THE ROLE OF CHROMATIC TEXTURE AND 3D SHAPE IN COLOUR DISCRIMINATION, MEMORY COLOUR, AND COLOUR CONSTANCY OF NATURAL OBJECTS

Milena Vurro

Newcastle University

Faculty of Medical Sciences

Institute of Neuroscience

Thesis submitted in partial fulfilment of the requirements of the regulations for the degree of Doctor of Philosophy

April 2010

Preface

The work presented in this thesis was entirely conducted during my PhD study. Reports on some of the work have been published or presented at international conferences. A list of publication and conference abstracts is shown as follows:

Milena Vurro, Yazhu Ling, and Anya Hurlbert. (2007). *The effect of shape on memory colour and colour constancy*. Perception 36 ECVP 2007 Abstract Supplement.

Anya Hurlbert, Yazhu Ling, Milena Vurro. (2007). *Polychromatic Colour Constancy***.** Perception 2008, 37(2): 308 – 316, AVA meeting Abstract (Birmingham, UK, Dec 2007).

Yazhu Ling, Milena Vurro, and Anya Hurlbert. (2008). *Surface chromaticity distributions of natural objects under changing illumination*. Fourth European Conference on Colour in Graphics, Imaging, and Vision (CGIV 2008, Barcellona, Spain). (Peer reviewed paper)

Hurlbert, A., Vurro, M., and Ling, Y. (2008). *Colour constancy of polychromatic surfaces*. [Abstract]. Journal of Vision, 8(6): 1101, 1101a.

Ling Y, Allen-Clarke L, Vurro M, Hurlbert A (2008). *The effect of object familiarity and changing illumination on colour categorization*. [Abstract]. Perception 37 ECVP 2008 Abstract Supplement, page 149.

Ling, Y., Pietta, I., Vurro, M., and Hurlbert, A. (2009). *The interaction of colour and texture in an object classification task*. [Abstract]. Journal of Vision, 9(8): 788, 788a.

Vurro, M., Ling, Y., and Hurlbert, A. (2009). *Memory colours of polychromatic objects*. [Abstract]. Journal of Vision, 9(8): 333, 333a

Hurlbert, A., Pietta, I., Vurro, M., and Ling, Y. (2009). *Surface discrimination of natural objects: When is a blue kiwi off-colour?* [Abstract]. Journal of Vision, 9(8): 330, 330a

Ling, Y., Vurro, M, and Hurlbert, A. (under review). *The effects of shape diagnosticity and depth cues on memory color of natural objects*. Journal of Vision. (Peer reviewed paper)

Vurro, M., and Hurlbert, A. (under submission) *Familiarity influences Color Appearance of Natural Objects*. Journal of Vision. (Peer revied paper)

Abstract

The primary goal of this work was to investigate colour perception in a natural environment and to contribute to the understanding of how cues to familiar object identity influence colour appearance. A large number of studies on colour appearance employ 2D uniformly coloured patches, *discarding* perceptual cues such as binocular disparity, 3D luminance shading, mutual reflection, and glossy highlights are integral part of a natural scene. Moreover, natural objects possess specific cues that help our recognition (shape, surface texture or colour distribution).

The aim of the first main experiment presented in this thesis was to understand the effect of shape on (1) memory colour under constant and varying illumination and on (2) colour constancy for uniformly coloured stimuli. The results demonstrated the existence of a range of memory colours associated with a familiar object, the size of which was strongly object-shape-dependent. For all objects, memory retrieval was significantly faster for object-diagnostic shape relative to generic shapes. Based on two successive controls, the author suggests that shape cues to the object identity affect the range of memory colour proportionally to the original object chromatic distribution.

The second experiment examined the subject's accuracy and precision in adjusting a stimulus colour to its typical appearance. Independently on the illuminant, results showed that memory colour accuracy and precision were enhanced by the presence of chromatic textures, diagnostic shapes, or 3D configurations with a strong interaction between diagnosticity and dimensionality of the shape. Hence, more cues to the object identity and more natural stimuli facilitate the observers in accessing their colour information from memory. A direct relationship was demonstrated between chromatic surface representation, object's physical properties, and identificability and dimensionality of shape on memory colour accuracy, suggesting high-level mechanisms. Chromatic textures facilitated colour constancy.

The third and fourth experiments tested the subject's ability to discriminate between two chromatic stimuli in a simultaneous and successive 2AFC task, respectively. Simultaneous discrimination threshold performances for polychromatic surfaces were only due to low-level mechanism of the stimulus, whereas in the successive discrimination, i.e. when memory is involved, high-level mechanisms were established. The effect of shape was strongly task- dependent and was modulate by the object memory colour. These findings together with the strong interaction between chromatic cues and shape cues to the object identity lead to the conclusion that high level mechanisms linked to object recognition facilitated both tasks.

Hence, the current thesis presents new findings on memory colour and colour constancy presented in a natural context and demonstrates the effect of high-level mechanisms in chromatic discrimination as a function of cues to the object identity such as shape and texture. This work contributes to a deeper understanding of colour perception and object recognition in the natural world.

Acknowledgment

Deepest gratitude is due to Prof. Anya Hurlbert, my supervisor, for her advice, support and guidance, and for her invaluable assistance in this study. Thanks to her vital encouragements this work was able to flourish and be successful. Her dedication to scientific research and her desire for knowledge and understanding inspiered me everyday during these three years and made me a passionate scientist. Special thanks goes also to Dr. Yazhu Ling, who showed me the ropes of the lab and was patient enough to introduce me to the "world of colour perception". The author also wishes to thank the Engineering and Physical Science Research Council for supporting this project.

Special thanks also to my colleagues and friends in the IoN, for the fruitful discussions and sharing of literature, but most of all for supporting me, encouraging or just listening in those days when everything went totally wrong, and rejoyed with me in those brilliant days of success. Immense thanks to Katja and Richie, for helping in printing this thesis. Eternal gratitude to my dearest best friends that from all over the world (in Italy and not) just called to say "You can do this!!" Thank you for always being there.

And finally but not lastly, thanks to my parents, my brother and my beloved boyfriend, Richard; for their unshakable love, understanding and care that has kept me warm even in the cold English winter faraway from home.

Disclaimer

The Author accepts liability for the content of this thesis. Any errors, mistakes or omissions presented in this thesis are solely those of the Author.

To my family

But it was all right,
everything was all right,
the struggle was finished.
George Orwell, Nineteen Eighty-Four (1948)

Table of contexts

Abbreviations

2DG: generic two-dimensional shape

3DG: generic three-dimensional shape

2DN: two-dimensional stimulus shaped as a natural object

3DN: three-dimensional stimulus shaped as a natural object

TEX: natural chromatic texture of a familiar object

RAND: randomized chromatic texture of a familiar object

MM: mean colour of a familiar object

MS: most saturated colour of a familiar object

CS: congruent stimulus; shape and texture of the stimulus belong to the same object

IS: incongruent stimulus; shape and texture of the stimulus belong to two different objects

GS: generic stimulus; the shape of the stimulus does not belong to any object

INTRODUCTION AND THESIS OVERVIEW

Through our senses, we are able to collect information about the world around us, but often there is no simple relationship between such information and the properties of the objects in the world. For example, the size of the retinal image of an object depends both on the size of the object and on the viewing distance; its shape depends both on the shape of the object and on the viewing angle; and its intensity depends both on the reflectance of the object and on the light illuminating it. Similarly the colour of an object depends on its surrounding. Then, how are we able to visually recognize an object? What computations are required to recover the information about the object's physical properties from the signals acquired by our senses? This thesis is part of the large number of ongoing efforts to answer these unsolved questions focusing on the colour properties of an object and its interaction with its shape and the illuminant changes.

The colour appearance of an object depends on a number of factors and no model developed has been able to represent fully the underlying mechanisms. A significant example is colour constancy, i.e. the ability to perceive object colours as almost constant despite changes in illumination that change the spectral properties of the light reaching the observer's eye. Most experiments designed to explain this phenomenon simulate a bi-dimensional world and simple, uniformly coloured shapes, far from the objects in the real natural world. Most likely, instead, colour constancy arises from a combination of sensory (e.g. chromatic adaptation) and cognitive (e.g. memory colour) factors operating at different levels in the visual system. Attributes of the object such as shape, texture or familiarity have been proven to influence colour appearance. For example, Hansen et al. (2006) have shown the effect of memory colour on colour appearance of 2D images. Furthermore, two studies of Ling and Hurlbert (2004, 2008) have demonstrated the effect of colour memory on colour constancy of paper patches (Ling&Hurlbert (2008)) and the interaction between the size of a 3D dome and its colour (Ling&Hurbert (2004)).

However, to the author's knowledge and as described in Chapter 1 of this thesis, no study has examined the combination or interaction of these different levels of complexity (three-dimensionality, shape, texture and familiarity) on colour appearance of natural objects. For example, how does the recognition of the object identity through its attributes, presented one at the time or in combination, affect the perception of the object colour? Moreover, if the memory colour of the object affects its appearance, is memory colour itself influenced by context of the scene, namely the shape of the object and the chromatic properties of the object's surface? Obviously, to answer these questions it is necessary to analyse the physical features of the natural object such as its natural texture or shape.

Therefore the first aim of this thesis was to develop a novel apparatus capable of representing the complexity of a natural object's appearance while allowing the free manipulation of its surface properties. The setup devised allows us to display a real three-dimensional object and a realistic natural texture, i.e. to present the observer with a more natural environment than a CRT monitor. Chapter 2 describes the hardware of this virtual environment, while Chapter 3 describes the software packages developed to simulate the object colour surface and algorithms required. This apparatus is the core of this thesis and this project, and it was employed to investigate several aspects of human colour perception in the context of memory colour and colour constancy.

The structure of this thesis' work should be seen as a building: Chapter 4 and Chapter 5 are the ground floor on which Chapter 6 and 7 are built. In Chapter 4, I examined the influence of naturalness of the object shape on memory colour introducing the concept of memory colour range (or distribution). This chapter presents two psychophysical experiments and an analysis of the object surface properties. The findings showed a robust effect of shape on memory colour properties and subject performance. Furthermore, Chapter 4 proposes a simple "model" of memory colour. This experiment introduces one of the main theories of this thesis: surface polychromaticity affects both immediate colour appearance and memory colour of natural objects.

Chapter 5 describes a computational analysis of natural object chromatic textures. The main aim of this analysis was to describe the difference between the chromatic distributions of natural textures under changes in illumination for objects belonging to the same or different categories. The results identify one specific illuminant-independent parameter and one within-category-independent parameter of the natural object chromatic distribution. The findings of this chapter were employed to design the following three psychophysical experiments described in Chapter 6 and Chapter 7.

Chapter 6 is the second "floor" of this thesis. This chapter had two aims: (1) investigate the effect of variegation of the chromaticity of the stimulus and shape on memory colour under a reference light and (2) evaluate the effect of change in illumination on memory colour and calculate the constancy of the perceived object colour in relation to the variegation of the chromaticity of the stimulus, and the 2D and 3D cues to shape. The results reveal that the naturalness of stimuli influences memory colour and its constancy. Chapter 6 proposes a direct correlation between cues to object identity, object natural surface properties and memory colour properties. This chapter introduces the concept of surface chromaticity variegation and its diagnosticity of the object. As the experiment in Chapter 6 examined diagnostic shape and texture only where they were congruent with each other in specifying object identity, I designed two experiments that evaluated cases with incongruent attributes.

In fact, Chapter 7 concludes my investigation on memory colour with two related experiments. In the first experiment, I studied the effect of natural texture and natural shape of a

familiar object in contrast with (1) synthesised textures, (2) incongruent shape, (3) uniform colours and (4) generic shapes on simultaneous discrimination. In the second experiment the same effects were contrasted in a successive discrimination task. Results showed that shape and surface chromatic properties strongly affect subject performance in both simultaneous and successive discrimination, and no simple relationship can describe the factors involved. Chapter 8 concludes this thesis drawing the final conclusions and pointing to the challenges for the future.

Chapter 1

Background and Related Work

In this chapter I will provide a concise reminder of the fundaments of colour perception considered in this thesis and I will outline the findings of previous studies on colour constancy and colour memory. In addition, I will summarize the available approaches for texture (and chromatic texture) analysis and provide a starting point for Chapter 3 and 7. To conclude I will review the finding son the effect of chromatic texture and of familiarity on colour perception.

1.1 THE PRINCIPLE OF TRICHROMACY AND BASIC COLOUR REPRESENTATION SYSTEMS

In most humans, colour vision is generated by the neural responses of three types of photoreceptors, termed cones. Each type of cone responds to a broad range of wavelengths, but has a different wavelength of peak sensitivity. Figure 1.1 shows the response of each cone type to the different wavelength. Cones labelled as long-wavelength sensitive, or L cones, possess a peak of sensitivity around the 570 nm (red line in figure). Cones labelled as short-wavelength sensitive, or S cones, possess a peak of sensitivity around the 420 nm (blue line in figure), whereas cones with a peak of sensitivity between this two responses are define middle-wavelength sensitive, or M cones (green line in figure). Hence any colour is encoded by exactly three distinct neural responses; this is the principle of trichromacy postulated for the first time by von Helmholtz (1896). Note that two physically different distributions of light will be indistinguishable if they result in the same three responses and they will be named "metameres".

