no signal was ever presented. Most claimed that the task was very hard and felt they were guessing in early sessions, but after a few sessions they felt more comfortable with the task and claimed to have found a strategy for it. No one ever doubted the presence of the target. Figure 2.3, 2.4, 2.5 shows example templates for each of the conditions tested, vertical sections and best-fit Gaussians for each participant. Gaussian fit results are reported on Table 1. Only half of the small bump templates revealed a light-from-above prior. Figure 2.6a plots distance from the centre of the screen in terms of degrees of visual angle, for each observer in both small bump and large bump conditions. In the small bump condition, 3 subjects (PS,TY,NK) had a positive offset: that is, they demonstrated templates consentient with a shaded bump lit from above; 3 (EL,JH,HB) had no offset and 1 (KG) had a negative offset (lighting-from-below). In the large bumps condition, we found that 3 subjects (PS,TY,EL) out of 4 showed a lighting-from-above prior while 1 (KG) maintained a lighting-from-below interpretation. Observers tended to keep the same preference for light-from-above or below despite changes in the target size. The four observers (PS, TY, KG, EL) that were tested in both Small and Large condition, averaged a shift of 1deg from one condition to the other; specifically, PS and EL had a shift towards the direction of their prior of 1.3deg,and 0.6deg respectively, while both KG and TY had a shift on the direction opposite to their prior of 2.1deg and

0.11deg respectively. The effect of target size on highlight location is better shown in Fig 2.6b, where mean, absolute peak offsets are shown. As expected, highlights were on average further from the centre in the large bump condition than in the small bump condition. We also conducted a one-sample t-test, to see if the mean highlight position relative to the image centre was different from zero. There was no significant difference for small bumps ($t(6)= 2.164$ $p=0.074$) but a significant difference for the large bump condition ($t(3)= 3.844$ $p=0.031$).

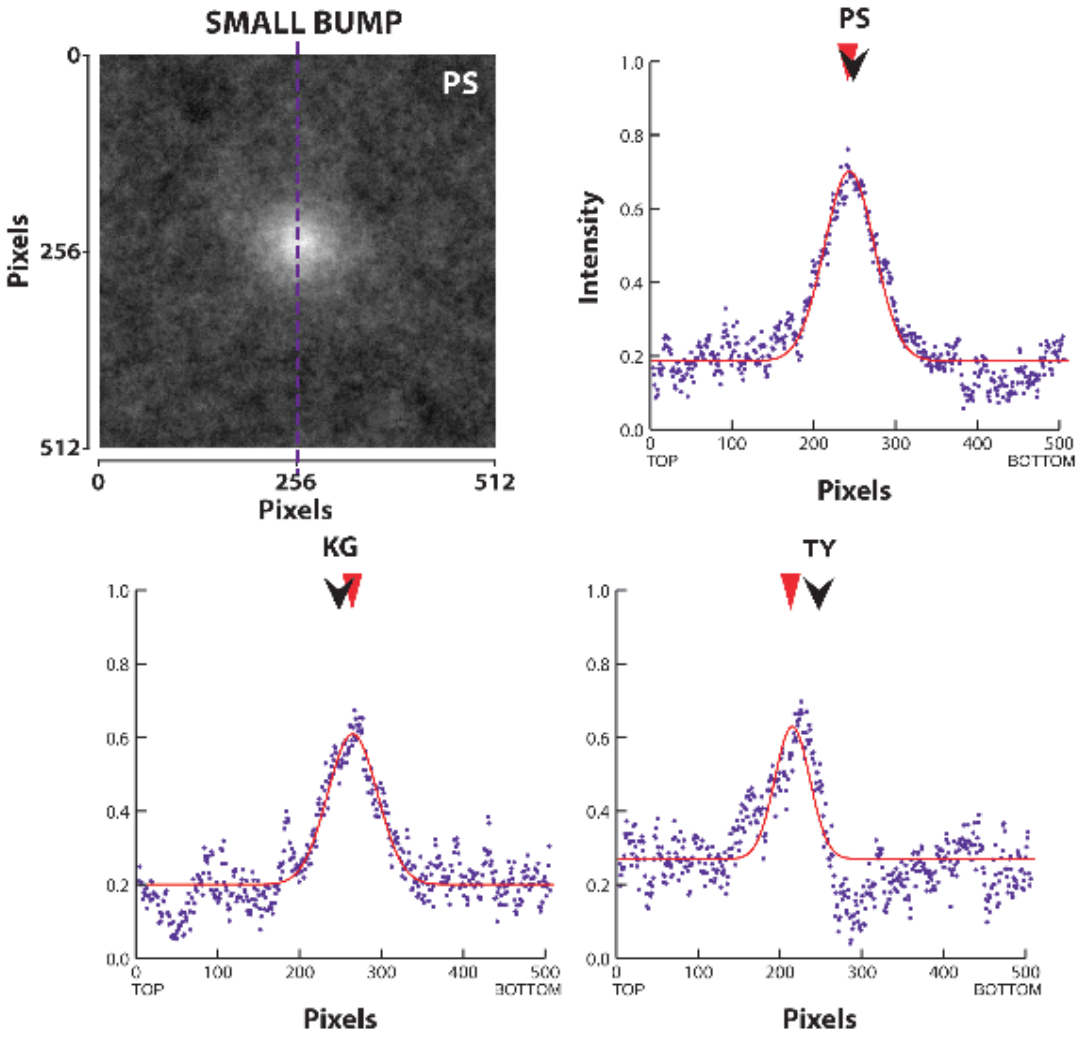

Fig 2.3. Example templates and cross sections for Small bump condition. The black arrow points at the centre of the image (Column 256) while the red arrow points at the peak of the Gaussian Fit. The sign of the difference between the two arrows indicates the direction of the lighting prior

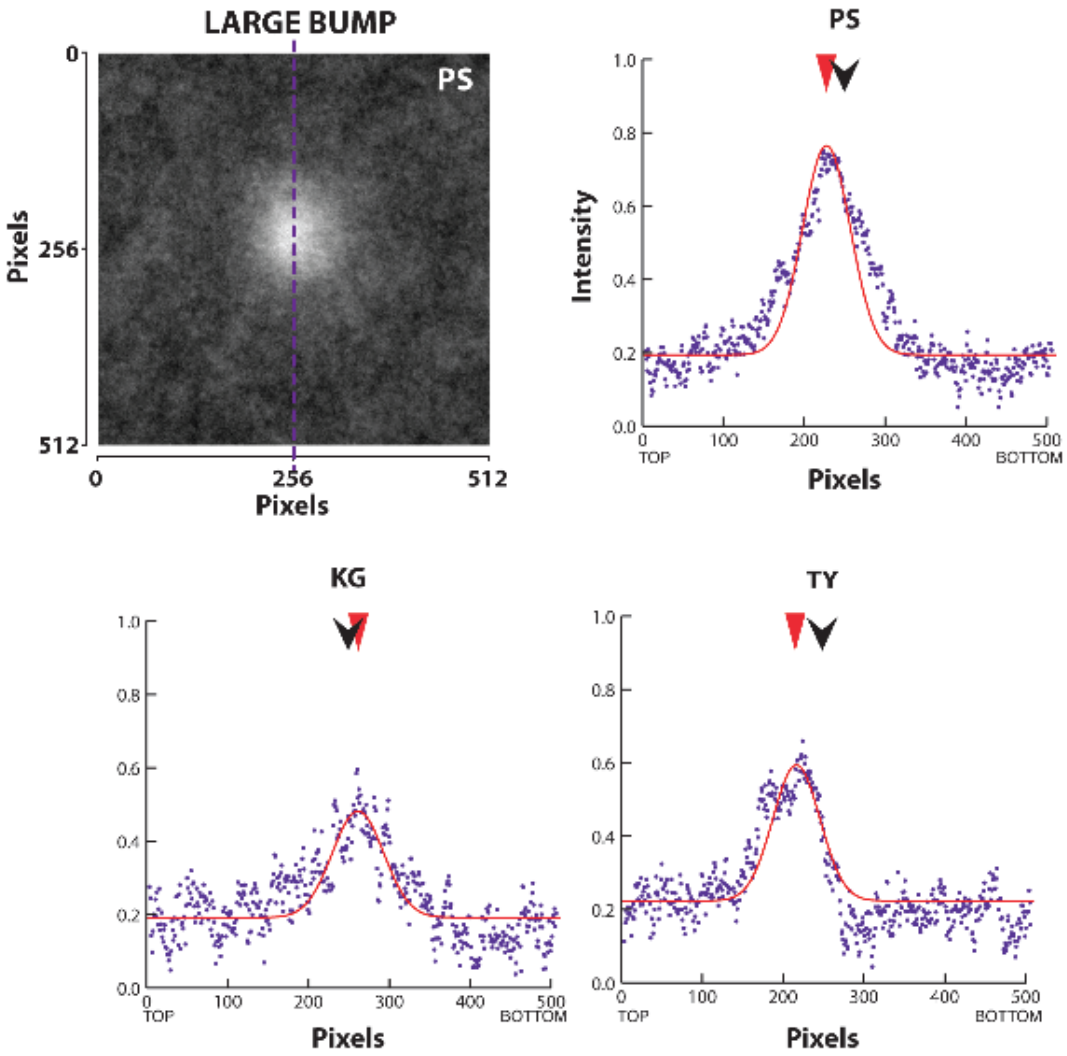

Fig 2.4. Large bump condition. The black arrow points at the centre of the image (Column 256) while the red arrow points at the peak of the Gaussian Fit. The sign of the difference between the two arrows indicates the direction of the lighting prior

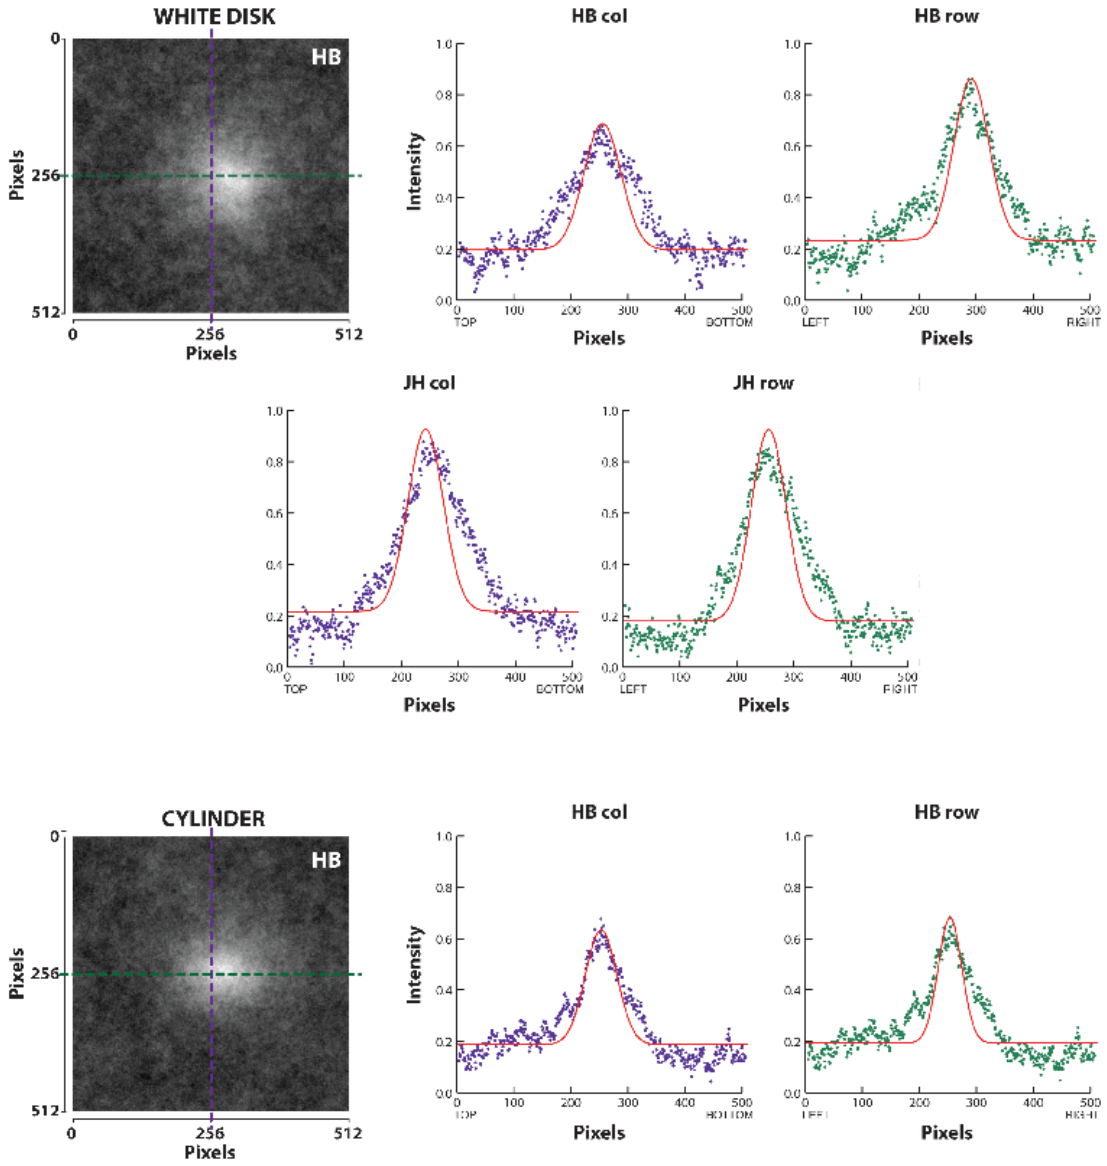

Fig 2.5. White disk (Top) and Cylinder (Bottom). The black arrow points at the centre of the image (Column 256) while the red arrow points at the peak of the Gaussian Fit. The sign of the difference between the two arrows indicates the direction of the lighting prior

Templates for the cylinder condition closely followed the small bump templates along the vertical axis but had a much wider profile on the horizontal axis. This reflects the elongated highlight one would expect from a cylindrical surface as compared to a spherical one. In the white disk condition, where the imagined object should not convey any lighting bias, template peaks were not offset from the centre of the image and templates were less well fit by Gaussian profiles being somewhat broader and flatter (that is, more disk like). These results confirm the task dependent nature of our templates and thus the validity of the peak offsets measured in the bump conditions. The good Gaussian fits to the bump templates suggest that people were expecting a concentrated highlight in the bump condition but the lack of consistency between observers in the location of the highlight suggests that the lighting-from-above prior is not universal. These results suggest that the lighting-from-above prior may be less robust and less prevalent than previously thought.

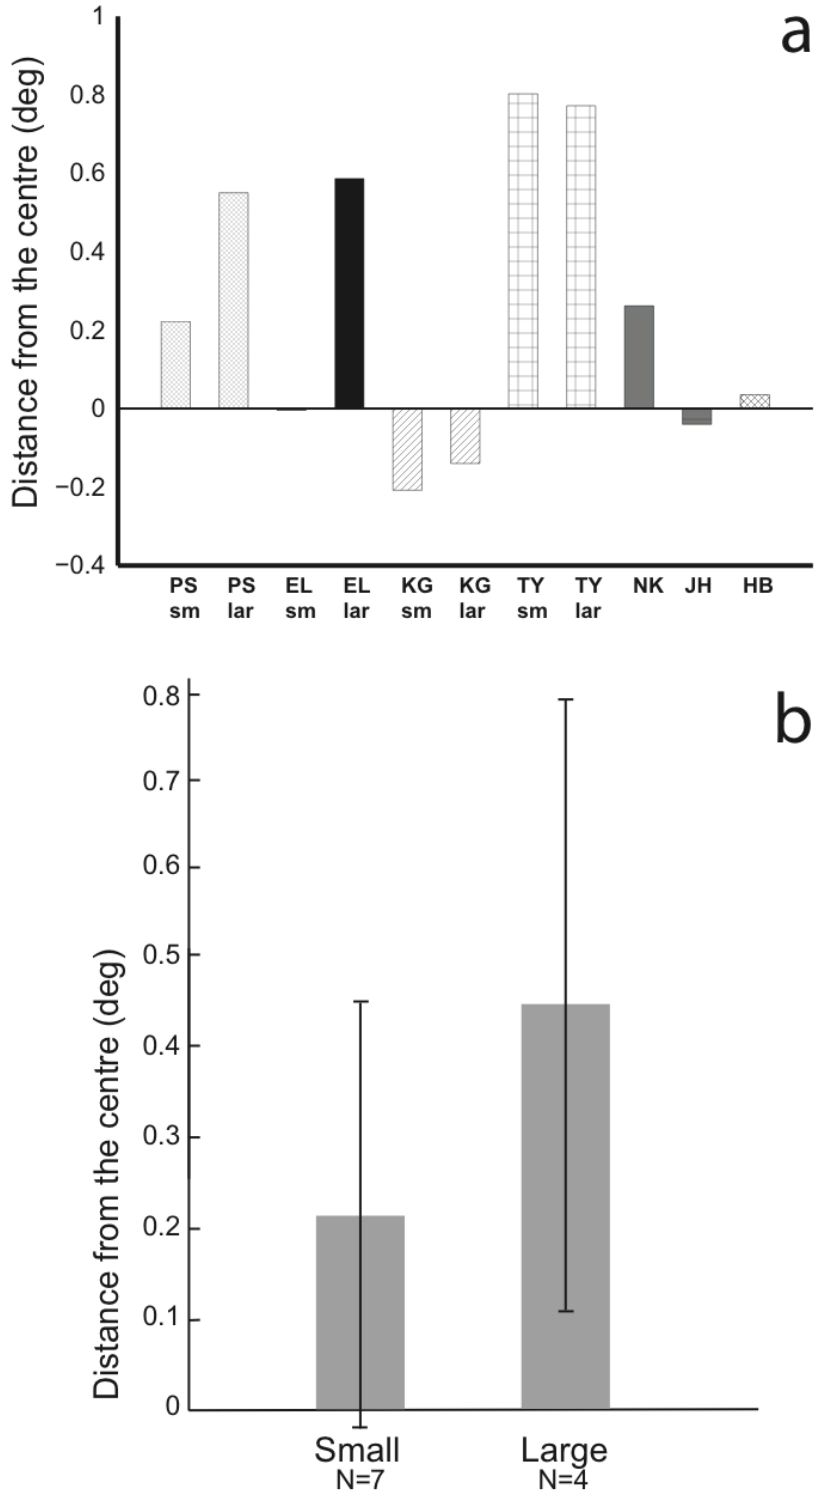

Fig 2.6 (a) all participants offset for both Small and Large bump condition (b). Lighting prior bar graph for small and large bump. The distance from the centre reported in terms of magnitude (no direction). Bars represent confidence intervals at 95%.

2.4 Conclusion

Assumptions made about the position of the light source play a crucial role in the perception of ambiguous stimuli (Rittenhouse, 1786; Brewster, 1826, Todd and Mingolla, 1983, Ramachandran,1988). In shape-from-shading, the use of ambiguous images has always been considered the best way to test observers' prior assumptions (Ramachandran,1988; Kleffner and Ramachandran, 1992; Sun and Perona, 1998; Mamassian and Goutcher, 2001; Geradin and Mamassian, 2007; Gerardin, Kourtzi and Mamassian, 2010; Morgenstern, Harris and Murray, 2011). Previous studies have found that a very high proportion of observers have a bias for seeing lighting-from-above at both long (Adams, Ernst, Graf, 2004; Sun and Perona, 1998; Ramachandran, 1988) and short (Mamassian and Goutcher, 2001) presentation time. However, these studies presented ambiguous stimuli which nonetheless depicted shaded images and may therefore have promoted the assumptions of a directional light source over, say, a diffuse interpretation. That is if the images were rendered under a directional light source (albeit one of unknown direction) there may be cues to this in the scene and such cues might lead the observer away from a diffuse lighting interpretation. We present an alternative method using noise only stimuli which – by definition – cannot bias observer responses in favor of a particular lighting type as the images contain no meaningful structure of any knd. A consequence of testing with no stimulus, as we have done, might be to reveal the latent robustness of the lighting-from-above prior. The large

variations we have found across participants suggest in fact that the light-from-above prior is fragile even in cases where a clear highlight is found. The type of noise used in our case produces a blurred template and hence object features such as the edges of the cylinder are absent from the templates and even the white disk results in a fuzzy template. Testing with simulated observers suggests that this is a property of the noise used not the underlying templates. Most participants produced a clear highlight in their bump templates as compared to the white disk template which was broader and the cylinder template which was elongated. This likely suggests that they were expecting to see a bright spot on the shaded bump that they were trying to detect and suggest that they were expecting directional illumination. We reason thus because diffuse illumination would produce images closer to the white disk condition. The fact that the distance of the highlight form the centre of the screen / template varied with expected bump size also supports the idea that observers expected a directional light source. This result futher suggests that they expected the light source to be directed at an angle that is not frontal to the object, as the latter illumination would produce a highlight in the centre of the screen. However the fact that several observers (a high proportion compared to what might be expected from previous studies) placed their highlight close to the midline or even below it suggests that the lighting from above assumption is not as common as would be expected from the majority of studies. These observers either expected a frontal illuminant or illumination from below although a prior expectation for diffuse illumination cannot be ruled out in such

cases. We can probably exclude the possibility that some observers expected a concave surface lit form above (which would produce highlights below the midline) as this would be a perverse interpretation of the instructions nonetheless the data from such observers do not preclude this geometric interpretation.

In summary, the adoption of a lighting prior is a logical strategy for the visual system to solve ambiguities in shaded stimuli such and the light-from-above assumption has ecological validity. Nevertheless, the large variation found here reveals its influence to be somewhat less strong than might be expected from the literature. Our results therefore contribute to the growing evidence that the light-from-above prior is weak and dependent on stimulus features.

Appendix - Observers Instructions

Thank you for taking part in this study.

In this study you will be asked to judge the presence or the absence of a Target in a noise image.

On each trial it will be shown two different noise images, in one of the two there will be present a Target which has constant size and shape and it is always in the same location.

At the beginning of each session you will be informed about the type of Target you will be shown. There are 4 types of Targets:

1. Small Bump – A half convex sphere with a diameter of about 5cm
2. Large Bump – A half convex sphere with a diameter of about 8cm
3. White Disk – A flat disk with a diameter of about 5cm
4. Cylinder – A cylinder horizontal side view of about 5cm height and 8cm width

Your task is to respond in which of the two intervals the Target signal is present. We will collect 10.000 trials per subject. We are not measuring reaction time.

As the Target is most of the time under detection thresholds, you may feel like you are guessing sometimes. Target presence is completely random so it is not guaranteed that it will be equally present across the two intervals.

Please do not hesitate to ask if you have any question.

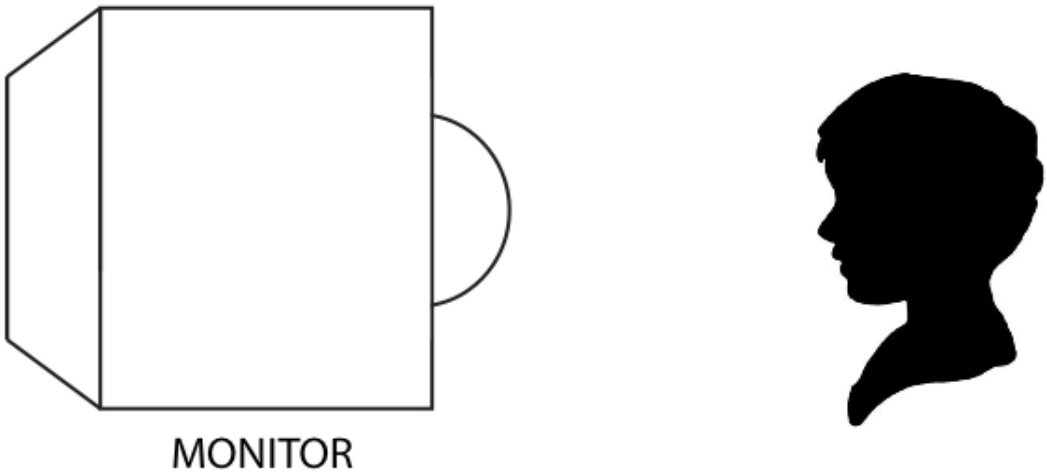

3. The effects of lighting direction and elevation on judgements of shape-from-shading.

The study reported in chapter 3 aims to investigate the effect of varying the declination (elevation) of the light source in recovering shape-from-shading. We generated images of three different surfaces (differing in their degree of anisotropy) illuminated with spotlights at 6 angles of declination (0 degrees frontal to 75 degrees oblique) with constant azimuth. We tested images at two orientations (upright and rotated by 90 degrees) and two contrasts (natural or normalised) since varying the declination of the light source affected the contrast of the resulting shading pattern. We found that shape judgements in isotropic shading patterns are more affected by the lighting-from-above prior than those for anisotropic surfaces and that the slant judgements depend on local luminance/contrast and are affected by changes in the declination of the light source because this affects the contrast of shading patterns.

3.1 Introduction

The knowledge of lighting direction is an important factor in determining the perception of shape-from-shading. Rittenhouse (1786) was the first to demonstrate changes in perceived surface shape with changes in actual - but

not perceived - lighting direction. More recent studies (Ramachandran,1988; Kleffner and Ramachandran, 1992; Sun and Perona, 1998; Mamassian and Goutcher, 2001; Gerardin et al., 2007; Gerardin, Kourtzi and Mamassian, 2010; Morgenstern, Harris and Murray, 2011) have confirmed that the direction of the light source, and the assumptions made about the position of the illuminant, play a crucial role on how people perceive ambiguous shading stimuli. There is good agreement that when the lighting direction is ambiguous people assume the dominant lighting direction to be from above with a small leftward bias (Mamassian and Goutcher, 2001; Sun and Perona, 1998). Most of the above studies used isotropic (specifically circular) stimuli that could be perceived as either a bump or a dent although Mamassian and Goutcher (2001) have shown that the lighting-from-above prior applies to non-isotropic stimuli.

Studies of shape-from-shading have also considered the effects of a known, or easily estimated, light source on perceived shape measure. Most studies use a viewer centered spherical coordinate system to measure observers' shape perception in these tasks and therefore, quantify their performance with two parameters: slant and tilt. Imagine a clock face. The illuminant tilt is the direction, around the clock face' from which the light is directed with zero=above, 90=right etc. The illuminant slant is the angle between the illuminant and the clock face and describes the degree to which the illuminant is directed from the direction of tilt versus being in-line with the observer. Curran and Johnston (1996) showed that variations in the tilt and slant of the illuminant can affect the perceived curvature of, isotropic surfaces. More specifically, they

found that as the tilt of the light source increases (from 0° -top- to 180°-bottom- in steps of 45°) the resulting shape perception would be flatter; on the other hand, perceived surface curvature increased with slant. Observers' curvature perception was veridical only when the light source tilt was set between 0° and 45° and the slant was set to 45°. These results support the idea that perceiving a 3-D object depends on the direction of the light source and on the observers' point of view as well as object geometry.

Christou and Koenderink (1997) demonstrated how shape constancy is weak in shape-from-shading since observers show systematic biases towards the position of the light source with particular emphasis to luminance gradients (the brightest region appeared to be closer to the viewer). Since this bias on the interpretation of shading was reduced when a contour was present, they suggest that shape-from-shading is heavily dependent on the location of the light source only when no outline is available. Similarly, Mamassian and Kersten (1996) have reported that changes in shading relative to variations in illumination direction do not alter observers' local surface orientation perception on croissant-like objects with well defined contours. They concluded that the occluding contour was overriding the shading cues in observers' responses, confirming their importance on the perception of shape-from-shading.

The study of lighting in shape-from-shading is not restricted to influences of light source position or default lighting assumptions: others have considered our

ability to perceive lighting direction from the evidence provided by shading stimuli. For iso-tropic surfaces directional lighting produces anisotropic shading patterns and the direction of the lighting can be inferred from the pattern of shading (Koenderink et al., 2004; Koenderink et al. 2007). We are quite good at judging the azimthual direction (or tilt) of a light source in such stimuli presuming it, correctly, to be orthogonal to the dominant orientation of the anisotropic visual texture produced by the scene.

Much of the above work has concerned variations in the azimuth or tilt of the light source. If one imagines a fronto-parallel surface as a clock face, tilt is the direction of the light source as measured around the perimeter of the clock. The aim of the current study is to analyse how the perception of surfaces is affected by the elevation (alternatively declination or slant) of the light source. Slant being the angle of the light source out of the plane of the clock face. Studies of shape-from-shading have typically used circular bumps or dips; we wanted to see how lighting priors would affect observers' shape judgments in presence of more patterned and anisotropic surfaces. Consequently we used surfaces varying in their level of isotropy to test the possibility that the level of anisotropy affects the application of prior assumptions. We used surfaces with no bounding contours to avoid the possibility that they would override the shading cues in observers' responses and to make sure that observers' responses relied on shading cues only. We also tested stimuli lit from two directions (tilts) lighting from above the observers head and lighting from the side. Finally, since the

declination of the light source affects the contrast of shading patterns resulting from a surface, we normalised the contrast of our stimuli in order to control for the effects of contrast on perceived shape.

3.3 Methods

3.3.1 Participants

Three participants with normal or corrected-to-normal vision took part in the experiment (age range 24-30 years old). We asked observers to adjust the 3D orientation of “gauge figure” probes (see Koenderink, 1992, and section 3.3.3) distributed across nine different locations for each surface.

3.3.2 Apparatus

Images were stored on a PC computer, which also controlled the experiment, and presented, via a VSG2/5 grapics card (CRS Ltd, Rochester, UK), at the centre of 20deg wide monitor CRT (Sony 520GDMF) with a refresh rate of 120Hz. The monitor was gamma corrected using look up tables in the VSG 2/5 generated from a 4 parameter monitor model (Brainard, Pelli and Robson, 2002) based on luminance readings taken with a CRS ColourCal photometer. The border between the images and the monitor limits was set to mean

luminance. The final image size was 24x24 cm on screen and therefore 12x12 deg at the viewing distance of 114cm (images occupied 512x512 pixels).

3.3.3 Stimuli

The stimuli (6 pictures, 3 levels of anisotropy x 2 levels of contrast just one orientation) were based on photographs taken from the BOLD database (www.bold.ac.uk; Jiang, Schofield & Wyatt, 2010). They were images of physical surfaces milled out of plastic blocks and generated from height maps which were in turn derived from Gabor noise textures. The maximum 'height' of the surface out of vertical plane (peak to trough distance) was 3.6 mm. We added an isotropic random noise pattern to each image (see Figure 3.1). We used images of three types of surface differing in their degree of anisotropy: High (orientation bandwidth for surface ripples = 40° full width and half height), Medium (orientation bandwidth=90°) and Low (orientation bandwidth=180°). Surface ripples had a dominant spatial frequency corresponding to 0.3c/deg co-ordinates allowing for all image transformations. The dominant spatial frequency of the shading profiles depended on the relative strength of the linear and quadratic shading components but was either 0.3 or 0.6 c/deg. Each surface was photographed six times under different light sources. The tilt of the light source was 90 degrees anti-clockwise (side lighting from the left). The slant of the light source varied from 0° (frontal) to 75° degrees, with constant tilt. Images were then presented in two orientations: upright (90 degree lighting) and rotated, as if both the surface and the illuminant has been rotated through 90

degrees, to simulate lighting from over head. Stimuli were presented with either their original contrast or with their Root Mean Squared (RMS) contrast normalised to a standard level (the mean contrast of the 30° and 45° declination conditions - see Equation 1). We thus tested 72 different Images (3 surface types, six lighting slants, two orientations / lighting tilts and two contrast settings).

$$S_{rms} = I_d * \left(\frac{X - \overline{X}}{std(X)} \right) + I_{\min} \quad (1)$$

Where I_d is the difference between the *Max* and the *Min* Intensity value of the image X, $\overline{X}$ is the average Intensity of images at 30 deg and 45 deg and $I_{\min}$ is the Min Intensity of the Image.

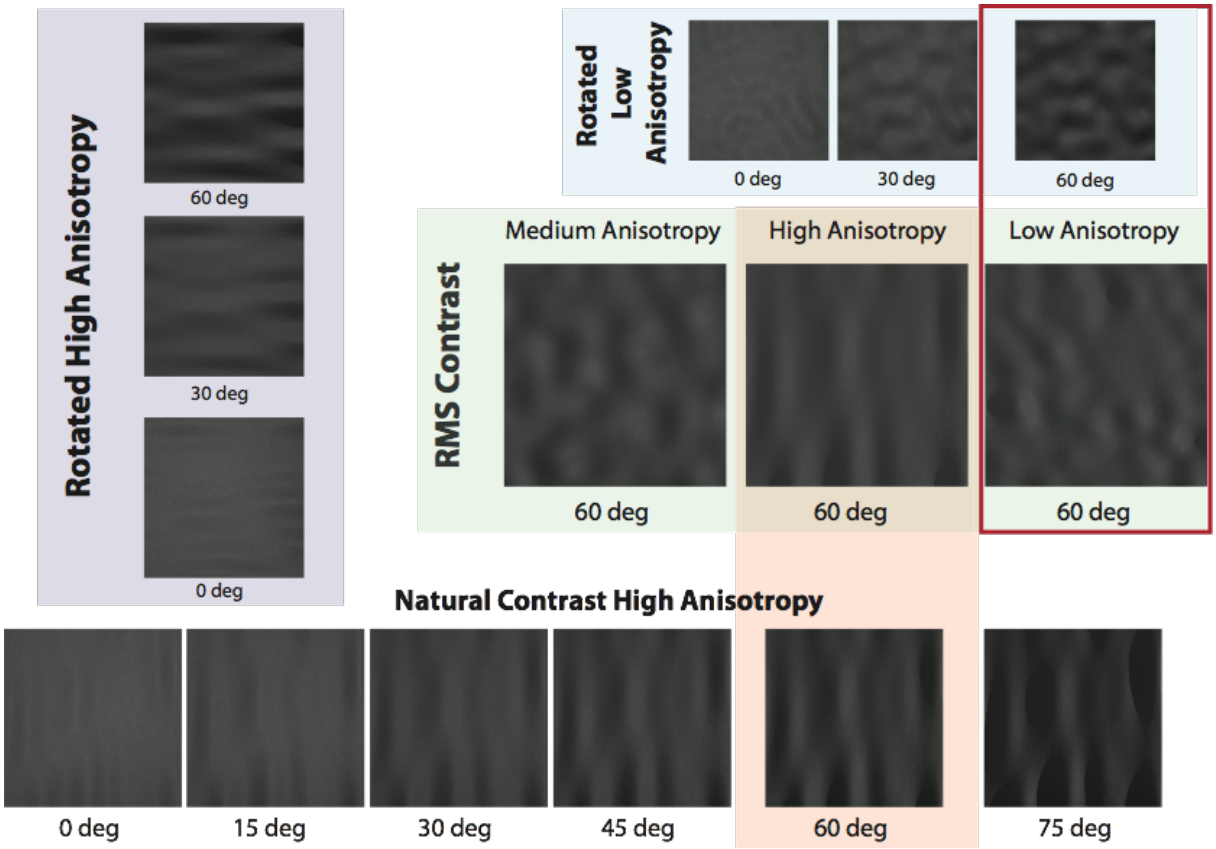

Figure 3.1. Few examples experiment images. images have natural contrast if not specified. On the left side column (Light Blue Box), Rotated High Anisotropy surfaces for different light elevation; central row (Green Box), RMS normalised contrast images for three levels of anisotropy and light elevation at 60 deg. Top row Rotated Low Anisotropy surfaces for different light elevation; bottom row High Anisotropy upright surfaces for all light elevation condition. Red Box area is highlighting similar conditions: High Anisotropy surfaces at 60 deg declination of the light source for Normalised (bottom) and RMS contrast (top); in the top right (Red frame), Low Anisotropy surfaces for RMS contrast (bottom) and Normalised contrast (top).

3.3.4 Procedure

Observers were asked to make judgements about the surface shape by setting the tilt and slant of a probe disk to match the perceived 3D orientation of the surface depicted at each probe location, in each stimulus (Figure 3.3). The probe disk, or "gauge figure", was defined by Koenderink (1992) as the "orthographic projection of a thin circular disk, pierced orthogonally through the centre with a straight line-shaped axle". The disk was 0.4 deg of visual angle; the size of the axle was equal to the disk radius and was protruding from only one side of the disk. The disk was black and the axle white. Adjustments to the probe disk settings were made via mouse movements (horizontally x-axis for tilt - range 0°-360° - , vertically y-axis for slant - range 0°-80°); participants were instructed to adjust the "gauge figure" in a way that looked tangential to the surface depicted in the scene – that is such that the disk appear to be lying on the perceived surface. The tilt and slant of the probe disk and thus our measurements are described as follows. Disk tilt is the direction in which the axel of the disk is pointed around the clock face with zero=up, 90=right, etc. Disk slant is the degree to which the disk is slanted out of the fronto-parallel plane. Zero slant means the disk is pointing back at the observer, 90 degrees of slant mean the disk is oriented orthogonal to the line of sight. There was no limitation on presentation time and observers were asked to be as precise as possible in their judgements. They pressed the "Enter" key on the PC's keyboard to record their answer and go to the next trial.

We tested 9 different probe disk locations for each of the 72 images. The location of the test points varied between surfaces in order to test the most informative locations for each stimulus.

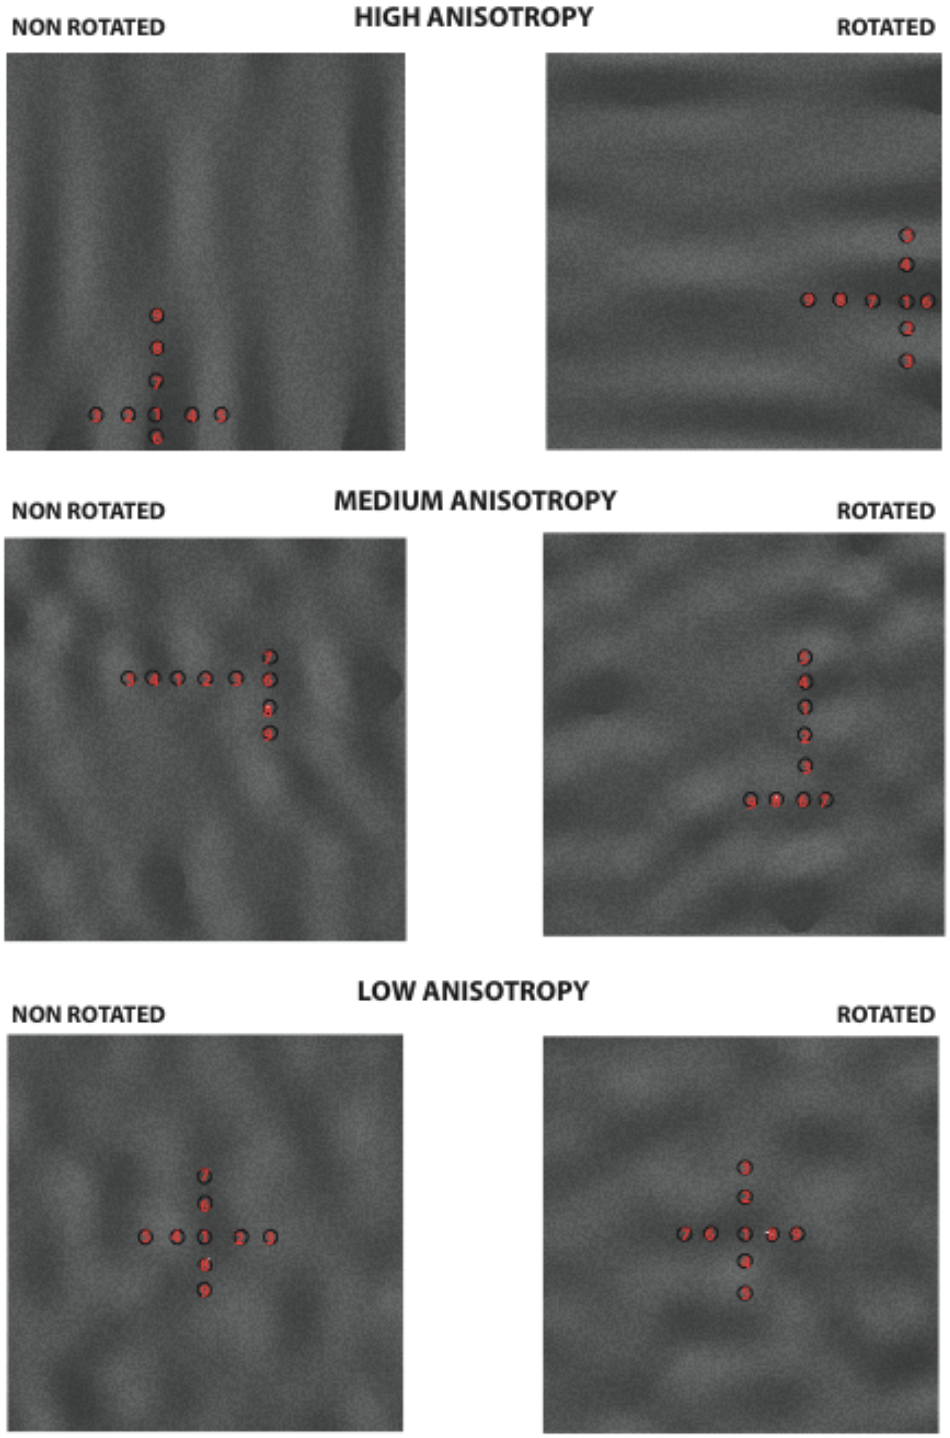

Figure 3.2. Illustrates the 9 different tested locations for each condition.

The same probe points (rotated as appropriate for the rotated images) were used on all repetitions of the same underlying surface. Participants undertook 3 repetitions of each measurement, a total of 1944 trials over 4 sessions (72 stimuli x 9 locations x 3 repetitions). The order of the tested orientations and contrast conditions was counterbalanced while trial presentation was randomised within each session. Counterbalancing was used as a method for controlling order effects in terms of RMS and NON RMS images so that the contrast levels of the tested images would not bias the results of the subsequent sessions. Observers were shown examples of images before the start of the experiment to familiarise themselves with the procedure.

3.3.5 Analysis

The data were analysed by averaging the settings made by each participant in each condition. Tilt estimates were averaged using the circular mean, slant with the arithmetic mean. For each surface we also extracted ground truth maps based on the grey levels of the height maps used to produce each physical surface. Ground truth is defined as the physical conformation of the surface in terms of depth (height). From these height maps we computed the ground truth tilt and slant for each location tested.

The variation in observers' tilt and slant settings seemed to vary systematically with surface anisotropy, image contrast and between the natural and normalised contrast conditions. To evaluate how contrast and anisotropy

affected the variation in both tilt and slant settings, we measured the degree of "spread" across participants (Figure 3.16) of the two surface metrics for each condition.

Equation 2 and 3 were used to measure the spread of tilt and slant respectively.

(2) $$1 - \frac{\sqrt{(\sum_n^1 \cos\theta)^2 \quad + \quad (\sum_n^1 \sin\theta)^2}}{n}$$

(3) $$\frac{\frac{std(\sum_n^1 \alpha)}{n}}{\max(std(\alpha))}$$

Where is the slant angle and α is the tilt angle.

3.4 Results

Figures from 3.4 to 3.16 show polar plots (12 plots: 3 anisotropy levels x 2 contrasts x 2 orientations) representing averaged tilt and slant settings for each participants in selected locations for each condition. Tables from 1 to 12 (supplementary data) show individual slant and tilt settings for each location and condition and their respective variances. The circular-mean perceived tilt is represented by the angle of the polar plot while the mean perceived slant is given by the radial distance from the origin; dot colours shows different declination conditions.

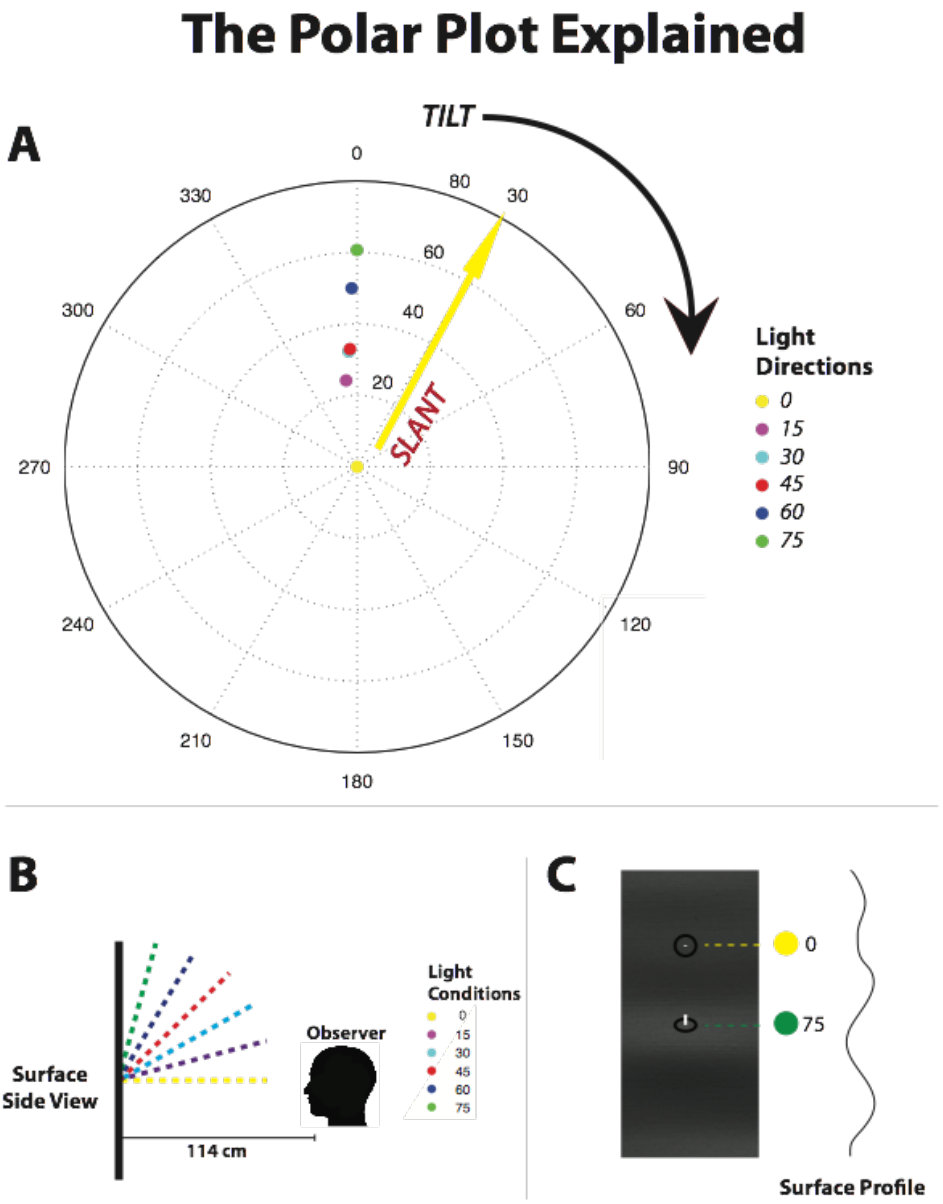

Figure 3.3. A) Polar plot examples: Tilt varies clockwise from 0 to 360 deg while Slant varies from 0 to 80 from the origin to the border of the plot (the more the dot is closer to the border, the higher the slant value). B) Side view illustration of the setup: observer's head is held at a constant distance from the monitor (114 cm using a chin rest); the dotted lines illustrate the direction of the illuminant for each light condition. C) Gauge figure setting explanation: on flat locations the gauge figure will result in setting similar to 0 deg light condition of plot in figure 3.3A (yellow dot: Tilt 0, Slant 0) while for an upwards inclination slant will increase as per 75deg condition of plot in figure 3.3A (green dot: Tilt 0, Slant 60).

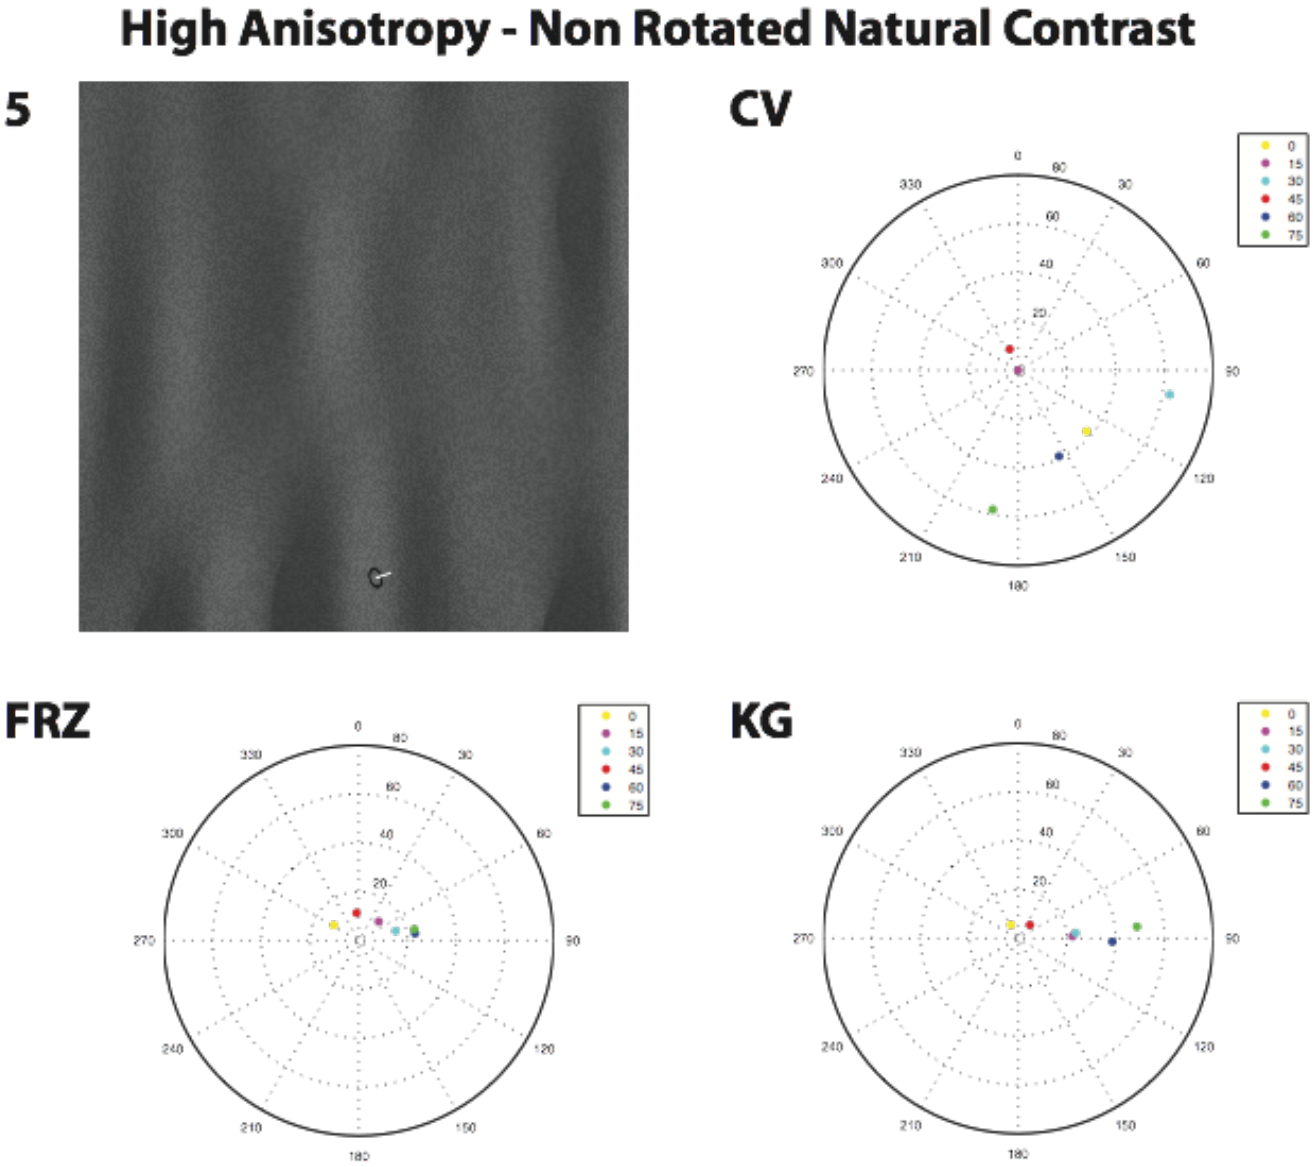

Figure 3.4. Individual data from 3 participants for location 5 (top left) Non Rotated High Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

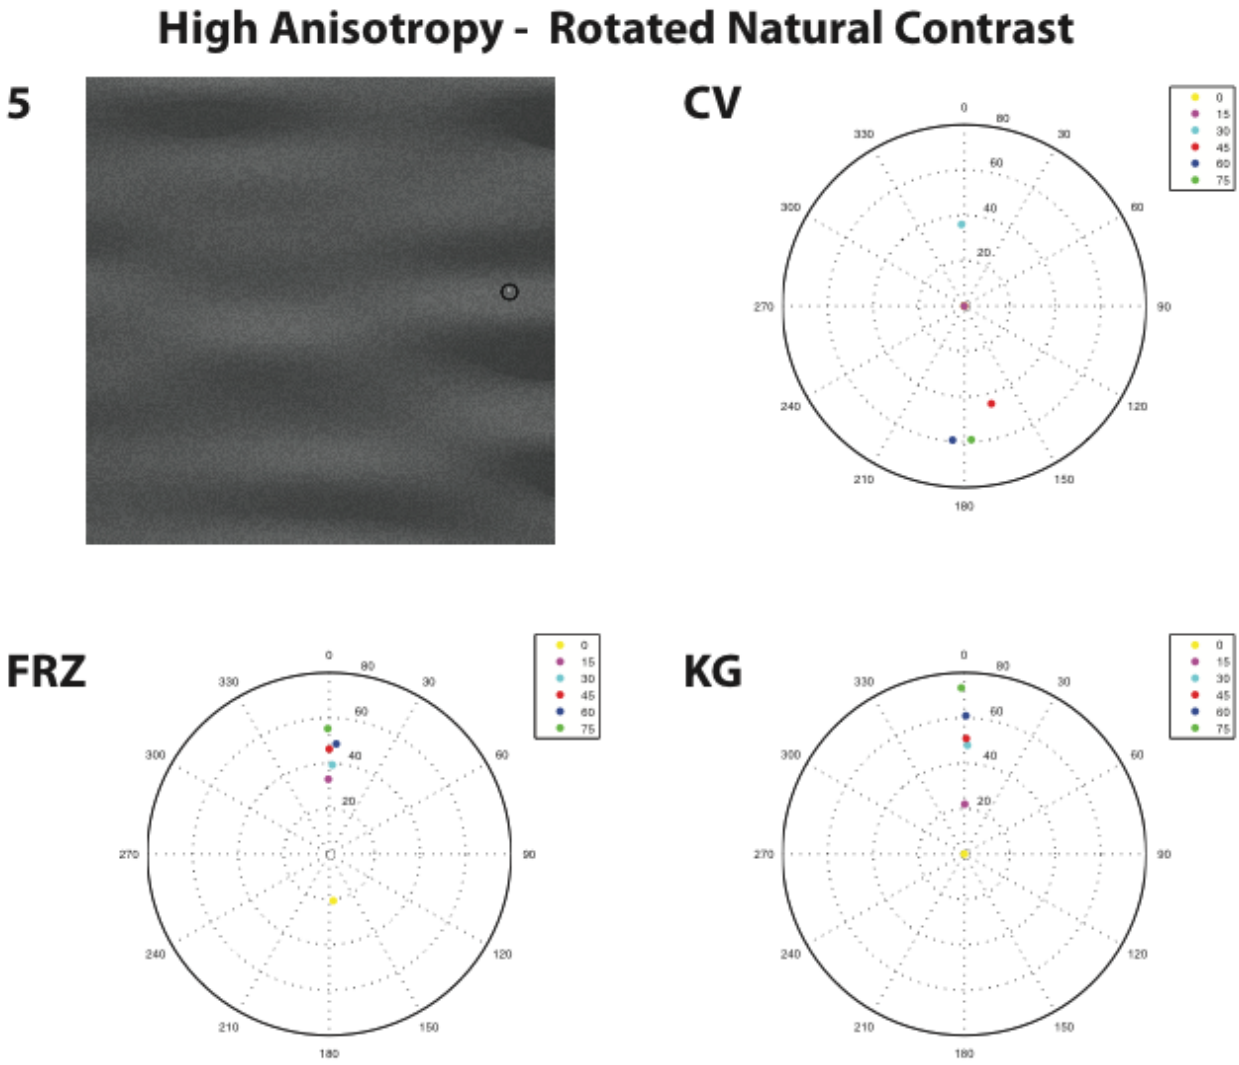

Figure 3.5. Individual data from 3 participants for location 5 (top left) Rotated High Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

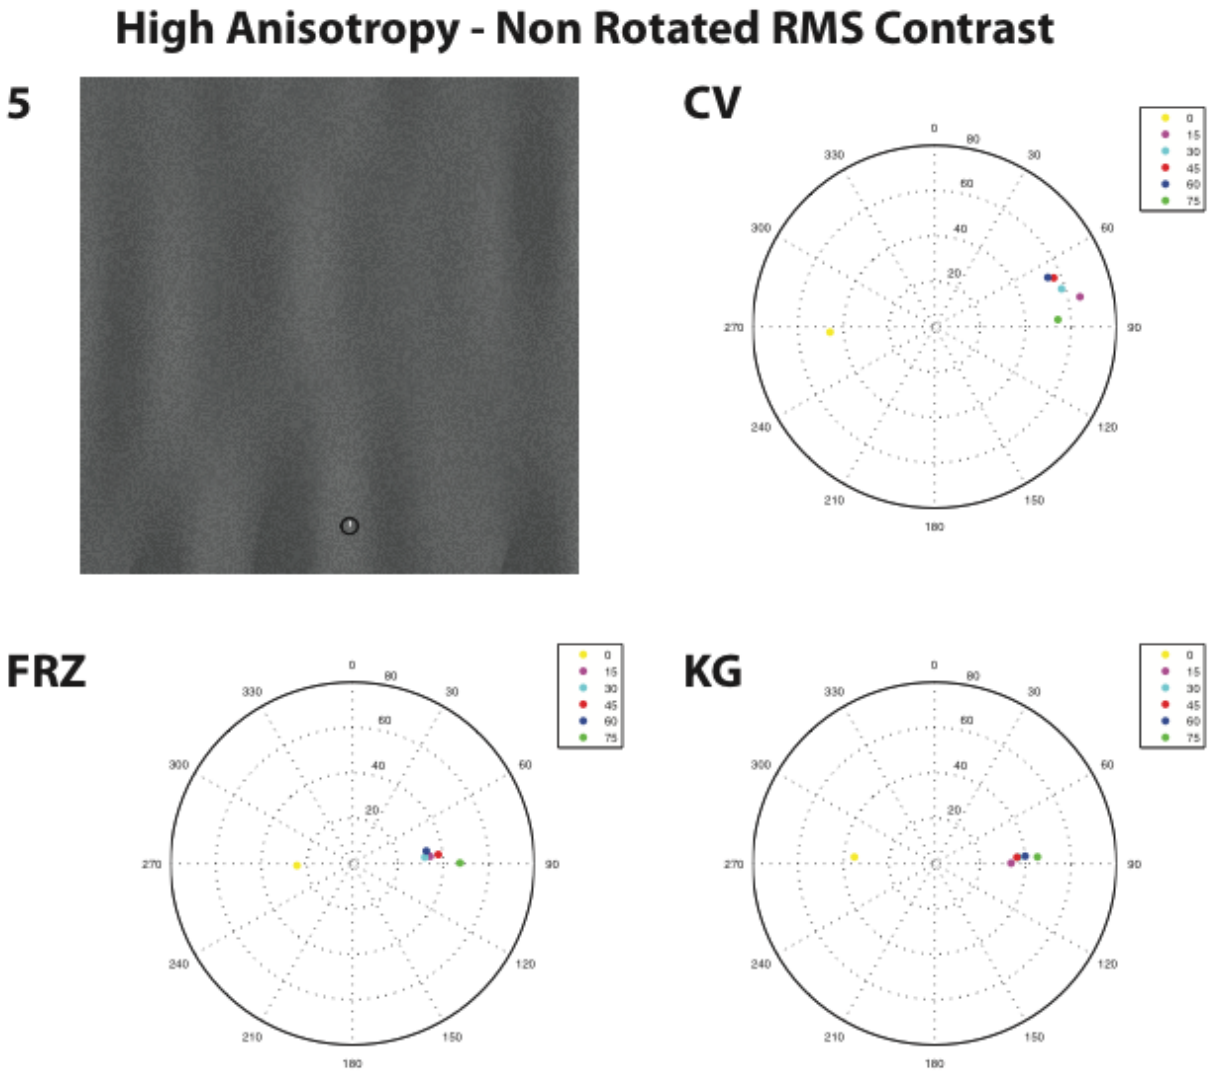

Figure 3.6. Individual data from 3 participants for location 5 (top left) Non Rotated High Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

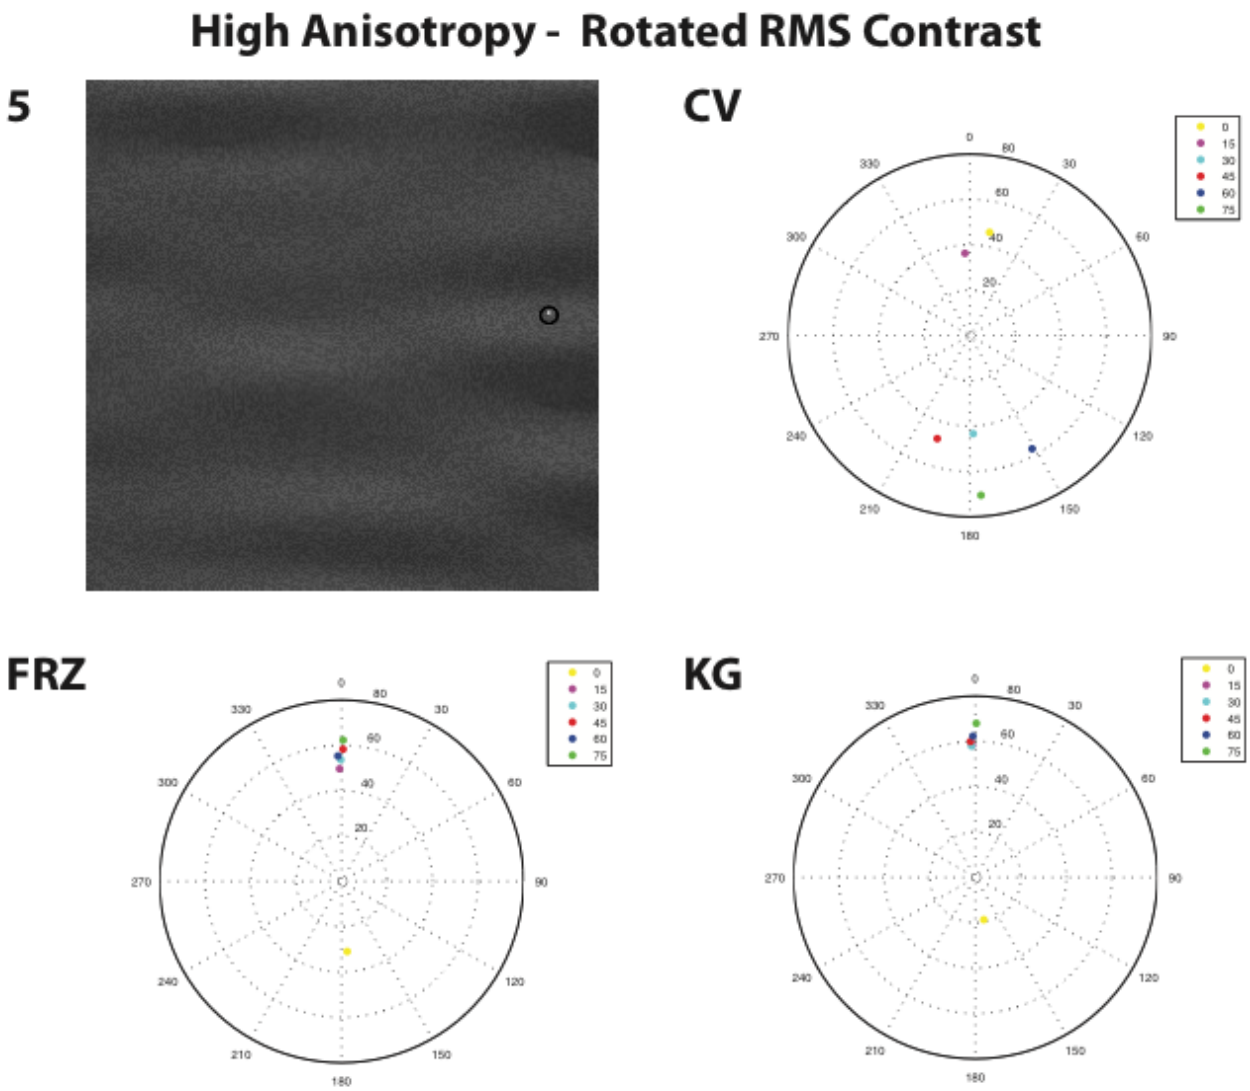

Figure 3.7. Individual data from 3 participants for location 5 (top left) Rotated High Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

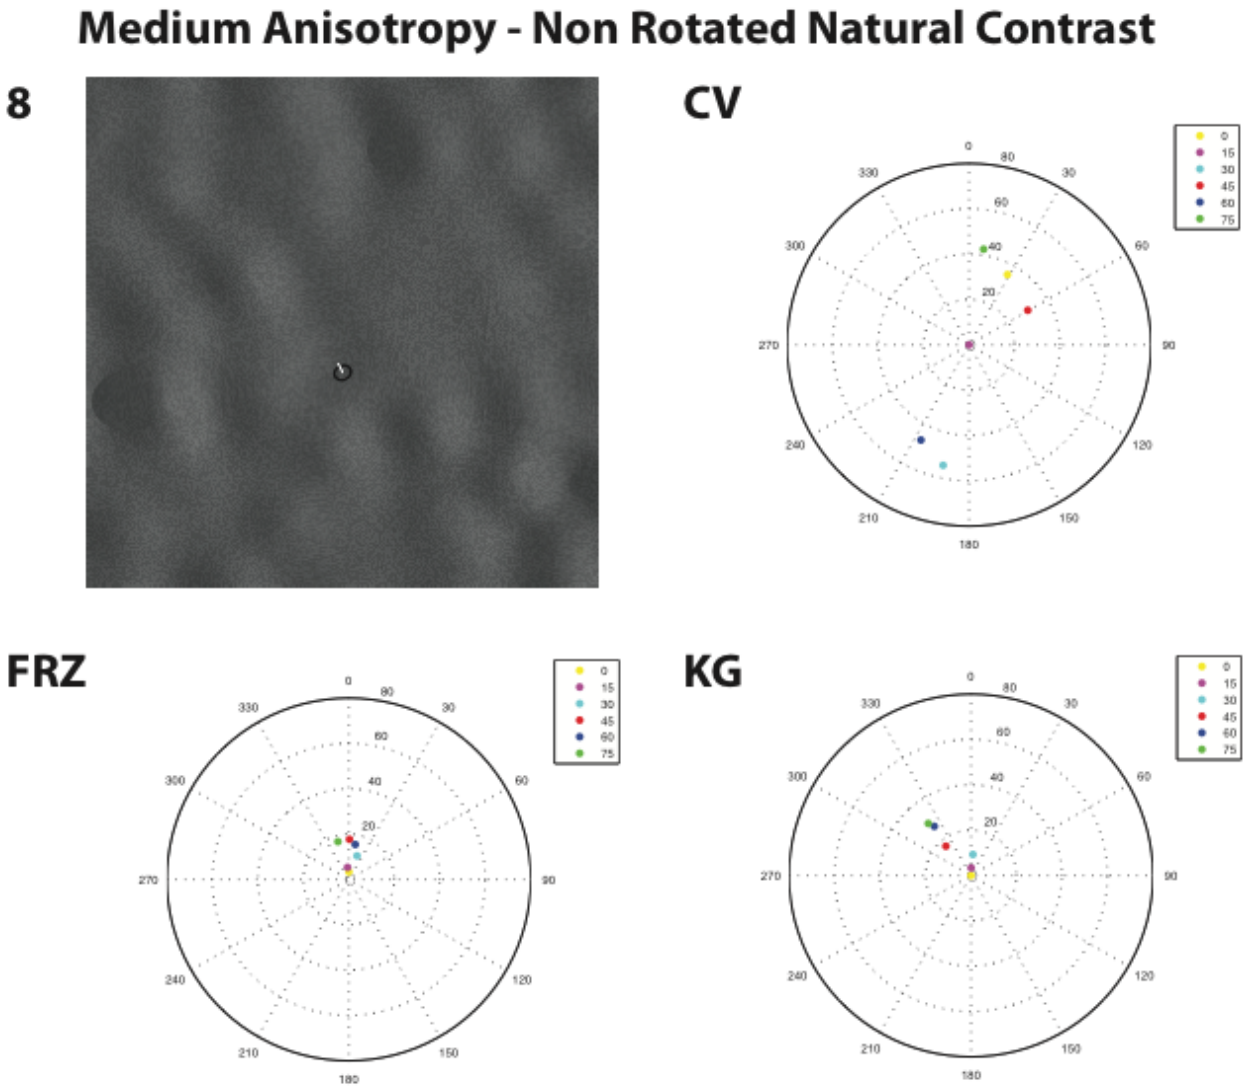

Figure 3.8. Individual data from 3 participants for location 5 (top left) Non Rotated Medium Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

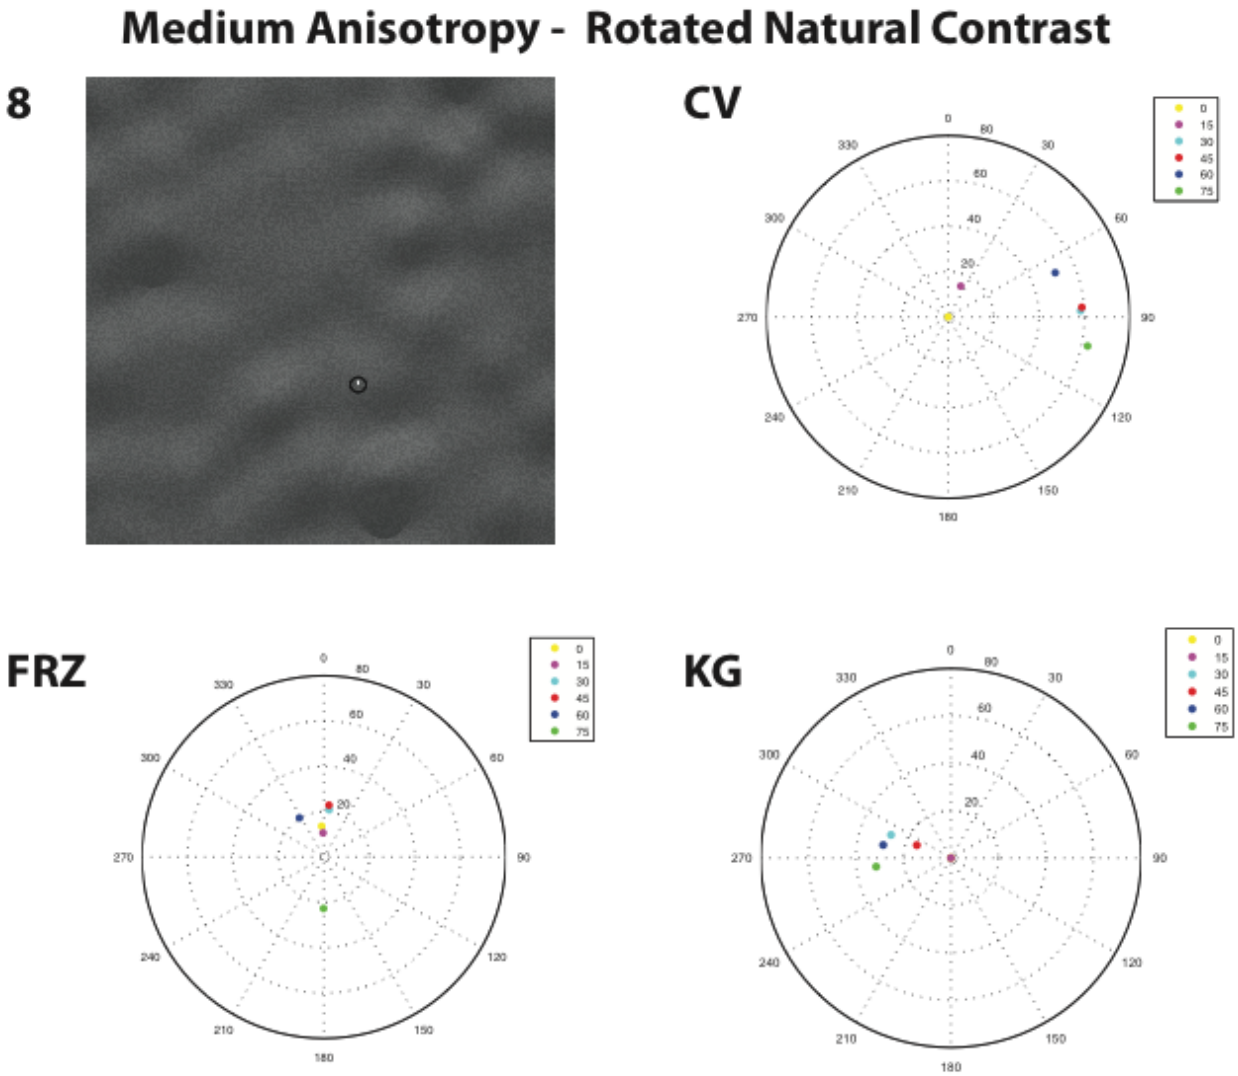

Figure 3.9. Individual data from 3 participants for location 5 (top left) Rotated Medium Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

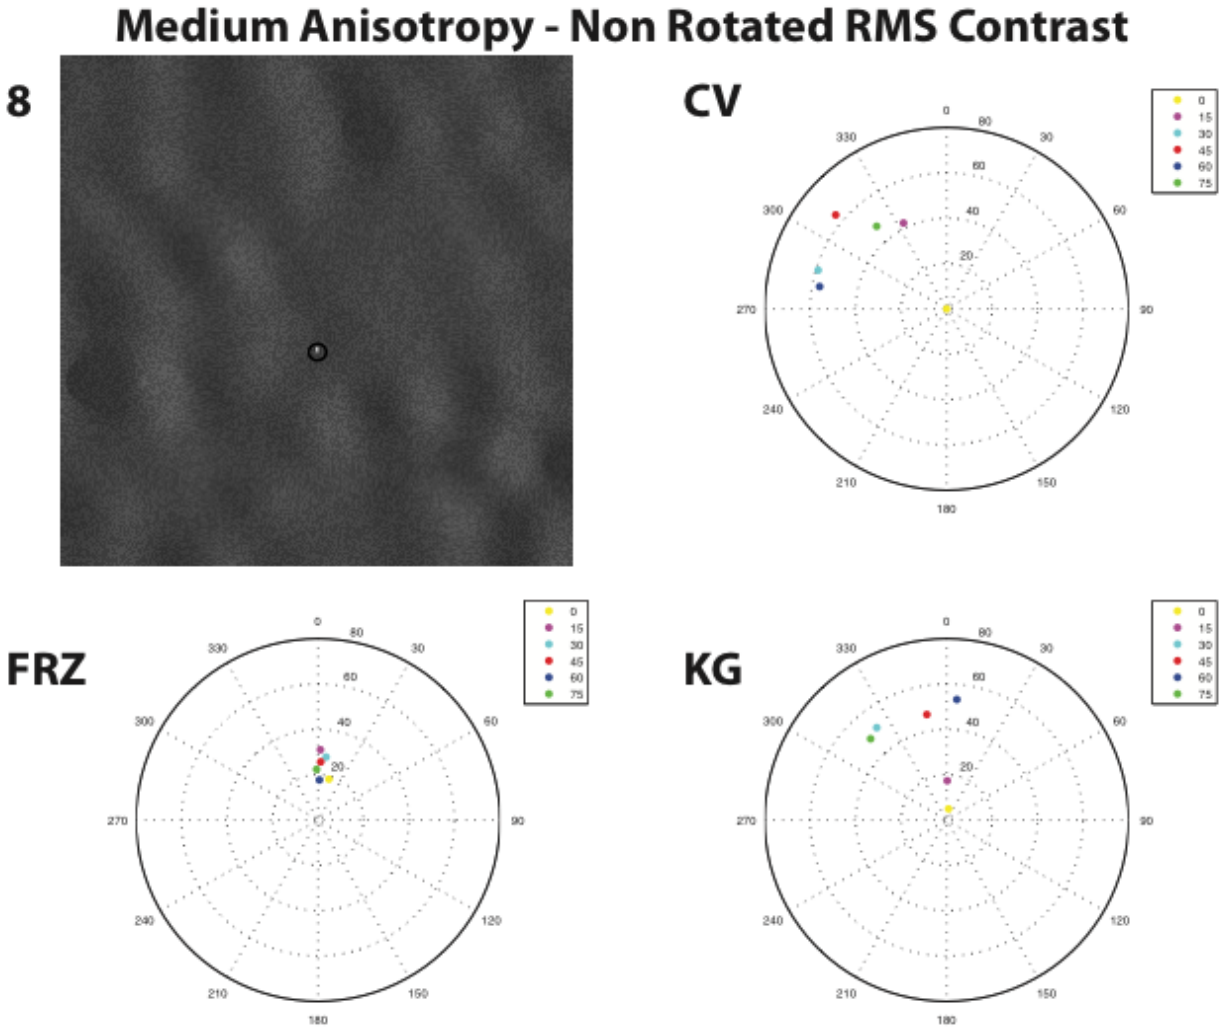

Figure 3.10. Individual data from 3 participants for location 5 (top left) Non Rotated Medium Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