However the cone response is not convenient method of representing a colour. In 1931, the CIE (Commission Internationale de l'Eclairage[1]) drew up a system of colour specifications (called CIE XYZ 1931 colour space) based on the colour matching functions obtained from the experiments of Wright (1929) and Guild (1931). In brief, in these experiments the observers were required to adjust a mixture of three beams of lights (red, green and blue), to match a test colour. The test and the matching colour were presented each in one half of the subject's field of view covering 2 degree of visual angle. The three coordinates XYZ of these CIE colour space define a specific perceived colour; even if the spectral components of two lights are different but their XYZ values are the same, then the "standard observer"[2] will perceive the two light identical. The CIE XYZ colour space is fundamental space from which the majority of other colour spaces are built, as the Yxy chromaticity coordinates, the La*b* space or the Lu*v* colour space.

---

1 www.cie.co.at

2 i.e. CIE 1931 Standard Colorimetric Observer.

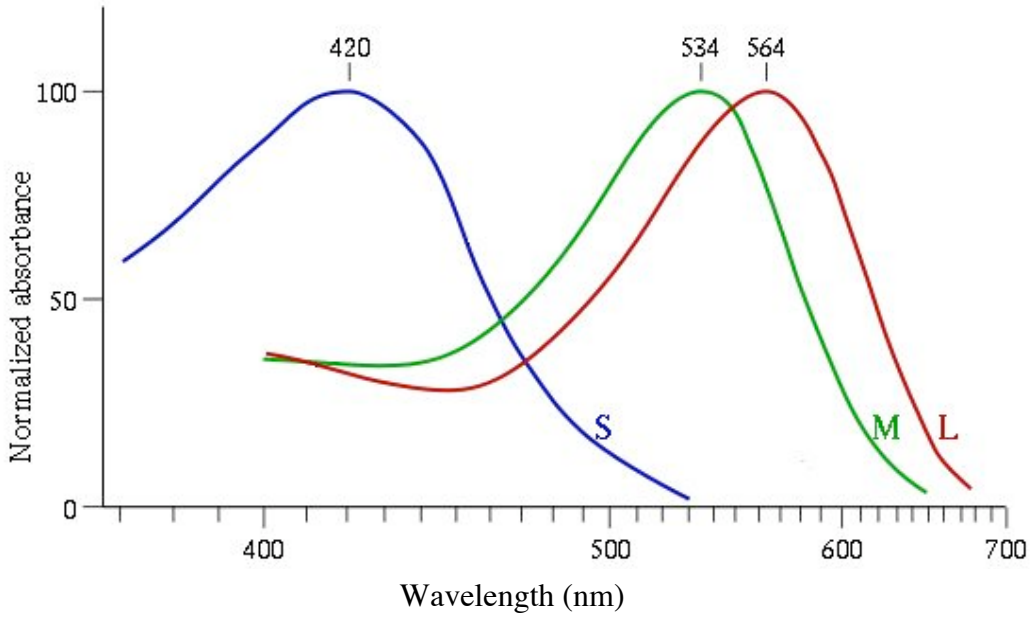

Figure 1.1 – Cone response a function of the light wavelength.

Since it will be frequently used in this thesis, it is worth reminding that the perceptual differences between two colours can be evaluated CIELAB colour difference ΔE^*_{ab} (or simply ΔE) defined as in the following formula:

$$\Delta E^*_{ab} = \sqrt{\Delta L^{*2} + \Delta a^{*2} + \Delta b^{*2}} \qquad \text{Equation 1-1}$$

Similarly we can define CIELUV colour difference as $\Delta E_{uv} = \sqrt{\Delta L^2 + \Delta u^2 + \Delta v^2}$

Finally it is worth reminding that the just noticeable difference (JND) between two colours is equal to about 1 ΔE unit.

1.1.1 THE EMG CONE-CONTRAST SPACE

Although very useful and widely used, tristimulus colour spaces as the CIE Yxy space fail to consider that: (1) cone responses actually depend on changes from a previous state of excitation and (2) their outputs are processed in opponency to generate a second set of signal outputs of the retina.

In fact, the photoreceptors actually respond proportionally to the amount of light at which they were exposed before and adapt to the average luminance and chromaticity of the scene. If this state is maintained the cone response can be described as the variation from its adapting point. Thus if we indicate with L, M and S the instantaneous response of the long-, middle- and short-wavelength sensitive cones to a test stimulus (as describe for example by Derrington, Krauskopf, Lennie 1984) and as L_0, M_0 and S_0 their response to the adapting field, then we can define the variation from the adaptation point as ΔL, ΔM and ΔS. Moreover, Noorlander&Koenderink (1983) described a colour space, termed cone contrast space, in which the variation ΔL, ΔM and ΔS were

divided by their correspondent response to the adapting field L_0, M_0 and S_0 mimicking "the sensitivity scaling effect of the adaptation occurring in cone-specific pathways" (Figure 1.2). Note that in this way we can obtain a three-dimensional space *in which the origin always represents the adapting condition*, and equal cone contrasts produce cone-contrast vectors along the equichromatic direction indicated by the dashed line in Figure 1.2.

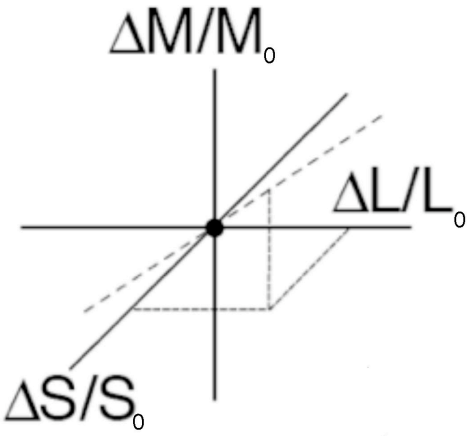

Figure 1.2 - Cone contrast colour space as described by Noorlander&Koenderink (1983). Figure modified from Eskew et al. (1999).

Furthermore, it is well established that the output of the retina is a function of the contrast between the cone responses (e.g. Hering (1964), Hurvich et al. (1957)). This theory, defined the opponent-colour theory, defines three combinations of the cone outputs: (1) the L- and M-cone signals are added to obtain the so called "luminance channel" (or L+M axis), (2) the L- and M-cone signals are subtracted to obtain the so called "red-green opponent" channel (or L-M axis), and (3) the S cone signal is subtracted from the luminance channel to obtain the so called "yellow-blue opponent" channel (or L+M-S axis) (for review Gegenfurtner (2003)). The result of these combinations is a colour space that is a linear combination of the previous cone space. For simplicity let us indicate the L+M channel as *L,* the L-M channel as *RG,* and the L+M-S channel as the *BY*.

Based on these two theories, Eskew et al. (1999) defined a cone contrast space that evaluates the cone signal modulations in relation to chromatic detection and discrimination. Firstly, let us model the opponent-colour mechanism as a linear combination of cone contrast factors as classically expressed by the following formula:

$$F_i = \gamma_F * \left(a_i \frac{\Delta L}{L_0} + b_i \frac{\Delta M}{M_0} + c_i \frac{\Delta S}{S_0} \right) + \beta(0, \sigma_i) \qquad \text{Equation 1-2}$$

where F_i is the mechanism response for the *i* colour opponent combination (or component) described above (i.e. RG, BY or L), a_i, b_i, and c_i are the weight given to each cone contrast variation (relative to the adaptation point) and $\beta(0, \sigma_i)$ is the probability density of a stochastic process, with zero mean and standard deviation σ_i different for each component and associated to the noise

sources that corrupt the its signal (e.g. neural noises). The second step is then to estimate the cone weights under neutral adaptation through psychophysical experiments. Based on the two studies of Cole, Hine and McIlhagga (1993) and Sankeralli &Mullen (1996) on six subjects, Eskew et al. (1999) computed such estimation of the a_i, b_i, and c_i weights for each of the three opponent mechanisms. The results of such computation are listed in Table 1.1; note that for the pourpose of this thesis the axis BY was inverted in direction (i.e. YB). The weights of RG and YB were chosen so that their vector had unit length and their sum was exactly zero. For a more detailed discussion on the above calculation refer to Eskew et al. (1999).

RG			YB			L		
a_{RG}	b_{RG}	c_{RG}	a_{YB}	b_{YB}	c_{YB}	a_L	b_L	c_L
0.70	0.72	-0.02	-0.55	0.25	-0.80	0.90	0.43	0.00

Table 1.1 – Cone contrast weights as computed by Eskew et al. (1999).

The cone contrast colour space proposed by Eskew et al. (1999) has been widely applied in the present thesis and will be referred as the EMG cone-contrast space. Keep in mind that, as said, in this colour space the origin (i.e. the point [0,0]) is always the cone response to the adaptation point (i.e. the chromaticity and luminance of such point) and all values are calculated relative to this point.

1.2 COLOUR CONSTANCY

> *At present, it is becoming increasingly common to work with very different colour reproduction media (e.g., paper, slides, and screens) and a great variety of environments, surrounds, and lighting conditions. It is desirable to reproduce the appearance of coloured stimuli with the greatest possible accuracy. (Fez et al. (2001))*

1.2.1 INTRODUCTION

Our ability to assign roughly constant colours to objects despite variation in the illumination allows us to identify objects correctly in a varying visual environment. This phenomenon, called colour constancy, has long been observed in psychophysical and physiological experiments. To understand colour constancy we must consider both the physical properties of the world, and the biological and psychological properties of the observer.

The light that reaches a local area of retina depends both on the spectral reflectance of the object or surface in view, and on the spectral composition of the illuminant. The human visual system response depends on the photons captured by just three classes of photoreceptors (cones),

and their absorption curves. Thus, we do not have direct access to full spectral information about the light that reaches the retina or, in general, direct information about the illuminant.

1.2.2 THE COMPUTATIONAL PROBLEM

From a computational point of view, colour constancy may be treated as a problem of transforming a sensory response matrix into an object's surface properties matrix.

Let us assume that the object surfaces in a two-dimensional scene illuminated by a source of light are flat and perfectly matte, and not affected by mutual reflections (i.e. the light reflected from each point has the same spectral composition in all directions). Furthermore, let us assume that s(λ), m(λ), and l(λ) are the spectral sensitivity functions for short-, middle-, long-wavelenght sensitive cones, respectively, while R(x,y,λ) and E(x,y,λ) the object's surface reflectance function and the spectral power distribution of the illuminant, respectively. The response for each cone for the point (x,y) will be:

$$S(x,y) = \int s(\lambda) R(x,y,\lambda) E(x,y,\lambda) d\lambda \qquad \text{Equation 1-3}$$

$$M(x,y) = \int m(\lambda) R(x,y,\lambda) E(x,y,\lambda) d\lambda \qquad \text{Equation 1-4}$$

$$L(x,y) = \int l(\lambda) R(x,y,\lambda) E(x,y,\lambda) d\lambda \qquad \text{Equation 1-5}$$

However, as said before, the spectral power distribution of the illuminant is unknown and is therefore impossible to recover, without further information given, the sensory response in which it is entangled with the surface reflectance function. So, given that colour constancy is an unconstrained problem, how does our visual system disentangle the information about the object from the light illuminating it?

An important finding of the last decade has been the empirical result that for environmental surfaces and illuminants, the spatial ratios of cone-photoreceptor excitations of the light reflected from pairs of surfaces under changing illumination is, to a great extent and for some surfaces, an invariant feature of the scene, a property that Foster and colleagues describe as the invariance of cone-excitation ratios (Foster and Nascimento 1994). In addition, they found that to a first approximation, a change in the spectrum of the illumination produces multiplicative changes in cone coordinates (Figure 1.3). Computationally, it consists of multiplying all cone coordinates by the same diagonal matrix, the elements of which are set by the illuminants cone coordinates. Furhtermore, Nascimento and colleagues (Nascimento, Ferreira et al. 2002) reported that spatial ratios of cone excitations of natural rural and urban images remain nearly invariant under daylight changes. They argue that this finding is a "general property of the visual environment (and …)

strengthen the argument for the basic role played by cone-excitation ratios in a range of color-constancy phenomena".

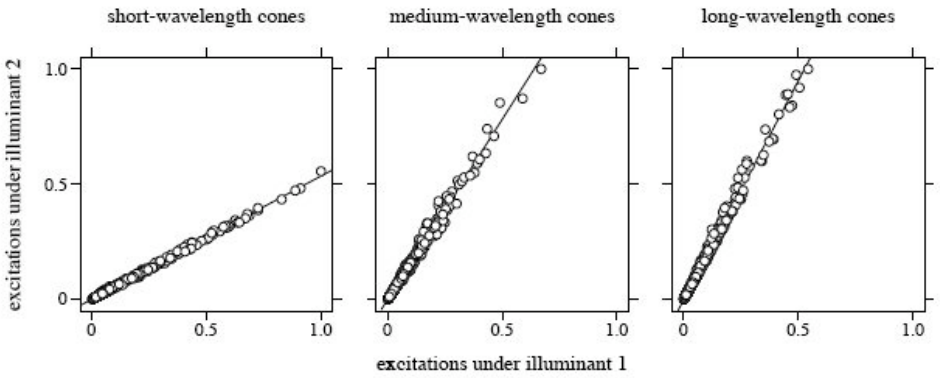

Figure 1.3 - Cone excitations for a set of rural scenes. Each point in the graphs represents a pair of excitations for a single pixel: The value on the x axis is for a daylight of correlated color temperature 25,000 K, and the value on the y axis is for a daylight of correlated color temperature 4300 K.(Nascimento, Ferreira et al. 2002)

Several models of colour constancy (Krantz 1968; Mamassian, Landy et al. 2002; Brainard 2004; Ling 2005), propose that regularities of the natural world may provide additional constraints to this problem. In the next section I concisely describe some theories to solve the colour constancy problem.

1.2.3 THEORIES OF COLOUR CONSTANCY

In 1902 Von Kries proposed a first theory on colour constancy. This theory, called the **adaptation theory** is based on the assumption that the eye possess three types of photoreceptors that decompose the light spectrum (Von Kries 1902, 1905). Since these photoreceptors, i.e. the cones, where not identified at the time, their existence was only postulated by Von Kries. He intuitively suggested that the sensitivity of these sensors were scaled to compensate for changes in illuminant. For example, when the retina was exposed to intense red light, the response amplitude of the "red" sensitive cones was reduced without a change in spectral shape, adapting to the light source. To apply his algorithm we have to make assumption on the illumination shining on the scene. There are two classic hypotheses to estimate the illuminant: (1) the 'grey world' assumption, which assumes that the average surface colour in a scene is grey, or (2) the 'brightest is white' assumption, which assumes that the brightest surface in a scene is white. However, the validity of both assumptions is limited. Nonetheless, the Von-Kries adaptation theory formalise an essential feature of colour constancy: the chromatic adaptation. This simple algorithm is been applied widely in many branches of vision science, e.g. the "white balance" calibration system in digital cameras.

A successive evolution of the Von-Kries model is the **retinex theory.** Developed by Land (1964), it suggests that the visual system determines colour based on the spatial distribution of received light signals in the scene rather than on the physical stimulus at a specific point (Land 1974). Based on this theory, each receptor records reflectance-related information about the object

colour: the‘lightness’ (McCann, McKee et al. 1976). The lightness signal is generated independently in each channel and compared in the retina, in the cortex, or even both, from which the term “Retinex”.

The key feature of retinex theory is the consideration of the effects of spatial distribution on colour perception. However, Lennie and D'zmura (1988) pointed out at least three key problems on this theory: (1) it does not account for the colour opponency mechanism of the cones, which can be considered if we modify the model applying instead of cone signals colour-opponent signals, (2) it is valid only for a two-dimensional “Mondrian” scene, i.e. uniformly coloured patches, under smooth spatial changes of illumination, and (3) requires a large receptive field to derive the lightness value. Nonetheless, the retinex theory provides a good theoretical foundation for the colour constancy problem solution and several retinex-based models have been proposed in the last decades (Horn (1974); Hurlbert (1986); Moore, Fox et al. (1990); Provenzi, Carli et al. (2005)).

An important objective that the above theories fail to achieve is to recover the spectral reflectance properties of surfaces from the incoming signals. The **linear basis algorithms** try to overcome this problem assuming that illuminants and reflectances can be characterized by three basic functions. A classic example is Maloney’s experiment in 1986 (Maloney and Wandell (1986)) which showed that three basic functions could account for 99% of the variance of 462 Munsell colour samples and 337 spectral reflectances of natural formations collected by Krinov. In other words, given a triplet of photoreceptor responses we can obtain a second triplet of illuminat-independent values through a 3x3 transformation matrix together with the set of three basic functions. They proposed that natural surfaces and illuminants are well represented by the weighted sum of a small number of basis functions, inherently stored within the brain or learned through past experience. As some studies have shown (Finlayson, Drew et al. 1994; Hurlbert 1998), the transformation is a simple diagonal transform. Alas, this theory is based on an assumption that the scene is illuminated by single spatially uniform illumination and that the surfaces within the scene are also restricted to flat, matt surfaces (Hurlbert 1998).

All the theories mentioned above can be considered as sensory models. They do not include the effects of past experience, complex judgments, language, context or content attributes. These are what we define as “higher perceptual or cognitive processes”. However there are numerous demonstrations of their contribution to perception from clinical, neuroanatomical and psychophysiological experiments, which have lead to the **cognitive theories** of colour constancy.

For example, from the neurophysiological point of view, visual information is conveyed to the cortex via two separate pathways, the Magnocellular (M) and Parvocellular (P) (Kandel, Schwartz et al. (2000)). Whereas the M pathway relays information about motion and depth, the P pathway transmits information about colour and form. Land (1986) tested the colour constancy of a

patient with complete resection of his corpus collosum. Since the patient's speech centre was in his left hemisphere, he was able verbally to describe things in his right visual field, but not in his left visual field. When a Mondrian was presented in his right visual field, his reports of the appearance of a centrally presented test-patch were consistent with those made by normal observers. However, when the Mondrian was presented in his left visual field, his reports of the same centrally presented test-patch paralleled those that a normal observer would give if the test-patch were seen in isolation. From this solid result, Land et al. argued that the constancy computations could not occur in the retina, but must occur in the cortex. Ruettinger presented further evidence that cortical computations are essential for colour constancy (Ruettiger, Braun et al. 1999). Five patients with circumscribed unilateral lesions in parieto-temporal cortex of the left or right hemisphere exhibited a selective loss of colour constancy, while their colour discrimination thresholds and colour associations for familiar objects were preserved. Furthermore, Zeki and colleagues, using Positron Emission Tomography (PET) scans, have demonstrated that colour is processed with information about context, specifically in the V4 region of the visual pathway (Zeki, Watson et al. 1991; Bartels and Zeki 1998; Kusunoki, Moutoussis et al. 2006).

In parallel, psychological experiments confirmed that contextual clues affect perception. Bodrogi and colleagues report that size alters the perception of hue and saturation (Kutas, Gocza et al. 2004). They found significant shifts of red, yellow and green unique hues when the stimulus size changed from 10° to 120°. Vice versa, Ling and Hurlbert (2004) prove that colour saturation significantly affects the perceived size. Some have proposed that after the recognition of a familiar object, its memory colour may help our visual system to perform colour constancy (Ling's thesis 2005). Arend & Reeves (1986) have presented a clear demonstration of the influence of experimental context on a colour-matching task. When observers were asked to make a match to 'look as if it were cut from the same piece of paper', they showed relatively good constancy compared with conditions where they were asked to match 'hue and saturation' (see also Arend et al. 1991). The appearance-based constancy obtained in the second case is a demonstration of phenomenal regression to the real object. Cornelissen and Brenner (Cornelissen and Brenner 1995) duplicated this experiment with concurrent measurements of subjects' eye movements and proved that observers spent more time looking at surrounds when making the "same paper" matches. They concluded that the change in eye movement, which led to a different state of adaptation of the fovea, was an effect of instruction.

Some recent colour constancy experiments attempted to create tasks in more natural viewing conditions obtaining higher degrees of colour constancy than traditional colour constancy experiments employing only 2D simple colour stimuli (Brainard and Freeman (1997); Brainard (1998); Kraft and Brainard (1999)). In addition, it is likely that other phenomena such as colours of

shadows, mutual reflections, texture attributes and highlights may aid us in determining the colour of the illumination (Lee (1986), Bloj, Kersten et al. (1999)).

Clearly, colour constancy is not a unitary process, achieved by all or nothing computation, but probably is achieved by transformations of visual information along specialized pathways that start in the retina, is enhanced in V1/V2 and continues in V4. Our current understanding of colour constancy is insufficient. A full understanding would require detailed knowledge of the physical world of illuminants and surfaces, and of the biological and psychological worlds of our sensory and cognitive processes. This work aims to enlighten a few of these aspects using a psychophysical approach.

1.2.4 APPROACHES FOR COLOUR CONSTANCY EXPERIMENT

While we have different methods to measure and evaluate human colour constancy, none of them are complete or have no disadvantages. In the last decades, all the studies have proven that different set ups or different instructions lead to entirely different results (Arend and Reeves 1986; Arend, Reeves et al. 1991).

On the other hand, the task itself may have different effects. One of the most popular approaches is to ask how well an observer can match the colour of a surface seen under one illuminant to the colour of a test surface seen under a second illuminant (asymmetric colour matching). The scenes may be real or artificial, and usually comprise multiple surfaces. The two illuminant-conditions may be presented side-by-side (simultaneous asymmetric matching; e.g. (Brainard, Brunt et al. 1997; Robilotto and Zaidi 2004), or one after the other (successive asymmetric matching; e.g. (Brainard and Wandell 1992; Hurlbert and Ling 2004), or to different eyes (haploscopic matching; e.g.(McCann, McKee et al. 1976)). Simultaneous matching has the drawback that adaptation to the two scenes will be determined by the pattern of eye movements across the two regions of the scene. Successive matching allows experimental control of adaptation to the two illuminants, but performance will additionally depend on the observer's ability to remember colours. Haploscopic matching allows separate adaptation states in the two eyes, but removes binocular cues to scene geometry. An interesting example of simultaneous matching is that set by Zaidi, in which he asks the observer to identify identical surfaces across illuminants (Brainard and Wandell 1992; Zaidi 1998; Zaidi 2001). This experiment has the diversity of showing four different objects coupled under two different lights, with only one different from the others. To choose the correct surface requires observers first to choose the illuminant condition in which the surfaces differ (a chromatic discrimination task), and then to identify which of those two surfaces is the same as the two standard surfaces under the second illuminant (the constancy task).

An alternative method is to ask an observer to adjust a test patch within a scene to appear white (achromatic setting; e.g. (Fairchild and Lennie 1992; Brainard 1998). However, the achromatic setting provides information about transformations of only a single point in perceptual colour space, and, therefore, cannot provide a general test of colour constancy mechanisms.

A further alternative is to ask the observer to assign to a surface a colour name under different illuminants (Troost and de Weert 1991; Hansen, Walter et al. 2007). However, our ability to discriminate surface colours far exceeds the number of categories represented in our vocabulary.

1.3 MEMORY COLOUR

The modeling of neuroanatomical and psychophysical data, coupled with colour vision theories, may help to provide an understanding of how the brain perceives colour and how colour perception interacts with memory processes.

Let us assume we are looking for a scarf to match our favourite hat and we enter a shop in our way home. To find the right colour we compare the "sampled" memory of the hat's colour with the actual colour of the scarf. Ones we have chosen the scarf we came back home and directly compare the two colours. If the viewing condition in the shop is the same of our home, the difference between the original colour and the chosen colour is completely due to memory effects. The memory of the hat's colour in my example is traditionally defined as "memory colour" and refers to colours that are recalled in association with familiar objects in long-term memory (Bartleson 1960).

Hering (1878) was the first to combine the terms object, colour, and memory. His work was a break through in colour vision also for his studies on opposite-colours theory. The early studies in memory colours used simple stimuli as paper cut shaped like familiar objects and their colours were representative of the typical colour of each object group (Duncker, 1939). Colour matching experiments were performed and matches were compared for figures with or without characteristic colours. The majority of early reports confirmed Hering's prediction in that two identically coloured patches were perceived to differ in colour if they were shaped as a familiar object or an unfamiliar object (Duncker (1939); Bruner et al. (1951)). A typical example is Duncker's experiment (1939) in which a leaf and a donkey were cut out of green paper and alternately illuminated with red light. The observers were asked to match the colour of the leaf and donkey using a colour wheel. The results showed that the leaf was judged greener than the donkey, providing evidence of a memory colour effect. Nonetheless successive experiments reported discrepant results probably due to different set up (Bolles et al. (1959); Delk and Fillenbaum (1965); Bruner et al. (1999)).

Other classic examples are Bartleson studies of the early ’60 (Bartleson (1960); Bartleson (1961); Bartleson&Bray (1962)). In his first experiment, fifty observers were required to select their memory colours for a previously named familiar object among a set of 931 Munsell patches. Ten natural objects were named: “red brick”, “green grass”, “dry grass”, “blue sky”, “skin”, “tanned skin”, “broad leaf summer foliage”, ”evergreen trees”, “inland soil”, and “beach sand”. Results showed that the mean memory colour of certain objects was more saturated than the object’s actual mean colour and its hue was shift toward dominant hues. However “red brick”, “skin”, “tanned skin”, and “beach sand” were not affected by memory. In the subsequent experiment (Bartleson (1961)) seven observers had to view a uniform patch for 15 second and the match its colour to the same 931 Munsell paper set as in the previous experiment. The subjects viewed in total four uniformly coloured patches of which colour was the mean of one familiar object (“Caucasian skin”, “blue sky”, “beach sand”, and “green foliage”). The object was named before the match. In this case the results showed no difference between hues with respect to the original colour, although there was an increase in the chroma. Hence he concluded that the effect of memory colour on colour appearance may depend on the specific object.

Pérez-Carpinell et al. (1998) investigated memory colour for a set of eight different familiar objects. The task of the observers in the experiment was to indicate individually the color of each of them. No additional information was given to observers about the objects to be examined except the name of an object. The observers pre-adapted to laboratory illumination for 5min before each session, then he/she examined binocularly ten comparison color samples distributed on the grey panel and chose the one that best represented the color of a familiar absent object. After completing the series of eight panels under one of the Macbeth illuminants, the subject carried out the same experiment again, but this time with the other illuminant after an adaptation phase (Perez-Carpinell, De Fez et al. 1998). The results obtained with one hundred observers and two illuminants indicate that: (1) the shifts that are produced in the dominant wavelength with memory depend on the familiar object considered; (2) colorimetric purity, as a measure of saturation, of the remembered objects is not the same as that of the familiar objects. With the second illuminant, red tomato was the best remembered color and yellow lemon the worst. The large number of subjects makes this work one of the most reliable in memory colour.