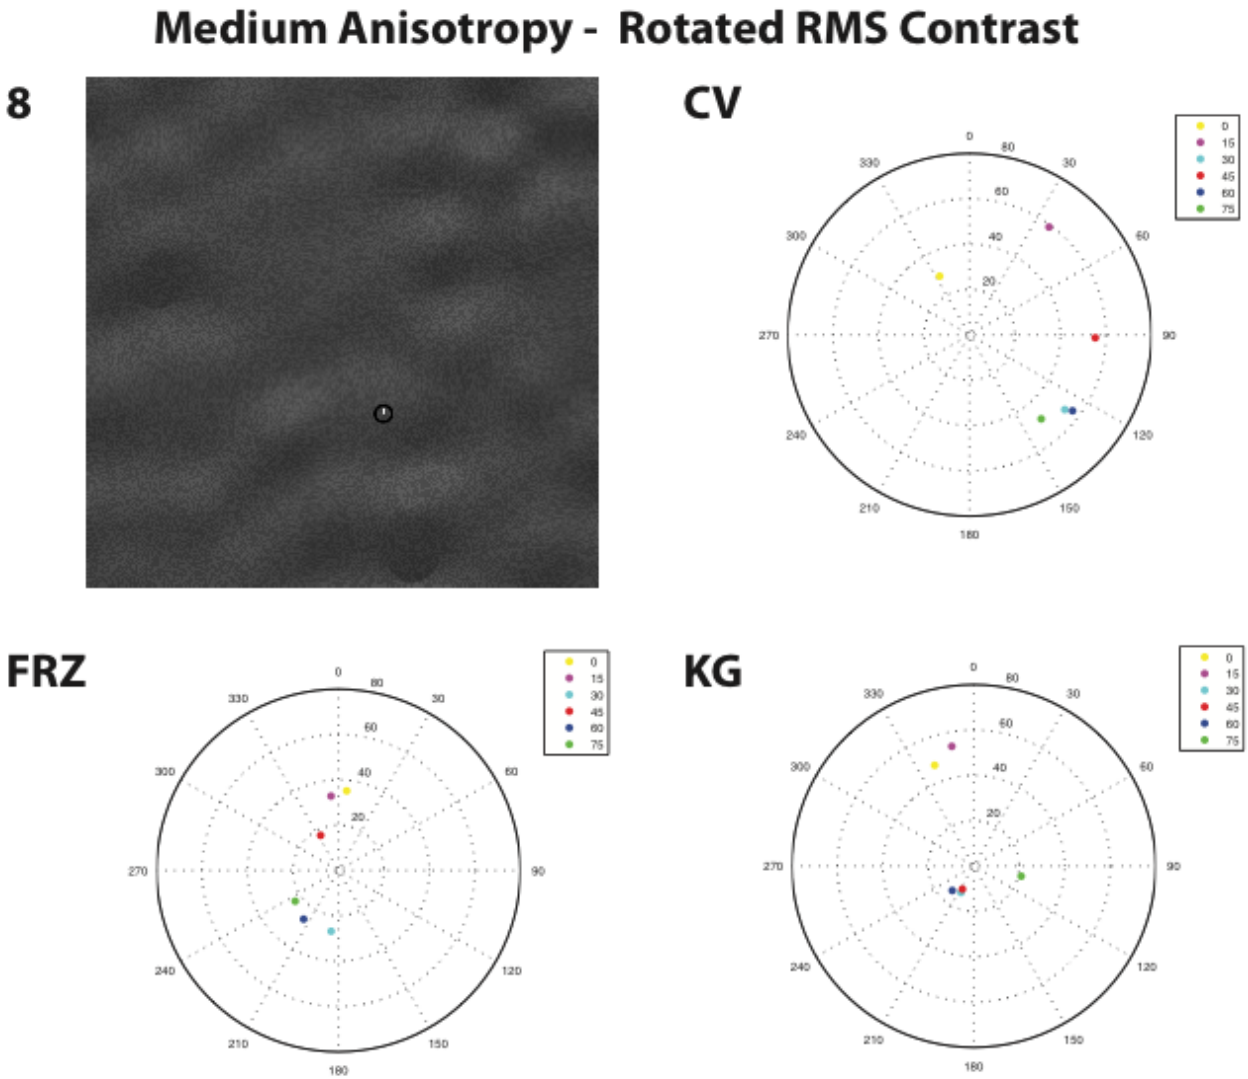

Figure 3.11. Individual data from 3 participants for location 5 (top left) Rotated Medium Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

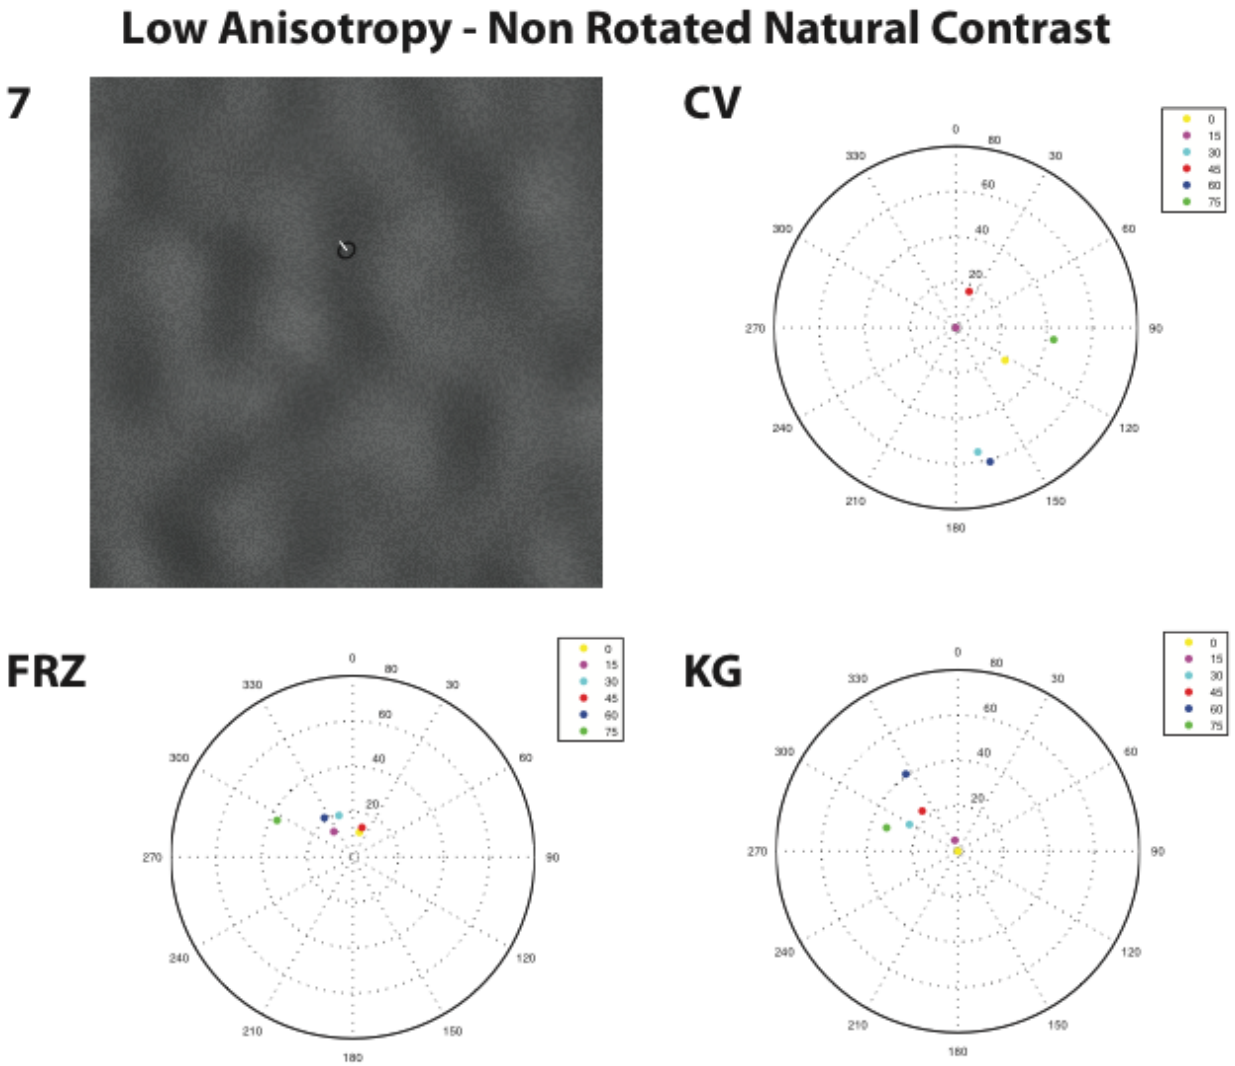

Figure 3.12. Individual data from 3 participants for location 5 (top left) Non Rotated Low Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

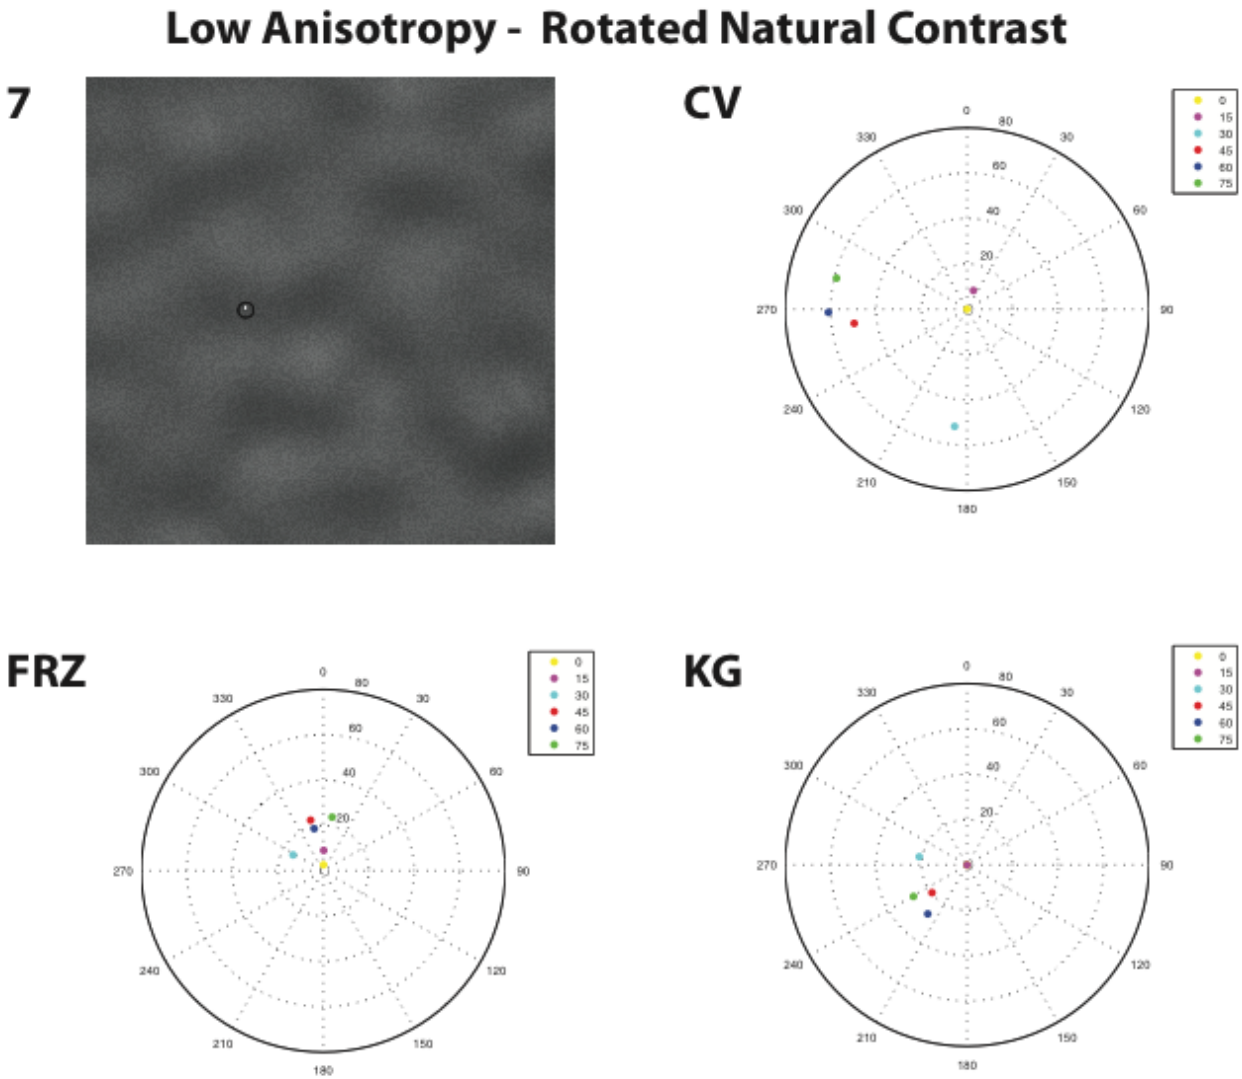

Figure 3.13. Individual data from 3 participants for location 5 (top left) Rotated Low Anisotropic surface with Natural contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

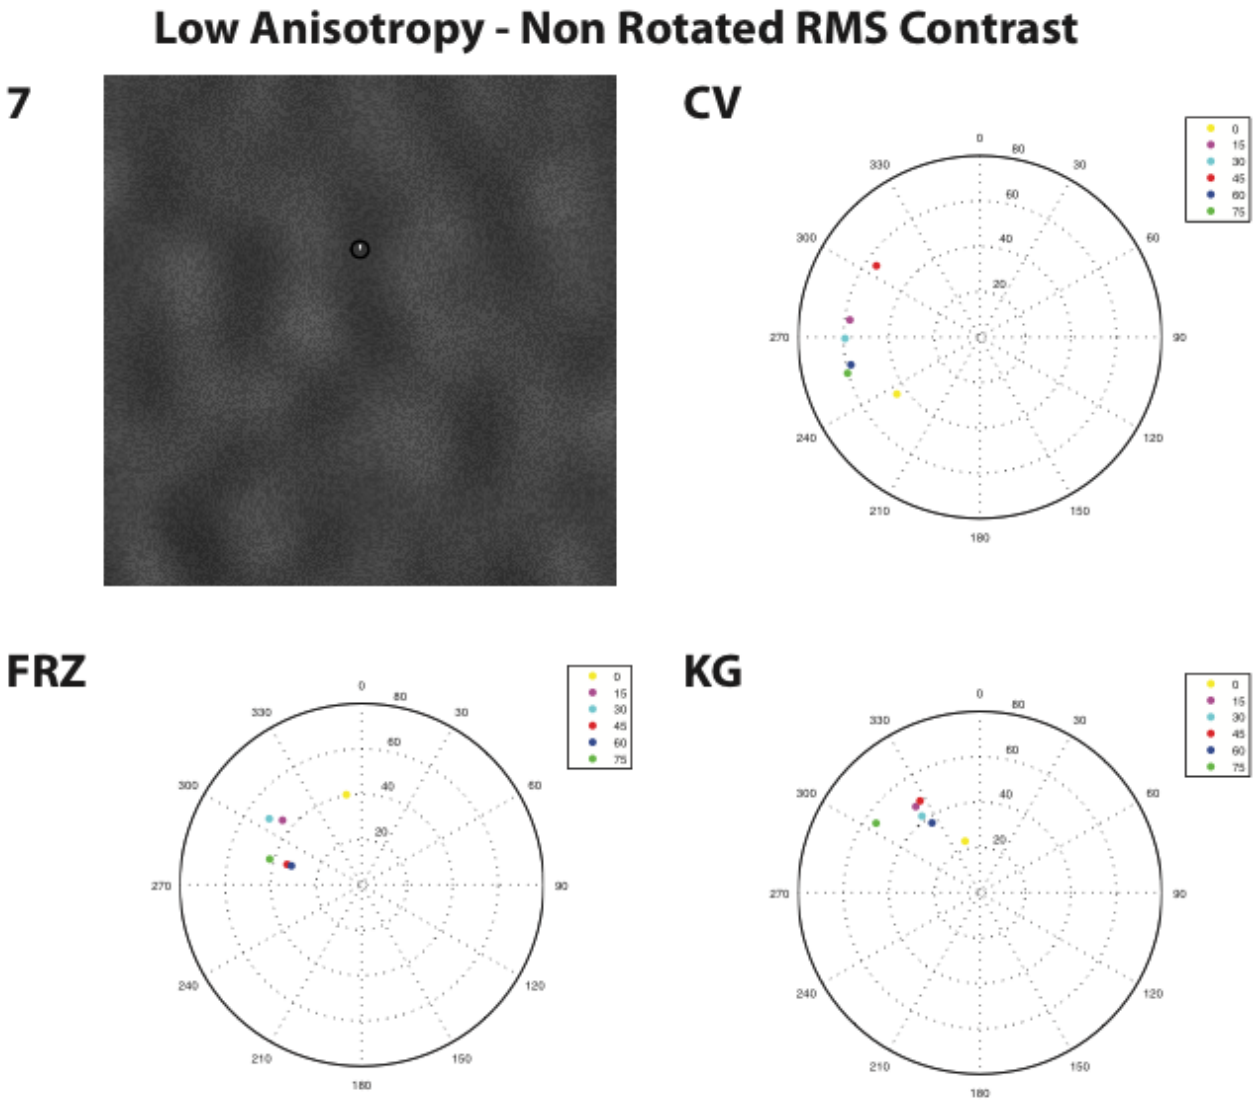

Figure 3.14. Individual data from 3 participants for location 5 (top left) Non Rotated Low Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

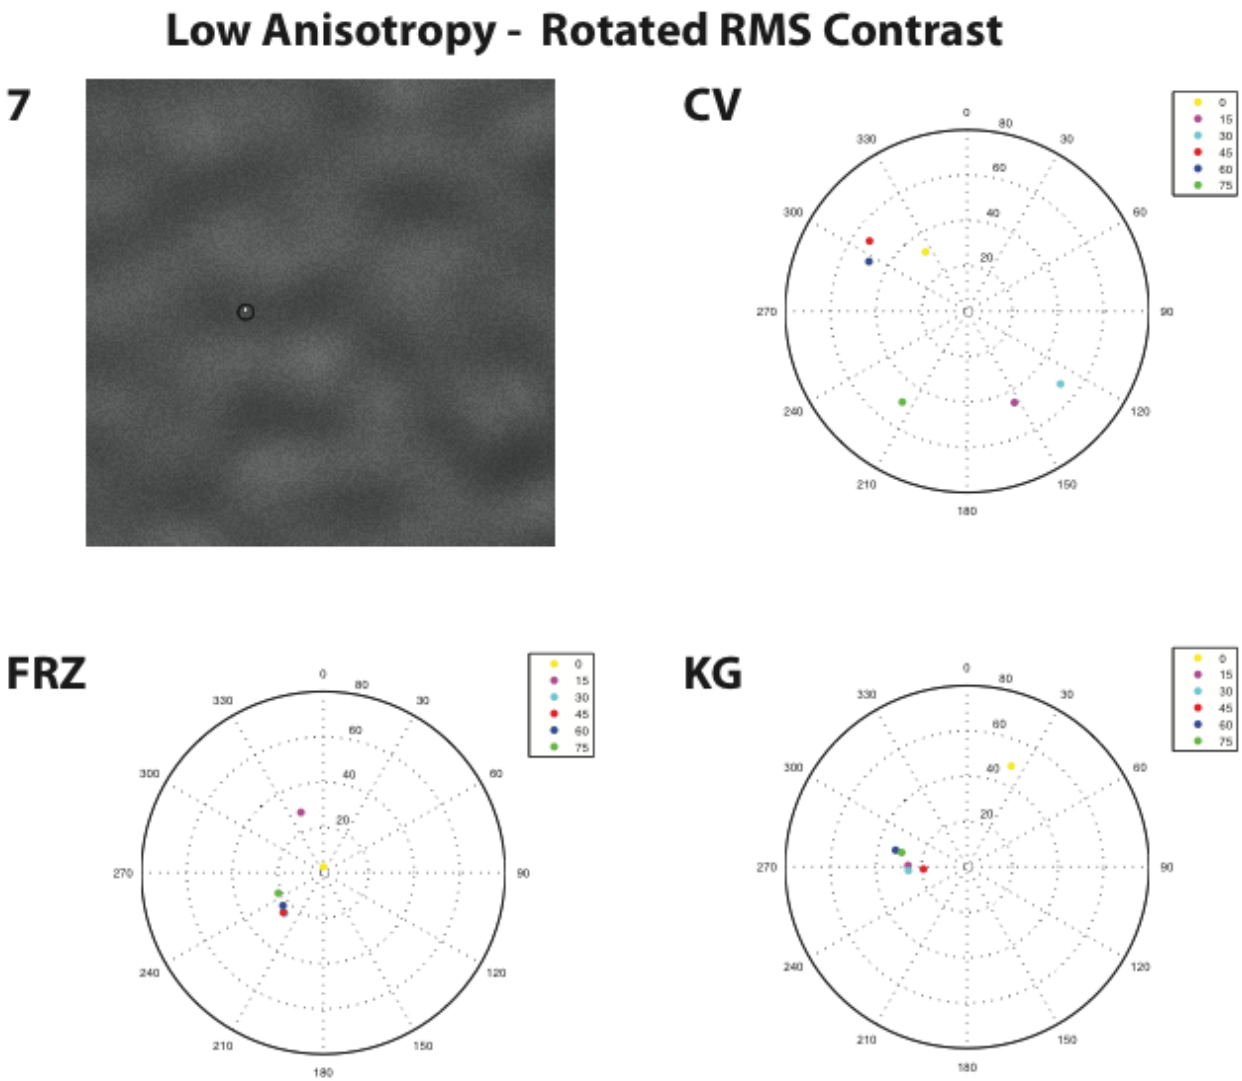

Figure 3.15. Individual data from 3 participants for location 5 (top left) Rotated Low Anisotropic surface with RMS contrast. The gauge figure is set in its starting position with randomised tilt and slant settings.

Surface anisotropy affects the degree to which observers rely on the light-from-above prior. The more isotropic the surface the more observers seem to rely on the lighting-from-above prior perceiving light regions of the image to be tilted

towards a notional light source placed above their own head. Thus, for the most isotropic surface (low anisotropy), tilt depends on local luminance and is largely unaffected by stimulus rotation even though this implies a shift in both the ground truth surface orientation and lighting direction; tilt deviated considerably from ground truth for this surface. For highly anisotropic surfaces tilt settings were dominated by the orientation of the visual texture formed by the shading pattern and thus shifted through 90 degrees when the image was rotated by 90 degrees; tilt settings remained close to the ground truth tilt despite the rotation. This trend is also consistent with a shift in the direction of the perceived light source.

To quantify how people are affected by the influence of lighting prior in respect to the surface isotropy, we measured the absolute difference of tilt settings from 0° and 90/270° which we then express a standardized differences by taking the mean difference across different surface positions and illumination directions divided by the standard deviation of these values. Because of the bas relief ambiguity we took the minimum of the difference from 90 and 270°. A small deviation from 0°, would mean that observers assume the lighting is coming from above; a small deviation from 90 or 270° would represent perceived lighting from the side as might be expected in the vertically oriented anisotropic case. Mean differences as multiples of the standard deviation are reported in table 3.1.

Non Rotated

	Non RMS		RMS	
	Mean difference from 0°	*Mean difference from 90°/270°*	*Mean difference from 0°*	*Mean difference from 90°/270°*
HIGH Anisotropy	1.615122107	1.360047045	2.916027684	1.098429161
MED Anisotropy	1.422105174	1.776088524	1.277095406	1.933304301
LOW Anisotropy	1.166745644	1.96686694	1.544587525	1.572214299

Rotated

	Non RMS		RMS	
	Mean difference from 0°	*Mean difference from 90°/270°*	*Mean difference from 0°*	*Mean difference from 90°/270°*
HIGH Anisotropy	0.926391202	3.17755028	1.016042678	3.456930739
MED Anisotropy	0.945759733	2.225571517	1.148944916	2.221339639
LOW Anisotropy	0.939490085	2.368415687	1.170250647	2.172102559

Table 3.1. Mean differences as multiples of the standard deviation 0° and 90°. Data are averaged across three observers.

Generally the deviation from 0° is smaller than that from 90/270° suggesting the operation of a lighting from above prior. The exception to this is the non-rotated Highly anisotropic surface where deviations from 90/270° are smaller suggesting that oblique lighting is perceived. When the anisotropic surface is rotated such that the luminance pattern runs horizontally, lighting from above is preferred but this result shows the dependence of the perceived light source position on the dominant orientation in anisotropic images. On the other hand, the deviations measured for Medium and Low anisotropic surfaces are more stable across orientation suggesting that lighting-from-above dominates here.

For images presented with natural contrast, slant settings (the degree to which the perceived surface slants away from the observer) seemed to increase with

increased declination in the light source (zero declination equals frontal lighting). We note however that image contrast also increases with light source declination: frontal lighting produces low contrast images while oblique lighting, which grazes the surface undulations, produces high contrast images. When contrast was normalised slant settings varied less systematically with light source declination even though the pattern of shading still varied with declination. This result suggests that human vision sets slant proportional to image contrast; largely ignoring more subtle variations in the luminance profile produced by changes in the declination of the light source.

The frontal lighting (zero degree declination, yellow dots) condition is something of a special case. For frontal lighting the quadratic (frequency doubled) component in the shading pattern dominates over the linear shading component (Pentland, 1989) thus shading patterns produced by frontal lighting differ considerably from the equivalent profiles for oblique lighting. Frontal lighting also produces very low contrast images. Thus slant settings tended to be very low for the frontal lighting conditions when contrast was not normalised and when contrast was naturally higher, and in the normalised contrast case, tilt setting for the frontal lighting condition tended to be at odds with those for more oblique lighting. This result highlights the dependence of human shape-from-shading on local luminance information when cues such as bounding contours are missing. Figure 3.16 shows the estimates spread for tilt and slant setting averaged across participants produced from equations 1 and 2 respectively. Panel A shows variance estimates for each surface type and contrast

adjustment separately but combining across rotations whereas panel B shows the data averaged across surface type.

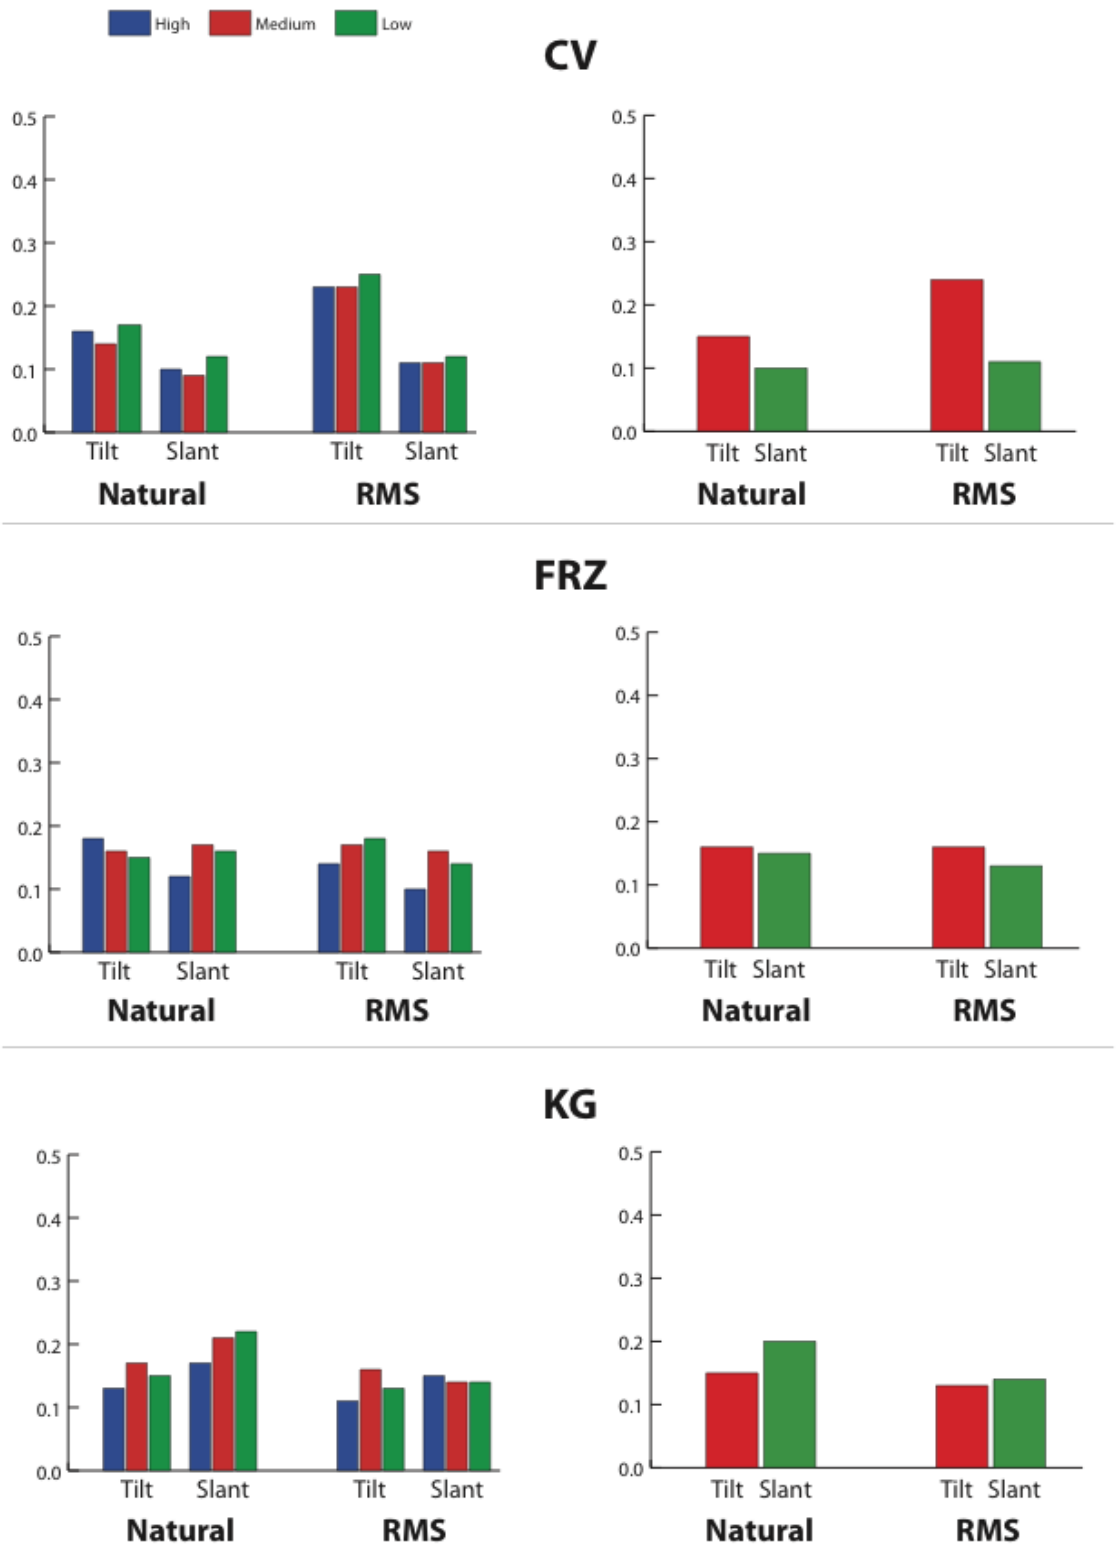

Figure 3.16. On the left: observers variance for tilt and slant adjustments divided by surface anisotropy (High, Medium and Low) and image contrast (Natural Vs RMS). On the right: individual variance tilt and slant adjustments averaged across surfaces.

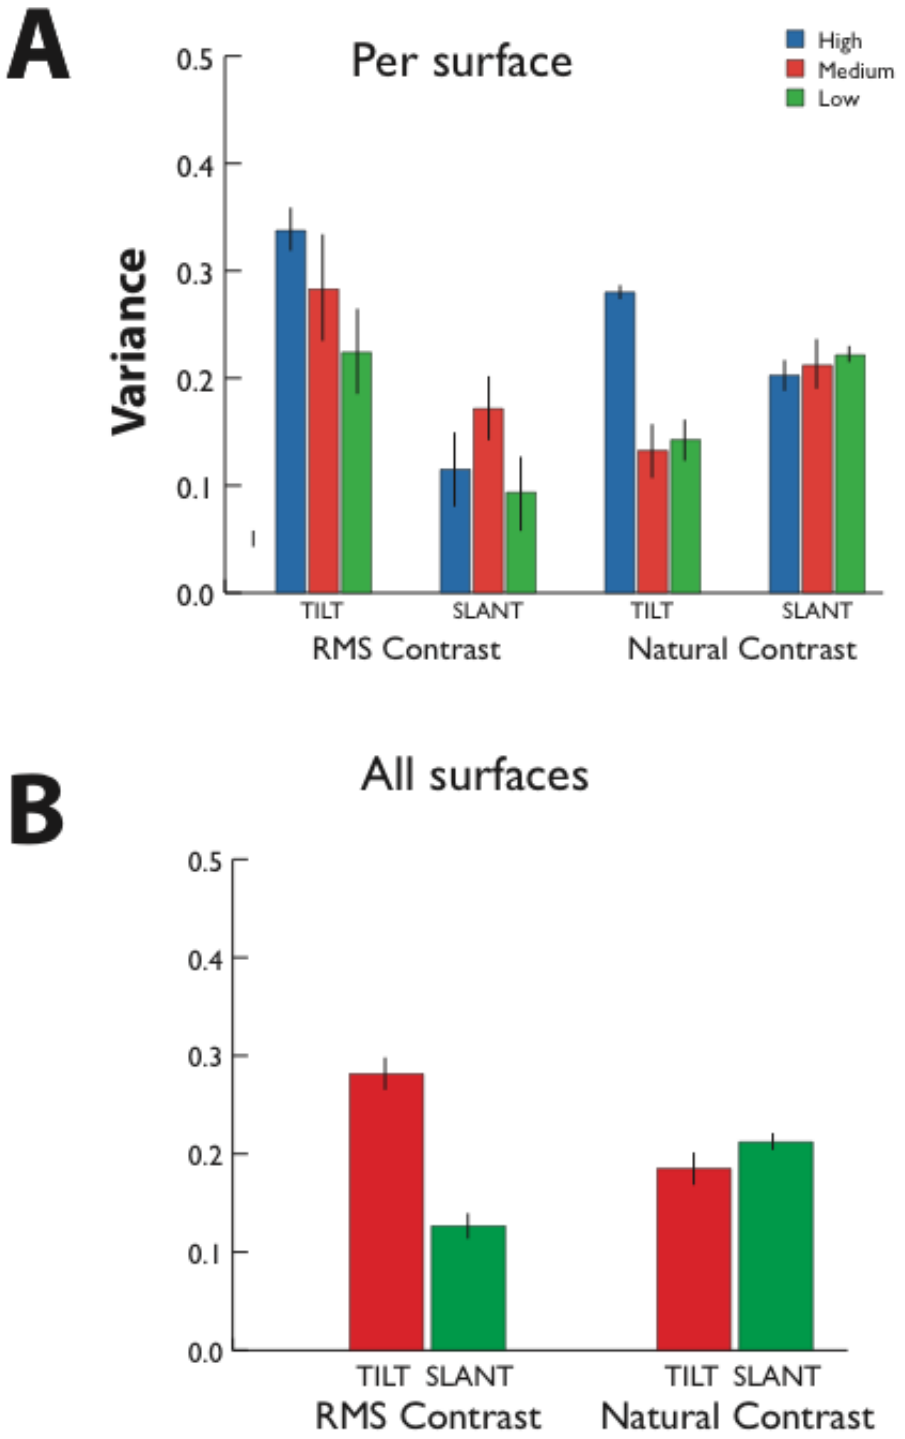

Figure 3.17. Variance estimates plot for each surface type (A) and averaged across surfaces (B), according to equations 1 and 2. Error bars are showing Standard Mean Error based on single observers variance data (A) or single surfaces variance data (B).

The variation in tilt settings reduces with decreasing anisotropy as is consistent with our observation that tilt depends on the dominant image orientation in anisotropic images (and thus varies with stimulus rotation) but is determined by the light-from-above prior in isotropic images (and thus largely fixed close to vertical). The variation in tilt also reduces somewhat when contrast is normalised. Slant setting are not systematically affected by anisotropy but are affected by contrast with less variation in slant being noted when contrast is normalised.

3.5 Modelling

3.5.1 Models

We constructed two simplified models for shape from shading based on the two most common approximations used when studying human vision. One model (Linear Shading model) set perceived slant proportional to image luminance. More specifically slant in this model is proportional to the deviation away from mid grey such that very light areas of the image are strongly slanted towards the perceived light source and very dark areas are strongly slanted away from the light source. Slant values were estimated by first smoothing each image with a Gaussian filter and then taking the absolute difference between mid-grey and each pixel value ($L=|(I*G)-g|$ where G is a Gaussian kernel * indicates convolution and $g=mean(I*G)$. There is no direct way to estimate the tilt direction in this model – it depends on the observers perceived lighting direction

we therefore used two defaults. For the high-anisotropy case we set the lighting direction orthogonal to the dominant orientation in the image. For the medium and low anisotropy cases where there is no dominant orientation we set the light direction as from above. This approach is consistent with the findings of (Erens, Kappers & Koenderink, 1993; Koenderink, Doorn, Pont, 2004; Khang, Koenderink, Kappers, 2007).

The second model is based on the alternative 'dark-is-deep' model of human shape from shading. In this model dark areas of the image are recessed (valleys) and light areas stand proud (ridges). Mid-gray regions have maximum slant as they mark the transitions from light to dark. Slant in this model is accordingly proportional to the magnitude of the luminance gradient local to the sampled point. In this model the light source is assumed to be diffuse (non-directional) but tilt can be estimated directly from the direction of the steepest luminance gradient measured locally. That is the surface if tilted towards light regions and away from dark regions. This model was implemented by filtering each image with two orthogonally oriented Gabor filters to produce two new images $H=I * G_h$ and $V=I * G_v$ where I is the image, G_h and G_v horizontal and vertical Gabor filters respectively and * indicates convolution. These filtered images are then combined to estimate the gradient directions, D=arctan(H,V), and gradient magnitudes, $M=\sqrt{H^2+V^2}$.

3.5.2 Comparison of models.

Models slants compared by correlating the model slant estimates (which are in arbitrary units) with the observers' slant settings. The preferred model in any condition being that with the more positive correlation. To compare tilts we had to first allow for the bas-relief ambiguity. This ambiguity applies only in the linear shading model and is due to the fact that any given point on the surface can be interpreted as a bump lit from one direction or a valley lit from the opposite direction. Thus we had to compare each response to two possible model outcomes and take the minimum. As a result the maximum difference between the model and observed settings in the linear shading model was only 90 degrees. The bas-relief ambiguity does not apply to the dark is deep model so here the maximum difference between model and observed tilt setting was 180 degrees. We allowed for this difference between the models by calculating the proportional difference between model and observed tilts dividing by the appropriate maximum difference. In each condition the model with the lower proportional difference is deemed the better model.

3.5.3 Model results

Tables 3.2 and 3.3 show the correlations between model and observed slants and the proportional difference between model and observed tilts in each case the mean is taken across all locations and physical lighting elevations. Cells highlighted in green show the better model in each case although sometimes the advantage was small. It should be noted that neither model provides an

especially good fit to the data however, for the high anisotropy images the linear shading model is clearly the better model. In this case observers tend to set slant according to luminance and tilt at right angles to the dominant orientation in the image. For the medium and low anisotropy cases the position is less clear. Slant settings seem to be better explained by the dark is deep model but tilt is better explained by the linear shading model with lighting from above (or below). This somewhat paradoxical result may reflect the dominance of the light from above prior in human vision; introducing a bias in tilt settings even when the image suggest a more diffuse light source and hence the dark-is-deep interpretation.

Luminance + Light Direction = Linear Shading

		nRMS			***RMS***	
	Tilt	*Mean angular differences as proportion of max*				
	CV	*FRZ*	*KG*	*CV*	*FRZ*	*KG*
High	0.46	0.42	0.39	0.49	0.18	0.21
Medium	0.49	0.34	0.33	0.55	0.38	0.39
Low	0.44	0.34	0.40	0.56	0.49	0.52

	Slant	*Mean correlations*				
	CV	*FRZ*	*KG*	*CV*	*FRZ*	*KG*
High	0.10	0.56	0.48	0.06	0.53	-0.10
Medium	-0.15	-0.08	-0.15	-0.28	-0.31	-0.21
Low	-0.06	-0.02	-0.15	-0.21	-0.17	0.02

Table 3.2 Linear Shading Model summary table. Green boxes highlight the better results compared with Dark is deep model (Table 3.3)

Gradient model = Dark is Deep

		nRMS			***RMS***	
	Tilt	*Mean angular differences as proportion of max*				
	CV	*FRZ*	*KG*	*CV*	*FRZ*	*KG*
High	0.47	0.46	0.36	0.47	0.51	0.40
Medium	0.48	0.51	0.49	0.36	0.52	0.47
Low	0.53	0.57	0.59	0.53	0.58	0.59
	Slant	*Mean correlations*				
	CV	*FRZ*	*KG*	*CV*	*FRZ*	*KG*
High	0.09	-0.20	0.27	-0.05	-0.37	0.23
Medium	0.32	0.42	0.55	0.44	0.57	0.35
Low	0.37	0.39	0.65	0.19	0.23	0.06

Table 3.3. Dark is Deep Model summary table. Green boxes highlight the better results compared with Linear Shading Model (Table 3.2)

3.6 Discussion

We varied the effective tilt and declination of a directional light source applied to undulating surfaces with different levels of anisotropy and asked observers to judge the perceived tilt and slant of the surfaces at a range of probe points. For anisotropic surfaces perceived surface tilt was most often orthogonal to the dominant orientation of the shading pattern and thus varied with planar rotations of the surface. Thus for these surfaces tilt setting tended to be veridical to the surface ground truth and also tended to vary with changes in the veridical direction of the light source, this being confounded with surface orientation in our stimuli. For isotropic surfaces perceived surface tilt did not follow the ground

truth tilt as images were rotated but was more consistent with a lighting-from-above interpretation with high luminance regions of each image being seen as titled upward regardless of the ground truth surface tilt and lighting direction.

Slant settings varied systematically with the declination of the light source and therefore were not always veridical; however light declination was confounded with image contrast due to the natural interaction between light source declination and surface geometry. When image contrast was normalised the variation **in** slant settings reduced and no longer varied systematically with the declination of the light source.

Frontal lighting is an interesting special case. Here image contrast would be naturally very low but, additionally, quadratic components dominate the shading profile and thus, when contrast is normalised, tilt settings for this condition can sometimes be radically different from those for the same surface under oblique lighting. This result supports the idea that, when the light source is ambiguous, the surfaces perceived via shape-from-shading are highly dependent on the luminance profile which suggest the operation of a rather simple and often not veridical computation whereby slant is proportional to local contrast (even luminance) with tilt dominated by the orientation of image features in conjunction with the lighting-from-above prior or some other estimate of light source direction.

In order to have a better understanding of the data, we constructed two simplified models for shape from shading based on Linear Shading and ‘dark-is-

deep' models. Our results suggest that neither model provides an especially good fit to the data. Nevertheless, we found that for the high anisotropy images the linear shading model is more accurate. Here, observers' slant results tend to rely on luminance and tilt to the dominant orientation in the image. The Linear Shading model accuracy decreases for the medium and low anisotropy cases in which. Here, slant settings are better explained by the dark is deep model while tilt is better explained by the linear shading model with lighting from above (or below). This anomaly on the models reliability can perhaps be related to the influence of lighting priors on the human visual system. More specifically, a bias in tilt settings may be introduced in those cases where the image suggest a diffuse light source and therefore a dark-is-deep interpretation.

In conclusion, we showed how the declination of the light source can influence the perception of shapes. More specifically, we found that shape judgments are affected by variation in the light source declination, confirming that shape constancy is weak in shape-from-shading in the absence of non-shading cues, as reported by Christou and Koenderink (1997). We found that the perception of isotropic surfaces tends to be influenced by the lighting priors with little regard to actual surface orientation. For highly anisotropic surfaces (or at least shading patterns) the affect of lighting priors is reduced and perceived tilt is more dependent on the dominant orientation of the shading pattern which in turn depends on the dominant orientation of the surface and the direction of the illuminating light source.

3.7 SUPPLEMENTARY INFO

Non Rotated Natural Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
High Anisotropy	**1**	*0*	335.58	44.33	0.57	0.07	357.48	4.33	0.50	0.39	52.71	19.00	0.32	0.29
		15	348.67	0.00	0.00	0.00	319.13	12.67	0.17	0.13	267.67	19.00	0.00	0.06
		30	354.67	51.33	0.00	0.02	313.41	17.33	0.33	0.13	284.97	25.00	0.01	0.05
		45	22.75	11.67	0.15	0.00	284.67	16.00	0.01	0.09	271.33	30.67	0.00	0.11
		60	277.19	57.67	0.22	0.07	287.33	29.00	0.00	0.10	286.19	24.67	0.04	0.06
		75	78.34	50.33	0.01	0.58	270.00	27.67	0.00	0.10	273.34	24.00	0.00	0.09
	2	*0*	310.75	53.67	0.08	0.06	9.09	11.00	0.33	0.32	343.00	0.00	0.01	0.00
		15	90.34	58.33	0.00	0.02	60.99	14.00	0.00	0.06	20.25	0.33	0.06	0.58
		30	231.17	46.67	0.06	0.03	76.33	15.33	0.00	0.03	55.32	10.33	0.27	0.58
		45	83.33	48.67	0.00	0.10	63.36	14.00	0.02	0.02	83.33	30.67	0.01	0.02
		60	322.57	62.33	0.13	0.16	81.35	21.33	0.02	0.04	75.18	18.00	0.04	0.20
		75	85.32	73.00	0.05	0.04	9.70	2.00	0.02	0.58	281.99	71.67	0.67	0.06
	3	*0*	78.59	58.33	0.03	0.05	63.08	15.33	0.05	0.05	26.30	7.67	0.22	0.58
		15	357.31	8.00	0.01	0.03	355.26	9.33	0.07	0.20	73.88	21.33	0.03	0.10
		30	82.01	57.33	0.01	0.05	32.95	8.67	0.10	0.15	34.08	17.33	0.14	0.33
		45	1.22	4.67	0.10	0.58	330.03	14.67	0.19	0.09	28.45	14.67	0.05	0.07
		60	99.34	53.67	0.01	0.11	341.23	15.33	0.21	0.18	25.28	16.00	0.36	0.29
		75	284.00	46.67	0.67	0.58	314.31	15.00	0.01	0.06	55.58	18.00	0.13	0.08
	4	*0*	110.59	54.67	0.24	0.09	69.52	11.33	0.06	0.14	351.55	0.00	0.04	0.00
		15	296.53	54.00	0.11	0.02	335.48	13.00	0.51	0.09	96.01	20.00	0.67	0.08
		30	103.04	54.00	0.01	0.05	292.89	18.00	0.66	0.14	305.57	13.00	0.10	0.16
		45	276.63	46.00	0.04	0.04	279.33	22.00	0.00	0.03	288.34	20.00	0.06	0.07
		60	170.21	66.67	0.66	0.03	277.33	15.67	0.00	0.05	298.28	17.67	0.08	0.04
		75	82.33	71.33	0.00	0.07	274.00	23.00	0.00	0.05	277.32	23.00	0.01	0.08
	5	*0*	131.67	37.67	0.00	0.06	301.68	12.00	0.05	0.19	331.95	6.33	0.21	0.58
		15	346.64	0.00	0.01	0.04	46.65	11.33	0.60	0.09	87.67	22.33	0.01	0.21
		30	99.06	63.00	0.03	0.15	75.34	15.67	0.01	0.02	84.67	23.67	0.00	0.19
		45	338.34	9.33	0.01	0.00	355.50	11.33	0.59	0.29	41.27	7.33	0.16	0.58
		60	154.53	39.00	0.59	0.08	82.69	23.33	0.01	0.09	92.00	38.67	0.00	0.02
		75	190.33	58.00	0.00	0.58	78.64	23.33	0.06	0.13	84.33	49.00	0.00	0.10
	6	*0*	76.00	57.33	0.01	0.29	27.26	10.00	0.31	0.33	357.38	10.67	0.56	0.30
		15	222.45	47.00	0.44	0.13	293.67	15.33	0.00	0.11	285.00	14.33	0.01	0.08
		30	109.05	55.33	0.33	0.05	349.50	14.00	0.31	0.13	272.33	26.67	0.00	0.06
		45	239.05	52.00	0.70	0.10	297.71	16.67	0.02	0.06	274.67	23.67	0.01	0.09
		60	117.30	60.00	0.26	0.04	277.67	31.33	0.00	0.05	282.67	25.00	0.00	0.12
		75	192.33	49.33	0.00	0.07	270.67	27.33	0.00	0.04	271.33	26.67	0.00	0.08
	7	*0*	57.96	47.67	0.04	0.04	304.99	5.33	0.13	0.29	6.33	0.00	0.01	0.00
		15	353.67	0.00	0.01	0.03	293.71	13.00	0.02	0.16	309.26	10.00	0.07	0.29
		30	311.59	53.67	0.49	0.01	250.45	16.33	0.57	0.09	289.40	17.33	0.02	0.11
		45	1.33	27.67	0.00	0.00	269.32	15.33	0.01	0.26	284.99	19.67	0.01	0.15
		60	307.75	37.67	0.23	0.01	132.77	8.00	0.65	0.22	321.31	40.33	0.46	0.17
		75	14.00	50.33	0.00	0.29	291.00	29.33	0.00	0.08	280.67	31.00	0.00	0.09
	8	*0*	110.27	64.00	0.29	0.29	8.29	9.00	0.02	0.10	348.00	0.00	0.00	0.00
		15	11.67	49.33	0.00	0.02	107.17	11.33	0.97	0.12	348.66	0.00	0.00	0.00
		30	282.00	59.67	0.67	0.02	169.50	10.67	0.59	0.12	280.66	16.67	0.00	0.05
		45	167.55	56.67	0.65	0.03	111.41	9.67	0.09	0.13	324.14	15.67	0.08	0.29
		60	142.98	64.67	0.30	0.01	141.42	12.00	0.05	0.06	102.67	25.00	0.00	0.04
		75	163.63	57.67	0.65	0.01	237.35	16.67	0.26	0.05	325.93	40.00	0.02	0.05
	9	*0*	267.00	46.00	0.70	0.04	49.68	6.67	0.25	0.31	357.00	0.00	0.00	0.00
		15	348.67	0.00	0.00	0.08	289.93	18.67	0.11	0.10	350.87	14.33	0.54	0.29
		30	78.33	65.33	0.00	0.12	1.43	6.67	0.31	0.13	34.51	0.00	0.22	0.00
		45	311.63	27.00	0.14	0.00	292.78	12.67	0.98	0.14	302.89	3.67	0.23	0.58
		60	73.33	58.33	0.00	0.02	112.24	9.33	0.04	0.05	9.68	13.00	0.31	0.29
		75	287.33	51.33	0.00	0.29	2.27	14.67	0.60	0.03	43.71	24.33	0.62	0.08

Table 1. Individual data from 3 participants in the Non Rotated High Anisotropy condition with Natural Contrast.

Rotated Natural Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
High Anisotropy	**1**	*0*	357.00	0.00	0.00	0.00	187.97	18.67	0.66	0.42	32.82	11.67	0.53	1.25
		15	358.33	30.33	0	0.29	177.35	29.00	0.01	0.15	170.74	12.33	0.02	0.97
		30	359.00	62.33	0.00	0.02	182.36	24.00	0.01	0.14	172.33	20.67	0.00	0.42
		45	357.33	58.67	0.00	0.01	176.67	22.67	0.00	0.05	177.33	20.33	0.00	0.32
		60	358.33	54.67	0.00	0.02	196.67	29.67	0.00	0.04	182.33	19.67	0.00	0.06
		75	357.33	61.33	0.00	0.04	184.67	41.67	0.00	0.06	168.67	17.33	0.00	0.03
	2	*0*	356.00	0.00	0.00	0.00	108.30	15.00	0.54	0.06	354.67	0.00	0.01	0.04
		15	352.67	10.33	0.00	0.58	7.00	35.00	0.00	0.02	350.03	13.00	0.02	0.08
		30	354.00	0.00	0.00	0.00	5.67	33.67	0.00	0.06	10.00	53.67	0.02	0.00
		45	353.33	0.00	0.00	0.00	7.33	30.00	0.00	0.03	2.00	50.33	0.00	0.58
		60	191.58	32.00	0.65	0.30	9.66	44.67	0.00	0.01	5.00	61.33	0.00	0.05
		75	359.33	67.67	0.00	0.03	36.03	3.33	0.69	0.07	6.00	80.00	0.00	0.05
	3	*0*	352.33	0.00	0.00	0.00	2.65	19.67	0.01	0.03	2.00	0.00	0.00	0.03
		15	14.15	13.33	0.53	0.58	4.00	33.33	0.00	0.58	3.66	0.00	0.01	0.00
		30	143.27	34.33	0.55	0.07	4.00	31.33	0.00	0.18	330.89	17.00	0.10	0.00
		45	95.33	48.67	0.00	0.02	4.00	26.33	0.00	0.10	320.96	29.67	0.12	0.00
		60	93.00	58.33	0.00	0.05	35.18	26.33	0.49	0.08	4.89	27.67	0.44	0.37
		75	185.29	61.33	0.67	0.08	156.93	25.00	0.70	0.10	303.00	17.67	0.67	0.13
	4	*0*	350.67	0.00	0.00	0.00	358.67	24.00	0.00	0.10	7.99	0.00	0.01	0.38
		15	184.14	48.67	0.79	0.10	360.00	21.33	0.00	0.17	15.29	35.00	0.62	0.14
		30	50.34	52.67	0.16	0.14	173.66	23.00	0.01	0.11	174.00	25.67	0.67	0.00
		45	307.13	68.00	0.61	0.03	172.34	25.33	0.01	0.09	332.81	22.33	0.72	0.29
		60	110.82	61.33	0.34	0.01	182.33	26.33	0.00	0.12	171.33	18.00	0.00	0.11
		75	105.08	69.00	0.32	0.05	173.00	27.33	0.00	0.03	178.00	20.00	0.00	0.21
	5	*0*	351.01	0.00	0.01	0.00	175.02	20.67	0.67	0.10	3.27	0.00	0.02	0.09
		15	350.69	0.00	0.01	0.00	359.33	33.00	0.00	0.08	0.67	22.00	0.00	0.17
		30	358.02	36.00	0.67	0.29	2.00	39.33	0.00	0.10	1.67	48.00	0.00	0.00
		45	164.52	44.67	0.65	0.28	0.00	46.33	0.00	0.10	1.00	51.00	0.00	0.33
		60	185.00	59.33	0.00	0.06	3.67	48.67	0.00	0.05	0.67	61.00	3.4E-05	0.08
		75	177.00	59.00	0.67	0.04	359.33	55.33	0.00	0.01	359.00	73.33	0.00	0.05
	6	*0*	358.33	0.00	0.00	0.00	223.00	15.33	0.67	0.02	5.33	0.00	0.00	0.03
		15	0.00	40.33	0.00	0.12	152.69	26.00	0.65	0.03	167.68	17.33	0.01	0.03
		30	358.67	63.33	0.00	0.04	184.00	27.33	0.00	0.10	169.33	19.33	0.00	0.00
		45	356.33	56.33	0.00	0.02	193.34	21.67	0.01	0.01	177.67	21.67	0.00	0.09
		60	357.33	58.67	0.00	0.06	184.66	31.67	0.01	0.10	178.99	19.33	0.00	0.12
		75	357.67	61.33	0.00	0.03	188.34	37.00	0.00	0.14	169.68	15.67	0.01	0.09
	7	*0*	356.66	0.00	0.00	0.00	17.09	7.00	0.62	0.09	9.61	13.33	0.65	0.09
		15	0.33	58.33	0.00	0.04	173.02	22.33	0.67	0.05	139.95	6.33	0.63	0.07
		30	0.67	58.33	0.00	0.04	158.33	22.67	0.00	0.58	192.32	18.00	0.01	0.31
		45	359.67	57.33	0.00	0.03	136.83	22.67	0.16	0.11	92.83	28.33	0.54	0.37
		60	357.33	63.33	3.4E-05	0.04	103.38	22.67	0.07	0.03	58.28	32.33	0.07	0.06
		75	0.67	66.33	0.00	0.02	183.69	33.33	0.01	0.07	156.83	17.67	0.05	0.14
	8	*0*	359.08	0.00	0.03	0.00	0.33	9.00	0.00	0.09	359.33	0.00	3.4E-05	0.07
		15	357.33	15.67	0.00	0.58	147.00	22.33	0.33	0.10	318.60	4.33	0.53	0.05
		30	356.37	49.33	0.02	0.02	105.86	20.67	0.16	0.29	23.31	19.00	0.01	0.00
		45	2.00	58.33	0.00	0.05	93.42	27.00	0.04	0.17	54.91	38.67	0.18	0.58
		60	357.34	62.33	0.00	0.05	66.21	25.67	0.04	0.11	66.34	32.00	0.00	0.11
		75	354.67	55.33	0.00	0.04	146.61	27.00	0.06	0.06	57.35	40.00	0.01	0.07
	9	*0*	353.33	0.00	0.00	0.00	358.67	7.67	0.00	0.01	1.67	0.00	0.00	0.15
		15	2.00	0.00	0.00	0.00	47.65	27.00	0.01	0.05	7.34	6.33	0.01	0.06
		30	336.40	30.00	0.07	0.32	40.00	18.33	0.00	0.50	47.30	17.00	0.20	0.00
		45	353.53	51.33	0.13	0.04	57.34	19.67	0.00	0.09	104.10	21.00	0.07	0.58
		60	353.07	63.67	0.24	0.03	229.73	12.67	0.69	0.02	20.59	21.67	0.60	0.32
		75	338.12	52.33	0.08	0.04	52.00	21.00	0.00	0.05	112.94	21.00	0.02	0.07

Table 2. Individual data from 3 participants in the Rotated High Anisotropy condition with Natural Contrast.

Non Rotated RMS Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
High Anisotropy	**1**	*0*	335.58	44.33	0.00	0.00	46.86	25.00	0.66	0.42	292.64	30.33	0.53	1.25
		15	354.67	51.33	0.00	0.29	271.33	39.00	0.01	0.15	279.00	35.33	0.02	0.97
		30	277.19	57.67	0.00	0.02	275.00	37.00	0.01	0.14	278.00	29.67	0.00	0.42
		45	310.75	53.67	0.00	0.01	272.37	34.33	0.00	0.05	271.33	30.33	0.00	0.32
		60	231.17	46.67	0.00	0.02	278.99	40.67	0.00	0.04	283.01	34.00	0.00	0.06
		75	322.57	62.33	0.00	0.04	270.33	38.00	0.00	0.06	285.35	29.67	0.00	0.03
	2	*0*	78.59	58.33	0.00	0.00	274.00	37.00	0.54	0.06	277.66	39.67	0.01	0.04
		15	82.01	57.33	0.00	0.58	87.33	31.00	0.00	0.02	31.97	21.67	0.02	0.08
		30	99.34	53.67	0.00	0.00	80.67	29.00	0.00	0.06	87.00	30.67	0.02	0.00
		45	110.59	54.67	0.00	0.00	89.00	42.67	0.00	0.03	84.11	37.00	0.00	0.58
		60	103.04	54.00	0.65	0.30	94.68	38.33	0.00	0.01	90.67	34.67	0.00	0.05
		75	170.21	66.67	0.00	0.03	18.38	8.67	0.69	0.07	276.00	70.00	0.00	0.05
	3	*0*	131.67	37.67	0.00	0.00	86.34	23.00	0.01	0.03	85.00	32.67	0.00	0.03
		15	99.06	63.00	0.53	0.58	73.64	37.33	0.00	0.58	83.00	35.33	0.01	0.00
		30	154.53	39.00	0.55	0.07	81.66	26.67	0.00	0.18	75.69	32.67	0.10	0.00
		45	76.00	57.33	0.00	0.02	81.67	23.67	0.00	0.10	38.35	35.67	0.12	0.00
		60	109.05	55.33	0.00	0.05	278.41	26.00	0.49	0.08	33.90	37.00	0.44	0.37
		75	117.30	60.00	0.67	0.08	78.04	32.67	0.70	0.10	64.95	37.00	0.67	0.13
	4	*0*	57.96	47.67	0.00	0.00	82.34	28.33	0.00	0.10	86.67	33.67	0.01	0.38
		15	311.59	53.67	0.79	0.10	272.33	23.33	0.00	0.17	293.20	28.00	0.62	0.14
		30	307.75	37.67	0.16	0.14	275.34	34.33	0.01	0.11	270.67	38.33	0.67	0.00
		45	110.27	64.00	0.61	0.03	281.87	32.67	0.01	0.09	274.66	31.00	0.72	0.29
		60	282.00	59.67	0.34	0.01	273.98	25.33	0.00	0.12	276.67	28.67	0.00	0.11
		75	142.98	64.67	0.32	0.05	268.00	29.33	0.00	0.03	276.97	25.00	0.00	0.21
	5	*0*	267.00	46.00	0.01	0.00	268.00	24.33	0.67	0.10	274.67	35.33	0.02	0.09
		15	78.33	65.33	0.01	0.00	84.67	34.33	0.00	0.08	89.67	33.67	0.00	0.17
		30	73.33	58.33	0.67	0.29	85.04	32.00	0.00	0.10	85.67	37.00	0.00	0.00
		45	67.67	56.67	0.65	0.28	84.00	38.00	0.00	0.10	85.67	36.33	0.00	0.33
		60	66.46	54.33	0.00	0.06	80.35	33.00	0.00	0.05	85.33	40.00	0.00	0.08
		75	86.69	54.33	0.67	0.04	89.67	47.33	0.00	0.01	86.33	45.33	0.00	0.05
	6	*0*	201.50	56.00	0.00	0.00	269.67	26.00	0.67	0.02	68.58	44.33	0.00	0.03
		15	193.00	68.33	0.00	0.12	268.67	36.67	0.65	0.03	277.67	33.67	0.01	0.03
		30	253.31	59.33	0.00	0.04	271.33	34.67	0.00	0.10	276.67	34.67	0.00	0.00
		45	247.19	58.67	0.00	0.02	272.67	33.67	0.01	0.01	271.66	34.67	0.00	0.09
		60	312.33	49.33	0.00	0.06	269.67	42.33	0.01	0.10	277.66	34.00	0.00	0.12
		75	9.78	72.33	0.00	0.03	272.00	42.33	0.00	0.14	273.67	27.00	0.01	0.09
	7	*0*	176.83	46.33	0.00	0.00	54.90	23.00	0.62	0.09	70.00	36.00	0.65	0.09
		15	205.92	58.67	0.00	0.04	278.92	24.00	0.67	0.05	292.88	27.00	0.63	0.07
		30	218.31	65.00	0.00	0.04	268.00	40.33	0.00	0.58	294.99	28.00	0.01	0.31
		45	232.47	61.33	0.00	0.03	254.47	26.00	0.16	0.11	286.95	37.00	0.54	0.37
		60	194.33	54.67	0.00	0.04	152.88	25.67	0.07	0.03	297.63	48.00	0.07	0.06
		75	333.49	64.67	0.00	0.02	281.33	36.33	0.01	0.07	290.98	30.67	0.05	0.14
	8	*0*	124.01	15.67	0.03	0.00	90.01	15.67	0.00	0.09	311.45	19.67	0.00	0.07
		15	193.30	61.00	0.00	0.58	192.67	25.33	0.33	0.10	299.99	21.33	0.53	0.05
		30	188.66	63.67	0.02	0.02	151.00	19.00	0.16	0.29	275.82	26.00	0.01	0.00
		45	205.00	60.33	0.00	0.05	181.77	17.33	0.04	0.17	285.14	28.67	0.18	0.58
		60	186.00	58.00	0.00	0.05	154.00	24.33	0.04	0.11	295.83	15.00	0.00	0.11
		75	239.33	66.00	0.00	0.04	284.21	35.00	0.06	0.06	330.01	50.33	0.01	0.07
	9	*0*	28.63	35.33	0.00	0.00	77.13	10.33	0.00	0.01	351.33	14.00	0.00	0.15
		15	192.09	54.33	0.00	0.00	246.74	17.00	0.01	0.05	260.17	15.33	0.01	0.06
		30	206.70	47.00	0.07	0.32	103.71	15.33	0.00	0.50	224.02	25.33	0.20	0.00
		45	188.67	55.00	0.13	0.04	101.42	13.67	0.00	0.09	249.50	27.67	0.07	0.58
		60	192.00	36.00	0.24	0.03	133.20	17.67	0.69	0.02	163.62	22.00	0.60	0.32
		75	200.33	53.00	0.08	0.04	134.14	28.33	0.00	0.05	131.33	23.33	0.02	0.07

Table 3. Individual data from 3 participants in the Non Rotated High Anisotropy condition with RMS Contrast.

Rotated RMS Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
High Anisotropy	**1**	*0*	357.00	42.67	0.00	0.08	1.00	32.67	0.67	0.12	11.88	59.67	0.03	0.05
		15	356.00	48.33	0.00	0.03	172.33	31.67	0.00	0.02	171.00	13.33	0.00	0.13
		30	1.00	58.00	0.00	0.05	183.00	36.67	0.00	0.02	168.34	23.33	0.01	0.10
		45	359.33	60.67	0.00	0.04	177.67	28.00	0.01	0.07	169.99	15.00	0.01	0.18
		60	349.16	48.00	0.04	0.03	192.00	35.67	0.00	0.01	158.00	18.67	0.03	0.04
		75	17.23	65.33	0.63	0.06	185.00	29.33	0.00	0.01	171.00	14.00	0.00	0.00
	2	*0*	358.33	15.00	0.00	0.58	183.33	31.00	0.00	0.10	166.66	18.33	0.01	0.12
		15	360.00	18.00	0.67	0.58	1.00	40.67	0.00	0.03	0.67	60.33	0.00	0.02
		30	13.20	15.33	0.63	0.56	7.67	39.00	0.00	0.06	0.33	61.00	0.00	0.01
		45	197.82	55.33	0.66	0.09	4.33	47.00	0.00	0.06	2.00	55.00	0.00	0.05
		60	170.00	58.33	0.67	0.04	10.67	55.33	0.00	0.02	4.00	62.00	0.00	0.04
		75	358.03	58.67	0.66	0.08	23.36	3.67	0.56	0.37	5.00	79.67	0.00	0.00
	3	*0*	94.20	47.33	0.51	0.06	1.33	33.33	0.00	0.14	356.00	61.00	0.00	0.02
		15	93.11	56.00	0.44	0.05	358.67	42.33	0.00	0.09	353.00	46.33	0.00	0.07
		30	82.08	65.33	0.59	0.06	1.00	43.33	0.00	0.10	358.33	53.67	0.00	0.03
		45	63.99	50.00	0.55	0.05	3.00	40.67	0.00	0.05	334.90	46.67	0.07	0.10
		60	242.25	58.67	0.49	0.05	190.50	21.67	0.06	0.09	313.80	42.00	0.54	0.15
		75	95.39	70.33	0.76	0.05	27.73	36.67	0.58	0.12	18.61	37.33	0.72	0.22
	4	*0*	209.98	65.00	0.19	0.06	356.00	44.67	0.67	0.13	359.67	59.00	0.00	0.01
		15	174.12	65.33	0.88	0.07	15.22	41.33	0.65	0.12	357.67	57.33	0.00	0.04
		30	209.92	58.33	0.66	0.04	179.67	28.33	0.00	0.04	358.33	58.33	0.00	0.04
		45	208.50	58.67	0.39	0.12	176.00	28.67	0.01	0.02	159.10	32.00	0.66	0.26
		60	210.78	73.67	0.46	0.05	178.00	32.00	0.00	0.08	181.67	17.00	0.00	0.09
		75	179.00	69.00	0.67	0.05	172.33	38.67	0.00	0.03	163.31	13.67	0.01	0.06
	5	*0*	10.87	46.33	0.04	0.09	175.33	31.00	0.00	0.04	168.67	19.00	0.00	0.11
		15	356.67	36.33	0.00	0.10	359.00	49.67	0.00	0.03	358.67	58.33	0.00	0.02
		30	177.99	43.33	0.67	0.07	359.67	53.67	0.00	0.05	358.33	58.00	0.00	0.03
		45	197.44	47.67	0.69	0.06	0.67	58.33	0.00	0.02	358.00	60.00	0.00	0.02
		60	151.28	57.00	0.68	0.11	356.33	55.33	0.00	0.04	359.00	62.33	0.00	0.01
		75	176.00	70.67	0.67	0.05	0.67	62.33	0.00	0.03	0.33	68.00	0.00	0.03
	6	*0*	335.42	38.00	0.66	0.06	183.35	31.33	0.01	0.04	358.00	43.00	0.67	0.16
		15	354.69	63.00	0.01	0.06	183.67	30.67	0.01	0.10	156.64	14.33	0.01	0.19
		30	0.67	52.67	0.00	0.04	188.67	32.00	0.00	0.07	170.36	24.00	0.02	0.06
		45	347.70	63.33	0.01	0.06	180.33	35.00	0.00	0.05	165.33	24.67	0.00	0.08
		60	359.00	52.33	0.00	0.04	187.67	38.33	0.00	0.05	182.67	21.00	0.00	0.06
		75	356.00	67.33	0.67	0.08	188.34	31.67	0.01	0.08	170.01	16.00	0.01	0.09
	7	*0*	352.68	24.33	0.01	0.90	183.00	20.33	0.67	0.18	13.09	49.67	0.08	0.11
		15	357.02	47.33	0.01	0.07	181.65	31.00	0.01	0.05	47.71	49.00	0.61	0.18
		30	350.06	56.67	0.02	0.07	200.32	26.33	0.01	0.08	199.09	21.00	0.08	0.04
		45	346.40	55.00	0.07	0.06	177.66	32.00	0.01	0.04	47.57	43.00	0.59	0.24
		60	337.19	53.33	0.18	0.07	170.66	23.00	0.01	0.12	27.33	65.33	0.01	0.08
		75	203.92	60.00	0.66	0.01	183.67	32.67	0.01	0.02	163.34	17.00	0.00	0.10
	8	*0*	0.34	0.00	0.01	0.00	14.97	20.33	0.64	0.15	10.76	16.00	0.06	0.58
		15	339.73	31.33	0.58	0.29	156.75	23.67	0.58	0.07	25.10	39.33	0.03	0.29
		30	328.07	53.67	0.06	0.04	120.25	18.67	0.19	0.05	40.32	47.67	0.01	0.13
		45	336.10	32.33	0.05	0.29	130.67	19.67	0.46	0.09	77.22	39.33	0.17	0.22
		60	32.59	57.33	0.38	0.05	81.87	25.00	0.29	0.11	50.69	45.33	0.02	0.07
		75	157.00	62.33	0.00	0.09	224.43	29.33	0.03	0.03	47.96	42.33	0.02	0.05
	9	*0*	1.67	0.00	0.01	0.00	358.00	1.33	0.00	0.29	14.52	16.00	0.04	0.58
		15	354.37	0.00	0.02	0.00	62.13	21.67	0.13	0.12	21.61	33.00	0.02	0.30
		30	347.29	35.67	0.10	0.29	94.24	21.67	0.11	0.10	21.25	12.00	0.05	0.58
		45	345.00	19.33	0.00	0.58	81.62	18.33	0.13	0.07	36.98	42.67	0.01	0.07
		60	302.82	39.00	0.52	0.31	52.67	21.00	0.02	0.14	27.33	53.00	0.01	0.02
		75	3.00	52.67	0.47	0.03	67.20	10.33	0.32	0.13	42.67	47.33	0.00	0.09

Table 4. Individual data from 3 participants in the Rotated High Anisotropy condition with RMS Contrast.

Non Rotated Natural Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Medium Anisotropy	**1**	*0*	67.67	56.67	0.01	0.05	351.00	2.33	0.00	0.46	358.67	0.00	0.00	0.00
		15	281.33	50.67	0.00	0.04	3.00	12.00	0.11	0.07	354.01	0.00	0.00	0.00
		30	66.46	54.33	0.04	0.00	346.59	14.00	0.03	0.06	8.99	3.67	0.01	0.58
		45	271.33	56.33	0.00	0.07	2.31	19.33	0.02	0.15	0.82	29.67	0.08	0.01
		60	86.69	54.33	0.06	0.03	72.67	11.67	0.24	0.04	350.35	44.33	0.01	0.13
		75	272.01	63.67	0.01	0.03	104.68	16.33	0.01	0.05	19.83	24.33	0.11	0.01
	2	*0*	201.50	58.00	0.65	0.11	349.00	0.00	0.01	0.00	330.97	5.33	0.53	0.58
		15	350.00	0.00	0.00	0.06	9.62	7.00	0.20	0.30	9.00	0.00	0.01	0.00
		30	193.00	68.33	0.00	0.06	52.70	15.33	0.02	0.21	38.83	20.00	0.50	0.15
		45	76.98	45.00	0.01	0.00	23.00	14.00	0.67	0.09	42.40	8.33	0.28	0.58
		60	253.31	59.33	0.01	0.02	53.19	17.67	0.08	0.04	80.88	25.33	0.61	0.14
		75	35.07	45.33	0.34	0.06	35.33	18.67	0.00	0.11	17.11	28.33	0.22	0.09
	3	*0*	247.19	58.67	0.08	0.03	13.67	14.33	0.01	0.16	329.44	6.33	0.07	0.58
		15	81.67	63.33	0.00	0.08	27.95	10.67	0.10	0.34	0.34	0.00	0.01	0.00
		30	312.33	49.33	0.00	0.01	22.85	11.33	0.58	0.12	307.63	9.00	0.32	0.29
		45	85.00	54.67	0.00	0.07	287.12	13.67	0.06	0.09	249.13	17.67	0.13	0.09
		60	9.78	72.33	0.62	0.01	226.69	16.00	0.03	0.07	226.00	14.00	0.00	0.11
		75	75.32	60.00	0.01	0.08	283.45	17.67	0.51	0.09	211.11	16.67	0.66	0.04
	4	*0*	176.83	46.33	0.06	0.05	14.71	8.67	0.02	0.30	13.29	0.33	0.02	0.58
		15	16.06	28.67	0.19	0.03	2.30	16.67	0.42	0.14	32.99	4.33	0.22	0.58
		30	205.92	58.67	0.10	0.06	114.66	18.33	0.01	0.05	295.58	10.67	0.49	0.30
		45	51.42	31.67	0.24	0.32	104.10	7.33	0.18	0.05	75.69	13.33	0.29	0.31
		60	218.31	65.00	0.07	0.08	119.65	12.67	0.23	0.11	96.00	21.67	0.06	0.10
		75	82.66	57.67	0.00	0.29	295.00	36.00	0.00	0.10	286.67	27.00	0.00	0.22
	5	*0*	232.47	61.33	0.09	0.06	4.65	4.33	0.04	0.58	5.00	0.00	0.00	0.00
		15	71.67	56.00	0.00	0.02	9.04	12.67	0.02	0.29	3.00	0.00	0.00	0.00
		30	194.33	54.67	0.00	0.03	351.93	16.00	0.10	0.13	305.62	17.67	0.45	0.12
		45	81.34	59.33	0.01	0.12	338.93	14.00	0.03	0.02	340.29	6.00	0.12	0.58
		60	333.49	64.67	0.17	0.05	336.97	11.33	0.01	0.09	323.99	22.33	0.02	0.15
		75	83.34	60.67	0.01	0.06	299.82	13.33	0.04	0.20	349.31	27.33	0.01	0.09
	6	*0*	124.01	15.67	0.53	0.05	342.37	4.00	0.04	0.58	5.00	0.00	0.00	0.00
		15	354.00	0.00	0.00	0.02	341.64	6.00	0.10	0.29	12.58	7.33	0.26	0.58
		30	193.30	61.00	0.01	0.58	285.26	8.33	0.38	0.09	17.47	1.67	0.09	0.58
		45	349.46	0.00	0.03	0.00	329.65	10.00	0.58	0.07	287.56	2.33	0.79	0.58
		60	188.66	63.67	0.00	0.02	225.60	13.67	0.21	0.10	345.03	8.33	0.66	0.58
		75	67.13	56.00	0.04	0.00	42.33	14.33	0.00	0.13	44.29	6.33	0.36	0.58
	7	*0*	205.00	60.33	0.01	0.02	341.16	17.33	0.66	0.14	9.84	0.00	0.04	0.00
		15	68.32	50.67	0.10	0.01	332.09	4.00	0.20	0.43	349.68	0.00	0.01	0.00
		30	186.00	58.00	0.01	0.10	3.13	11.00	0.22	0.14	29.37	3.67	0.12	0.58
		45	51.20	56.67	0.36	0.08	325.33	16.67	0.00	0.05	352.00	5.00	0.67	0.58
		60	239.33	66.00	0.00	0.06	350.12	14.67	0.54	0.23	34.67	11.00	0.91	0.32
		75	91.00	65.67	0.00	0.04	322.67	13.67	0.01	0.05	220.14	16.67	0.05	0.30
	8	*0*	28.63	35.33	0.66	0.04	0.33	3.33	0.00	0.58	352.02	0.00	0.01	0.00
		15	351.41	0.00	0.03	0.03	354.01	5.33	0.01	0.58	1.32	3.33	0.01	0.58
		30	192.09	54.33	0.06	0.34	19.19	11.00	0.23	0.31	5.33	9.33	0.01	0.58
		45	59.51	30.00	0.27	0.00	1.11	17.67	0.03	0.01	319.52	17.00	0.28	0.31
		60	206.70	47.00	0.15	0.04	10.47	15.67	0.43	0.10	323.38	27.00	0.03	0.04
		75	8.67	42.67	0.03	0.29	344.03	17.33	0.19	0.14	320.67	29.67	0.00	0.03
	9	*0*	188.67	55.00	0.01	0.05	357.00	0.00	0.00	0.00	0.99	0.00	0.01	0.00
		15	21.97	49.67	0.02	0.02	35.83	10.33	0.15	0.31	342.74	4.67	0.07	0.58
		30	192.00	36.00	0.67	0.08	343.59	14.33	0.42	0.10	338.30	7.33	0.15	0.58
		45	22.33	52.67	0.00	0.09	4.97	9.67	0.01	0.20	318.59	7.00	0.06	0.54
		60	200.33	53.00	0.00	0.29	307.41	12.67	0.18	0.29	224.35	21.33	0.01	0.03
		75	54.66	55.00	0.32	0.05	287.12	24.33	0.50	0.13	245.29	18.67	0.04	0.17

Table 5. Individual data from 3 participants in the Non Rotated Medium Anisotropy condition with Natural Contrast.