As outlined in the previous section context plays an important role in colour constancy, hence on colour appearance. It is then logical to ask if it also influence colour appearance relative to the memory colour of an object. Siple and Springer (1983) investigated memory shifts in different context. They selected six fruit and vegetable objects: carrot, maize, lettuce, lime, orange, and peanut. Each object was presented in three conditions: (1) without context (uniformly coloured disk), (2) with shape context (homogeneous filled contour of the object or homogeneous

"silhouette"), and (3) with texture context (textured filled contour of the object or textured "silhouette"). All stimuli were obtained from photographs of these objects. Eighteen observers were asked to select the colours that they typically associate with the fruit or vegetable named, while in a second part of the experiment they had to select the colour they "preferred" for it. The results showed that change in context produced no significant hue or lightness shift, although memory chroma and preferred chroma were higher than the original. These results seemed assert the independence of memory colour on shape and textural information.

However, more recently, Hansen and colleagues (Hansen et al. (2006)) have carried out an interesting experiment showing the effect of memory colour on the colour appearance of familiar object. Fourteen observers were asked to adjust the colour of digitalized photographs of natural fruit objects until they appeared achromatic. Additionally, two control stimuli were employed: (1) a uniformly coloured disk and (2) a random noise disk resembling the chromatic characteristics of a banana. All stimuli were placed on a neutral grey background. The subject could, in real-time, scale and rotate the entire distribution of pixels around the adaptation point of the scene in the DKL isoluminant plane (i.e. a cone constrast color space of chromatic axis L – M and (L + M) – S). Figure 1.4 show an example of such manipulation. Results showed that the photographs of the familiar objects were generally perceived to be achromatic when their colour was shifted away from the observers' neutral point in a direction opposite to the typical colour of the fruit. Such effect was not found for the control stimuli. Based on these results they suggested that the memory colour of the familiar object prompted a bias on the perception of its displayed colour.

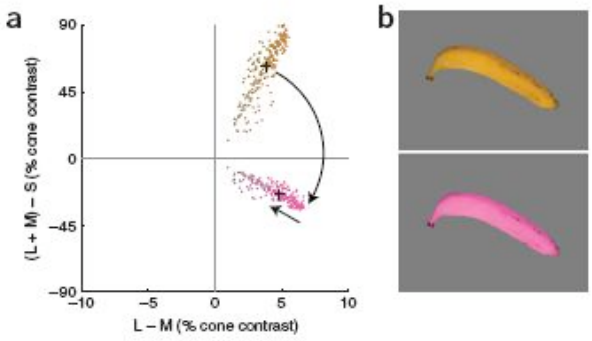

Figure 1.4 - The chromatic adjustment method. Back cross: mean of the distribution and how it changes when it is rotated. Hansen et al. (2006)

In a following up paper, Olkkonen et al. (2008) asked the observers to perform a similar task using the same stimuli under various simulated illumination. Two additional control stimuli were introduced: photographs of 3D fruit shapes without texture and 2D outline shapes. The subject had to adjust the simulus colour to the achromatic point, as in the previous study, or to its typical colour. Together with their previous findings, i.e. the biasing effect of natural textures, this experiment showed that the strength of such effect depended on the degree of naturalness of the stimuli, and it

was present under all tested illuminations. Furthermore, although colour constancy for textured stimuli was found lower than for uniformly coloured stimuli, colour constancy indeces resulted significantly higher for natural textures that synthesised texture, suggests an additional advantage for stimuli with a familiar texture. Finally, they concluded that these results indicated the potential significance of visual identity (memory colour) as an additional mechanism for colour constancy.

1.4 COLOUR CONSTANCY AND MEMORY

Since the changes in illumination occur over a period of time, the role of memory in colour constancy must be considered.

In general an individual views an object under one chromatic light and *then* under a second light. Hence it is possible to argue that because colour constancy is defined in terms of a change in illumination, it implies that the phenomenon occurs over some period of time. Based on this consideration, Jin and Shevell (1996) postulate a relationship between colour memory and colour constancy. They asked the observers to memorize a central patch surrounded by either: (1) a complex pattern composed of several coloured patches (1st condition), (2) a uniform grey field at the chromaticity of the illuminant (2nd condition), or (2) a dark background (3rd condition – control). After the training period, the subjects had to adjust a test colour under test illumination until it matched the remembered reference colour. They tested two hypotheses of colour memory: (1) colour is remembered as the light absorbed by each photoreceptor (photoreceptor hypothesis) or (2) as its surface reflectance (surface-reflectance hypothesis). Results showed that the colour was remembered as the light absorbed by the photoreceptors when there is no background light or only "neutral" background light. Conversely, the colour was remembered as the spectral reflectance of the surface when it is seen against a complex background. Hence in a complex environment the observers tend to remember the 'reflectance' (or at least an illuminant-invariant description) of surfaces presented under different illuminants, rather than the light that reaches the eye. They concluded that "a possible explanation for larger memory shifts in the 1st condition concerns the representation of colour in memory". Ling&Hurlbert (2008) carried out an experiment to evaluate the contribution of memory in colour constancy. Seven observers were asked to memorize a coloured paper sample under one illumination (reference sample). After adapting for 60 sec to a second illuminant while performing a distracting task, they were required to identify the reference sample among 15 other alternatives. The second illuminant was alternatively the reference illuminant, presented in the memorizing phase or a different illuminant. Therefore they were able to pool two "memory" shift: one obtained for the same illumination and one obtained for the change in illimination. Using such information, they defined a new colour constancy index that additionally

accounts for the memory shift in colour constancy, and proves that memory is involved in the constancy mechanism. Furthermore, if memory modulates colour constancy, does also memory colour influence colour constancy? As mention in the previous section, the work of Okkonen et al. (2008) has demonstrated such relationship.

1.5 TEXTURE

Texture is one of the most complex visual cues and at present there is no unified theory for its contribution to perception.

Texture is the visual cue derived from non-homogeneous surfaces in the scenes. Depending on the surface reflectance, viewing angle of the observer and lighting conditions, we may obtain different texture images from the same surface. Although there are some recent studies dealing with the recovery of the physical reflectance properties of a texture (Dana, Nayar et al. 1997; van Ginneken and Koenderink 1999), the most traditional approach in texture segmentation, classification or synthesis has been the analysis of the texture images without taking consideration of the image formation process. It is the author's opinion that a "real" computational representation of texture perception must have its basis on a psychophysical or neuroanatomical theory. In addition, despite the large number of works, there is still a lack of a standard texture definition nor a widely accepted texture representation space (for review Reed and Hans Du Buf (1993)).

In psychophysics, the aim has been to understand how the human visual system represents textures and which are the mechanisms used for texture segregation. Two basic categories of approaches are prominent at the moment: one just based on first order statistics and the second on higher order statistics.

The first category has as its father the work of Julesz and Bergen (1983): the Julesz's texton theory. They were the first to define as "texton" the primitive element (micro-structure) in natural images that forms the texture. For example, texton attributes can be the size and contrast between elements. Julesz's texton theory states that differences between two textures are due to differences in the first order statistics (e.g. mean), or densities, of the texton attributes, ignoring the positional relationships between adjacent textons (Julesz and Bergen 1983).

The second category, led by J. Beck, advocates that differences on first order statistics of local properties independently of the element arrangement is not enough to be able to capture the segregation of textures, since in a wide range of cases, differences are due to patterns emerging from the distinctive arrangements of image elements (Beck, Sutter et al. 1987). In these cases a global spatial-frequency analysis is needed in order to represent different textures.

The first class is clearly the simplest and many studies argue that the visual system may look for the simplest solution in resolving visual problems, although it can be easily demonstrated that this is not enough. In Figure 1.5, while the textures (a) and (b) can be differentiated using Julesz's texton theory (e.g diverse luminance mean), textures (b) and (c) are equals from this theory. Differences between textures (b) and (c) can be instead derived in the frame of a global frequency analysis considering differences in orientation.

Within these categories we can distinguish three different approaches: (1) geometrical, (2) model-based, and (3) filtering methods.

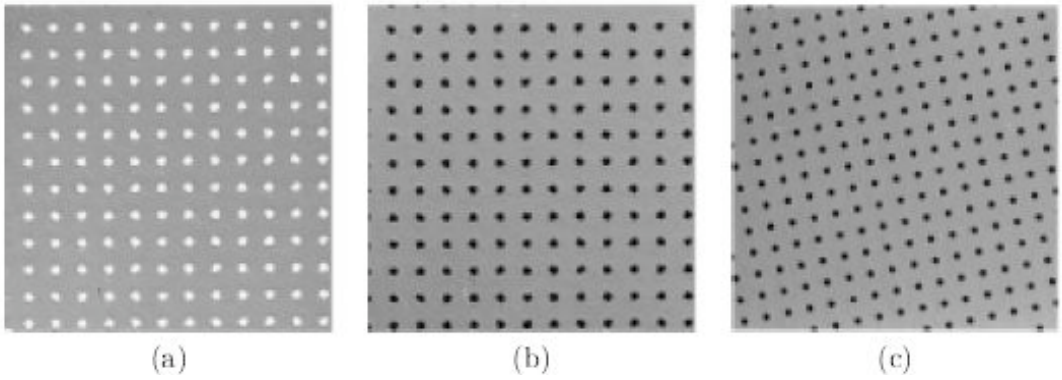

Figure 1.5 – Difference between the two categories of texture analysis (Caselles 2001).

In the **geometrical approach** texture is described by the set of textural primitives that compose the image (therefore a texton isolation step is always needed). Once the basic elements have been extracted we can: (1) compute the statistical properties of the extracted elements and their attributes [e.g. (Voorhees and Poggio 1988)]; (2) re-constructed by structural rules, i.e. primitives forming repetitive patterns are what define the image and such patterns can be described in terms of their construction rules.

In the **model-based approach** the texture is considered as the realization of a concrete mathematical model, hence it is defined by the model parameters. From a methodological point of view this is the most well defined solution, but, in general, there is not a unique model that can represent any natural texture. We can have a statistical model, of which the co-occurrence matrix is the most popular and representative method [e.g. (Haralick, Shanmugam et al. 1973)], or a stochastical model, of which the Markov random fields are probably the most widespread [e.g. (Cross and Jain 1983)].

In the **filtering approach** the texture is described by the responses of convolving a set of filters with the image. Spatial-frequency models use the fact that most textures, due to their repetitive behavior, are easier to represent in the spectral domain than in the spatial one. Spatial-frequency approach to texture analysis have provide the most successful solutions to the problems of texture characterization. Their success is reflected in the large number publications in this field. Here we found techniques in which features are obtained studying new representations as spectral

ones or conjoint representations of space and frequency. Fourier, Gabor, Wavelets are some of the most representative techniques of this group. For example, Malik and Perona proposed a global pre-attentive texture perception model based on neuro-physiological and psychophysical considerations (Malik and Perona. 1990). A global Fourier-based analysis of textures has been instead proposed by Liu and Picard (Liu and Picard 1996), while Manjunath and Ma suggested a Gabor transformation (Manjunath and Ma 1996). More recently, more complex filtering framework have been developed: for example, Gaussian filters based on bayesian assumptions (Lanterman, Grenander et al. 2000) or using independent component analysis (ICA) of Gabor features (ICAG) (Yang and Runsheng 2006).

1.6 COLOUR AND TEXTURE

Modeling of colour and texture in combination, rather than colour or texture alone, may provide better understanding of many visual mechanisms such as colour constancy.

It is well-known that colour and texture are visual cues in the distinction and recognition of objects, and thus colour texture representation has become a current topic in any field of vision research. Although both are properties of surfaces that have been studied separately. In fact, most of the existing studies on texture analysis have only focused on gray-level texture images. Figure 1.6 shows an example of the difference in appearance in three stimuli possessing the same geometrical texture and mean colour. Therefore, considering each individual cues separately is not enough when studying object recognition.

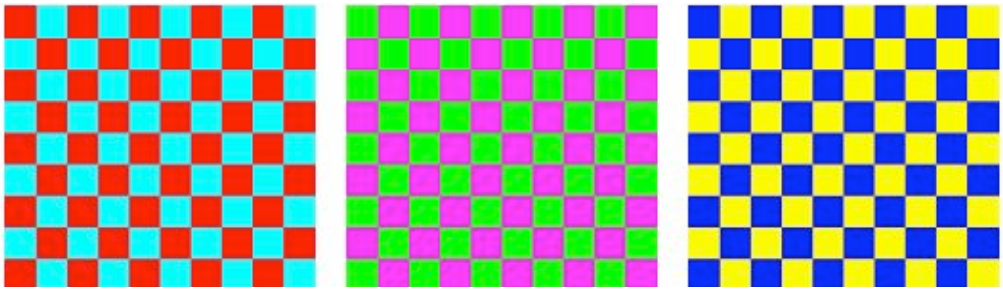

Figure 1.6 – Image with different appearance but with the same mean colour and pattern/texture.

The objective of many researchers has been to find conjoint representations of spatial and chromatic information that capture the spatial dependence (in particular, correlation) within and among spectral bands. Alas, no single method is comprehensive or exhaustive.