Rotated Natural Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Medium Anisotropy	**1**	*0*	354.33	0.00	0.00	0.00	358.33	3.00	0.00	0.22	0.33	0.00	0.00	0.13
		15	354.33	0.00	0.00	0.00	356.52	25.00	0.04	0.05	320.82	22.00	0.11	0.18
		30	99.25	43.33	0.63	0.02	354.67	30.33	0.00	0.58	327.60	24.00	0.08	0.00
		45	320.30	52.33	0.65	0.04	7.33	26.33	0.00	0.16	286.40	32.33	0.03	0.31
		60	117.32	55.33	0.01	0.08	346.09	42.33	0.20	0.17	293.00	29.33	0.00	0.29
		75	316.19	55.67	0.65	0.04	18.99	35.67	0.01	0.07	314.16	46.33	0.24	0.01
	2	*0*	359.34	0.00	0.00	0.00	13.99	4.67	0.01	0.06	356.00	0.00	0.00	0.08
		15	359.67	0.00	0.00	0.00	344.73	15.00	0.04	0.12	356.00	12.33	0.01	0.02
		30	18.67	55.00	0.01	0.04	357.58	25.33	0.05	0.58	343.83	23.67	0.24	0.00
		45	38.00	64.67	0.00	0.05	352.08	22.33	0.28	0.17	314.30	35.00	0.55	0.58
		60	31.02	62.00	0.01	0.01	351.74	23.00	0.13	0.21	345.48	48.33	0.06	0.35
		75	53.43	63.00	0.18	0.05	313.65	20.67	0.01	0.08	294.17	32.67	0.35	0.17
	3	*0*	354.66	0.00	0.01	0.00	344.68	0.00	0.01	0.25	359.33	0.00	0.00	0.04
		15	331.39	0.00	0.13	0.00	26.79	14.00	0.10	0.07	356.66	0.00	0.01	0.10
		30	304.60	24.00	0.28	0.33	133.21	26.67	0.18	0.00	62.59	15.67	0.17	0.00
		45	320.09	52.67	0.22	0.04	31.81	19.00	0.05	0.15	77.57	27.00	0.03	0.00
		60	284.76	56.33	0.03	0.07	123.67	22.33	0.04	0.03	85.02	29.00	0.01	0.29
		75	340.98	56.00	0.16	0.01	116.00	22.33	0.00	0.14	82.27	27.00	0.02	0.06
	4	*0*	352.00	0.00	0.01	0.00	357.34	8.00	0.01	0.05	358.33	0.00	0.00	0.04
		15	353.00	24.33	0.00	0.29	81.20	14.67	0.44	0.15	2.30	0.00	0.01	0.08
		30	356.00	32.33	0.00	0.29	24.67	20.67	0.00	0.58	10.51	49.67	0.13	0.00
		45	356.67	33.00	0.00	0.29	100.22	22.33	0.11	0.17	41.98	43.67	0.01	0.00
		60	358.33	50.00	0.00	0.01	77.36	20.33	0.19	0.07	38.97	20.33	0.01	0.04
		75	360.00	60.67	0.00	0.01	193.26	30.33	0.07	0.17	65.62	46.00	0.02	0.05
	5	*0*	352.34	0.00	0.00	0.00	354.68	0.00	0.01	0.10	2.63	0.00	0.02	0.29
		15	23.78	18.00	0.18	0.58	352.34	18.00	0.01	0.07	334.80	12.00	0.14	0.07
		30	91.71	49.67	0.02	0.05	27.78	30.00	0.49	0.00	222.95	23.67	0.47	0.00
		45	74.32	47.33	0.01	0.04	165.34	20.67	0.00	0.10	253.63	19.33	0.02	0.58
		60	95.27	55.00	0.41	0.04	159.81	33.00	0.68	0.16	229.04	21.67	0.08	0.11
		75	43.13	50.00	0.12	0.12	154.84	24.00	0.04	0.01	194.47	16.33	0.39	0.10
	6	*0*	355.67	0.00	0.00	0.00	351.43	13.00	0.03	0.12	351.41	12.33	0.03	0.10
		15	356.67	0.00	0.00	0.00	4.98	9.33	0.01	0.08	4.67	0.33	0.00	0.11
		30	333.10	16.67	0.16	0.58	8.60	27.33	0.03	0.30	0.03	20.67	0.28	0.58
		45	328.59	48.33	0.17	0.03	343.85	18.33	0.16	0.30	358.39	19.00	0.07	0.58
		60	314.88	53.67	0.14	0.07	358.06	19.00	0.10	0.17	310.62	49.67	0.66	0.29
		75	334.68	59.67	0.01	0.01	7.90	32.67	0.18	0.18	336.02	49.00	0.41	0.58
	7	*0*	353.33	0.00	0.01	0.00	356.33	0.00	0.00	0.10	4.67	0.00	0.00	0.07
		15	351.67	1.00	0.00	0.58	25.35	9.00	0.23	0.04	358.34	0.00	0.00	0.07
		30	17.06	40.33	0.03	0.29	357.36	17.67	0.03	0.00	340.00	5.33	0.51	0.00
		45	14.23	36.33	0.08	0.29	281.51	18.67	0.37	0.40	169.30	20.67	0.40	0.00
		60	21.51	52.67	0.05	0.04	244.46	21.00	0.13	0.19	25.07	27.33	0.40	0.58
		75	33.05	60.00	0.02	0.08	131.17	22.33	0.65	0.11	26.70	33.00	0.64	0.04
	8	*0*	355.67	0.00	0.00	0.00	355.97	13.67	0.01	0.12	353.66	0.00	0.01	0.12
		15	22.34	14.67	0.31	0.58	358.00	10.67	0.00	0.11	4.99	0.00	0.01	0.12
		30	87.34	58.33	0.01	0.04	6.33	21.00	0.00	0.33	291.22	27.00	0.10	0.00
		45	85.86	59.00	0.04	0.04	5.64	23.00	0.01	0.29	290.66	15.33	0.32	0.00
		60	67.50	51.00	0.64	0.02	328.18	20.33	0.51	0.10	280.89	29.00	0.45	0.13
		75	101.83	62.67	0.35	0.02	180.35	22.67	0.68	0.05	263.35	31.67	0.41	0.14
	9	*0*	357.33	0.00	0.00	0.00	356.68	8.33	0.01	0.05	358.33	0.00	0.00	0.07
		15	250.66	31.00	0.75	0.29	358.67	13.33	0.00	0.08	350.00	15.67	0.00	0.18
		30	61.84	36.33	0.24	0.29	17.62	23.67	0.21	0.58	354.67	0.00	0.00	0.00
		45	227.10	49.67	0.64	0.03	356.00	13.67	0.67	0.07	63.54	19.67	0.64	0.58
		60	76.98	59.33	0.67	0.04	1.75	17.00	0.12	0.13	259.76	31.33	0.21	0.00
		75	261.81	69.00	0.71	0.06	127.00	21.67	0.73	0.07	282.71	23.33	0.26	0.08

Table 6. Individual data from 3 participants in the Rotated Medium Anisotropy condition with Natural Contrast.

Non Rotated RMS Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Medium Anisotropy	**1**	*0*	334.18	18.33	0.00	0.00	358.00	14.67	0.00	0.22	33.50	7.67	0.00	0.13
		15	112.60	52.00	0.00	0.00	13.78	20.00	0.04	0.05	41.25	23.67	0.11	0.18
		30	91.00	54.33	0.63	0.02	17.11	21.33	0.00	0.58	13.63	33.33	0.08	0.00
		45	101.35	58.67	0.65	0.04	346.90	48.33	0.00	0.16	349.77	38.33	0.03	0.31
		60	100.36	65.00	0.01	0.08	21.01	30.67	0.20	0.17	35.82	42.00	0.00	0.29
		75	106.31	71.00	0.65	0.04	96.67	24.00	0.01	0.07	358.69	52.67	0.24	0.01
	2	*0*	298.68	30.00	0.00	0.00	10.14	17.33	0.01	0.06	351.99	12.67	0.00	0.08
		15	11.49	34.67	0.00	0.00	33.00	24.33	0.04	0.12	29.67	43.00	0.01	0.02
		30	25.96	71.00	0.01	0.04	25.71	19.00	0.05	0.58	61.60	46.33	0.24	0.00
		45	63.50	60.67	0.00	0.05	70.92	32.33	0.28	0.17	75.62	40.33	0.55	0.58
		60	25.00	53.67	0.01	0.01	42.22	28.33	0.13	0.21	67.14	39.33	0.06	0.35
		75	45.66	57.33	0.18	0.05	27.66	23.33	0.01	0.08	58.93	44.00	0.35	0.17
	3	*0*	42.41	38.67	0.01	0.00	34.33	5.67	0.01	0.25	1.36	0.00	0.00	0.04
		15	224.14	50.67	0.13	0.00	224.09	20.00	0.10	0.07	188.95	19.00	0.01	0.10
		30	69.91	55.67	0.28	0.33	236.34	25.33	0.18	0.00	213.00	27.33	0.17	0.00
		45	91.69	52.33	0.22	0.04	250.81	33.33	0.05	0.15	222.95	20.67	0.03	0.00
		60	100.78	47.67	0.03	0.07	245.90	24.00	0.04	0.03	91.08	15.00	0.01	0.29
		75	230.38	54.33	0.16	0.01	242.68	21.67	0.00	0.14	189.63	24.00	0.02	0.06
	4	*0*	123.55	26.00	0.01	0.00	359.33	6.33	0.01	0.05	41.82	26.33	0.00	0.04
		15	169.99	55.67	0.00	0.29	143.33	26.67	0.44	0.15	120.00	32.33	0.01	0.08
		30	165.67	61.00	0.00	0.29	137.02	18.33	0.00	0.58	148.99	25.00	0.13	0.00
		45	179.01	57.00	0.00	0.29	172.66	27.67	0.11	0.17	249.34	33.67	0.01	0.00
		60	182.65	49.33	0.00	0.01	209.64	27.67	0.19	0.07	116.64	40.33	0.01	0.04
		75	246.99	37.67	0.00	0.01	283.52	39.33	0.07	0.17	294.73	38.00	0.02	0.05
	5	*0*	1.33	0.00	0.00	0.00	348.80	5.67	0.01	0.10	2.33	8.00	0.02	0.29
		15	240.21	51.00	0.18	0.58	305.97	20.67	0.01	0.07	16.57	29.67	0.14	0.07
		30	261.81	62.00	0.02	0.05	18.28	30.67	0.49	0.00	358.34	27.33	0.47	0.00
		45	301.50	52.67	0.01	0.04	297.31	23.33	0.00	0.10	331.10	33.33	0.02	0.58
		60	269.91	51.67	0.41	0.04	272.27	28.33	0.68	0.16	336.26	33.00	0.08	0.11
		75	80.76	52.00	0.12	0.12	269.68	29.00	0.04	0.01	1.00	19.33	0.39	0.10
	6	*0*	356.67	11.33	0.00	0.00	359.02	18.33	0.03	0.12	344.15	8.67	0.03	0.10
		15	44.03	52.67	0.00	0.00	28.49	18.00	0.01	0.08	88.43	15.00	0.00	0.11
		30	42.76	62.00	0.16	0.58	17.15	21.00	0.03	0.30	9.04	18.00	0.28	0.58
		45	56.28	52.67	0.17	0.03	52.14	24.67	0.16	0.30	16.46	19.33	0.07	0.58
		60	55.67	54.33	0.14	0.07	54.98	28.33	0.10	0.17	2.51	26.67	0.66	0.29
		75	54.66	62.33	0.01	0.01	55.00	34.67	0.18	0.18	7.33	12.00	0.41	0.58
	7	*0*	330.57	29.33	0.01	0.00	344.88	10.67	0.00	0.10	341.65	23.67	0.00	0.07
		15	294.11	62.33	0.00	0.58	320.00	20.67	0.23	0.04	46.22	29.67	0.00	0.07
		30	115.05	55.33	0.03	0.29	313.66	23.00	0.03	0.00	341.81	23.33	0.51	0.00
		45	122.92	59.00	0.08	0.29	326.03	32.67	0.37	0.40	348.54	31.67	0.40	0.00
		60	141.76	57.33	0.05	0.04	316.00	26.00	0.13	0.19	52.70	23.67	0.40	0.58
		75	48.48	59.33	0.02	0.08	306.01	18.67	0.65	0.11	237.43	19.33	0.64	0.04
	8	*0*	6.33	0.00	0.00	0.00	14.99	18.67	0.01	0.12	9.99	5.00	0.01	0.12
		15	333.39	42.33	0.31	0.58	1.99	31.00	0.00	0.11	1.00	17.33	0.01	0.12
		30	286.80	59.00	0.01	0.04	7.24	28.00	0.00	0.33	322.99	51.00	0.10	0.00
		45	310.33	64.00	0.04	0.04	2.68	25.67	0.01	0.29	349.41	47.33	0.32	0.00
		60	279.98	56.67	0.64	0.02	2.00	17.67	0.51	0.10	4.84	53.33	0.45	0.13
		75	319.84	47.67	0.35	0.02	358.33	22.33	0.68	0.05	317.00	49.00	0.41	0.14
	9	*0*	0.33	0.00	0.00	0.00	10.18	10.00	0.01	0.05	24.26	9.67	0.00	0.07
		15	296.78	35.33	0.75	0.29	333.55	24.67	0.00	0.08	258.57	29.67	0.00	0.18
		30	206.56	55.67	0.24	0.29	334.97	19.00	0.21	0.58	269.99	24.67	0.00	0.00
		45	241.00	56.33	0.64	0.03	310.83	26.00	0.67	0.07	289.00	37.33	0.64	0.58
		60	259.17	53.00	0.67	0.04	183.14	17.00	0.12	0.13	237.36	30.33	0.21	0.00
		75	227.93	58.67	0.71	0.06	307.87	11.33	0.73	0.07	256.23	23.67	0.26	0.08

Table 7. Individual data from 3 participants in the Non Rotated Medium Anisotropy condition with RMS Contrast.

Rotated RMS Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Medium Anisotropy	**1**	*0*	4.66	16.00	0.01	0.58	358.01	3.00	0.01	0.58	22.96	23.67	0.04	0.30
		15	108.88	49.33	0.42	0.09	358.34	13.00	0.01	0.39	347.03	41.00	0.24	0.09
		30	130.61	47.00	0.10	0.05	2.00	44.33	0.01	0.06	290.67	35.33	0.00	0.10
		45	303.11	40.67	0.59	0.09	354.10	32.33	0.04	0.13	293.31	38.33	0.01	0.08
		60	95.70	50.67	0.65	0.08	357.33	32.67	0.00	0.18	309.65	46.00	0.01	0.01
		75	87.51	65.67	0.41	0.10	8.67	37.00	0.00	0.05	325.01	45.67	0.00	0.02
	2	*0*	326.05	20.67	0.50	0.58	354.01	15.00	0.01	0.10	349.79	48.00	0.04	0.01
		15	3.82	51.00	0.05	0.07	337.34	42.00	0.01	0.15	355.31	61.00	0.02	0.05
		30	357.51	57.33	0.05	0.04	330.33	22.00	0.00	0.22	8.62	61.67	0.02	0.02
		45	24.98	55.67	0.02	0.09	346.32	34.33	0.02	0.20	316.80	37.33	0.51	0.18
		60	342.00	57.33	0.67	0.02	11.70	31.00	0.42	0.14	16.37	57.33	0.02	0.04
		75	168.48	52.33	0.08	0.04	331.99	16.00	0.01	0.08	332.12	57.00	0.05	0.03
	3	*0*	356.61	22.33	0.13	0.29	352.70	18.00	0.01	0.39	281.66	25.33	0.61	0.36
		15	90.37	67.00	0.54	0.05	347.48	26.67	0.04	0.17	46.70	25.00	0.17	0.25
		30	205.52	56.33	0.35	0.05	106.24	28.33	0.35	0.17	77.57	28.33	0.22	0.08
		45	219.37	45.33	0.39	0.10	114.16	27.33	0.04	0.13	77.31	25.00	0.01	0.06
		60	217.40	29.00	0.84	0.29	117.66	25.67	0.03	0.13	92.28	28.33	0.02	0.09
		75	208.91	65.00	0.36	0.07	132.66	30.67	0.02	0.13	94.33	27.00	0.01	0.13
	4	*0*	29.75	18.67	0.29	0.58	356.35	31.33	0.01	0.12	2.71	36.00	0.03	0.30
		15	0.33	61.33	0.01	0.19	143.03	19.33	0.59	0.09	10.20	59.33	0.06	0.03
		30	252.03	49.00	0.69	0.04	180.79	29.00	0.04	0.01	7.61	57.33	0.11	0.06
		45	344.32	47.67	0.02	0.03	166.58	27.00	0.65	0.07	15.95	66.67	0.03	0.05
		60	338.58	43.33	0.03	0.04	70.08	23.33	0.60	0.14	34.00	51.00	0.00	0.05
		75	255.66	50.67	0.50	0.07	192.18	36.67	0.05	0.06	44.82	52.67	0.66	0.13
	5	*0*	360.00	0.00	0.01	0.00	354.00	14.00	0.00	0.04	328.62	28.67	0.07	0.31
		15	72.83	47.00	0.41	0.04	322.03	23.67	0.66	0.14	317.25	43.67	0.08	0.17
		30	100.96	57.67	0.36	0.07	208.34	23.00	0.65	0.04	237.60	24.00	0.02	0.07
		45	183.25	60.67	0.73	0.09	149.01	26.33	0.01	0.04	235.31	16.67	0.10	0.18
		60	328.96	55.00	0.07	0.04	154.66	26.67	0.01	0.05	257.75	26.00	0.34	0.21
		75	308.67	67.67	0.56	0.05	155.68	25.33	0.01	0.17	215.62	17.33	0.21	0.10
	6	*0*	357.33	0.00	0.00	0.00	354.33	7.33	0.00	0.21	19.67	44.67	0.01	0.08
		15	255.74	58.33	0.31	0.05	3.01	16.67	0.01	0.19	32.16	41.67	0.50	0.04
		30	261.38	47.00	0.27	0.07	3.99	15.67	0.00	0.18	56.40	35.67	0.06	0.13
		45	345.29	52.00	0.38	0.04	354.06	26.00	0.03	0.21	25.29	42.67	0.19	0.15
		60	342.57	45.67	0.26	0.12	357.00	38.33	0.00	0.12	340.05	56.33	0.18	0.06
		75	165.53	42.33	0.52	0.06	334.63	10.67	0.02	0.07	38.35	59.33	0.01	0.04
	7	*0*	359.01	0.00	0.01	0.00	1.67	2.67	0.00	0.38	27.99	28.33	0.00	0.29
		15	4.71	49.00	0.21	0.02	152.00	13.67	0.69	0.14	93.00	25.00	0.00	0.13
		30	9.36	55.33	0.01	0.02	262.70	15.00	0.34	0.12	112.29	22.33	0.04	0.11
		45	203.12	64.33	0.65	0.08	314.61	13.67	0.55	0.32	269.42	27.67	0.89	0.25
		60	10.83	50.33	0.05	0.09	323.50	22.00	0.60	0.32	145.11	32.33	0.61	0.18
		75	353.00	48.67	0.09	0.06	297.90	14.00	0.51	0.42	329.35	36.00	0.60	0.11
	8	*0*	333.07	29.00	0.09	0.29	5.93	35.33	0.02	0.13	338.87	47.67	0.04	0.04
		15	36.36	59.00	0.32	0.09	354.36	33.00	0.01	0.24	349.67	53.67	0.01	0.02
		30	128.23	53.33	0.26	0.11	186.84	27.00	0.04	0.02	205.58	13.00	0.13	0.12
		45	91.35	55.33	0.01	0.02	333.27	17.33	0.58	0.08	205.97	11.33	0.21	0.10
		60	126.49	56.33	0.28	0.04	215.50	26.33	0.69	0.10	221.02	14.33	0.42	0.13
		75	139.62	48.67	0.24	0.06	234.51	23.33	0.66	0.10	102.12	21.33	0.43	0.05
	9	*0*	265.18	34.00	0.47	0.29	356.33	2.33	0.00	0.33	7.26	44.67	0.02	0.04
		15	76.32	56.33	0.13	0.04	353.33	34.33	0.00	0.02	343.66	51.33	0.00	0.06
		30	207.82	57.33	0.34	0.04	339.14	40.67	0.64	0.12	333.09	33.00	0.32	0.09
		45	90.09	57.00	0.63	0.07	182.61	17.67	0.21	0.08	7.24	33.00	0.16	0.18
		60	69.04	41.67	0.13	0.04	144.09	23.67	0.72	0.31	62.01	36.00	0.00	0.12
		75	312.39	50.00	0.77	0.01	131.90	16.33	0.05	0.06	103.82	20.33	0.11	0.06

Table 8. Individual data from 3 participants in the Rotated Medium Anisotropy condition with RMS Contrast.

Non Rotated Natural Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Low Anisotropy	**1**	*0*	334.18	18.33	0.59	0.05	337.09	0.67	0.05	0.58	24.20	4.33	0.47	0.58
		15	353.06	0.00	0.02	0.05	14.36	5.33	0.18	0.31	345.47	6.67	0.13	0.58
		30	112.60	52.00	0.16	0.58	267.84	11.33	0.70	0.12	308.46	13.00	0.13	0.33
		45	351.00	6.67	0.00	0.00	270.08	17.67	0.65	0.27	282.93	16.00	0.13	0.08
		60	91.00	54.33	0.01	0.03	11.00	19.67	0.39	0.21	280.54	22.33	0.13	0.04
		75	354.41	43.67	0.56	0.58	262.72	23.33	0.03	0.07	280.33	21.67	0.00	0.14
	2	*0*	101.35	58.67	0.02	0.02	21.98	9.33	0.02	0.29	338.33	2.33	0.50	0.58
		15	241.01	47.00	0.01	0.04	357.72	13.00	0.23	0.18	29.57	4.67	0.41	0.58
		30	100.36	65.00	0.06	0.06	351.21	13.00	0.31	0.16	303.31	7.33	0.54	0.30
		45	291.96	52.33	0.13	0.10	319.31	14.00	0.20	0.11	223.30	6.33	0.62	0.29
		60	106.31	71.00	0.01	0.02	315.71	11.00	0.20	0.10	226.26	15.67	0.35	0.14
		75	221.65	53.00	0.07	0.10	295.36	15.00	0.07	0.18	195.15	9.00	0.71	0.29
	3	*0*	298.68	30.00	0.47	0.04	335.93	6.67	0.10	0.35	2.66	0.00	0.01	0.00
		15	348.00	0.00	0.00	0.10	345.80	7.00	0.38	0.42	358.45	0.00	0.05	0.00
		30	11.49	34.67	0.05	0.29	8.75	3.33	0.03	0.58	15.59	6.33	0.08	0.58
		45	22.41	17.00	0.33	0.00	26.67	15.33	0.00	0.03	22.89	2.33	0.29	0.58
		60	25.96	71.00	0.02	0.29	32.97	9.00	0.01	0.23	41.67	25.33	0.00	0.07
		75	133.46	46.67	0.64	0.58	18.54	16.67	0.03	0.10	41.72	15.67	0.09	0.29
	4	*0*	63.50	60.67	0.57	0.06	348.33	4.00	0.01	0.58	355.67	0.33	0.00	0.58
		15	333.16	20.33	0.52	0.04	2.23	13.67	0.05	0.04	18.11	6.33	0.12	0.53
		30	25.00	53.67	0.00	0.06	1.27	7.67	0.14	0.31	26.39	8.00	0.23	0.58
		45	301.95	50.33	0.02	0.30	24.31	12.67	0.03	0.04	84.36	23.67	0.01	0.01
		60	45.66	57.33	0.01	0.04	47.62	21.00	0.03	0.02	77.82	19.67	0.14	0.29
		75	266.01	50.00	0.03	0.07	71.66	20.00	0.02	0.08	94.67	27.00	0.00	0.04
	5	*0*	42.41	38.67	0.19	0.04	16.99	16.67	0.23	0.13	0.34	0.00	0.00	0.00
		15	352.67	0.00	0.00	0.05	57.98	14.00	0.06	0.11	9.24	10.67	0.25	0.58
		30	224.14	50.67	0.29	0.29	0.70	11.33	0.08	0.07	334.36	4.67	0.64	0.58
		45	357.67	12.67	0.00	0.00	317.18	17.33	0.04	0.23	109.21	16.33	0.25	0.16
		60	69.91	55.67	0.03	0.02	356.44	13.67	0.08	0.04	300.82	9.33	0.50	0.29
		75	353.98	42.33	0.01	0.58	62.95	17.67	0.02	0.05	158.00	22.00	0.24	0.07
	6	*0*	91.69	52.33	0.60	0.03	348.66	15.67	0.07	0.07	331.23	0.33	0.25	0.58
		15	357.02	51.33	0.01	0.02	304.24	16.67	0.27	0.13	354.33	0.33	0.01	0.58
		30	100.78	47.67	0.90	0.03	302.04	16.00	0.15	0.10	273.00	19.67	0.00	0.14
		45	360.00	57.67	0.00	0.06	312.04	13.67	0.22	0.20	277.34	29.67	0.00	0.07
		60	230.38	54.33	0.44	0.10	285.59	14.33	0.04	0.08	280.32	25.33	0.01	0.04
		75	353.68	58.33	0.01	0.04	271.33	26.00	0.00	0.03	274.67	25.67	0.00	0.01
	7	*0*	123.55	26.00	0.60	0.04	15.08	11.33	0.20	0.09	354.67	0.00	0.00	0.00
		15	0.34	0.00	0.00	0.03	323.68	14.00	0.01	0.06	345.29	5.00	0.06	0.58
		30	169.99	55.67	0.01	0.35	341.72	19.33	0.10	0.12	299.00	24.33	0.00	0.02
		45	20.09	17.00	0.34	0.00	17.34	13.67	0.61	0.08	318.67	23.67	0.00	0.10
		60	165.67	61.00	0.00	0.04	324.24	21.33	0.16	0.27	326.08	41.00	0.08	0.12
		75	97.00	43.33	0.01	0.58	295.99	37.00	0.01	0.05	288.33	33.00	0.00	0.07
	8	*0*	179.01	57.00	0.01	0.01	7.00	9.33	0.03	0.29	14.44	0.00	0.29	0.00
		15	60.87	33.67	0.39	0.06	15.00	11.67	0.00	0.09	31.67	2.67	0.68	0.58
		30	182.65	49.33	0.01	0.01	7.18	15.00	0.12	0.00	337.32	13.33	0.01	0.29
		45	64.19	47.67	0.14	0.29	14.99	19.33	0.00	0.04	352.15	23.33	0.05	0.29
		60	246.99	37.67	0.01	0.02	2.63	11.67	0.11	0.03	338.30	39.33	0.10	0.11
		75	80.97	70.00	0.01	0.06	335.56	34.00	0.03	0.13	332.29	35.33	0.02	0.01
	9	*0*	1.33	0.00	0.00	0.07	1.00	3.00	0.00	0.58	357.67	0.00	0.00	0.00
		15	352.33	0.00	0.00	0.05	327.57	12.67	0.52	0.13	2.00	0.00	0.00	0.00
		30	240.21	51.00	0.07	0.00	336.00	12.33	0.00	0.16	351.39	12.33	0.08	0.29
		45	348.33	0.00	0.00	0.00	340.00	12.00	0.00	0.08	345.53	11.33	0.05	0.35
		60	261.81	62.00	0.07	0.02	339.00	14.00	0.00	0.04	334.40	18.67	0.06	0.03
		75	325.67	38.33	0.00	0.00	19.48	15.33	0.31	0.05	5.23	24.00	0.33	0.12

Table 9. Individual data from 3 participants in the Non Rotated Low Anisotropy condition with Natural Contrast.

Rotated Natural Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Low Anisotropy	**1**	*0*	344.78	0.00	0.03	0.00	339.24	14.00	0.63	0.07	356.67	0.00	0.00	0.12
		15	355.35	11.00	0.01	0.58	106.50	19.00	0.39	0.06	3.87	4.00	0.64	0.33
		30	346.55	51.33	0.05	0.10	156.01	17.33	0.02	0.09	91.13	26.67	0.45	0.00
		45	356.33	51.67	0.00	0.05	148.00	20.33	0.00	0.03	145.22	14.33	0.14	0.58
		60	356.33	60.00	0.00	0.02	126.10	27.00	0.56	0.13	172.94	18.00	0.02	0.11
		75	357.67	59.33	0.00	0.04	179.34	31.00	0.01	0.05	149.36	21.33	0.08	0.11
	2	*0*	352.34	0.00	0.01	0.00	6.58	7.00	0.63	0.11	355.34	0.00	0.00	0.16
		15	17.36	9.33	0.56	0.58	14.92	9.67	0.08	0.03	326.93	19.00	0.07	0.07
		30	93.00	34.33	0.50	0.30	8.00	34.00	0.00	0.33	354.70	40.33	0.03	0.00
		45	160.93	60.00	0.06	0.04	4.38	29.33	0.07	0.46	2.01	51.00	0.01	0.29
		60	178.00	50.67	0.00	0.02	3.33	49.33	0.00	0.11	12.00	40.67	0.00	0.15
		75	341.02	56.00	0.64	0.04	0.33	50.67	0.00	0.15	13.98	47.33	0.01	0.06
	3	*0*	357.33	0.00	0.00	0.00	16.60	19.00	0.04	0.04	353.33	0.00	0.00	0.14
		15	333.81	16.33	0.37	0.58	9.65	16.33	0.01	0.03	1.67	10.33	0.03	0.11
		30	226.50	45.33	0.17	0.09	29.38	18.67	0.03	0.13	35.20	26.67	0.18	0.00
		45	241.98	60.00	0.23	0.07	20.15	19.67	0.10	0.16	18.50	16.00	0.34	0.58
		60	254.36	63.33	0.32	0.08	25.35	18.00	0.64	0.18	85.04	17.00	0.13	0.34
		75	266.55	57.67	0.12	0.04	5.67	38.00	0.00	0.12	46.27	9.00	0.33	0.31
	4	*0*	354.67	0.33	0.00	0.58	357.67	10.33	0.01	0.11	358.33	0.00	0.00	0.22
		15	192.58	31.00	0.81	0.29	26.66	11.67	0.01	0.04	5.64	0.00	0.01	0.58
		30	37.53	44.33	0.56	0.11	63.65	14.00	0.35	0.36	356.67	0.00	0.00	0.00
		45	20.40	63.67	0.56	0.01	109.47	20.33	0.41	0.58	20.62	4.00	0.63	0.00
		60	227.82	60.00	0.79	0.07	146.06	19.00	0.04	0.04	77.13	24.67	0.44	0.00
		75	255.53	62.33	0.38	0.09	163.65	27.67	0.01	0.23	139.70	10.33	0.60	0.58
	5	*0*	358.33	0.00	0.00	0.00	3.99	7.00	0.00	0.09	2.00	0.00	0.00	0.23
		15	0.67	0.00	0.00	0.00	356.44	27.67	0.05	0.06	348.16	12.00	0.04	0.34
		30	180.35	51.67	0.04	0.07	327.64	24.67	0.01	0.58	333.86	24.00	0.04	0.00
		45	125.64	35.67	0.69	0.30	347.74	27.33	0.07	0.18	324.66	43.67	0.01	0.58
		60	181.33	64.00	0.02	0.08	341.16	25.67	0.24	0.06	314.32	41.67	0.10	0.32
		75	175.33	50.33	0.00	0.06	324.20	16.67	0.05	0.29	328.74	44.00	0.04	0.11
	6	*0*	1.67	0.00	0.00	0.00	2.67	7.33	0.00	0.05	0.67	0.00	0.00	0.08
		15	356.66	0.00	0.01	0.00	337.00	20.67	0.00	0.23	354.00	0.00	0.00	0.02
		30	109.62	53.33	0.57	0.02	324.69	18.33	0.02	0.32	280.33	38.33	0.00	0.00
		45	98.76	56.67	0.11	0.10	327.96	14.00	0.02	0.18	296.06	24.67	0.15	0.00
		60	81.29	52.33	0.02	0.03	285.01	28.33	0.14	0.09	282.00	29.67	0.00	0.08
		75	346.45	66.33	0.04	0.03	269.61	43.33	0.08	0.11	274.16	41.33	0.06	0.10
	7	*0*	356.67	0.00	0.00	0.00	2.31	2.67	0.01	0.19	350.21	12.00	0.05	0.03
		15	18.77	8.67	0.51	0.58	1.15	9.00	0.64	0.06	343.61	6.67	0.12	0.18
		30	186.03	52.00	0.63	0.01	297.87	15.00	0.65	0.58	345.99	11.33	0.10	0.58
		45	262.79	50.00	0.09	0.05	346.00	23.00	0.58	0.29	256.67	21.00	0.00	0.53
		60	268.68	61.00	0.42	0.03	347.95	19.00	0.41	0.06	268.17	23.33	0.25	0.58
		75	283.40	59.00	0.20	0.10	9.33	24.00	0.00	0.03	304.85	23.33	0.14	0.09
	8	*0*	352.00	0.00	0.00	0.00	358.00	11.00	0.00	0.05	357.33	0.00	0.00	0.08
		15	358.67	16.67	0.00	0.56	158.99	17.33	0.67	0.04	356.00	0.00	0.00	0.17
		30	351.00	37.00	0.00	0.30	122.44	27.67	0.37	0.35	144.88	9.33	0.67	0.00
		45	358.67	60.33	0.00	0.03	160.36	30.33	0.01	0.13	147.25	16.33	0.24	0.00
		60	358.33	68.00	0.00	0.02	166.33	30.33	0.01	0.03	136.34	17.33	0.13	0.31
		75	357.33	59.00	0.00	0.07	176.33	27.67	0.00	0.04	149.85	23.67	0.14	0.11
	9	*0*	355.33	0.00	0.00	0.00	355.67	0.00	0.01	0.11	348.02	0.00	0.01	0.10
		15	0.82	18.67	0.05	0.58	314.78	10.33	0.77	0.01	359.67	0.00	0.00	0.05
		30	14.81	45.33	0.05	0.08	222.98	24.00	0.07	0.00	279.80	21.33	0.29	0.00
		45	2.66	51.67	0.00	0.02	216.21	24.33	0.31	0.13	231.55	19.67	0.45	0.00
		60	3.33	64.00	0.00	0.06	217.57	30.67	0.03	0.04	218.65	27.67	0.02	0.36
		75	1.00	55.00	0.00	0.06	224.08	33.33	0.14	0.07	239.47	27.33	0.14	0.11

Table 10. Individual data from 3 participants in the Rotated Low Anisotropy condition with Natural Contrast.

Non Rotated RMS Contrast

			CV				FRZ				KG			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Low Anisotropy	**1**	*0*	212.66	54.67	0.03	0.00	19.32	33.00	0.63	0.07	33.79	36.33	0.00	0.12
		15	228.91	65.00	0.01	0.58	251.95	36.67	0.39	0.06	293.59	30.67	0.64	0.33
		30	270.18	63.00	0.05	0.10	31.55	29.67	0.02	0.09	325.05	39.00	0.45	0.00
		45	236.58	47.33	0.00	0.05	220.84	33.00	0.00	0.03	261.19	24.00	0.14	0.58
		60	294.34	46.67	0.00	0.02	271.46	31.67	0.56	0.13	292.97	45.67	0.02	0.11
		75	280.35	57.67	0.00	0.04	247.00	39.33	0.01	0.05	262.02	41.33	0.08	0.11
	2	*0*	191.73	50.33	0.01	0.00	35.67	25.33	0.63	0.11	69.83	17.33	0.00	0.16
		15	301.28	56.00	0.56	0.58	310.77	25.67	0.08	0.03	86.36	32.00	0.07	0.07
		30	264.65	62.67	0.50	0.30	278.67	35.33	0.00	0.33	260.64	27.33	0.03	0.00
		45	241.24	56.67	0.06	0.04	271.00	22.67	0.07	0.46	296.00	34.00	0.01	0.29
		60	243.36	57.00	0.00	0.02	284.36	25.00	0.00	0.11	315.14	40.67	0.00	0.15
		75	220.67	61.00	0.64	0.04	282.00	31.00	0.00	0.15	289.48	37.00	0.01	0.06
	3	*0*	145.09	37.00	0.00	0.00	340.69	16.00	0.04	0.04	7.67	31.67	0.00	0.14
		15	8.01	37.67	0.37	0.58	23.33	19.33	0.01	0.03	44.64	29.00	0.03	0.11
		30	2.62	41.67	0.17	0.09	359.16	18.67	0.03	0.13	30.79	42.33	0.18	0.00
		45	74.77	48.33	0.23	0.07	358.22	19.67	0.10	0.16	50.33	28.67	0.34	0.58
		60	163.28	44.33	0.32	0.08	28.01	19.67	0.64	0.18	27.54	35.33	0.13	0.34
		75	6.44	62.33	0.12	0.04	31.33	23.67	0.00	0.12	41.02	42.33	0.33	0.31
	4	*0*	303.71	59.67	0.00	0.58	337.89	22.33	0.01	0.11	344.68	39.00	0.00	0.22
		15	90.31	62.67	0.81	0.29	33.35	29.33	0.01	0.04	27.92	39.67	0.01	0.58
		30	80.67	63.67	0.56	0.11	24.74	21.33	0.35	0.36	125.87	34.00	0.00	0.00
		45	50.57	67.33	0.56	0.01	31.33	21.33	0.41	0.58	88.00	31.33	0.63	0.00
		60	57.53	60.67	0.79	0.07	47.01	17.67	0.04	0.04	87.65	33.67	0.44	0.00
		75	88.76	52.33	0.38	0.09	76.67	35.33	0.01	0.23	82.33	35.33	0.60	0.58
	5	*0*	168.52	45.00	0.00	0.00	74.84	28.00	0.00	0.09	96.66	31.67	0.00	0.23
		15	154.81	55.33	0.00	0.00	70.28	27.67	0.05	0.06	98.33	32.00	0.04	0.34
		30	103.81	43.00	0.04	0.07	326.24	16.00	0.01	0.58	209.13	13.33	0.04	0.00
		45	31.33	64.00	0.69	0.30	304.64	25.67	0.07	0.18	6.65	21.33	0.01	0.58
		60	75.65	57.67	0.02	0.08	20.36	29.00	0.24	0.06	352.85	10.33	0.10	0.32
		75	140.72	58.33	0.00	0.06	112.62	28.00	0.05	0.29	93.16	20.67	0.04	0.11
	6	*0*	227.93	58.33	0.00	0.00	302.74	16.00	0.00	0.05	301.08	33.33	0.00	0.08
		15	267.74	51.67	0.01	0.00	249.64	29.67	0.00	0.23	267.63	25.67	0.00	0.02
		30	239.07	57.33	0.57	0.02	280.06	23.67	0.02	0.32	276.33	32.33	0.00	0.00
		45	263.99	54.00	0.11	0.10	240.62	36.67	0.02	0.18	279.79	29.67	0.15	0.00
		60	231.91	58.00	0.02	0.03	257.01	34.67	0.14	0.09	280.66	35.33	0.00	0.08
		75	237.10	60.67	0.04	0.03	267.67	38.00	0.08	0.11	291.00	38.00	0.06	0.10
	7	*0*	235.46	44.33	0.00	0.00	350.43	40.33	0.01	0.19	344.00	23.67	0.05	0.03
		15	277.54	57.67	0.51	0.58	309.34	45.00	0.64	0.06	323.43	47.33	0.12	0.18
		30	269.42	59.33	0.63	0.01	305.67	50.00	0.65	0.58	323.00	42.33	0.10	0.58
		45	304.67	55.33	0.09	0.05	285.35	34.00	0.58	0.29	327.00	48.33	0.00	0.53
		60	257.82	58.00	0.42	0.03	285.26	32.00	0.41	0.06	325.67	37.33	0.25	0.58
		75	254.69	60.33	0.20	0.10	285.67	42.00	0.00	0.03	303.94	55.00	0.14	0.09
	8	*0*	260.64	56.67	0.00	0.00	353.32	19.33	0.00	0.05	2.47	53.67	0.00	0.08
		15	325.53	48.33	0.00	0.56	347.45	34.00	0.67	0.04	326.28	51.33	0.00	0.17
		30	321.06	62.67	0.00	0.30	21.66	43.00	0.37	0.35	328.67	54.33	0.67	0.00
		45	335.00	54.00	0.00	0.03	2.66	41.67	0.01	0.13	335.01	46.67	0.24	0.00
		60	345.31	54.00	0.00	0.02	13.00	35.00	0.01	0.03	356.65	52.67	0.13	0.31
		75	321.00	60.00	0.00	0.07	353.44	40.00	0.00	0.04	338.05	63.33	0.14	0.11
	9	*0*	352.02	42.33	0.00	0.00	355.45	16.33	0.01	0.11	155.00	19.00	0.01	0.10
		15	317.46	55.33	0.05	0.58	320.67	28.67	0.77	0.01	339.66	39.00	0.00	0.05
		30	19.41	55.00	0.05	0.08	312.58	22.33	0.07	0.00	346.30	32.67	0.29	0.00
		45	12.26	54.33	0.00	0.02	324.65	28.33	0.31	0.13	351.00	25.33	0.45	0.00
		60	325.05	50.00	0.00	0.06	332.69	23.67	0.03	0.04	340.30	35.67	0.02	0.36
		75	314.53	65.67	0.00	0.06	346.68	31.33	0.14	0.07	338.89	40.00	0.14	0.11

Table 11. Individual data from 3 participants in the Non Rotated Low Anisotropy condition with RMS Contrast.

Rotated RMS Contrast

			CV				*FRZ*				*KG*			
Surface	*Location*	*Lighting*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*	*AVG Tilt*	*AVG Slant*	*Tilt variance*	*Slant variance*
Low Anisotropy	**1**	*0*	106.85	29.67	0.51	0.29	350.14	26.33	0.68	0.10	345.02	66.00	0.01	0.01
		15	304.79	63.00	0.56	0.08	156.34	26.00	0.01	0.13	355.96	57.33	0.02	0.03
		30	15.57	56.00	0.05	0.28	166.82	32.00	0.04	0.06	134.75	35.67	0.70	0.15
		45	353.01	55.67	0.01	0.04	158.21	32.00	0.05	0.03	104.62	39.00	0.10	0.03
		60	7.15	58.00	0.05	0.07	154.95	25.00	0.67	0.06	132.53	24.33	0.55	0.16
		75	350.33	52.00	0.00	0.05	175.66	30.67	0.01	0.05	158.61	15.33	0.02	0.14
	2	*0*	109.96	49.67	0.15	0.04	213.83	12.33	0.55	0.20	316.99	54.00	0.01	0.02
		15	98.41	46.33	0.31	0.03	354.66	36.33	0.00	0.15	329.60	49.00	0.05	0.01
		30	137.84	64.67	0.70	0.07	349.70	41.33	0.03	0.07	11.33	53.67	0.01	0.03
		45	74.08	64.33	0.60	0.05	357.33	20.33	0.00	0.16	2.67	57.67	0.00	0.03
		60	134.43	53.00	0.22	0.11	346.12	37.33	0.03	0.20	3.33	62.00	0.00	0.04
		75	340.23	52.00	0.66	0.01	360.00	54.33	0.00	0.02	355.47	64.67	0.05	0.04
	3	*0*	37.43	37.67	0.60	0.30	14.09	38.33	0.04	0.13	16.32	56.33	0.01	0.03
		15	162.00	56.67	0.67	0.11	348.33	33.67	0.01	0.12	53.93	29.67	0.18	0.25
		30	68.92	62.67	0.38	0.05	67.55	17.33	0.50	0.06	70.04	21.00	0.10	0.13
		45	37.66	61.00	0.01	0.07	23.12	20.00	0.54	0.06	82.77	16.00	0.04	0.06
		60	30.00	63.00	0.01	0.10	130.14	21.33	0.16	0.03	98.93	20.67	0.08	0.06
		75	118.08	35.33	0.58	0.29	341.05	30.33	0.57	0.15	50.75	34.67	0.09	0.11
	4	*0*	44.79	25.00	0.51	0.29	348.85	16.33	0.05	0.30	360.00	50.67	0.00	0.05
		15	39.59	61.33	0.65	0.05	146.87	23.00	0.68	0.17	350.34	52.00	0.01	0.06
		30	166.91	72.67	0.84	0.05	107.58	23.67	0.63	0.04	332.66	39.67	0.02	0.10
		45	125.67	69.67	0.38	0.05	161.86	21.33	0.07	0.07	95.68	23.67	0.01	0.15
		60	200.91	63.00	0.25	0.08	153.68	20.00	0.01	0.13	84.74	18.00	0.08	0.12
		75	168.67	66.00	0.01	0.08	182.67	28.67	0.00	0.02	76.42	25.67	0.47	0.11
	5	*0*	352.56	32.67	0.29	0.33	330.64	17.67	0.24	0.33	348.48	42.67	0.11	0.11
		15	128.30	52.00	0.42	0.08	331.67	21.33	0.00	0.21	330.39	53.67	0.02	0.07
		30	331.04	49.00	0.67	0.04	13.69	20.00	0.24	0.18	338.99	30.00	0.03	0.29
		45	185.55	51.67	0.61	0.02	340.34	24.67	0.02	0.18	318.33	42.67	0.00	0.04
		60	168.34	58.33	0.74	0.07	336.82	39.67	0.04	0.14	340.99	46.67	0.01	0.04
		75	98.09	59.00	0.36	0.10	317.32	22.67	0.02	0.09	341.96	46.67	0.09	0.07
	6	*0*	131.42	21.67	0.63	0.30	13.24	13.67	0.12	0.17	332.91	26.67	0.08	0.30
		15	30.58	61.33	0.20	0.06	305.75	14.33	0.11	0.07	301.32	44.00	0.05	0.08
		30	352.33	55.00	0.00	0.04	299.84	38.00	0.06	0.07	296.95	49.33	0.07	0.05
		45	71.46	53.67	0.48	0.08	280.02	40.33	0.15	0.02	308.41	48.33	0.03	0.09
		60	183.28	55.33	0.18	0.08	292.22	24.67	0.11	0.22	288.32	44.33	0.01	0.05
		75	154.23	63.67	0.64	0.08	273.19	50.33	0.22	0.12	308.00	58.33	0.01	0.03
	7	*0*	326.33	34.33	0.52	0.29	3.99	16.33	0.00	0.26	27.90	48.00	0.04	0.05
		15	152.70	49.33	0.01	0.04	341.89	27.00	0.12	0.12	270.08	26.00	0.02	0.06
		30	128.77	55.00	0.68	0.02	226.72	23.67	0.08	0.03	263.14	26.33	0.25	0.21
		45	304.57	54.33	0.65	0.19	228.01	24.33	0.02	0.03	268.29	19.00	0.07	0.12
		60	298.66	47.33	0.01	0.05	238.66	24.33	0.01	0.06	284.99	33.33	0.25	0.20
		75	217.69	45.00	0.21	0.01	248.00	21.67	0.09	0.10	282.06	27.00	0.08	0.12
	8	*0*	1.33	18.67	0.00	0.58	19.64	4.67	0.50	0.30	25.02	49.33	0.17	0.11
		15	25.68	46.33	0.51	0.08	149.96	33.00	0.02	0.06	93.05	16.00	0.28	0.30
		30	358.62	58.33	0.11	0.04	154.03	33.00	0.06	0.02	137.66	26.33	0.01	0.04
		45	346.00	61.33	0.02	0.01	159.05	27.67	0.04	0.04	126.72	24.33	0.03	0.15
		60	347.52	62.33	0.04	0.06	158.33	28.00	0.00	0.14	140.75	16.67	0.12	0.06
		75	0.33	50.33	0.00	0.03	172.63	36.00	0.02	0.02	152.83	18.00	0.06	0.07
	9	*0*	354.70	24.67	0.03	0.30	337.23	15.67	0.60	0.41	2.65	35.00	0.33	0.07
		15	4.33	50.33	0.00	0.05	197.45	19.67	0.39	0.16	291.58	19.33	0.71	0.29
		30	357.38	63.00	0.02	0.01	193.74	20.33	0.60	0.20	208.23	19.67	0.66	0.09
		45	22.08	52.67	0.03	0.04	249.30	18.67	0.50	0.29	218.56	22.33	0.03	0.10
		60	28.11	55.00	0.67	0.03	216.67	26.67	0.08	0.10	185.68	20.33	0.06	0.12
		75	174.37	56.67	0.09	0.03	187.85	32.67	0.36	0.02	199.24	20.67	0.13	0.02

Table 12. Individual data from 3 participants in the Rotated Low Anisotropy condition with RMS Contrast.

4. Is there only one sun?

In this chapter we explore the human visual system's ability to cope with lighting ambiguities; testing whether the brain adopts a global or local approach in the presence of bi-directional illumination. Observers made shape judgements for shaded disks embedded in complex scenes containing shading and shadows providing cues to the direction of the illuminant. Scenes were lit by two spotlights such that local lighting cues depended on the test location but indicated a single source whereas global cues suggested two opposing light sources. We found that strong local lighting cues tend to override priors and therefore to guide observers' perception of shape-from-shading. Specifically, where two light sources illuminate the scene, the one illuminating the part of the scene under test determines the percept; on the other hand, where the two sources illuminate the local area equally, giving an impression of diffuse lighting, the observer's prior dominates as is also the case when local lighting cues were absent or weak.

4.1 Introduction

The human visual system is able to reconstruct the shape of objects based on limited sources of information such as patterns of shading. For recovering shape-from-shading, the role of prior assumptions and bounding constraints is crucial, given that shading itself is highly ambiguous (Ramachandran, 1988, Ramachandran and Kleffner, 1992; Mamassian and Goutcher, 2001). Each shading pattern is a result of a combination of material properties, object shape,

and light source position, and any given pattern can result from several scene and lighting geometries. However, the human visual system consistently produces useful interpretations of such highly ambiguous stimuli. For example, the ability to perceive shape-from-shading implies that the visual system can estimate the light field, locate boundary points in the image to separate each surface from the other and discriminate shading from shadow. In order to complete all of these tasks and to have a stable percept of the environment, the visual system has to make use of prior knowledge such as assuming the direction of the dominant light source for the scene. The influence and the strength of these assumptions can be revealed by visual illusions. A typical example is the crater illusion where we interpret a crater lit from below as a mound lit from above based on the assumption that typically light comes from above. This effect, first reported by Rittenhouse (1786), has been studied for almost two centuries, and experimental versions of the illusion have been used to clarify the precise position of the assumed light source. Early studies (Ramachandran, 1988; Sun and Perona, 1998) suggested it to be located above the observer's head, but others have found it to be slightly shifted to the left (Mamassian and Goutcher, 2001). The light-from-above prior is a reasonable assumption that could help us perceiving ambiguous images in an environment where the sun shines above the line of horizon, but it can be altered by experience (Adams, Graf & Ernst, 2004) and overridden by relatively weak lighting cues (Morgenstern, Murray, Harris, 2011). These results suggest that estimates of the light source direction have a stronger influence than prior

assumptions on perceived shape. In particular, Morgenstern et al. (2011) found that the observer's percept is influenced by light direction cues in a way that is proportional to their strength. They showed how relatively strong local lighting cues could override the observer's lighting priors even in case where local lighting cues are removed or noticeably reduced. The role of lighting is so important in shape-from-shading that some studies have focused on its relationship with shape judgments, in particular, whether shape perception and lighting estimation are separate or linked processes. Todd and Mingolla (1983) first found that for Lambertian and glossy surfaces, the perceived curvature of a cylinder varied with surface treatment while the estimated lighting direction did not. In a follow up study Mingolla and Todd (1986), found that errors made by observers in a shape judgments task did not correlate well with estimates of the lighting direction, advancing the idea that these two processes may be independent. In a more recent study, Gerardin, Kourtzi and Mamassian (2010) investigated whether information about light directions and 3D shape were processed in the same or different brain areas. By combining psychophysics and fMRI measurements in a shape-from-shading task, they placed light estimation in early retinotopic areas and shape processing in higher cortical areas. Their results suggest a two-stage process of shape-from-shading where shading/luminance patterns are first analysed to estimate the lighting direction; then, presumably given this estimate, analysed to extract surface shape.

In chapters 2 and 3 we considered the relationship between lighting priors and the perception of shape. More specifically, we found that shape perception is

highly influenced by the type of object present, the orientation of the light source and the regularities in the shape of the surface itself. Chapter 2 examined lighting priors in the absence of any influence from a physical source of illumination; chapter 3, examined changes in observers' perception, related to the declination of the light source. The aim of the experiment in this chapter 4 is to investigate the relationship between lighting direction and shape perception in the case of multiple, directional lighting cues, in order to understand if and how the visual system resolves contradictory illumination cues. Most studies of shape-from-shading have used either scenes with no explicit (or articulated) lights source thus allowing observers to adopt their own lighting prior (Ramachandran, 1988, Ramachandran and Kleffner, 1992; Mamassian and Goutcher, 2001) or have used a simple, but nonetheless articulated within the image, light source model comprising a single light source which is either directional or diffuse (Langer and Bulthoff, 2001). Morgenstern et al (2011) extended this to two well-articulated light sources, one diffuse one, directional, in order to observe the strength of the lighting from above prior in the presence of local and global cues. We now further extend this line of enquiry to consider the case of two directional light sources. Such a lighting model can produce strong localised cues to lighting direction that are different in different parts of the image. Thus allowing us to directly assess the contribution of local and distal lighting cues.

One might ask about the ecological relevance of our two directional light sources. How often is this encountered in the real world where outdoor scenes

are illuminated by a sun in sky model and indoor scenes typically by a single relatively diffuse room light with the possible addition of an more directional ‘work light’. We agree that scenes illuminated by two distinct directional sources are relatively rare but we would argue that there is still a need to test such cases because it may allow the design of more advance lighting for specific situations. For example, museum exhibits often have special – multiple point source – lighting to highlight artefacts, could these be confusing to the observer? We would also argue that real world diffuse sources are more directional than many model illumination set ups suppose. Sky and room lighting is predominantly from above whereas work lighting is often side on. Windows create their own – directional – light source in a room. Further, computer screens create directional light sources that may illuminate a scene in addition to the room light and work-lamp. Generally the typical light fields are often more complex than is assumed in studies of shape from shading and we address some of that complexity here.

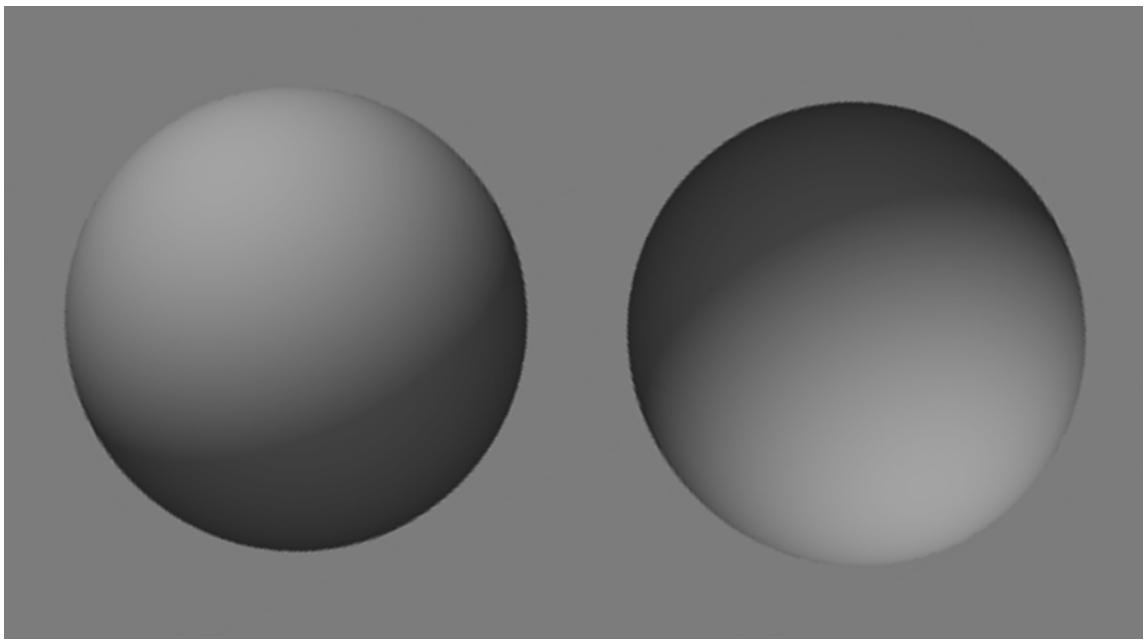

Figure 4.1. A rendered images of a convexity and a concavity lit from above left. Observers tend to perceive both as convexities and assume that there are two light sources instead.

To assess the observers' performance, we used a visual task in which an ambiguous stimulus, that could be seen either as a bump or a dent depending on the estimated (or assumed) lighting direction, was embedded in a complex scene providing information about the light field.

Stimulus ambiguities can lead to great inaccuracies in shape-from-shading, as shown in Figure 4.1, where a bias for convexity is overriding any lighting cues to perceive the two objects depicted in the scene. For instance, the visual system prefers to assume that two convex objects are lit by two opposite light sources instead of assuming that one of the two objects is concave. Questions about the type of light source involved in shapes-from-shading has led to some

contradictory results. In fact, early studies reported that the human visual system assumes there is only one directional light source (Ramachandran, 1988; Ramachandran and Kleffner, 1992); while others found that either diffuse lighting is the default assumption for SFS tasks (Tyler, 1997), or that the brain adopts a common depth assumption strategy (Zaidi, 2005). To explore how the brain interprets inconsistencies in lighting cues, we used stimuli lit by multiple light sources: two opposite ones, each providing illumination direction to half of the scene and a feeble diffuse one to test observers' bias in the absence of strong local lighting cues. This extends the work of Morgenstern at al (2011) who used only one directional light source in addition to a diffuse source.

We tested observers on a simple scene with no structure present, and thus no lighting cues, in order to measure their assumed lighting direction. These results were used as a reference for the following two conditions: first, stimuli were embedded on a structure composed of Euclidean polygons giving cues to light source direction (stimuli were modified version of those used by Morgenstern, et al, 2011); second, stimuli were embedded in the same structure but local lighting cues were eliminated by removing portions of the structure close to the location under test.

4.3 Methods

4.3.1 Participants

Twenty-one observers with normal or corrected-to-normal vision took part in the abridged version of the experiment, while three observers took part in the full experiment. Individual results from the abridged experiment where grouped for analysis whereas the data for the full experiment are displayed on an individual basis. For the abridged experiment, participants were assigned to different groups depending on individual lighting priors in the No Structure condition. In this segregation process, only two groups (*Light-from-above* prior and *Concavity* prior) reached the minimum number of participants needed (9).

4.3.2 Apparatus

Stimuli were presented on a Sony 520GDMF monitor using a VSG2/5 graphic card (CRS Ltd, Rochester, UK) and the screen resolution was set to 1024x768. Image size was 768x768 pixels (15 deg x 15 deg) at the viewing distance of 114cm, the background was set to black. Observers responded via a CRS CB3 response box. The monitor's gamma non-linearity was corrected using lookup tables in the VSG using a 4 parameters monitor model (Brainard, Pelli and Robson, 2002) fit to luminance values obtained with a CRS ColourCal luminance meter. The scenes were rendered for a linear gamma.

4.3.3 Stimuli

The stimuli (figure 4.2) were images of matte objects rendered in PovRay (Persistence of Vision Pty. Ltd., Williamstown, Victoria, Australia. http://www.povray.org/). They comprised a 'space station' design with platonic shapes made from cylinders, connected to one another with cylindrical rods. Test disks were placed on the flat frontal surface of each of 7 platonic shapes and on each trail one test location was marked with a small pentagon placed to one side of the relevant structural element. Additional platonic shapes surrounding the main structure served to provide further articulation of the light field. The simulated lighting consisted of three different light sources: one source was in the direction of the observer's point of view contributing 20% of the total illuminance of the scene while the other two distant point sources contributed equally to the overall light field but were concentrated onto different regions of the scene. The two point sources always shone from opposite directions and, being offset slightly from the centre lines of the scene, each illuminated one half the scene. Modelled as spotlights, the spread of these lights was set such that they each contributed 40% of the total illumination at the centre of the scene such that between the three light sources this location appeared as if illuminated by a diffuse source. Test disks had a diameter of approximately 0.8 deg and the whole scene measured 15 deg x 20 deg.

4.3.4 Procedure

The aim of the experiment was to test the robustness of observers' lighting priors under different types of lighting in terms of strength and direction and also to compare the effects of local and more distant light sources. Observers had to report (pressing one of two button on a CRS response box) whether test disks looked more like a bump or a dent. Scenes comprised 7 platonic solids connected by a tubular structure rich in local shading gradients to give observers information about the light source. In each trial, test disks were placed on top of each platonic solid and observers were informed about the target by a marker placed next to one of the 7 locations. Presentation time was unlimited as we wanted to give observers the time to look at the entire scene and they were instructed to be as accurate as possible.

We tested three types of scene: No structure, Full Structure, No Local cues. In the 'Full Structure' condition (figures 4.2 and 4.4) all the test disks were mounted on a platonic shape and all the shapes were connected. In the 'No Local Cues' condition the platonic shape and its connecting rods were removed from the location under test (see figure 4.5). In the 'No Structure' condition only the test disks were presented each on a plain circular background patch (see figure 4.3). The No Structure condition was used to estimate observers' lighting priors in absence of any lighting cues: in this particular case each shaded disk was lying on top of a flat grey disk with a diameter of 2.5°. The full structure

condition gave participants access to both local and distal lighting cues and enabled us to test the influence of local lighting cues on the perceived lighting direction relative to the observers' prior. The no local cues condition removed local lighting cues but retained the more distal cues enabling us to judge the effects of global lighting cues on perceived lighting direction. The use of two light sources is critical to enable us to explore how the visual system weighs local and global lighting cues. In fact, when local cues are strongly reduced, observers could only rely on strong distant light sources information or on weak (diffuse like) local cues.

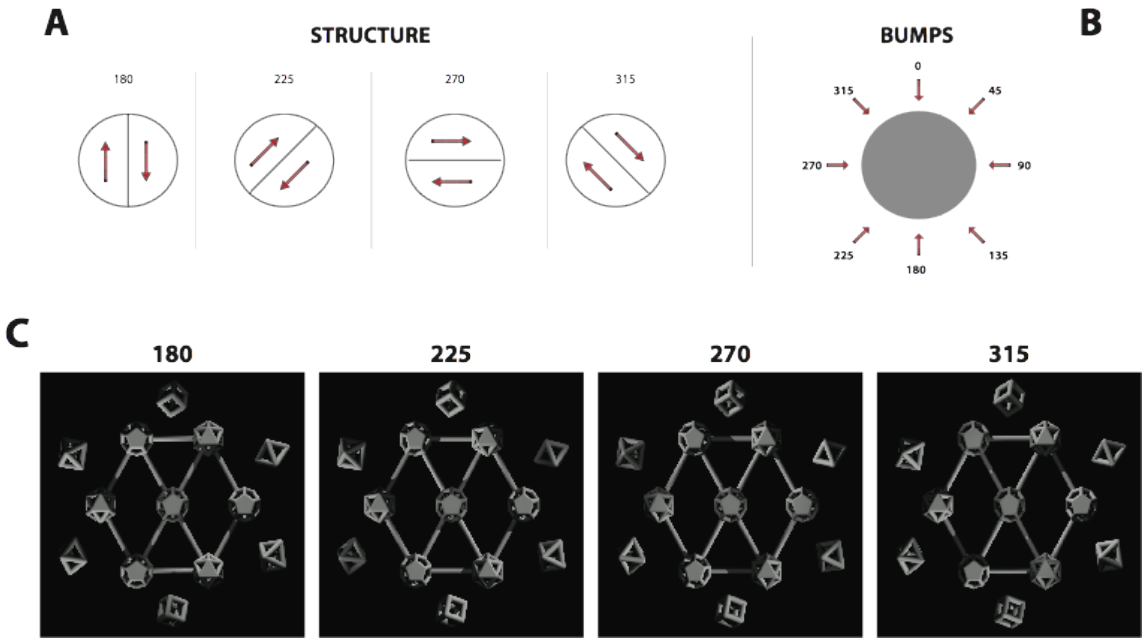

Figure 4.2. A) Lighting composition for the four different conditions. B) Lighting directions for the shaded disks. C) Four examples for the Full Structure condition without the additional bump stimuli.

Images were illuminated by two spot-lights (see Stimuli for details). We tested four different spot lighting combinations for both basic and no local cues condition: 0-180°, 45-225°, 90-270°, 135-315° (see figure 4.2 for references). The test disks (shown in figures 4.3-4.5) were bumps rendered as if lit by one of 8 evenly spaced directions in the range 0-360°. The illumination direction for the test disk placed at the location under test (marked by the pentagon) was chosen systematically (but in randomised order) as part of the experiment. The illumination direction for the remaining 6 disks was chosen at random on each trial**.** Test disks were rendered separately from the main scene and 'pasted' into each scene on a trial by trial basis. In total**,** participants who participated in the full experiment undertook 9 sessions of 560 trials each. A total of 4480 trials for the Full and No Local Cues condition (4 lighting x 7 locations x 8 bumps x 10 repetition) and 560 for the No Structure condition (7 locations x 8 bumps x 10 repetition). For the abridged experiment, we tested participants on 56 trials for the No Structure condition (7 locations x 8 bumps) and 96 for the Full Structure one (4 lighting x 3 locations x 8 bumps).

4.4 Results and Discussion

4.4.1 No Structure Condition

To determine observers' lighting priors in absence of any cue to light source direction, we tested participants in the No Structure condition (figure 4.3). This setup, composed of six flat grey disks with no local or global lighting cues, was explicitly designed to measure observers' biases for lighting direction. In this case, the six flat disks were lit by a distant diffuse light source, resulting in a uniform and homogeneous light field. Participants were asked to judge if the target disk appeared concave or convex. It is well established that in the absence of directional lighting cues observer will make such judgements based on the shading on the surface under test in relation their internal bias or prior assumption about the light source. When the light parts of the target disk are oriented towards the assumed light source observers will perceive a convex shape. When the darker side of the shaded disk is oriented towards the assumed light source participants will perceive a concavity. Thus the orientation(s) at which convexity is seen most often can be used to infer the assumed lighting prior for any given observer. See for example Adams, Ernst and Graf (2004) for a further example of this procedure. Using a similar method, previous studies have found the observers' lighting prior to be above the observers' head (Adams,Ernst and Graf, 2004; Sun and Perona, 1998; Morgenstern, Murray and Harris, 2011).

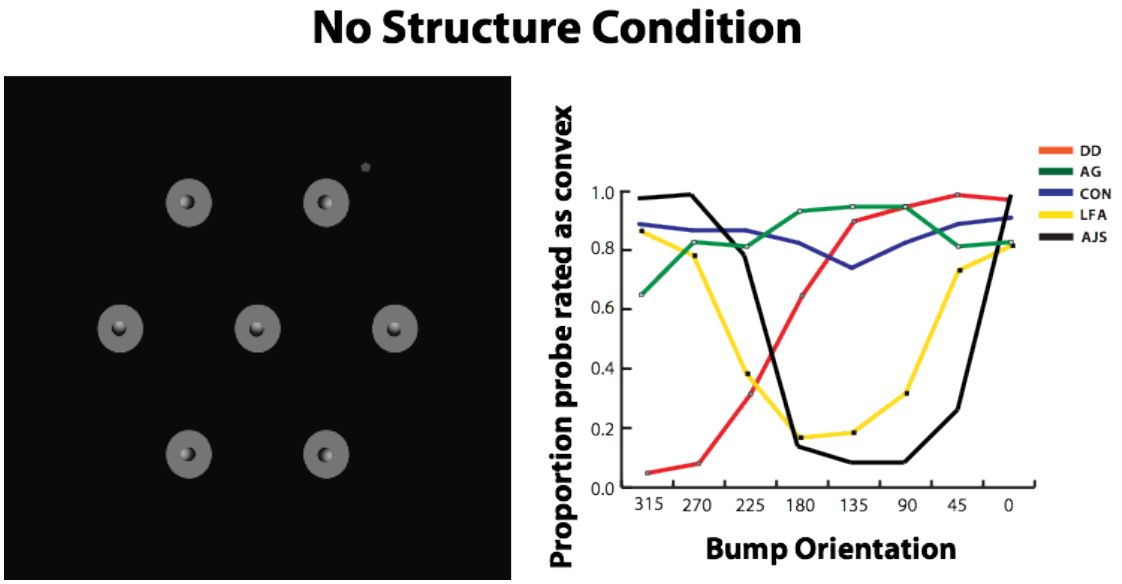

Figure 4.3. Observers' Lighting biases as measured in the No Structure condition. Three participants (AJS, DD, AG) and two groups (CON - preference for convexity; LFA - light-from-above bias).

The results (figure 4.3) were used not only to reveal the observers' lighting priors, but also as a reference to see the influence of global and local lighting fields on the other two conditions (i.e. Full Structure and No Local Cues). Observers AJS and those in the *LFA* group reported a light-from-above preference that was slightly shifted towards the left side while observer DD reported a light-from-above right preference. Observers AG and those in the *CON* group had a preference for convexity under all lighting condition as their bumpiness judgements failed to fall below 65% for any of the possible test patch orientations. For observers AJS, DD and the *LFA* group, the results are consistent with previous studies showing that human observers assume light is coming from above (Ramachandran, 1988; Ramachandran and Kleffner,1992; Sun and Perona, 1998; Mamassian and Goutcher, 2001; Adams, Graf and

Ernst, 2004; Morgenstern, Murray and Harris, 2011). On the other hand, data from participants AG and the *CON* group are consistent with Liu and Todd (2004) showing that some observers tend to have a bias for perceiving objects as convex rather than concave.

4.4.2 Full Structure Condition

In order to investigate how the strength of local and global lighting cues could affect the perception of ambiguous shaded disks, we tested observers in the full structure condition (figure 4.4). In this case, images were depicting a more complex structure that was lit by two strongly articulated light sources (see Stimuli and Materials for more details). Each trial was composed as follows: the central location was lit by a diffuse weak light source as a combination of two fading opposite light sources; three locations lit by an illuminant that could be either from 180°, 225°, 270°, 315° and the remaining three locations lit by an opposite light source (respectively 0°, 45°, 90°, 135°) of the same strength.

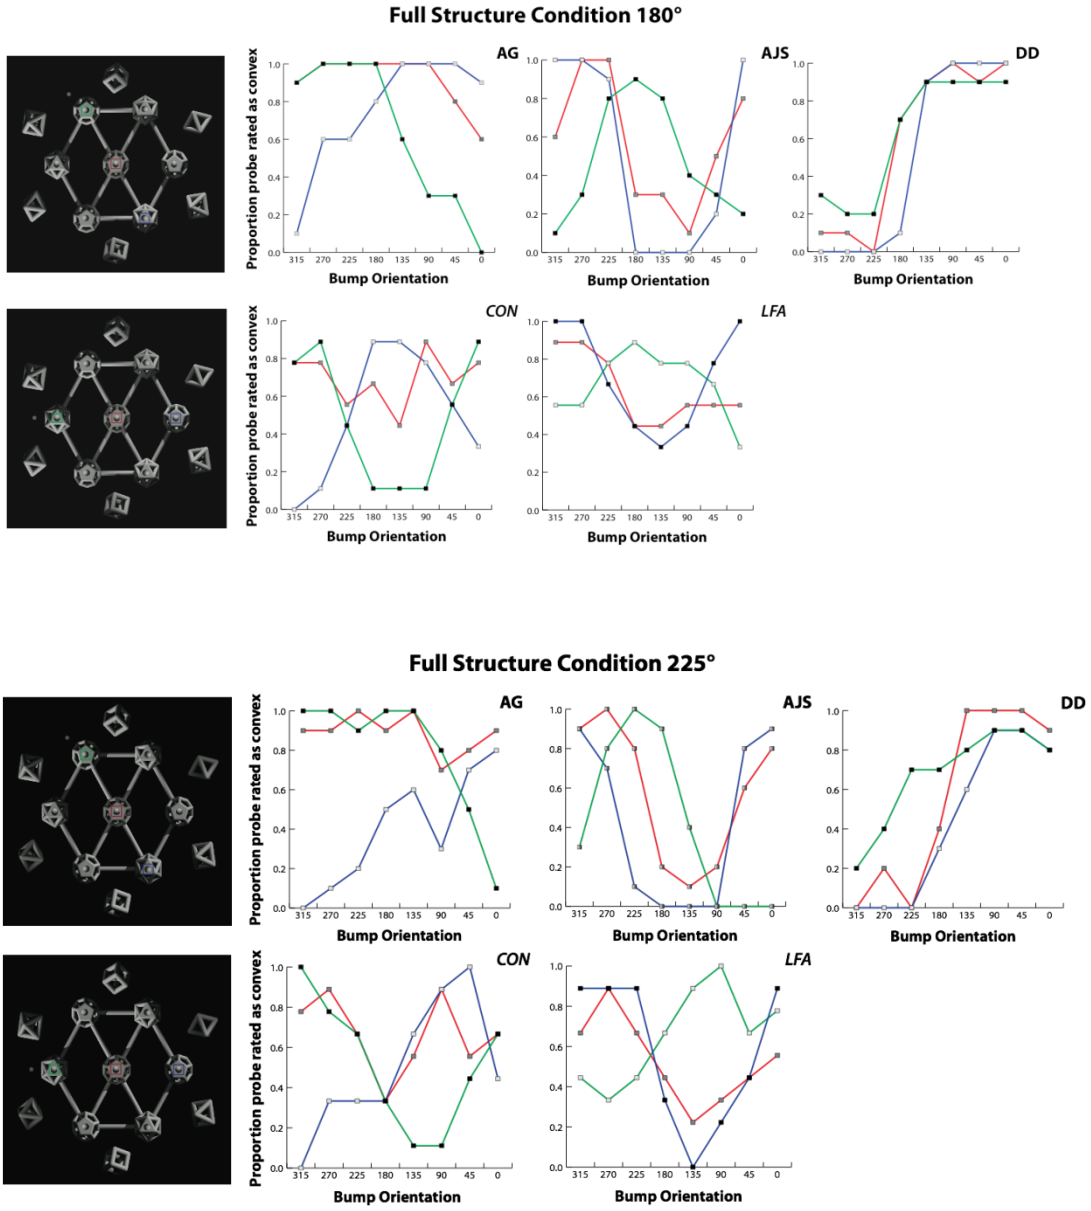

Figure 4.4a. Results for the Full Structure condition 180° and 225° for 3 subjects (DD, AG, AJS) and 2 groups (CON, LFA). Data reported here are from three locations highlighted with Green, Red and Blue squares. Data from the remaining 4 locations are showed in supplementary figure 1.

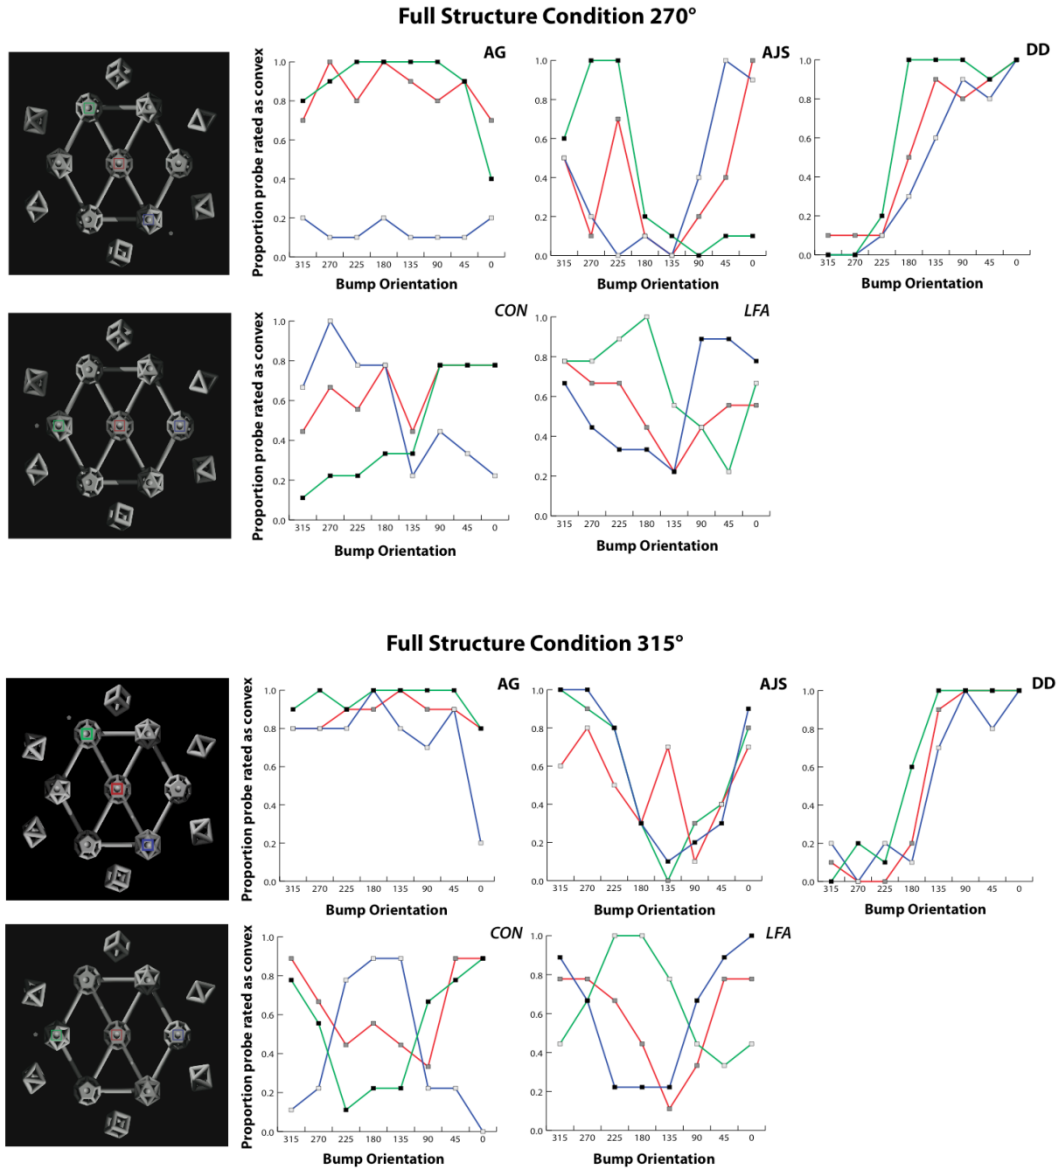

Figure 4.4b. Results for the Full Structure condition 270° and 315° for 3 subjects (DD, AG, AJS) and 2 groups (CON, LFA). Data reported here are from three locations highlighted with Green, Red and Blue squares. Data from the remaining 4 locations are showed in supplementary figure 1.