Two of the most frequent approaches in texture synthesis are (1) the construction of a feature vector mixing grey level texture features and colour features (see for review (Cremens, Rousson et al. 2007)) or (2) to extend classical texture methods, such as image quilting, Markov Random fields, the autocorrelation function (e.g. (Panjwani and Healey 1995)), or the structural tensor (Weickert 1999), in order to deal with multichannel images. An additional approach, as in

Gagalowicz and colleagues' work, consists of converting RGB values into a single code from which texture measurements are computed as if it were a grey scale image (Gagalowicz, De Ma et al. 1986). Filtering methods are staring to generate a great interest. Caelli and Reye have proposed the use of three multiscale isotropic filters to extract features from three colour spectral channels, but this method does not consider correlation between spectral bands (Caelli and Reye 1993). In the last few years the interest has grown and more complex approaches have been proposed; for example: Millán and Valencia combine the Laplacian of Gaussian (LoG) operator with spatial filters that approximate the contrast sensitivity functions of human visual systems (Millán and Valencia 2006), Bannai and colleagues employ a 3D colour-texture fusion technique (Bannai, Fisher et al. 2007), and Wu and Li have developed a complex tree-structured wavelet decomposition algorithm (Wu, Li et al. 2007).

In this thesis I will employ a modified image quilting technique applied in parallel to all the channels (Efros&Freeman (2001), Efros&Leung (1999)).

1.6.1 CHROMATIC TEXTURE AND FAMILIAR OBJECTS

Among the published works on the effect of chromatic texture on colour appearance (some recent works are for example Hurlbert&Wolf (2004), Giesel et al. (2009)), there are very few examining the effect of natural chromatic texture on colour appearance. A first interesting finding on natural chromatic textures came from the study of Hurlbert and Ling (2006). They showed that the distribution of within-surface cone contrasts, due to chromatic texture, for a given object forms an elongated cluster in three-dimensional cone-contrast space that can be represented by a "vector". For many surfaces, the direction of the cone contrast 'vector' remains constant under changes in illumination, if the cone contrasts are calculated with respect to the illumination chromaticity. They therefore argued that surface colour constancy might be better for objects with natural chromatic texture than for the homogeneous surfaces typically used in laboratory measurements of constancy

In a recent study Hansen et al. (2008) explored the effect on discrimination of chromatic natural and synthesised textures under different sates of adaptation. Three subjects were asked to perform a 4AFC and select the stimuli, which differ in colour to the other. Discrimination thresholds were measured at the adaptation point, where the mean chromaticity of the test stimuli and of the adaptation point are the same, or away from it. The discrimination thresholds were measured along 8 directions and the results were analysed using discrimination ellipses. Their results showed an effect of the chromatic distribution independent from the familiarity of the texture. Specifically the discrimination ellipses of the chromatic texture (natural or synthetic) were elongated in the direction of maximum variation of the chromatic distribution of the stimuli. Although in this experiment this effect was found only when the stimulus was centered with the

adaptation point in a subsequent experiment (Giesel et al. (2009)), same findings were reported also for distribution centred away from the adaptation point.

The experiments presented in this thesis are an evolution of the findings described in this chapter and aim to answer the some unresolved questions on the influence of texture, shape and familiarity on colour constancy, colour memory and chromatic discrimination.

Chapter 2

THE APPARATUS

2.1 INTRODUCTION

Contrary to previous studies that have investigated colour memory and colour constancy in monochromatic bi-dimensional (2D) spaces, this work aims to analyze interaction of colour and shape in a real three-dimensional environment. The final goal of the apparatus developed for this thesis is to create a virtual reality environment that gives the illusion of a solid (3D) object changing surface colour. The main idea consists in using a data projector to display an image onto a real solid object, which is previously been painted white, so that its surface colour depends directly on the colour projected. Since the data projector is connected to a computer, the object' surface colours can be adjusted by changing the colour of pixels overlapping the visible surface of the object in *interactively*. To determine the geometry of the visible surface, and therefore the area in the image corresponding to it, we have developed a geometrical calibration technique of the system that will be explained in the next chapter. This technique employs a digital camera as recording device, described in section 2.3.1.3.

While for the first experiment of this thesis I employed a setup constructed by Y. Ling (2005), for all the subsequent experiments I have designed a new apparatus to produce a better naturalness of the scene and adaptation. A brief description of the first is given in the next section (for further detail refer to Ling's thesis (2005)). A full description of the new setup design is detailed in section 2.3.

2.2 FIRST EXPERIMENTAL SET UP: THE BOX

Figure 2.1 show the main design of the experimental set up used for the first experiment in this thesis. The box maximum dimensions are: 800mm x 600mm x 1000mm. Its interior and exterior are entirely painted with matt black paint to eliminate undesired light reflection, and a matt white paper is placed on the bottom. The subject is allowed to look inside the box through a rectangular viewing hole of 85mm x 95mm and his/her view is restricted by a truncated pyramidal tunnel attached internally at the borders of the hole (Figure 2.2).

A projector is located at the top of the box, hidden from the outside observer. The projected light is reflected by a first mirror (350mm x 350mm; Galvoptics float glass), and transmitted through a one-way mirror (Pilkington Mirropane one-way observation mirror; 38% transmission, 800mm x 800mm). Then the light hits the bottom of the box and is reflected back towards the reflective side of the one-way mirror, reaching the subject eyes. The viewing hole is aligned perpendicularly with the centre of the observation one-way mirror; as a result the subject will see

the object located in front, in the vertical plane. The truncate pyramid hides both the chamber walls and the mirrors borders (Figure 2.2).

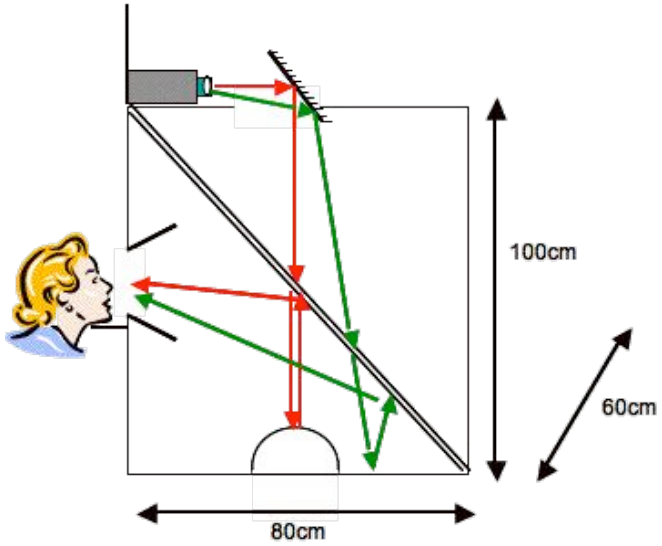

Figure 2.1 - Schematic illustration of the hardware set up (Ling 2005)

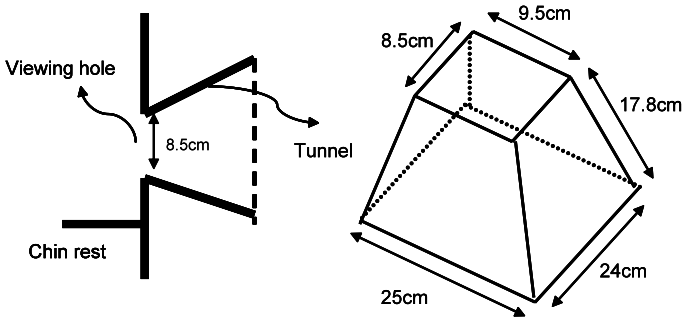

Figure 2.2 -Diagram of the geometry of the view aperture (Ling 2005)

2.3 SECOND EXPERIMENTAL SET UP: THE CHAMBER

Figure 2.3 shows a simplified diagram of the experimental apparatus used for the remaing experiments in this thesis (3 main experiments). It consists in a chamber, mainly constructed in wood, placed on a wheeled steel platform (not in figure). The chamber's maximum dimensions are 1400mm wide, 1500 mm deep and 1050 mm high and it was designed using a CAD-like software. The roof of the chamber is sloped such that it is highest on the viewer side (1050 mm – Figure 2.3), and lowest at the opposite of the viewer (far end of the chamber; 610 mm). The chamber can be conceptually divided in two compartments: a first compartment or "adaptation" compartment and a second compartment or "display" compartment. The exsterior of the chamber is completely painted with deep-black matt paint.

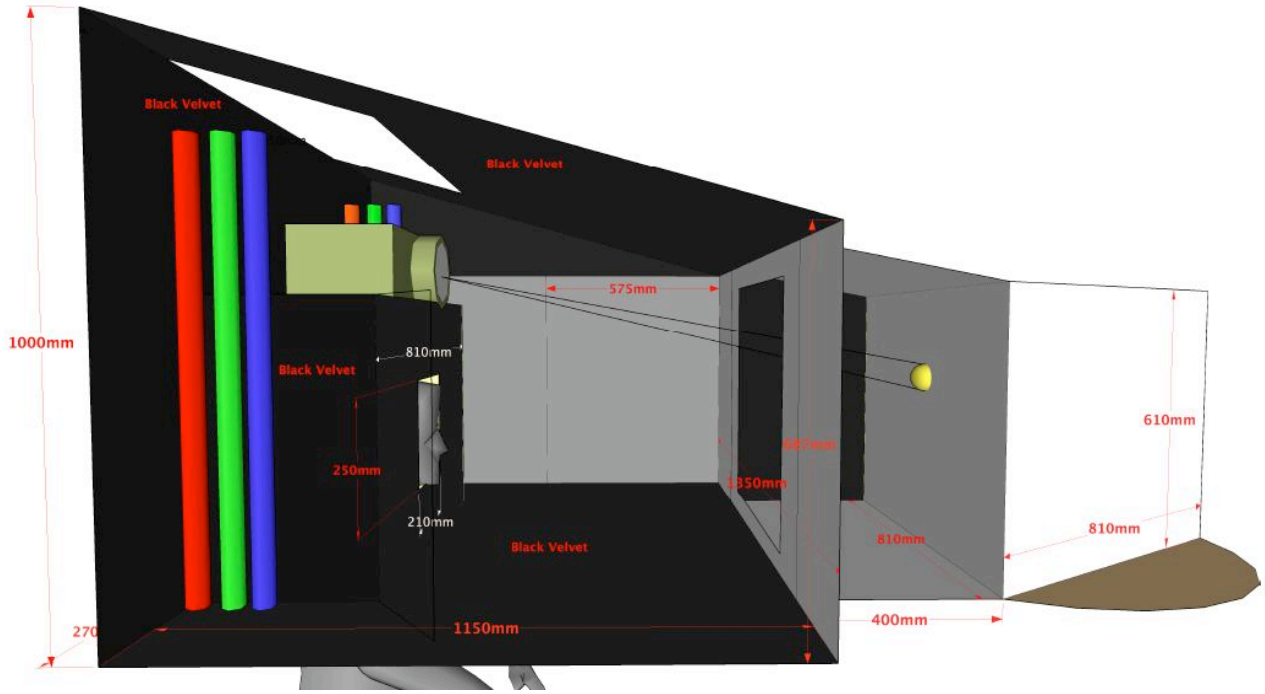

Figure 2.3 – Schematic illustration of the experimental set up.

2.3.1 ADAPTATION COMPARTMENT

The adaptation compartment contains: (1) a lighting system consisting in two sets of three fluorescent lights (red, green, and blue); (2) a data projector; (3) a digital camera; (4) a removable frame; (5) a set of movable panels; and (6) a pair of sliding doors. This wooden compartment has a U-shaped bottom (Figure 2.4). The subject sits into a rectangular inlet (400mm deep, 810mm wide and 650mm high) as in Figure 2.4. In the centre of the vertical front wall of inlet, a rectangular viewing hole (dimension 400x250mm) is cut at 150mm from the bottom of the chamber and 200mm from the sides; the inlet ceiling was covered with black sponge for the subject safety. Figure 2.5 shows a photograph of the front of the chamber; a black-painted ergonomic chin-rest is placed centrally to the viewing hole and a black thin rubber and velvet cushion is glued to its top. Top and bottom of the compartment are covered with deep black matt velvet. The front walls and the visible part (from the subject point-of-view) of the sidewalls are covered in ultra-white matt paper. The parts of the sidewalls that are not visible from the subject point-of-view are covered with deep-black velvet.

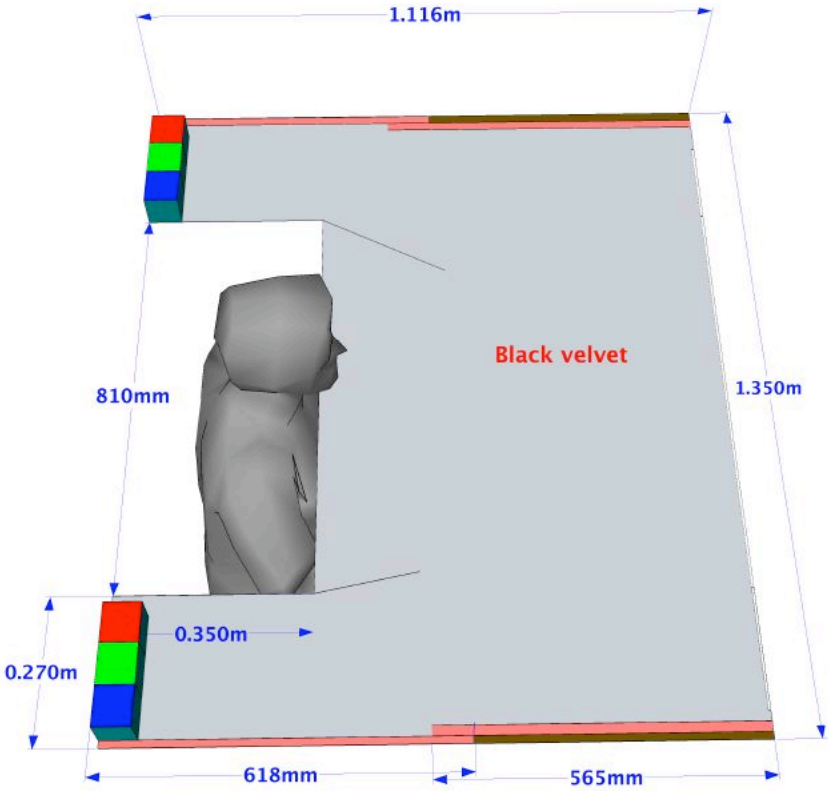

Figure 2.4 – Ground plan of the first compartment: viewed from top looking down. Red, green, and blue cubes: three coloured fluorescents lights. The pink rectangles: sliding door (or sliding doors tracks). Black lines: movable panels.