Morgenstern et al. (2011) show that observers' lighting biases (i.e. light-from-above prior) are weak and their influence is inversely correlated to local lighting

cues' strength. Thus we should find that for centrally located targets, under any light source pairing, observers should provide results similar to the No Structure condition as here the directional lighting cues are weak. In this location, the two light sources are balanced out to give a diffuse lighting impression, therefore any local lighting cue will not be strong enough to override the observer's prior. For the other two tested locations, where local lighting cues were relatively strong, we expect lighting cues to override personal biases according to the position of the illuminant. Results from all participants (figure 4.4a, 4.4b), for the central test location (red line) are very close to the performance they each provided in the No Structure condition, confirming our hypothesis that in case of weak lighting cues the observer's personal biases would dominate perception. Moreover, for each observer, the bumpiness rate is stable across variations of the light source setup, further emphasising the pivotal role of personal biases. On the other hand, as expected, the influence of local lighting cues becomes remarkable for locations were bumps have strong local lighting cues, hence observers' lighting biases are overridden by local lighting cues. Results for all observers (figure 4.4a, 4.4b), show that the peripheral locations (green and blue lines) produce opposite trends consistent with their being bathed by opposing light sources. The influence of strong local lighting cues is reflected on observers' performances depending on the light sources setup for each tested location. In particular, AJS and the *LFA* group - that is those observers with a light from above-left bias - have very similar distributions for all lighting condition. There is a shift in the apparent illuminant direction, indicated by the

position of the bumpiness peak, toward the direction of the local illuminant. For participant DD, local lighting cues enhanced performance in the case where the lights setup is similar to the observer's bias (i.e. blue line for 180° and green for 270°), but they negatively affect the performance when the illuminant position is opposite to the observer's prior (i.e. green for 180° and blue for 270°). Figures 4b show special cases for single observers' condition 315° and groups condition 270°. In particular, for both conditions, the hexagonal shape of the "space station" let the three locations reported in the graphs fall in a spot where the two directional lights are fading out. These circumstances let the test location have weaker local lighting cues that are more similar to diffuse lighting. The data reported in the graphs are strongly affected by the type of local lighting cues since blue, green and red lines for single observers tend to overlap. The same does not happen for groups data. We believe that this may be due to inconsistencies within participants in each group. In other words, the data could be very noisy due to the internal variance of the group. However, the results from the Full Structure condition do not completely clarify how the visual system interprets bi-directional (ambiguous) lighting cues. The results here reported suggest that observers tend not to perceive the scene as a whole but they rather prefer to rely on local lighting cues. In presence of strong local lighting cues, in fact, the information provided is salient enough to drive the observers' perception and at least partially override their internal biases. To test whether global lighting cues are still dominating in cases where strong local cues are removed, we tested participants in the No Local cues condition.

4.4.3 No Local cues condition

To better understand how the visual system interprets bi-directional ambiguous lighting information, we tested observers in the No Local cues condition. In this case the scene was deprived of local lighting cues (see figure 4.5) while the observer's task was identical to the previous conditions. Previously, in a similar setting, Morgenstern et al. (2011) showed that deleting local cues to lighting direction had little effect in a shape-from-shading task. For instance, they found no statistical difference on observers' performances in a scene with or without strong local lighting cues. They conclude that global lighting cues were overriding observers' lighting biases.

As explained before, the use of two illuminants is crucial to our experiment. By depriving observers of local lighting cues, we left them with a very ambiguous and unnatural lighting setup: the solids next to the tested location were providing weak cues to a light source with a given direction while the one illuminating the rest of the scene was stronger but distant and from the opposite direction. Based on Morgenstern et al's (2011) results, the human visual system is able to make use of the lighting information available even if they are locally corrupted. We predict that when the strong local cues are removed the remaining but distant lighting cues will dominate the observers' percept.

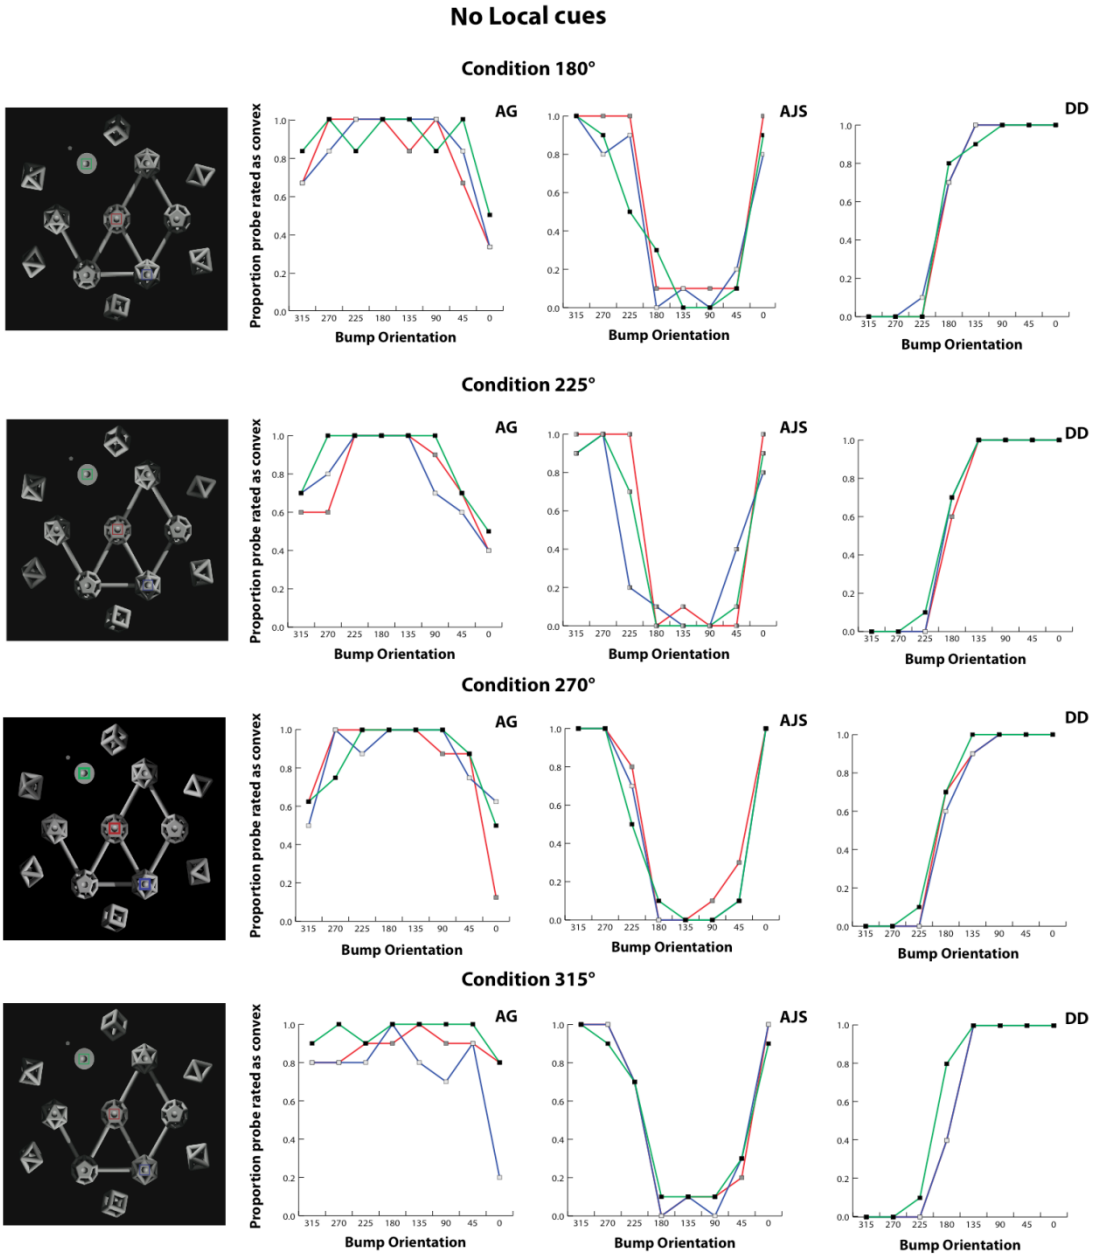

Figure 4.5. Results from the No Local cues condition for 3 subjects (DD, AG, AJS). Data reported here are from three locations highlighted with Green, Red and Blue squares. Data from the remaining 4 locations are showed in supplementary figure 2.

The results obtained from all participants in the No Local cues condition (figure 5) show that observers have a constant percept across lighting direction variations for all the tested locations. For three observers data distribution

shows very little variance: blue, red and green lines are strictly overlapping. Since changing the direction of the light-sources does not affect the performance, it seems like the visual system either ignores any global lighting cue or that the two lighting cues (locally weak and distally strong) are averaged out to form the impression of a diffuse source which is then not strong / directional enough to override any internal bias. In conclusion, when multiple light sources are present but local lighting cues are weak or absent, the lighting prior dominates even if other parts of the scene present strong lighting cues. Ecologically this might occur when a workspace is illuminated by a relatively overhead light and a directional work-lamp. Areas lit from overhead only will have weak directional lighting cues as this light is distal and in any case quite diffuse. Areas close to the work lamp will have a directional lighting profile but this profile will not affect the perception of objects in the diffusely lit area.

Taken together our results suggest that when lighting cues are fully articulated locally (full structure condition) people tend to use those lighting cues to aid shape from shading. Conversely when local cues are removed (no local cues condition) observers are strongly influenced by their prior assumptions for lighting and convexity. However comparison between the full-structure and no local cues conditions is made difficult by large individual differences in behaviour. For example AJS, DD (and the LFA subject group) are strongly influenced by a lighting from above prior whereas AG (and the CON subject group) are more influenced by the convexity prior. In terms of variations in the weight given to such priors AJS places a lower weight on his prior when good

lighting cues are present whereas DD places a strong weight on her prior at all times. Nonetheless it is clear from the figures presented above that shape judgments are more consistent and closer to the individual's prior (as established in the no cues condition) when local cues are absent than they are in the full structure condition; suggesting that prior assumptions are less dominant in the latter case. We sought to quantify these differences in two ways. First we calculated the correlation coefficients between the 'proportion bump' scores for each of the lighting conditions test positions correlated with the no cues conditions. That is we compared the individual results curves of figures 4.5 and 4.4 with those of figure 4.3. If observers are responding according to their prior assumptions then we would expect such correlations to be strong. If they are influenced by image cues we would expect weaker correlations. We then compared the mean correlations across all lighting positions and text positions for each observer. For the Full structure condition mean correlations were: AJS=0.47; AG=0.25; and DD = 0.96. These results reflect the use of images based lighting cues by AJS and AG and a greater reliance on prior assumptions by DD. For the no local cues condition mean correlations were: ASJ = 0.97; AG=0.49; and DD=0.98. These increased correlations show that AJS and AG clearly on prior assumptions more in the no local cues condition than the full structure condition (DD was already close to ceiling).

To further quantify the changes between the full structure and no local cues conditions we calculated the absolute difference in 'proportion bump' scores between each person's scores under their prior assumptions (given by the no

cues results) and those under the full structure and no local cues conditions. Such difference scores we calculated for all disk orientations at each test location and for each lighting position. Grand means (see table 4.1.) were then calculated collapsing across disk orientation, test position and lighting position to yield a single mean difference score for each person in each of the two structure conditions (full structure and no local cues). Predicating that difference scores would be higher in the lower in the no local cues condition (where people respond in line with their prior assumptions) and higher in the full structure condition were peoples results are more affected by image cues we found that the difference between conditions approached significance (one tailed paried t-test, t=2.14, df=3, p=0.08). People do seem to respond according to their prior assumptions more in the no local cues condition than they do in the full structure condition.

	Structure Condition	
Participant	***Full Structure***	***No Local Cues***
AJS	0.28	0.08
AG	0.26	0.12
DD	0.11	0.1

Table 4.1. Mean differences between the 'proportion bump' scores for 'structure' vs the 'no cues' conditions averages across lighting, test position and disk orientation.

4.5 Conclusions

The role of lighting priors in shape-from-shading studies is crucial to understanding the strategies used by the visual system to interpret ambiguity. In the past, many studies have used bas-relief ambiguities in the stimuli itself to reveal the assumptions made by observers about the position of the light source – the typical finding being that observers assume a directional light source from above their own head. More recent studies have discovered that light-from-above prior is weak and can be easily overridden by relatively weak lighting cues. This study furthers our understanding of visual processing in ambiguous bi-directional lighting. Specifically, what happens when the brain has to perceive stimuli embedded in an ambiguous light field for example when there are multiple spot lights in a room, or competing directional light sources in the form of a window, a desk-lamp and a computer screen. Morgenstern et al. (2011) suggested that global lighting cues have an influence on perception although this effect is smaller than that induced by local cues. What happens with two strongly articulated but opposing light sources? Will the percept be influenced by the strong local one or will be the bi-directional global one (which is ambiguous) dominate, leaving the prior unaffected? Our results show that in presence of two well articulated opposite light sources (Full structure condition), observers are influenced by local lighting cues only and they ignore more distal cues. In cases where scenes are deprived of local lighting cues (No Local cues

condition), we found that observers tend to perceive the stimuli as if there was no cue to light direction, that is in accordance to their personal biases. A possible explanation is that given the ambiguity of lighting cues, the visual system tends to ignore them and therefore relies heavily on lighting priors. Alternatively, it could be that the visual system merely averages out the remaining, opposing lighting cues (weak local cues from one direction and strong distant cues from the other) and consequently the No Local cues condition is treated like a diffuse lighting case. Given the results from Morgenstern et al (2011) we would favour the second explanation but we have no direct evidence to for it. Future studies should therefore consider in more detail what the role of global cues is, and by which means the human visual system weighs the influence of local cues to light direction and lets them dominate the observer's percept. In summary we have extended the work of Morgenstern et al (2011) by adding a second directional light source. In the full scene condition observers' shape judgements were governed by the local lighting cues and more distal cues – even tough strong and well articulated were largely ignored. When the local cues were removed – by removing scene structure close to the probed location – observers reverted to their default lighting assumption rather than use the strong but more distant cues present in the image. This is interesting since in this case there were strong directional lighting cues present in the image and we might assume that these would override the lighting prior as this has been shown to be weak (Morgenstern et al 2011). Of course our non-local cues condition retains cues to both light sources,

its just that cues to the more proximal light have been weaken by removing the local scene elements, this residual ambiguity may be enough to allow the lighting prior to dominate. We know for example that when directional lights are equally balances as in the central probe location the prior dominates and we assume this to be the case whenever lighting cues are ambiguous.

4.6 Supplementary figures

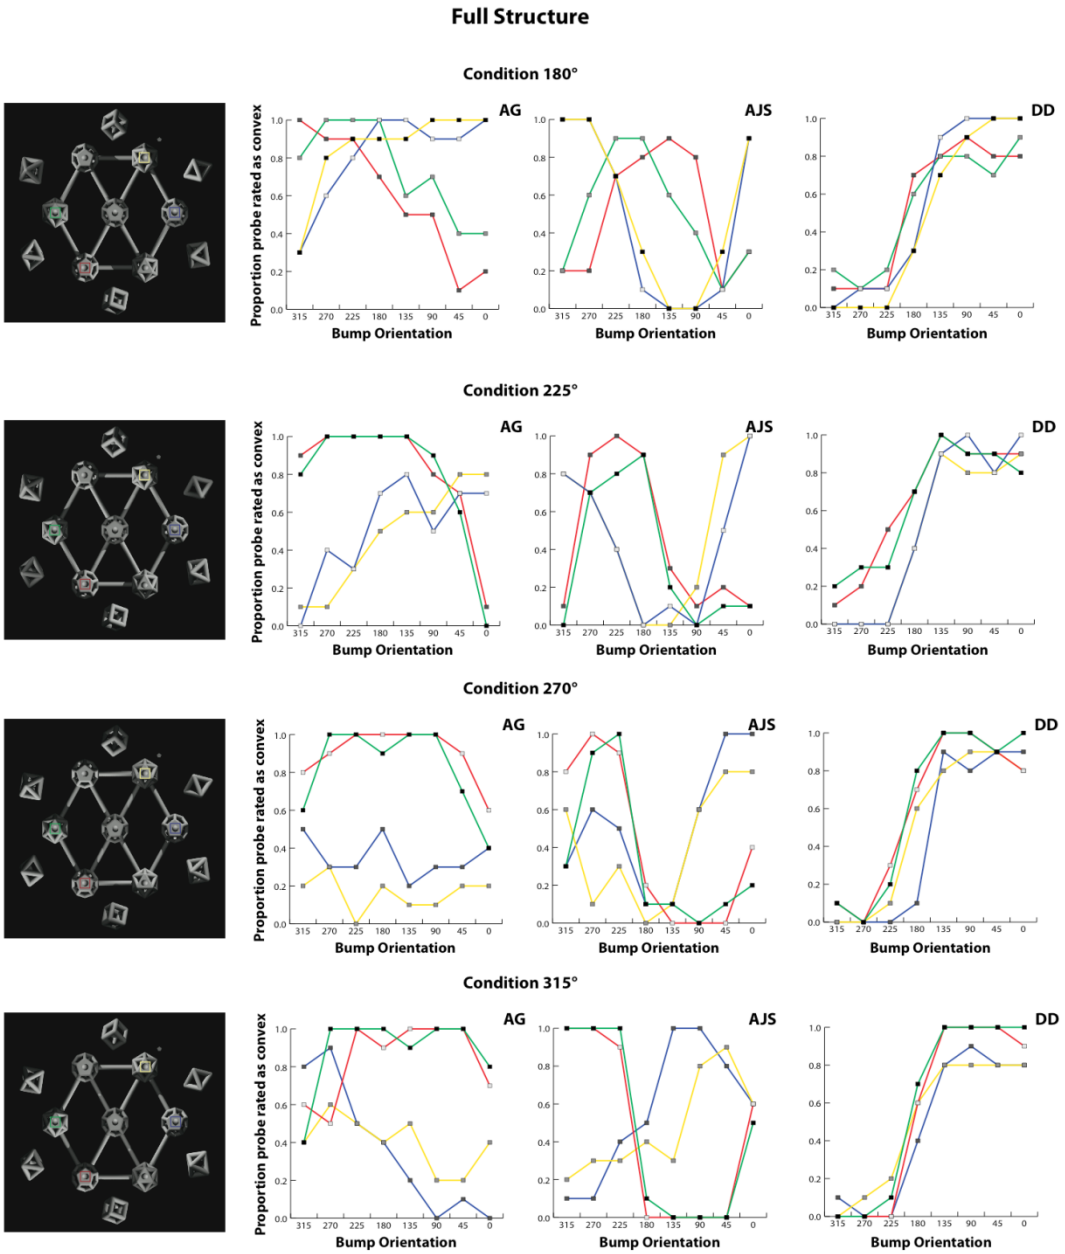

Figure 4.1S. Results from the Full Structure condition for 3 subjects (DD, AG, AJS). Data are relative to locations missing in the main text figures.

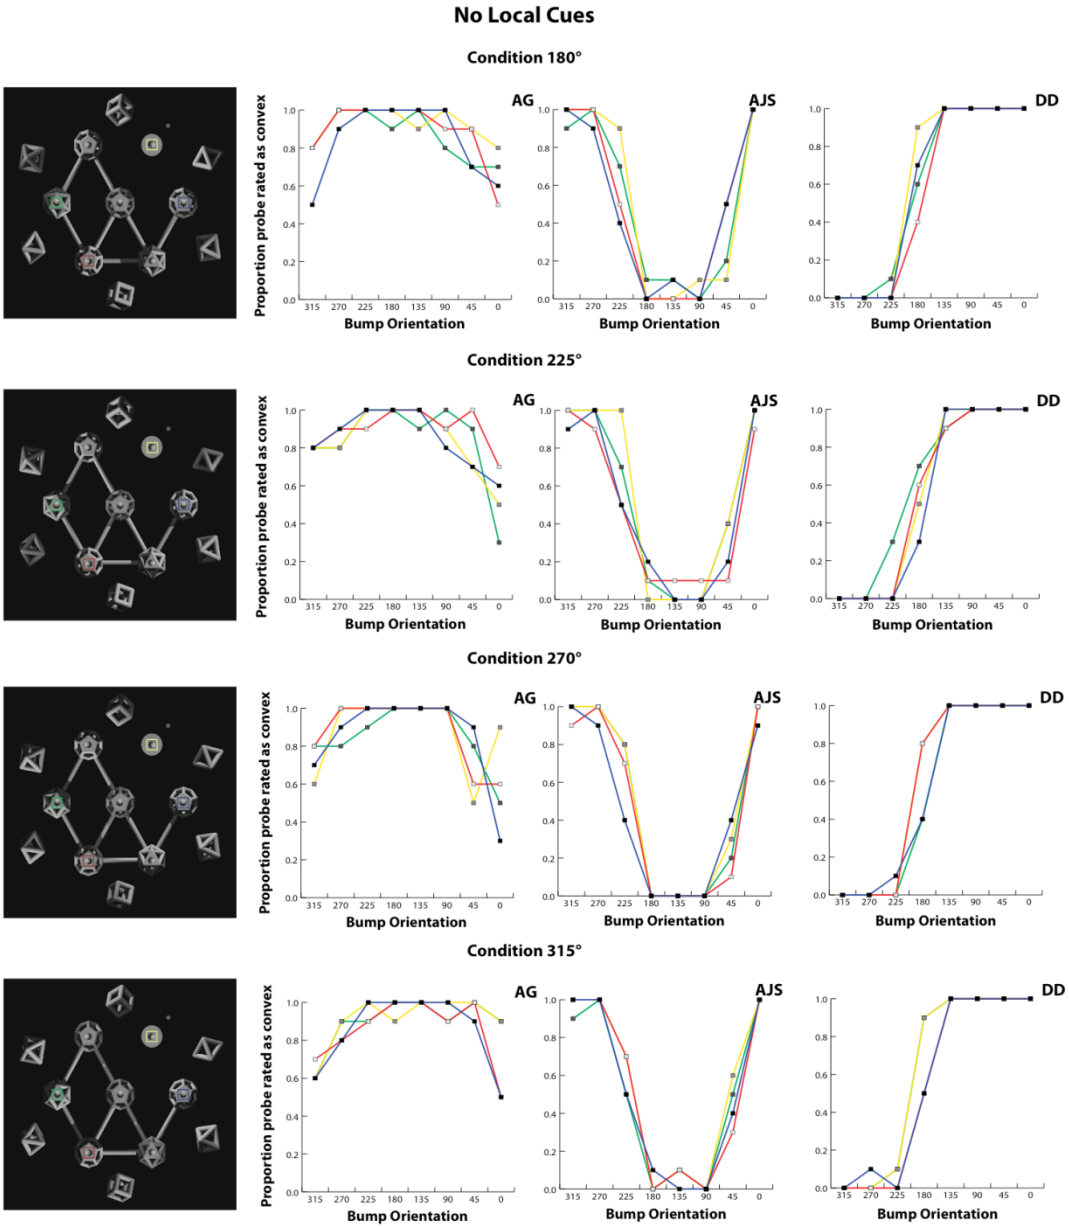

Figure 4.2S. Results from the No Local cues condition for 3 subjects (DD, AG, AJS). Data are relative to locations missing in the main text figures.

5. The Role of Shadows on Visual Search

So far this thesis has considered three aspects of the human visual systems response to illumination. In chapters 2 we considered the strength and prevalence of the lighting from above prior in human vision. In chapter 3 we considered how this prior interacts with lighting cues provided by a range of surface types, and, in Chapter 4, we considered the role of local and distal cues to illumination and there separate influence over the lighting from above prior. In this chapter we consider a rather different aspect of illumination, specifically the treatment of shadows within the visual system. Some studies have found that shadows are treated as objects while others found that they are ignored by the visual processing system. This latter position is intuitive in that our everyday experience of shadows is that we don't process them as objects or combine them into the objects that they may happen to be attached to. Attention paradigms have been used to clarify the role of shadows by comparing the observers' reaction times between within object and within shadow conditions. Here we replicate and extend the work of de-Wit, Milner, Kentridge (2012) using a new condition of object-like shadows. Our results are consistent with the idea that shadows are not treated as objects nor segmented as distinct regions in object-based attention paradigms.

5.1 Introduction

The human visual system makes use of many sources of information to discriminate between objects and shadows. In particular, small changes on the illuminant position, contours or shape may affect the perception of objects or surfaces. For instance, early level features such as surface convexity (Kleffner and Ramachandran, 1992; Ramachandran 1988), position of the light source (Sun and Perona, 1998; Mamassian and Goutcher, 2001) and contour edges (Knill, 1992; Sun and Schofield, 2012) are important cues to shape. The role of shadows in object recognition is somewhat ambiguous. Shadows seem to play an important role in the discrimination of object movement (Kersten, Knill, Mamassian, Bulthoff,1996), they can contribute to correctly estimate objects' distances (Allen, 1999) and provide useful information to perceive the 3-dimensional layout of a visual scene (Madison, Thompson, Kersten, Shirley, Smits, 2001). Nevertheless, shadows show little influence in certain contexts such as the identification (Bonfiglioli, Pavani, Castiello, 2004) and recognition (Braje, 2003) of real 3-dimensional objects.

Previous studies have tried to determine whether or not shadows are treated as objects by measuring their influence on visual search tasks (Rensink and Cavanagh, 2004; Lovell, Gilchrist, Tolhurst, Troscianko, 2009). In a series of experiments Rensink and Cavanagh (2004), investigated the role of shadows and shadow-casting objects by changing their borders, texture and background. They found that the visual system can rapidly identify a region as a shadow but

then relegates its status to some degree relative to objects. For instance, a region that is not interpreted as a shadow can still be available for rapid search while shadow regions slow down any search task. Lovell et al (2009) replicated their experiments for visual search with a variation that was consisting in an increment of the angle between target and distractors (from 30deg to 90deg). Their findings suggest that the discrepancies in terms of angular distance between target and distractors may be a crucial parameter to define the ability of the visual system to process shadows and objects in visual search tasks. More specifically, they found that reaction times decreased proportionally to the enlargement of the angle between target and distractors. They argued that the differences found in respect to Rensink and Cavanagh (2004) results were mainly due to shadows being represented by the visual system at a more coarse scale. The quickness and the reliability of this spatially coarse shadow discrimination mechanism, by definition, is largely affected by the ability to detect shadows (ie bigger angles between shadows and objects implies faster reaction times in discrimination). However, Lovell et al. (2009) affirm that the visual system simply “cannot use shadow information for spatially fine tasks”, therefore, they agree with the notion that shadows are coarsely represented.

Few studies in object recognition have shown that the perceptual organisation at the base of object representation is mainly due to the way boundaries are interpreted rather than by the contrast of those boundaries themselves (Anstis, 1990; Naber, Carlson, Verstrate, Einhauser, 2011, Scholl, 2001; Tadin, 2002; Albrecth, 2008). Object representation is important in attention paradigms since

it has been demonstrated that visual search tasks are object-based rather than spaced-based (Egly, Driver, Rafal, 1994). In their study, Egly et al (1994) asked observers to fixate a display containing two rectangles and a fixation cross in which: first a cue was shown to inform participants about the location where a target could pop up, then a target was presented. They then compared reaction times for detecting targets appearing at the cue location (valid), at a different location within the same rectangle (invalid-within) and a different location in the other rectangle (invalid-between). They found that observers were faster at detecting an invalidly-cued target when they appeared *within* the same object rather than on different objects even when are both at the same distance from the cued location. This advantage does not occur when the 'objects' are perceived to be holes cut into a surface rather than objects on top of the surface (Albrecth, List, Robertson, 2008).

To sum up, shadows have an influential role in shape and object perception (see Mamassian, Knill, Kersten, 1998), but their interpretation by the visual system is as yet not clear. Some studies claim that they are interpreted as object (Lovell et al 2009; de-Wit et al., 2012) and processed accordingly; others have shown that their processing is somehow downgraded relative to objects (Rensink and Cavanagh, 2004).

Our experiment is based on de-Wit et al. (2012) approach. They tested stimuli in which shadows were presented with or without a casting object. They tested conditions were a target was presented in one of four possible location and it was preceded by a location cue that did not always correctly indicate the

location of the subsequent target. Their results confirmed that reaction times were not influenced by the presence of casting objects but only by the location of the preceding cue. Specifically they found that targets presented within the same shadow as an invalid cue were detected faster than targets presented in a non-cued shadow region. These results are in agreement with object-based attention (Egly et al, 1994). De Wit et al argued that in the absence of a casting object the 'shadow' region should appear less shadow-like and more object-like. Thus they concluded that objects and shadows are processed similarly because the presence of casting objects made little difference to their results. However, they always used soft edged luminance decrements as 'shadows' and such edges provide a strong cue to shadow-hood. Anecdotally, humans can recognise shadows as such even when casting objects are absent. We wondered if sharp edged luminance decrements, which are more object like, would be treated in the same way as soft-edges regions and clearly cast shadows. Thus we added a new condition to de Wit et al's (2012) experiment in which 'shadow' regions were presented without casting objects and with sharp (well-defined) contours that gave a stronger impression of object-hood. Further, we presented our shadows on a textured background where the texture was still visible within the shadow regions. This manipulation enhances the perception of shadows. In the sharp edged 'shadow' condition we set the shaded region to uniform luminance (no texture) and thus further promoted the impression that these regions were objects rather than shadows.

5.2 Materials and Methods

5.2.1 Participants

Twenty-one participants (9 female; age range 18-56, mean age 30) with normal or corrected-to-normal vision took park in the experiment.

5.2.2 Apparatus

Stimuli were presented on 20inch Sony ViewSonic CRT monitor from a viewing distance of 57cm. The screen resolution was set to 800 by 600 pixels and image size was 30deg x40deg. The experiment was created in E-Prime (Psychology Software Tools, Sharpsburg, PA).

5.2.3 Stimuli

The stimuli were images rendered in PovRay (Persistence of Vision Pty. Ltd., Williamstown, Victoria, Australia. http://www.povray.org/). They depicted a 3D scene comprised of two rectangular objects lit by a distant light source from above and behind the objects. The shadow casting objects could be present (Basic condition) or absent (No Object and Sharp) while the shadows could have natural faded borders (Basic and No Object) or sharp contours. For the Sharp condition depicted shadows were also deprived of any background texture.

5.2.4 Procedure

In each trial, observers had to report the presence of a letter that could be either X or N. A cue appeared before the target-letter at one of four possible locations. There were three types of cue: Valid, where the cue appeared at the same locations as the following target-letter; invalid-Within, where the cue appeared within the same object as the letter but in a different location; and Invalid-between where the cue appeared in a different object to the letter. We tested observers in three conditions differentiated on the basis of Shadows type (figure 1): 'Basic', shadows were presented with casting objects; 'No Object', shadows were presented without casting objects but with soft edges; and 'Sharp' shadows were presented without casting objects, a sharp border and with no background texture in the shaded region.

Figure 5.1. From left to right, stimuli examples for Basic, No Object and Sharp shadows type.

On each trial (figure 2), a cue (a small place blue square) was presented in one of four possible locations for 250ms, subsequently a grey background was shown for 200ms before the observer was shown the target image.

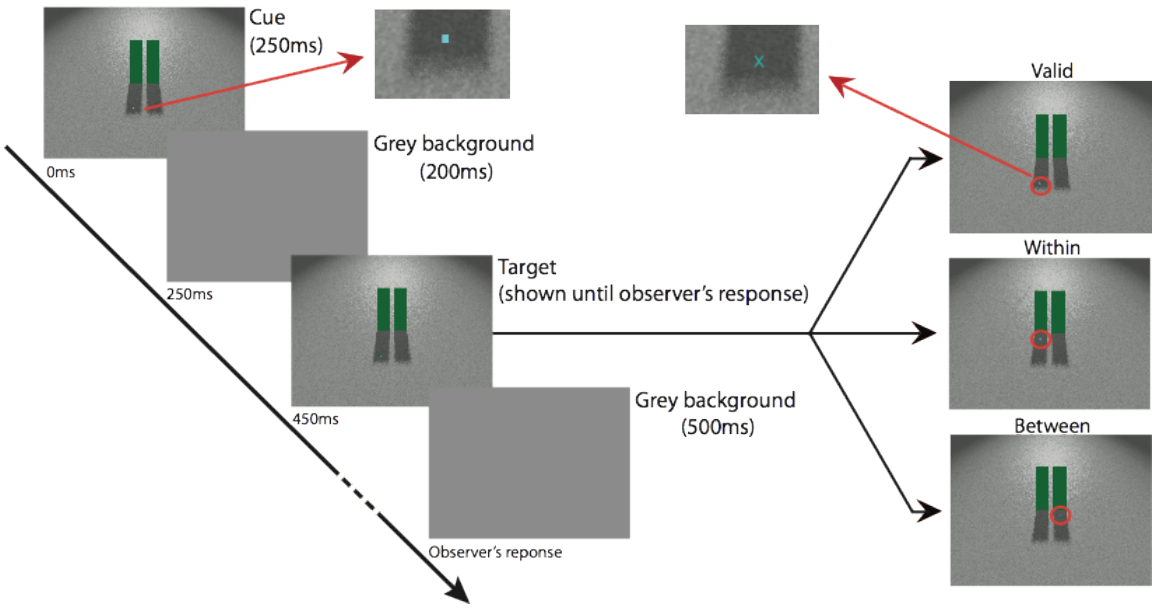

Fig 2. Experiment Design and Timeline

There were two types of target images (letter X or N) and cue locations were valid 33% of the time. Observers were asked to pay attention to the location of the cue and to report which target had been presented by pressing a key (X or N) as accurately as possible. They were aware that reaction times were being recorded. After the observer's response, a grey background was shown for 500ms before the next trail commenced.

Each participant undertook two sessions of 216 trials each (3 Shadows type x 4 Cue locations x 3 Cue Type x 2 Letters x 3 repetitions). Before the start of the experiment they undertook practice trials until confident with the task.

5.3 Results

Reaction times data were at first filtered to avoid anticipation error data and outliers; both correctly and incorrectly answered trails were used. RT's below 250ms and RT's above 2.5 Standard Deviations from the condition's mean were discarded.

A two-factor repeated measure ANOVA revealed a clear effect of Cue type ($F(2,20)=22.053$, $p < 0.0001$) but no effect of shadow type ($F(2,20)=0.305$, $p = 0.739$) on reaction times. The effect of cue type was dominated by the comparison between valid and invalid cues. The interaction of Cue type Vs Shadow type was not significant ($F(4,17)= 1.359$, $p =.256$).

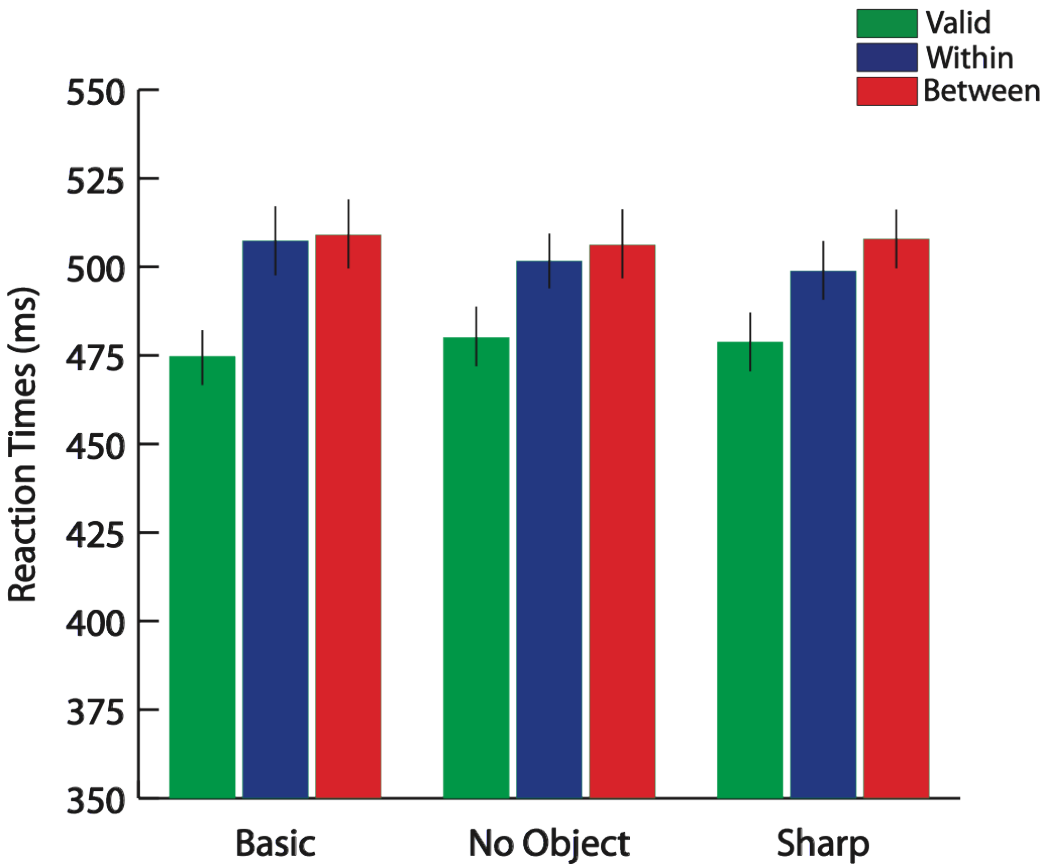

Figure 3. Reaction times results averaged across participants for Basic (shadows with casting object), No Object (shadow without casting object) and Sharp (shadows without casting object and with sharp edges). Each column colour corresponds to a different cue type and error bars represents standard mean error.

Accuracy data revealed no significant influence of cue type ($F(2,20)=1.268$, $p = 0.292$) although there was a non-significant trend for shadow type ($F(2,20)=2.590$, $p = 0.088$). The accuracy rates were: 95.9%, 95.2% and 94.8% for Valid, *Within* and *Between* cues respectively. For Shadow type, the accuracy

rates were respectively 95.5%, 94.6% and 95.8% (Basic, No Object, Sharp). Overall we conclude that there was no speed accuracy trade off.

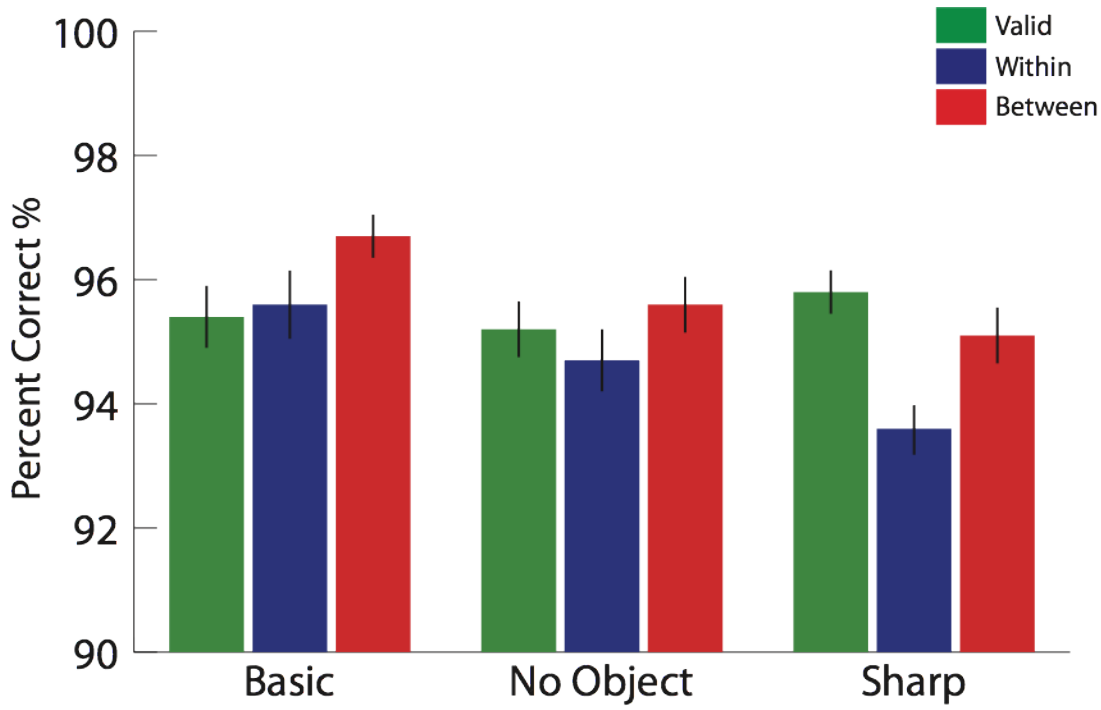

Figure 4. Accuracy results averaged across participants for Basic (shadows with casting object), No Object (shadow without casting object) and Sharp (shadows without casting object and with sharp edges). Each column colour corresponds to a different cue type and error bars represents standard mean error.

5.4 Discussion

We used an object-based attention paradigm (after Egly et al 1994) to investigate the role of shadows. The aim of our experiment was to use object-based attention paradigms in order to assess whether shadowed regions (pure shadows, shadows without casting objects and shadows with sharp contours and solid fills) are represented in a way similar to objects by the visual system. Previous studies have found divergent results: some claim that shadows are discounted to some extent (Rensik and Cavanagh, 2004); others claim that they are not discounted but rather processed at a distinct course spatial scale (Lovell, 2009; de-Wit, 2012).

Our results from both accuracy and reaction times confirmed no overall influence of shadows type (Basic, No Object and Sharp) on observers' perception, suggesting that all the tested conditions were treated similarly. On the other hand, Cue locations (Valid, Within, Between) influenced reaction times significantly and proportionally to their distance from the cue location: Valid cues produced faster reaction times, while Non-valid cues (Within and Between) produced slower responses. The comparison of within against between object cues did not produce any significant effect in conditions where shadows were clear, namely Basic and No Object. Our results are in Contrast to studies reporting a difference in reaction times between within- and between-object cueing for shadows (Egly 1994; de-Wit, 2012). We found no such within object benefit and only a slight hint of such a difference even in our sharp edges

condition. Our observers did not seem to treat shadows as objects and thus did not show a within object benefit. Following to Rensink and Cavanagh (2004), these results suggest that shadows are discounted even in cases where shadows are presented without casting objects.

Our study presents 2 rather curious results. First, the basic and no-object conditions are essentially identical to the similar conditions use by de Wit et al (2012) yet we failed to replicate their results. Second, our additional sharp edges condition – which was meant to represent object hood more clearly only produced a hint of the expected within object bias – and one which we must regard as unreliable due to the lack of an overall interaction between our cue and shadow types. Thus we have failed to replicate either the expected within object bias or its translation to shadows. Nonetheless our finding of a valid cue benefit suggests that observers understood the task and were influenced by exogenous cues to attention.

One main difference between our study and previous ones is the 3-dimensional nature of the depicted scenes: while both Rensink and Egly papers presented stimuli on top of flat grey background, our stimuli were presented in a scene with strong perspective cues (figure 1) and a background texture for both Basic and No Object conditions. It could be possible that observers were not misled in the perception of target regions as they were always placed on the ground, as a consequence, they were always considered shadows (see Albrecth, 2008) this impression might have extended to the sharp conditional also. However our stimuli were similar to those of de Wit et al so our failure to replicate their basis

finding remains a mystery. Had we found a strong within object benefit for the sharp condition we might have concluded that the introduction of this more object like condition encouraged the observers to see the other conditions as more shadow like and thus removed the within object benefit from those conditions. As it is however we cannot make that assertion.

Replications are important in science. Since we can see no methodological error in our approach – which was in all respects similar to that of de Wit et al (2012) we must conclude that their result is less robust than might be supposed. Perhaps there were differences in the observers used – ours had very little experience of the paradigm at all and no prior expectation in favour of a within object benefit. It would seem that further replication attempts are necessary for the de Wit et al study. Once the reliability of that result has been established perhaps then a better, more object-like condition can be added for comparison as was our intention here.

Further studies are also necessary to specify the role of shadows and in particular to determine the conditions in which they are treated as objects. A study in which the within object benefit is measured for clearly defined object regions as well as clearly defined shadow regions and more ambiguous shadows is required. A consideration of spatial scale would also be beneficial. Such manipulations would be useful not only to model how the visual system processes shadows but also to reveal how vision interprets objects and surfaces that stand in the shade.

6.Discussion

Perceiving 3-dimensional shapes is one of the most challenging computations for the human visual system, yet our brain is able to form a stable perception of the environment making full use of the information at his disposal. One of the major problems is the need to disambiguate the pattern of light that is projected into the retina in order to have a veridical (or at least usable) perception of the world.

The central aim of this thesis was to clarify the role of illumination in the process of Shape-from-Shading (SFS). The role of illumination can be considered from three perspectives. The influence of in built lighting assumption on perceived shape, the indirect influence of such assumptions on perceived lighting and hence shape, and the role of stimulus properties in determining perceived lighting and hence shape. From a psychological point of view, the first and most important step was to reveal the influence of prior assumptions for light source direction on visual processing. In particular, we wanted to determine how the position of the light source and the type of illumination contributed to the interpretation of ambiguous shaded stimuli. In doing so, we also considered the role of shadows and the physical properties of the object itself such as its isotropicity, orientation and contrast.

6.1 The prevalence of the lighting from above prior.

There is considerable evidence that when viewing ambiguously illuminated stimuli, humans assume a light source that is above their own head (Adams, Graf and Ernst, 2004; Brewster, 1826; Gerardin, Kourtzi and Mamassian, 2010; Mamassian and Goutcher, 2001; Ramachandran, 1988; Rittenhouse, 1786; Sun and Perona, 1998; Todd, 2004). Although it has been shown that the lighting form above prior can be slightly but significantly modified by experience (Adams, Graf and Ernst, 2004) the bulk of the literature suggests that most individuals have a lighting from above prior. However, these experiments have typically used stimuli that could bias the participant towards a directional lighting assumption. A few studies have shown that the lighting from above prior may be weak and easily overridden by stimulus cues or indeed other more potent prior assumptions. For example, Tyler (1997) suggested that the default illuminant may be diffuse rather than directional. Although he used a stimulus that could not support a point source interpretation, these results suggest that the stimuli used for testing lighting priors is important. In a similar contest, Schofield et al (2011) used a relatively neutral non-rendered sinewave stimulus and found that the default lighting assumption is a mix of diffuse and directional components. Morgenstern et al. (2011) shows that this prediction is very close match to the degree of diffuse vs. directional illumination in the natural environment. If we take an overall look at studies of the lighting prior, we can see that they have often used stimuli which could bias the participant to interpret the light source as

directional rather than diffuse. It may be then that under these constraints, the 'above' direction dominates even though observers' lighting prior as a whole is not especially directional.

In Chapter 2 we sought to test the validity of the lighting from above prior using a novel stimulus free design in which participants were asked to 'detect' the presence of a bump in random noise stimuli. However no bumps were ever actually presented. Nonetheless our accumulation images (classification images) seemed to contain a bump stimuli even though none were present. The final Template, or accumulation image, represented the observers' classification image or expected template for a bump stimulus. Such images revealed that most observers do expect bumps to have distinct shading highlight as predicted by a directional light source, but the location of this highlight relative to fixation varied considerably across observers, and was not always consistent with lighting from above. The validity of this result was affirmed by asking participants to detect other stimulus features. Highlights in the classification image moved further from fixation for when observers were asked to detect larger bumps; they became elongated in one direction only when observers were asked to detect a cylinder, and they became enlarged and somewhat flatter when observers were asked to detect a flat white disk.

We thus conclude that the apparent dominance of the lighting from above prior may be less prevalent in the population than might be supposed from the literature. We speculate that stimulus factors contribute to the impression of a directional light source and that this in turn promotes the lighting from above

prior. There could be no stimulus induced bias in our experiments: we tested the lighting prior in the most basic sense by deriving each participant's internal bump template from cue free random noise images. To determine the stimuli, and consequently the final template, observers had to rely on their internal template as supported by their lighting prior. This is not to say that we think the lighting from above prior is instantiated in such templates or that such templates are fixed. Rather we suppose merely that so far as templates are establish for complex detection tasks, those templates should be conditioned by ones prior assumptions about the appearance of stimuli. Our results suggest that people expect bumps to have a shading highlight (consistent with directional light source) but that the location of this highlight is highly idiosyncratic suggesting that frontal lighting and lighting from below maybe as common as the lighting from above assumption.

These results add to the growing literature suggesting that the lighting from above prior may be less prevalent and more easily manipulated or overridden than has previously been supposed.

6.2 The effects of light source elevation

Changes in illumination can determine shape perception systematically, as a result, recovering shape-from-shading is not independent of the lighting conditions used in forming the image. Observers can sometimes attribute alterations in the perceived shape to changes in the surface profile rather than

changes in the light source position. Most studies have considered the direction of the light source in terms of its azimuth – the position as if around a clock face from which the light is directed (Koenderink, Pont, vanDoorn, Kappers & Todd, 2007; Koenderink, vanDoorn & Pont, 2004). Typical findings are that when the surface under study is relatively isotropic, luminance features and surface shape vary considerably with the direction of the illumination. On the other hand, for anisotropic surfaces the orientation of the undulations dominates producing surface interpretations that are more stable (up to an inversion of the relief direction).

Relatively few studies have considered the elevation of the light source - the degree to which is it tilted out of the stimulus plane. Previous studies (Johnston and Passmore, 1994; Curran and Johnston, 1996) have reported that varying the elevation of the light source can affect surfaces' curvature perception, specifically decreasing slant of the light position (towards the viewing direction) caused the surfaces to be perceived as more flat than when lit obliquely.

In chapter 3 we tested shape from shading for surfaces illuminated at different elevations using three surfaces with varying degrees of surface anisotropy. In order to establish the role of image contrast, we tested shape from shading when image contrast was allowed to vary as it did in the original photographs and, also, with normalised image contrast (such that overall contrast was the

same for all illuminant elevations). Surface shape was tested at a wide range of surface locations.

We have found that when contrast was allowed to vary with illuminant elevation, grazing lights, which produce stronger luminance gradients, tended to produce larger slant settings whereas more frontal lighting produced smaller slant settings. In the cases when contrast was normalised slant settings converged towards the average value, highlighting the role of image contrast in determining slant. Given that the physical slants of the original surfaces did not change with light direction, and image features such as the presence of shadow edges were not removed by the normalisation process, these results highlight the dependency of shape-from-shading on basic image properties such as luminance and contrast.

Modelling the data with two of the most common models of humans shape-from-shading – the linear shading model and the dark is deep model – revealed that for isotropic surfaces slant settings are best approximated by a diffuse – dark-is-deep interpretation but tilt settings are better explained by lighting from above. This is significant because the dark-is-deep rule is associated with diffuse illumination. Hence our slant results suggest that observers assumed a diffuse illuminant for the surfaces presented despite the fact that they were photographed under a largely directional illumination. There were of course no direct cues to illumination type or direction in the image to bias any assumption.

For anisotropic surfaces tilt and slant settings were most consistent with the linear shading model which is in turn consistent with a directional light source directed orthogonally to the dominant orientation structure in the image. Thus a directional lighting hypothesis prevails but the direction itself is determined by the image not the default illuminant. Here, anisotropic stimuli have a tendency to force a directional interpretation with the apparent direction of the light source to be orthogonal to the ripples of the surface. In these circumstances, the position of the actual light source has no influence and the image will always be perceived as if illuminated from the same direction – this is analogous to the aperture problem in motion; if the surface has oriented ripples those ripples will determine the perceived lighting direction.

We conclude that human SFS is dependent on shading patterns in the image and it is insensitive to small changes in light source directions. While this is consistent with previous discoveries that human SFS is independent of light source estimation (Mingolla & Todd, 1986; Mamassian et al., 1996), humans still need a rough estimation of the position of the light source. One cue could be based on edge boundaries as reported by few studies (Koenderink et al.,2007; Sun and Schofield, 2012) who demonstrated that distributions of edges were decisive in the compositions of illumination field processed by humans. The presence of oriented features in the shading profile itself is another (if sometimes false) cue to lighting orientation.

6.3 The effect of multiple light sources

The relatively weak influence of lighting priors in SFS that has been shown on Chapter 2 is related to the type of stimuli used. In a recent study, Morgenstern et al. (2011) have confirmed how the reliability of assumptions about the light source position could be affected by the strength of the lighting cues provided in the stimuli used by the experimenter. Relatively weak lighting cues were overridden by lighting priors while strong lighting cues guided the perception of SFS. Interestingly, strong lighting cues were less effective when local cues were removed from the stimuli and the observers' performance resulted in a trade-off between lighting cues strength and observers' lighting priors.

Chapter 4 considered the issue from an extreme case that provided a full-articulated, rendered, geometric stimulus with twin light sources and used bump-like disks as a probe for lighting direction. The aim of the experiment is to evaluate how twin lights would affect the perception of SFS when local cues are absent or very weak. When first testing observers lighting priors, we found two very different classes of prior: light-from-above (LFA) and surface convexity. The fact that we tested a large number of experimentally very naïve participants and that many of these favored the convexity prior suggests that LFA is not as common as might be suggested in the literature.

Our two main conditions demonstrated that local cues can override the prior to some extent but not completely – as demonstrated by DD's results; while distal lighting cues are not used at all. In the absence of local cues, these distal lights do not show any effect on observers' performance. This result differs to that found by Morgenstern's et al (2011). It seems like when local cues are taken off the image, people switch to their prior assumption (whatever that might be) and not to the lighting cues provided by the more distant light source. Our light field interpretation is therefore to be considered as very local and strongly limited to the part of the image we are trying to interpret at any given time. Once again, human lighting priors in SFS strongly rely on the type of stimulus provided and not the strength of the local lighting cues provided.

6.4 The interpretation of shadows

We considered the issue of shadows in Chapter 5. While previous chapters explored how shading cues determine the perceived shape for surface undulations and how such interpretations are influenced by prior assumptions and cues to the illuminant type, chapter 5 dealt with how and when the visual system might discount illumination cues entirely.

Previous results have suggested different (opposing) treatments for shadows: some have found that shadows are treated like any other object (Lovell et al.

2009, de Wit et al., 2012) others that they are not treated as objects (Rensink and Cavanagh, 2004). We tried to replicate de Wit et al. (2012) study and to add a stronger object-like condition to test their findings – but without success. Our vain attempt to replicate even the basis of the within object benefit for shadows suggest that there is no benefit for unambiguous shadows. Note that we did find a trend towards a within object benefit for some observers but that this did not extend to shadow like stimuli. We would therefore support the notion that object processing ignores shadows.

7. Conclusions

The aim of this thesis was to examine and explain the role of illumination in human SFS. As previously documented, the role of illumination is very important to solve any ambiguity in the perception of shape. Thus, the first question we tried to address is whether the assumptions made by the human visual system about the position of the light source may be related to the type of stimuli used. Furthermore, we wanted to test if this assumption is consistent with the notion that light comes from above.

We found that in stimuli composed of noise only and therefore deprived of any feature that could drive observers' response, the human visual system tends to show a slight bias towards a directional light source. Even though the majority of participants were producing templates that showed a light from above preference, still this proportion was not comparable to that found in previous studies. One explanation is that these results were biased by the type of noise used in the classification images task, however, observers produced different templates in different conditions suggesting that a change in their template was occurring when they were asked to search for different targets.

In order to investigate the case where lighting priors could be weakened by the physical features of the stimuli used in relation to the actual illumination present we showed how the declination of the light source could affect the perception of

SFS (chapter 3), and how the strength of local lighting cues could affect the reliability of lighting priors (chapter 4) in the extreme case of twin, opposite lights.

Previous studies found that observers tend to underestimate the curvature of isotropic stimuli that are lit from a frontal light while they tend to overestimate the curvature when the light is oblique (Curran and Johnston, 1996). We tested isotropic and anisotropic surfaces to see whether this effect could be extended to both types of surfaces, furthermore, we manipulated image contrast to see whether this effect was related to contrast. The differences found in our study suggest that lighting priors are active in the human visual system in cases where the regularities of the surfaces (Isotropy) or the low contrast do not provide enough information to the visual system and observers cannot entirely rely on physical features of the image. This is confirmed by the fact that the human visual system process anisotropic surfaces with natural contrast in a way that is more predictable by the linear shading model, with illuminant direction determined by the image, but switch to a more assumption led interpretation for isotropic stimuli.

With regard to the effects of local and global lighting, we tried to test the idea that lighting cues could override the lighting priors even in cases where ambiguous twin lights were illuminating the stimuli. Our results showed that the interpretation of local ambiguous stimuli is influenced mostly by local cues and that, again, in cases where the stimulus features are ambiguous, observers rely

on lighting priors. The lighting priors were overridden by strong local lighting cues when present as demonstrated by Morgenstern et al. (2011).

Finally consideration of why global lighting cues might be ignored led us to consider the role of shadows – which we might expect to be ignored in certain circumstance. The results of chapter 5 seem to suggest that, in terms of object processing at least, the human visual system seems to ignore shadows as if they were removed from the image.

To conclude, in this thesis we examined the role of illumination in relation to the processing of SFS in the visual system. The results show that assumptions about the light source can play a role only in cases where stimuli are highly ambiguous or in cases were lighting cues are absent or weak. Nevertheless, the role of lighting priors is strongly determined by the type of stimuli used in respect to their isotropy, orientation and contrast. Also, the human visual system seems to adopt a very local interpretation of light-field and tends to discount shadows in object processing.

In all we find the relationship between lighting assumptions and stimulus based cues to lighting to be very fluid. Lighting assumptions are themselves highly idiosyncratic. We find evidence for the lighting-from-above prior in a relatively small proportion of experimentally naïve participants. Frontal, below and diffuse lighting priors and a tendency to assume surface convexity are also common assumptions. None the less we do find a consistent bias for directional lighting.

However, such priors are easily overridden by strong, high contrast, local image cues that can suggest a lighting direction even when this is not the direction of the actual light source that illuminated the scene.

Bibliography

Abbey, C. K., Eckstein, M. P., & Bochud, F. O. (1999). Estimation of human-observer templates in two- alternative forced-choice experiments. *Proceedings of SPIE*, **3663**, 284–295.

Abbey, C.K., Eckstein, M.P. (2000). Estimates of human-observer templates for a simple detection task in correlated noise. *Proc. SPIE 3981, Medical Imaging 2000: Image Perception and Performance*, 70

Abbey,C.K., & Eckstein,M.P. (2002). Optimal shifted estimate

s of human-observer templates in two-alternative forced-choice experiments. *IEEE Transactions on Medical Imaging*, **21**, 429–440.

Abbey, C. K., & Eckstein, M. P. (2009). Frequency tuning of perceptual templates changes with noise magnitude. *Journal of the Optical Society of America A*, **26**, B72–B83.

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nature Neuroscience,* **7**(10), 1057-1058.

Adams,W.J. (2007). A common light prior for visual search, shape and reflectance judgments. *Journal of Vision*, **7**, 1-7.

Adelson, E. H., & Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In M. Landy & J. A. Movshon (Eds.), *Computational models of visula processing* (pp. 3–20). Cambridge, MA: MIT Press.

Ahumada, A. J. (1996). Perceptual classification images from Vernier acuity masked by noise. *Perception,* **25**, 18-18.

Ahumada, A. J. (2002). Classificaation images weights and internal noise level estimation. *Journal of Vision,* **2**(1):8, 121-131.

Ahumada, A. J., & Beard, B. L. (1996). *Object detection in a noisy scene* (Vol. 2657).

Ahumada, A. J., Jr., & Lovell, J. (1971). Stimulus features in signal detection. *Journal of the Acoustical Society of America*, **49**, 1751–1756.

Albrecht, A., List A., Robertson L. (2008). Attentional selection and the representation of holes and objects. *Journal of Vision,* **8**, 13–13.

Allen, B.P. (1999). Shadows as sources of cues for distance of shadow-casting objects. *Perceptual and Motor Skills*, **89**, 571-584.

Anderson,B.L. & Winawer,J. (2005). Image segmentation and lightness perception. *Nature*, **434**, 79-83.

Anstis, S. (1990). Imperceptible intersections: The chopstick illusion. In *AI and the Eye* Eds A. Blake, T. Troscianko (Chichester, UK: JohnWiley & Sons).

Ban, H., Preston, T. J., Meeson, A., & Welchman, A. E. (2012). The integration of motion and disparity cues to depth in dorsal visual cortex. *Nature Neuroscience*, **15**, 636–643.

Beard, B. L., & Ahumada, A. J., Jr. (1998). A technique to extract relevant image features for visual tasks. *Proceedings of SPIE*, **3299**, 79–85.

Blinn, J.F. (1977). Models of light reflection for computer synthesized pictures. *Proceeding SIGGRAPH '77*, 192-198.

Bonfiglioli,C., Pavani,F., Castiello,U. (2004). Differential effects of cast shadows on perception and action. *Perception,* **33**, 1291–1304.

Brainard,D.H., Pelli,D.G. & Bobson,T. (2002). Display characterization. In *Encyclopedia of Imaging Science and Technology*, ed. Hornak,J., pp. 172-188.

Braje, W.L. (2003). Illumination encoding in face recognition: effect of position shift. *Journal of Vision*, **3**, 161-170.

Brewster,D. (1826). On the optical illusion of the conversion of cameos into intaglios, and intaglios into cameos, with and account of other analogous phenomena. *Edinburgh Journal of Science*, **4**, 99-108.

Bulthoff,H.H. & Mallot,H.A. (1988). Integration of depth cues: stereo and shading. *Journal of the Optical Society of America*, **5**, 1749-1758.

Cavanagh, P., & Leclerc, Y. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance,* **15**, 3-27

Christou,C. & Koenderink,J.J. (1997). Light source dependency in shape from shading. *Vision Research*, **37**, 1441-1449.

Curran,W. & Johnston,A. (1994). Integration of shading and texture cues: testing the linear model. *Vision Research*, **34**, 1863-1874.

Curran,W. & Johnston,A. (1996). The effect of illuminant position on perceived curvature. *Vision Research*, **36**, 1399-1410.

Debevec, P. (1998). Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In M. Cohen (Ed.), *Computer graphics (proceedings of SIGGRAPH 1998)* (pp. 189–198). New York: ACM Press.

De Monasterio, F.M., and Gouras, P. (1975). Functional properties of ganglion cells of the rhesus monkey retina. *Journal of Physiology*, **251**, 167-195.

De Valois, R.L., Yund, E.W., and Hepler, N. (1982). The orientation and direction selectivity o cells in macaque visual cortez. *Vision Research*, **22**, 531-544.

de-Wit, L., Milner,D., Kentridge,R. (2012). Shadows remain segmented as selectable regions in object-based attention paradigms. *i-Perception*, **3**, 150-158.

Dovencioglu,D., Ban,H., Schofield,A.J., Welchman,A.E. (2013). Perceptual integration for qualitatively different 3-D cues in the human brain. *Journal of Cognitive Neuroscience*, **25**(9), 1527-1541.

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. *Journal of Vision*, **4**(9):11, 821–837.

Egly,R., Driver,J., Rafal,R.D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology General,* **123**, 161–177.

Erens,R.G.F., Kappers,A.M.L. & Koenderink,J.J. (1993a). Perception of local shape from shading. *Perception and Psychophysics*, **54**, 145-156.

Erens,R.G.F., Kappers,A.M.L. & Koenderink,J.J. (1993b). Estimating local shape from shading in the presence of global shading. *Perception and Psychophysics*, **54**, 334-342.

Gerardin, P., de Montalembert, M., & Mamassian, P. (2007). Shape from shading: New perspectives from the polo mint stimulus. *Journal of Vision*, **7**(11):13, 1–11.

Gerardin, P., Kourtzi, Z., & Mamassian, P. (2010). Prior knowledge of illumination for 3D perception in the human brain. *Proceedings of the National Academy of Sciences of the United States of America,* **107**(37), 16309-16314.

Gershun, A. (1939). The light field (P. Moon & G. Timoshenko, Trans.). *Journal of Mathematics and Physics*, **18**, 51–151. (Original work published in 1936)

Gibson,J.J. (1950). The perception of visual surfaces. *The American Journal of Psychology*, **100**, 646-664.

Gilbert, C.D. (1977). Laminar differences in receptive field properties of cells in cat primary visual cortex. *Journal of Physiology*, **268**, 391-421.

Gilchrist,A.L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science*, **195**, 185-187.

Gilchrist,A.L. (1988). Lightness contrast and failures of constancy: A common explanation. *Perception and Psychophysics*, **43**, 415-424.

Gilchrist, A. L. (2006). Seeing black and white. New York: Oxford University Press.

Gold, J. M., Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2000). Deriving behavioural receptive fields for visually completed contours. *Current Biology,* **10**(11), 663-666.

Gosselin, F., Bacon, B. A. and Mamassian, P. (2004). Internal surface representations approximated by reverse correlation. *Vision Research*, **44**, 2515-2520.

Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science,* **14**(5), 505-509.

Hendrickson, A.E. , Wilson, J.R., and Ogren, M.P. (1978). The neuroanatomical organization of pathways between the dorsal lateral geniculate nucleus and visual cortex in old world and new world primates. *Journal of Comparative Neurology*, **182**,123-136.

Hill, H., & Bruce, V. (1994). A comparison between the hollow-face and ''hollow-potato'' illusions. *Perception*, **23**, 1335–1337.

Horn,B.K.P. (1975). Obtaining Shape from Shading Information. In *The Psychology of Computer Vision*, ed. Winston,P.H., pp. 115-155. New York: McGrawHill.

Horn,B.K.P. (1977). Understanding image intensities. *Artificial Intelligence*, **8**, 201- 231.

Horn,B.K.P. & Brooks,M.J. (1989). *Shape from Shading*. MIT Press.

Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex*. Journal of Physiology*. **195**, 215-243.

Johnston,A. & Passmore,P.J. (1994a). Shape from shading: Surface curvature and orientation. *Perception*, **23**, 169-189.

Jiang, X., Schofield, A.J., Wyatt, J.L., (2010) Correlation-Based Intrinsic Image Extraction from a Single Image, in K. Daniilidis, P. Maragos, N. Paragios (Eds.) ECCV 2010, Part IV, LNCS 6314, pp. 58–71, Springer-Verlag, Berlin.

Johnston,A. & Passmore,P.J. (1994). Independent encoding of surface orientation and surface curvature. *Vision Research*, **34**, 3005-3012.

Kanizsa, G. (1976). Subjective contours. *Scientific American*, **234**, 48–52.

Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Praege.

Kanizsa, G.,Gerbino, W. (1982). Amodal Completion: Seeing and Thinking?. In *Organization and Representation in Perception*, ed. J. Beck. Hillsdale, NJ: LEA, pp. 167-190.

Keane, B. P., Lu, H., & Kellman, P. J. (2007). Classification images reveal spatiotemporal contour interpolation. *Vision Research*, **47**, 3460–3475.

Kersten,D., Knill,D.C., Mamassian,P., Bulthoff,I. (1996). Illusory motion from shadows. *Nature,* **379**, 31–31.

Khang,B.G., Koenderink,J.J. & Kappers,A.M.L. (2007). Shape from shading from images rendered with various surface types and light fields. *Perception*, **36**, 1191- 1213.

Kingdom,F.A. (2008). Perceiving light versus material. *Vision Research*, **48**, 2090- 2105.

Kleffner,D. & Ramachandran,V.S. (1992). On the perception of shape from shading. *Perception and Psychophysics*, **52**, 18-36.

Knill,D.C. (1992). Perception of surface contours and surface shape: from computation to psychophysics. *Journal of the Optical Society of America*, **9**, 1449- 1464.

Knill, D.C., Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature*, **351**, 228-230.

Koenderink,J.J., van Doorn,A.J. & Kappers,A.M.L. (1992). Surface perception in pictures. *Perception and Psychophysics*, **52**, 487-496.

Koenderink,J.J., van Doorn,A.J., Christou,C. & Lappin,J.S. (1996a). Shape constancy in pictorial relief. *Perception*, **25**, 155-164.

Koenderink,J.J., van Doorn,A.J., Christou,C. & Lappin,J.S. (1996b). Perturbation study of shading in pictures. *Perception*, **25**, 1009-1026.

Koenderink,J.J., van Doorn,A.J. & Kappers,A.M.L. (2001). Ambiguity and the 'mental eye' in pictorial relief. *Perception*, **30**, 431-448.

Koenderink,J.J., van Doorn,A.J. & Pont,S.C. (2004). Light direction from shad(ow)ed random Gaussian surfaces. *Perception*, **33**, 1405-1420.

Koenderink,J.J. & van Doorn,A.J. (2004). Shape and Shading. In *The Visual Neurosciences*, eds. Chalupa,L.M. & Werner,J.S., pp. 1090-1105.The MIT Press.

Koenderink,J.J., van Doorn,A.J. & Pont,S.C. (2007). Perception of illuminance flow in the case of anisotropic rough surfaces. *Perception and Psychophysics*, **69**, 895-903.

Langer,M.S. & Bulthoff,H.H. (2000). Depth discrimination from shading under diffuse lighting. *Perception*, **29**, 649-660.

Langer,M.S. & Bulthoff,H.H. (2001). A prior for global convexity in local shape-from- shading. *Perception*, **30**, 403-410.

Lee, T.S., Mumford, D., Romero, R.D., Lamme, V.A.F.(1998). The role of the primary visual cortex in higher level vision. *Vision Research*, **38**, 2429–2454.

Lennie,P. (1998). Single units and visual cortical organization. *Perception*, **27**, 889-935.

Levi, D. M., & Klein, S. A. (2002). Classification images for detection and position discrimination in the fovea and parafovea. *Journal of Vision*, 2(1):4, 46–65.

Liu,B. & Todd,J.T. (2004). Perceptual biases in the interpretation of 3D shape from shading. *Vision Research*, **44**, 2135-2145.

Livingstone, M., and Hubel, D.H. (1988). Segregation of from, color, movement, and depth: Anatomy, physiology, and perception. *Science*, **240**, 740-749.

Lovell,P.G., Gilchrist,I.D., Tolhurst,D.J., Troscianko,T. (2009). Search for gross illumination discrepancies in images of natural objects. *Journal of Vision,* **9**, 1–1.

Madison, C., Thompson, W., Kersten, D., Shirley, P., Smits, B. (2001) Use of interreflection and shadow for surface contact. *Perception & Psychophysics*, **63**, 187-194.

Mamassian,P., Kersten,D. & Knill,D.C. (1996). Categorical local shape perception. *Perception*, **25**, 95-107.

Mamassian,P. & Kersten,D. (1996). Illumination, shading and the perception of local orientation. *Vision Research*, **36**, 2351-2367.

Mamassian,P., Knill,D.C., Kersten,D. (1998). The perception of cast shadows. *Trends in Cognitive Science*, **2**(8), 288-295.

Mamassian,P. & Goutcher,R. (2001). Prior knowledge on the illumination position. *Cognition*, **81**, B1-B9.

Marr, D.. (1982). Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. New York: W. H. Freeman.

Mazzilli,G., Schofield,A.J. (2013). A cue-free method to probe human lighting biases. *Perception*, **42**, 932-940.

Mingolla,E. & Todd,J.T. (1986). Perception of solid shape from shading. *Biological* Cybernetics, **53**, 137-151.

Moon, P., & Spencer, D. E. (1981). The photic field. Cambridge, MA: MIT Press.

Morgenstern, Y., Geisler, W.S., & Murray, R.F. (2014). Human vision is attuned to the diffuseness of natural light. *Journal of Vision*, **14**(9):15, 1–17.

Morgenstern, Y., Murray, R.F., & Harris, L.R. (2011). The human visual system's assumption that light comes from above is weak. *Proceedings of the National Academy of Sciences of the United States of America,* **108**(30), 12551-12553.

Murray, R. F. (2011). Classification images: A review. *Journal of Vision,* **11**(5).

Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2002). Optimal methods for calculating classification images: Weighted sums. *Journal of Vision*, **2**(1):6, 79–104.

Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2005). Classification images predict absolute efficiency. *Journal of Vision*, **5**(2):5, 139–149.

Naber,M., Carlson,T.A., Verstraten,F.A.J., Einhäuser,W. (2011). Perceptual benefits of objecthood, *Journal of Vision*, **11**, 4–4.

Nefs,H.T., Koenderink,J.J. & Kappers,A.M.L. (2005). The influence of illumination direction on the pictorial reliefs of Lambertian surfaces. *Perception*, **34**, 275-287.

Nefs,H.T., Koenderink,J.J. & Kappers,A.M.L. (2006). Shape-from-shading for matte and glossy objects. *Acta Psychologica*, **121**, 297-316.

Nefs,H.T. (2008). Three-dimensional object shape from shading and contour disparities. *Journal of Vision*, **8**, 1-16.

Nicodemus, F. (1965). Directional reflectance and emissivity of an opaque surface (abstract). *Applied Optics*, **4**(7), 767-775.

Penrose,L.S., Penrose,R. (1958), Impossible objects: a special type of illusion. *British Journal of Psychology*, **49**, 31–33.

Pentland,A. (1989). Shape information from shading: a theory about human perception. *Spatial Vision*, **4** , 165-182.

Phong, B.T. (1975). Illumination for computer generated pictures, *Communications of the ACM,* **18**(6), 311-317

Pizlo,Z. (2008). *3D Shape: Its Unique Place in Visual Perception*. The MIT Press.

Pont,S.C. & Koenderink,J.J. (2007). Matching illumination of solid objects. *Perception and Psychophysics*, **69**, 659-668.

Radonjic, A., Todorovic, D., Gilchrist, A.L. (2010). Adjacency and surroundedness in the depth effect of lightness. *Journal of Vision*, **10**(9), 12.

Ramachandran,V.S. (1988). Perception of shape from shading. *Nature*, **331**, 163- 166.

Rao, R.P., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, **2**(1), 79-87

Rensink,R.A., Cavanagh,P. (2004). The influence of cast shadows on visual search, *Perception,* **33**, 1339–1358.

Rittenhouse,D. (1786). Explanation of an optical deception. *Transactions of the American Philosophical Society*, **2**, 37-42.

Schiller, P.H., Finlay, B.L., and Volman, S.F. (1976). Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specificity and ocular dominance. *Journal of Neurophysiology*, **39**, 1320-1333.

Schofield, A. J., Rock, P. B., & Georgeson, M. A. (2011). Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading? *Vision Research,* **51**(21-22), 2317-2330.

Scholl,B. (2001). Objects and attention: the state of the art. *Cognition*, **80**, 1–46.

Solomon, J. A. (2002). Noise reveals visual mechanisms of detection and discrimination. *Journal of Vision*, **2**(1):7, 105–120.

Sun,J. & Perona,P. (1996). Early computation of shape and reflectance in the visual system. *Nature*, **379**, 165-168.

Sun,J. & Perona,P. (1998). Where is the sun? *Nature Neuroscience*, **1**, 183-184.

Sun, P., & Schofield, A.J. (2012). Two operational modes in the perception of shape from shading revealed by the effects of edge information in slant settings. *Journal of Vision*, **12**(1):12.

Svetlana, S.G, Todd, J.T., Peeters, R., Orban, G.A. (2008). The extraction of 3D shape from texture and shading in the human brain, *Cerebral Cortex*, **18**(10), 2416-2438.

Tadin,D., Lappin,J.S., Blake,R., Grossman,E.D. (2002). What constitutes an efficient reference frame for vision? *Nature Neuroscience*, **5**, 1010–1015.

Todd,J.T. & Mingolla,E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance*, **9**, 583-595.

Todd,J.T., Koenderink,J.J., van Doorn,A.J. & Kappers,A.M.L. (1996). Effect of changing viewing conditions on the perceived structure of smoothly curved surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, **22**, 695- 706.

Todd,J.T. (2004). The visual perception of 3D shape. *TRENDS in Cognitive Science*, **8**, 116-121.

Tyler,C.W. (1997). Diffuse illumination as a default assumption for shape-from-shading in the absence of shadows. *Journal of Imaging Science and Technology*, **42**, 319-325.

Tyler, C. W., Likova, L. T., Kontsevich, L. L., & Wade, A. R. (2006). The specificity of cortical region KO to depth structure. *Neuroimage*, **30**, 228–238.

Vuong,Q.C., Domini,F. & Caudek,C. (2006). Disparity and shading cues cooperate for surface interpolation. *Perception*, **35**, 145-155.

Watson, A. B. (1998). Multi-category Classification: Template models and classification Images. *Investigative Ophtalmology and Visual Science*, 39(S912).

Wiener, N. (1958). *Nonlinear problems in random theory.* New York: John Wiley & Sons.

Zaidi, Q. (1998). Identification of illuminant and object colors: heuristic-based algorithms. *Journal of the Optical Society of America*, **15**(7), 1767-1776.

Zaidi,Q. (2005). A sculpture technique for rendering complex impossible objects. *Perception*, **34**, 121-132.

Zeki, S. and Shipp, S. (1988) The functional logic of cortical connections. *Nature*, **335**, 311-317.