A)  B) 

Figure 2.5 – Photographs of the chamber: A) front view; B) left side view; red arrow: side panels.

Drawn in dark pink in Figure 2.4 are the tracks of two pairs of sliding doors (550 mm wide and 610mm high). The doors give complete access to any part of the first compartment, but, once closed, seal the interior from any undesired light. The parts of the doors visible from the viewing hole are covered in ultra-white matt paper, while the doors not visible and nearer to the light bulbs

are covered in deep black velvet. Any surface close to the lights is also covered with the same black velvet to avoid the reflection on other surface and control the direction of the light. To perfectly shut the door to intrusive undesired light a plastic gasket was glued on the borders.

2.3.1.1 LIGHTING SYSTEM AND SIDE PANELS

In Figure 2.4 is sketched the ground plant of the first compartment seen from the top looking down. Attached internally on the front walls, at the two sides of the inlet, three pairs of coloured fluorescent lamps (Osram T5-FQ39W-220 red, -230 green, and -240 blue) are placed behind a diffusing filter film (416 Lee uniform three-quarter white diffusion filter film; 25% light intensity loss at a distance of 175 mm), a set of three per side of the chamber. Each pair of coloured bulb is digitally controlled by double dimmable ballast (Osram Quicktronick - FQ 2x39/230-240 DIM) and connected through a custom-made interface to a terminal block with 68-pin SCSI female connector (DIN-68S). This in turn is connected to an 8-channel digital/analog output PCI multifunctionboard (DAQ 2502).

Two adjustable panels, placed inside the chamber at the two side of the viewing hole at a horizontal distance of 250 mm, restrict the direction of the light (black lines in Figure 2.4 and indicated by the red arrow in the photograph in Figure 2.5B). The area of the panels exposed to the observer view is covered with ultra-white matt paper while the rest is covered with deep black velvet. A frame is placed between the first and second compartment to avoid interaction between the light coming from the lighting system and the projector. Hence, in practice, the fluorescent lights illuminate only the first compartment. The side of the frame exposed to the subject view (front side) is covered with ultra-white matt paper, while the other side (back side) is covered with deep black paper to avoid mutual reflection. The light placements, panels' location, type of bulbs and filters, and, in general, the lighting system calculations were computed using DIALux software (DIAL GmbH, www.dialux.com). Ventilation holes were created to avoid the light bulbs overheating.

2.3.1.2 DATA PROJECTOR AND PLATFORM

A computer-driven LCOS data projector (Canon Xeed SX6) is placed on a platform above the subjects' seat, as shown in Figure 2.6. The platform is screwed inside the chamber and hidden from the observer's view. To allow total flexibility of the projector position, the stand has three degree of freedom, but the projector is solidly anchored to its center. The projector is placed upside-down to use the best angle for the chamber. To reiterate, the projector and the lighting system locations, as well as the chamber geometry, are designed to create the illusion of a solid object changing surface colour and so that there is no mutual reflection or interaction between the images

projected on the screen and the illumination from the lighting system. The projector and light system can be controlled independently. Note that the top of the projector is covered by a grey cardboard during the experiment to block possible undesired light.

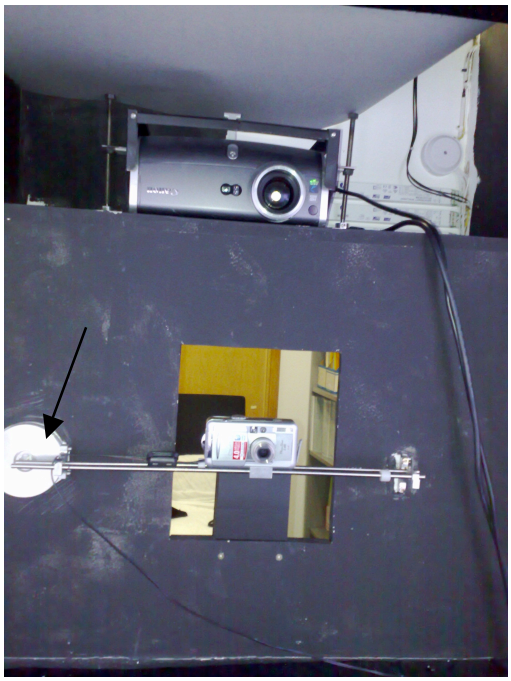

Figure 2.6 – Photograph of the data projector, digital camera, and camera platform (black arrow indicates the rotating arm block) taken from the inside of the chamber.

2.3.1.3 LATERAL PANELS AND CAMERA PLATFORM

As said in the introduction, to execute the calibration of the projected image geometry, a digital camera was employed. As will be explained later, a series of images had to be recorded from the same viewing point as the observers, every time at exactly the same point. Therefore the camera has to be placed at the observer eye level in front of the viewing hole. However, since we do not want to obstacle the subject view, it was necessary to design a platform that allowed perfect reallocation of the camera each time a picture was taken, but that was easy to remove before the experiment.

Figure 2.6 shows the location of the camera and its platform. A thin steel track is mounted as the arm of a rotating block fixed inside the chamber, on the right side of the viewing hole. Such block consists of two parallel, overlapping steel disks, of which one is free to rotate around their common centre and the other is solidly attached to the chamber internal wall. The camera is secured to a slider inserted in the track and can easily moved left and right if necessary. Two screws are used to block the slider on a fix position. When the camera is not in use the arm is lifted and anchored to a steel hook, and it is not visible by the observer. Conversely when needed the arm is rotated in horizontal position and fasten to a second steel hook.

2.3.2 DISPLAY COMPARTMENT

This second compartment consists mainly in a transition room and a rotating system. A closer schematic plant of the second compartment is illustrated in Figure 2.7.

2.3.2.1 TRANSITION ROOM

Photographs of the transition room are shown in Figure 2.8. Top and bottom of this room are painted with ultra-white matt paint, while the lateral sides were covered with black matt paint. In the same photograph it is also possible to see the back of the frame described in 2.3.1.

To enable the experimenter to arrange the objects before or during the experiment, a door was built on one side of the transition room (Figure 2.8A). In addition, a blind was build into the ceiling of the transition room, behind the frame and hidden to the observer during the experiment, to block the subject view during objects' re-arrangement (and eventually change of the scene by the rotating platform) without interfering with his/her adaptation to the illumination.

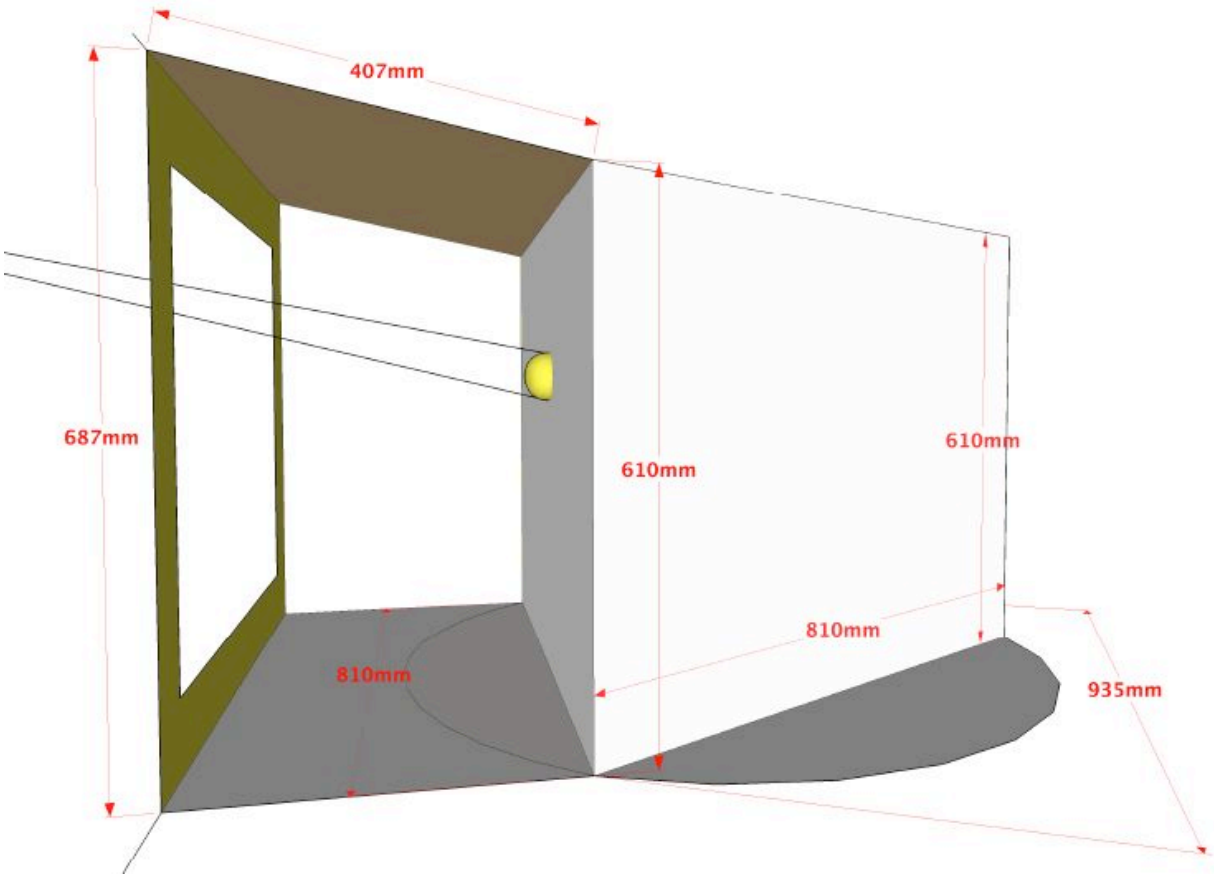

Figure 2.7 – Sketch of the second compartment or display compartment.

2.3.2.2 ROTATING SYSTEM

A rotating mechanical system was developed to allow rapid changes in the scene during the experiment (Figure 2.8B). The platform of this system consists of a wood disc of ~935mm diameter and 16.5mm thickness attached through a 100.5mm steel disc (7mm thick). The steel disc is welded to a steel cylinder of 30mm diameter and 50mm length. The latter is inserted coaxially to a second steel cylinder of 33.5mm and 60mm length (detail of this block in Figure 2.8C). Three steel boards are placed on the wood platform in vertical position, forming a triangle as show in Figure 2.8B). Each board is 810mm by 610mm and 7mm thick. To support the platform, reduce the friction and allow free rotation of the disc, four wheels are attached to a structure coaxial to the disc centre and located under the wood disc. Additionally, a blocking system is welded to this structure; it is designed to be set off by specific indents in the wood disc. The automatic block indents are placed at three distinct locations of the disc corresponding to the three locations in which the three boards were perfectly in front of the subject and projector. In this way the steel boards (or screens) are set at the exact same location every time, i.e. the image projected by the projector falls always exactly on the same area of the screen. The screens were covered with ultra-white matt paper and the rotating table was painted with ultra-white matt paint. To seal the box from any undesired light coming from the outside, the borders of the box in contact with the screen were overlaid with black plastic gaskets.

Small magnetic buttons (Electronics- bonded Neodymium-Iron-Boron discs: 2mm diameter and 1mm thickness) glued to the experimental objects (3 to 10 per object) hold them in place on the boards during the experiment and, eventually, during the rotation of the platform (change of scene).

A)  B) 

C) 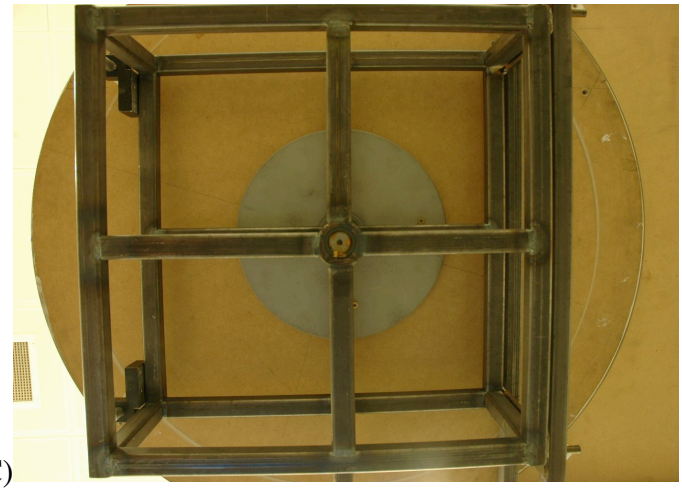

Figure 2.8 – Photographs of the second compartment. A) Transition room. B) Top-left view of the rotating system. C) Rotating block detail.

Chapter 3

SOFTWARE DESIGN AND TEXTURE MANIPULATION

3.1 INTRODUCTION

This chapter describes the general technique used to generate the images for the experiments described in this thesis. The global idea is to extrapolate the information of a chromatic texture from a real, natural three-dimensional object and generate a second 3D scene. The latter presents one or more solid real objects; the surface appearance of these objects is a manipulation of the object's original texture. The surface colour distribution of the object can be digitally altered in real time. The schematic flowchart of the entire basic framework is illustrated in Figure 3.1, and can be divided into 5 blocks: (1) texture acquisition and extrapolation; (2) texture manipulation; (3) geometry calculations and shape morphing; (4) final image conversion; (5) texture remapping for display. Texture manipulation algorithms were specifically designed for each different experimental purpose, and described in the methods section of each chapter. Thus, this chapter will not cover point (2).

In Chapter 2, I described the set-ups employed in this thesis' experiments and the projector used to illuminate the screen in front of the observer. The projected image consists of an object O placed within a uniform area. The geometry of the contour of the object was calculated so that it perfectly overlaps the contour of a solid object OS placed on the screen in front of the observers. The area within the contour of the object O projected to the solid object OS is defined as the ROI (region of interest). The ROI can either be: (1) the remapped chromatic texture of the natural object, obtained through the above algorithm and described in detail in this chapter, or (2) simply a uniformly coloured filled outline of the object. In the second case, blocks 1, 2, and 5 are unnecessary and only block 3 and 4 are partially performed. In describing the framework, it will be presumed that the object of which the texture belongs and the object presented in the chamber to the observer belong to the same "family" or "category", e.g banana texture on banana shape. Conversely, in Chapter 7, texture and shape of the displayed object could belong to two different categories, for example, a carrot texture on an apple shape. Therefore a sixth intermediate block was introduced; its description can be found in the methods section of Chapter 7. All software packages were written in Matlab 7.4 (the Mathworks, Inc.).

3.2 MAIN TEXTURE SYNTHESIS DESIGN

The synthesis of real object texture on a 3D real solid object naturally deals with more information than its 2D counterpart. In fact, 3D texture synthesis either implicitly or explicitly requires generation of surface geometric information and reflectance properties of the original

object from which the texture is obtained, as well as the geometry of the final object upon which it will be projected.

For 3D texture synthesis, we have to employ multiple images of the 3D object and collect enough information regarding geometric and reflectance properties as the input data (section 3.3). Once this knowledge is acquired, we can exstract the "pure" texture of the object from the scene (section 3.4) and obtain a pixel-by-pixel representation of it (block 1). In parallel, it is necessary to determine the geometry of the visible surface of a particular 3D object placed in the chamber and then to design the object's "stencil" (section 3.5). A "stencil" is a contour that, once projected on the screen in the chamber, aligns exactly with the solid object's outline (block 3). Then the flattened texture, output of the first block, is morphed to fill the stencil (section 3.6) and remapped (section 3.7) to be projected in the chamber (block 4 and 5). Each stage is divided in different steps, or algorithms, described in the following sections. It is worth notincing that *only* the geometric calibration, projector characterization, and conversion are performed to generate the uniform stimuli. The general framework is shown in Figure 3.1.

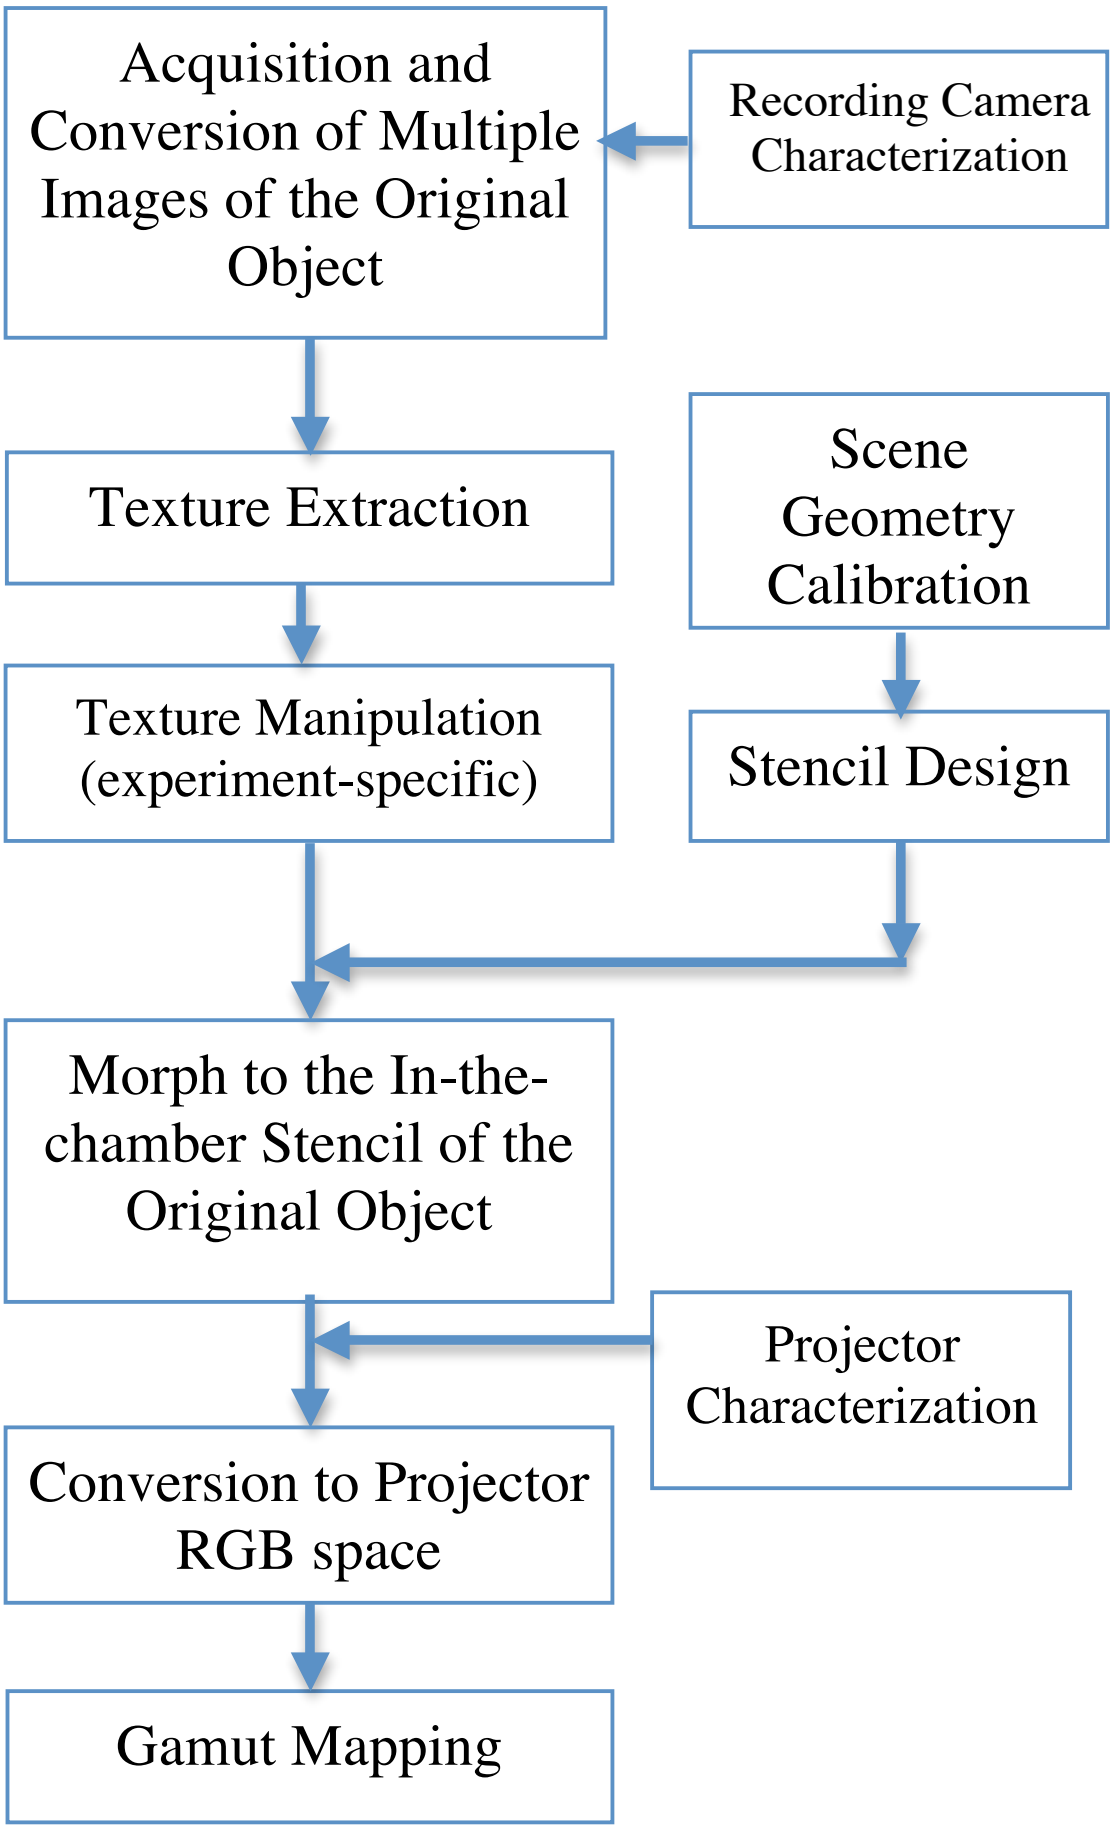

Figure 3.1 – Flowchart of the general framework of the image generation.

3.3 IMAGE ACQUISITION

According to the proposed framework, first we need to extract a suitable representation of the sample texture from a set of pre-recorded images taken by a camera. Each image recording depends on three factors: the physical content of the scene, the illumination incident on the scene, and the characteristics of the camera. Since the camera is an integral part of the resulting image, we require a camera model. The final goal of this model is to predict, from device dependent RGB values of each pixel of the image, its original CIE XYZ tristimulus values. In the next paragraph I discuss the image formation, the camera characteristics and the method used to model the camera in the specific conditions in which the images are recorded.

3.3.1 CAMERA CHARACTERIZATION

The recordings for the camera characterization were performed in a Verivide Colour Assessment cabinet containing 3 independent, stable illumination sources. The lights were metameric to: (1) standard daylight at correlated colour temperature 6,500 K (D65); (2) standard cool white fluorescent light (CWF); (3) standard CIE F light (F). Figure 3.2 illustrates for CWF, for example, the stability of the lights on two distinct days with an interval of three weeks.

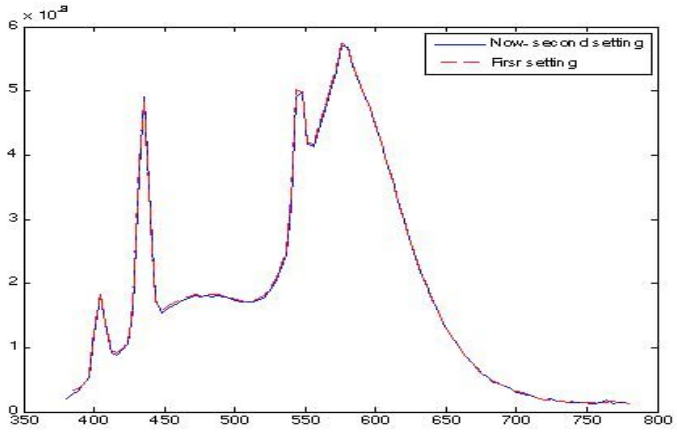

Figure 3.2 – CWF spectra measured on two distinct days. Red line: fist measurement. Blue line: second measurement (3 weeks later). Note that the two curves practically completely overlap.

The camera used is a Nikon D70 SLR camera with an 18-70mm kit lens, and was characterized to capture the chromaticity values of objects under each illumination in a specific location of the cabinet, as described later. The general idea is to measure the corresponding CIE XYZ tristimulus values and the camera RGB values of each group of paper patches and obtain the coefficients of a third order polynomial that maps one set into the other using a third order regression model. The final formula of the model is expressed in Equation 3-1; R, G, and B are the RGB camera values from which have been subtracted the camera black point; the latter was obtained by recording an image with the lens covered.

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 & . & . & \alpha_{18} & \alpha_{19} \\ \beta_1 & \beta_2 & . & . & \beta_{18} & \beta_{19} \\ \gamma_1 & \gamma_2 & . & . & \gamma_{18} & \gamma_{19} \end{bmatrix} \bullet \begin{bmatrix} R \\ G \\ B \\ R^2 \\ . \\ RG \\ . \\ R^3 \\ . \\ G^2B \\ RGB \end{bmatrix} + \begin{bmatrix} \alpha_0 \\ \beta_0 \\ \gamma_0 \end{bmatrix}$$

Equation 3-1

Initially a set of 87 Munsell paper patches, 10 Macbeth Digital ColorChecker SG chips, 7 coloured art paper pieces, and 1 Ocean Optics WS-1 Diffuse Reflectance Standard, which is >98% reflective from 250-1500nm, were chosen in a way to approximately uniformly sample the colour space both in saturation, hue and luminance. Two examples of the ordinary papers are illustrated in Figure 3.3. This was our training set for the camera model. Their Yxy values and spectrum were measured by a PR-650 spectroradiometer; the spectrum was acquired from 380nm to 780nm with a sampling step of 4nm. The angle of view was about 10 deg. The image resolution was 3008 by 2000. Successively a second set of 37 Munsell chips was measured and used to validate the model (testing set). Table 3.1 describes both sets, while Table 3.2 illustrates the tristimulus values of the art papers.

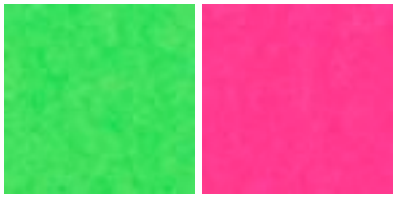

Figure 3.3 – Two examples of non-standardized paper (1G and 2O)

The main criterion of the setup is to ensure that the camera and the spectraradiometer measure the same signal, i.e the camera data for each patch had to be from the same central image location measured by the spectroradiometer. Therefore a stand was designed to place the chip always in the exact same location, as shown in Figure 3.4. The sample was manually placed behind a steel board aligned with a small rectangular aperture in the middle of the board, exposing the area of interest, but as little of the rest of the paper as was practical. The sample was attached to the steel board by magnetic buttons. The steel board was painted using Munsell N7 grey paint the same as the inside of the viewing cabinet. Some of the periphery of the pages of the Munsell book was exposed to aid alignment. This also reduced any mutual reflection with other patches in the same page. The effect of flare was small due to the smallness of the area recorded and its distance from the light source.

Traning Set: 105 sample
Munsell 5R: 3/2 – 3/6 – 5/2 – 5/6 – 5/10 – 5/14 – 7/2 – 7/6 – 7/10
Munsell 5YR: 3/2 – 5/2 – 5/6 – 7/4 – 7/8 – 7/12 – 8/8
Munsell 5Y: 3/2 – 5/2 – 5/6 – 7/4 – 7/8 – 7/12 – 8.5/4 – 8.5/8 – 8.5/12 – 9/8
Munsell 5GY: 3/2 – 5/4 – 5/8 – 7/4 – 7/8 – 8.5/2 – 8.5/6 – 8.5/10
Munsell 5G: 3/2 – 5/4 – 5/8 – 7/2 – 7/6 – 7/10 – 8/6
Munsell 5BG: 3/2 – 5/4 – 5/8 – 7/4 – 7/8 – 8/4
Munsell 5B: 3/2 – 3/6 – 5/4 – 5/8 – 7/4 – 7/8 – 8/4
Munsell 5PB: 3/4 - 3/8 – 5/4 – 5/8 – 5/12 – 7/4 – 7/8 – 8/6
Munsell 5P: 3/4 – 3/8 – 5/2 – 5/6 – 5/10 – 7/4 – 7/8 – 8/4
Munsell 5RP: 3/2 – 3/6 – 5/4 – 5/8 – 5/12 – 7/4 – 7/8 – 8/6
Munsell N: 2 – 3 – 4 – 5 – 6 – 7 – 8 – 9 – 9.5
COLOUR CHECKER: 2A - 5B – 6B – 7C – 2D – 5E – 4H – 6L – 5M – 6M
NON-STANDARD PAPERS: 2B – 3B – 1G – 1O – 2O – 3V – 1Y
REFLECTANCE STANDARD

Testing set: 37 samples
Munsell 10R: 2.5/2 – 5/8 – 8/6
Munsell 10YR: 3/2 – 6/8 – 8/10
Munsell 10Y: 4/2 – 7/4 – 8.5/10
Munsell 10GY: 5/6 – 8/6 – 9/4
Munsell 10G: 4/8 – 6/4 – 9/2
Munsell 10BG: 3/4 - 6/6 – 8/2
Munsell 10B: 3/6 – 6/10 – 9/2
Munsell 10PB: 2.5/6 – 4/6 – 7/8
Munsell 10P: 3/4 - 5/10 – 7/6
Munsell 10RP: 5/8 – 7/4 – 9/2
Munsell N: 2.5 – 3.5 – 4.5 – 5.5 – 6.5 – 7.5 – 8.5

Table 3.1 – Training (A) and testing (B) data sets description for the camera characterization

	X	Y	Z
2B	240.22	283.37	363.04
3B	107.86	140.83	290.13
1G	420.57	231.47	221.68
1O	197.62	382.66	70.917
2O	439.09	260.74	40.517
3V	508.1	342.91	59.992
1Y	402.81	567.22	90.781

Table 3.2 – CIE XYZ values of the 7 coloured art paper pieces used to characterize the Nikon D70 SLR camera (illum=D65)

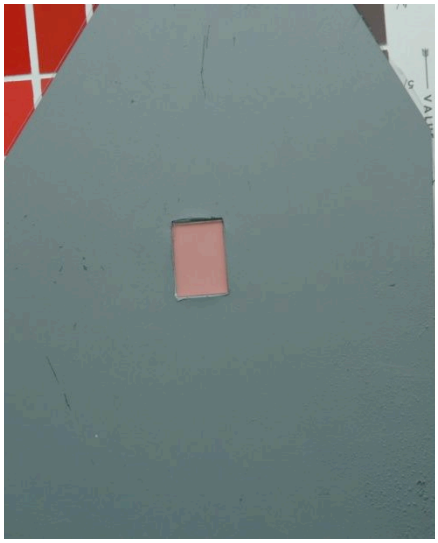

Figure 3.4 – Stand for colour chips alignment.

The camera and spectroradiometer were successively mounted on the same tripod. Rather than aim them simultaneously at the target, it was decided instead to set the optical axes to be parallel. Thus, first the entire group of camera data was recorded before capturing all the spectroradiometer measurements. Additional steps were taken to reduce the error as follows. As indicated above, it is important that the camera and the spectroradiometer are exposed to the same signal. In order to minimize the effect of misalignment, the illumination was made as uniform as possible and a window from the image was extracted, which corresponded as closely as possible to the area used by the spectroradiometer. The pixels in this window were averaged.

Both a tilt standard white sample and Macbeth colour checker were used to decide the exposure time and the white balance of the camera for each illuminant, while the focus was manually adjusted before each picture. The colour patches were recorded under the 3 illuminants (D65, CWF, and F), although not all of them were used in all the experiments. Figure 3.5 illustrates the chromaticities of the training and testing sets under the three lights in u' v' isoilluminant space. Once the XYZ and the RGB values of the training data sets are known, a characterization model can be fitted using a third order polynomial regression model.

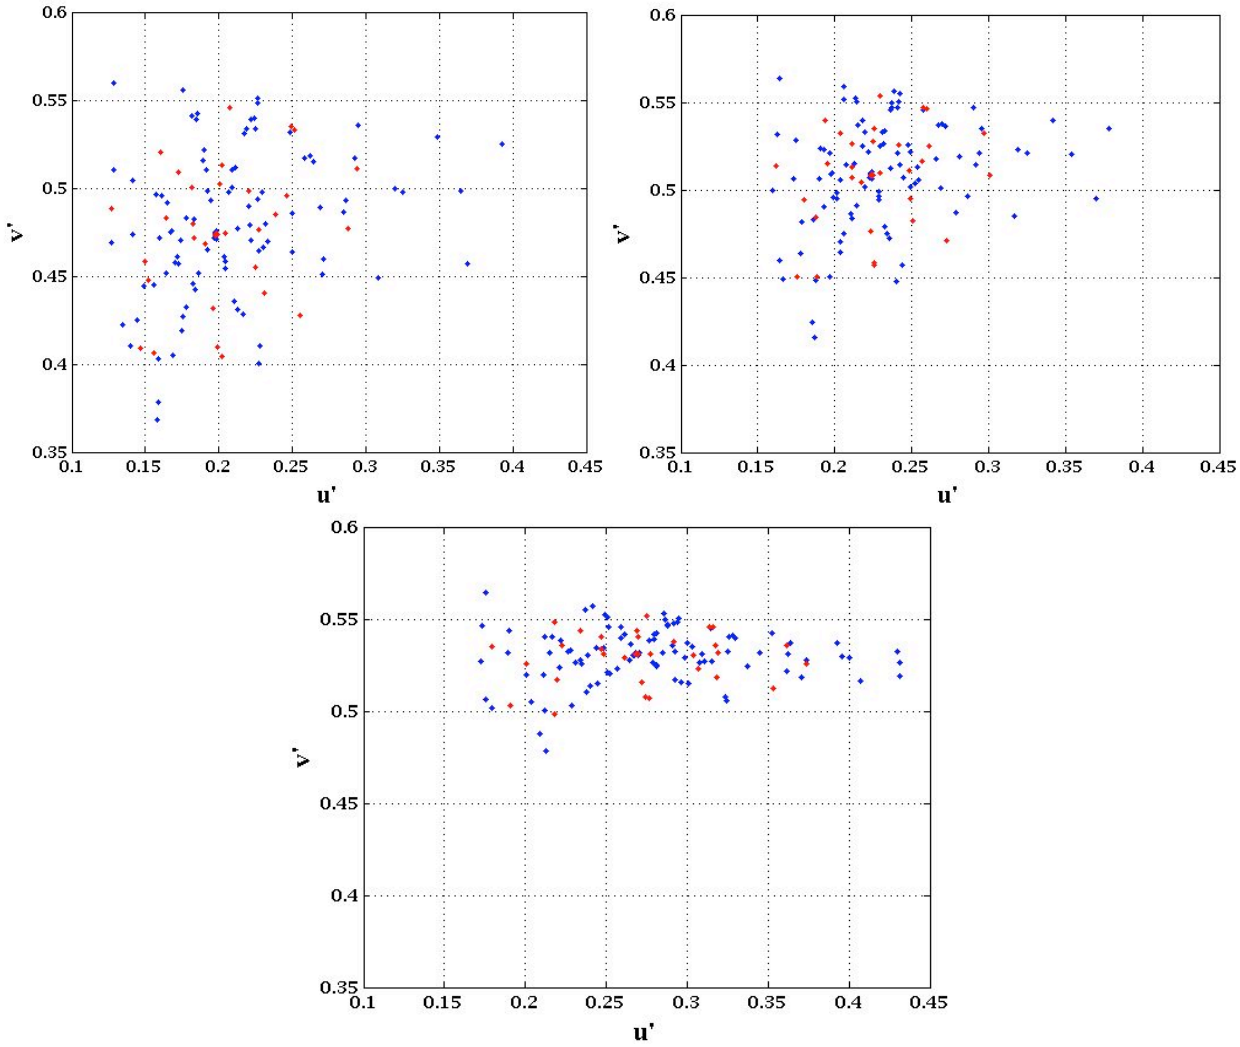

Figure 3.5 - Chromaticities of the training and testing sets projected on the u' v' isoluminant plane. Top left: D65. Top right: CWF. Bottom: F. Blue dots: training set. Red dots: testing set. Note the compression of the sets for F along the CIE v' axis.

To assess the model performance, "model error" was computed, defined as the mean deviation between the measured and the predicted tristimulus value over the set of testing or training patches. It was calculated in perceptually uniform space Lu'v'- yielding an average ΔE; this defines the accuracy of the prediction. For both sets, ΔEs are low, comparable to 3 JND[3], and consistent among illuminants, revealing good performance of the model. A summary of the model deviation from the real values (i.e. the model error) measured in perceptual difference, ΔE, for the testing and training sets are listed in Table 3.3. Alternatively, for each illuminant, the measured XYZ value of the training and testing sets were transformed in Lab colour space, and the Lab and RGB training values were used to fit the regression model. For each illuminant, the XYZ values of the reference white employed for the conversion to Lab space were measured using an Ocean Optics WS-1 Diffuse Reflectance Standard. Applying this alternative method I re-calculated the accuracy of the prediction of the model; result showed no difference between the two alternatives (overall mean difference 0.01).

[3] As said in Chapter 1, one just noticeable difference (JND) between colours is classically quantified by 1 ΔE.

	D65 (ΔE)	CWF (ΔE)	F (ΔE)
Training set (ΔE)	mean = 2.840 min = 0.454 max = 10.07	mean = 3.604 min = 0.306 max = 14.86	mean = 4.227 min = 0.613 max = 28.73
Testing set (ΔE)	mean = 2.988 min = 0.451 max = 8.518	mean = 3.638 min = 1.283 max = 8.001	mean = 3.714 min = 0.304 max = 9.645

Table 3.3 - Mean, max and min ΔE values of the model performance for the training and testing set.

Once the camera characteristics have been determined, it is possible to photograph a real textured object, segment it from the background and convert each pixel RGB value into its corresponding CIE XYZ tristimulus values.

3.4 TEXTURE EXTRAPOLATION

After the segmented image is transformed into XYZ colour space, the object's texture can be analysed and extrapolated. The aim of this procedure is ultimately to extract veridical surface texture of the object due solely to intrinsic variation in surface reflectance, and specifically not to record variation in chromaticity or luminance due to geometric factor such as illumination, object 3D shape or scene configuration. To this end, ***before*** photographing the natural object, I avoided recording its glossiness spraying its surface with full-matte totally transparent spray paint (Vanish Matt MARABU finishing). Nonetheless in some cases a further correction was necessary as described in the next section.

Therefore, at this point, we have a 2D image of a 3D object and the luminance changes on it is due to both the intrinsic characteristics of pigment of the object's surface and the shading gradient dependent on the position and geometrical properties of the light source (for example a linear tubular lamp above the object). Shading removal is the necessary step to subsequently process the image. Several techniques are available to remove shading from an image; the most common are briefly presented in 3.4.1. This thesis' approach is described in 3.4.1.3.

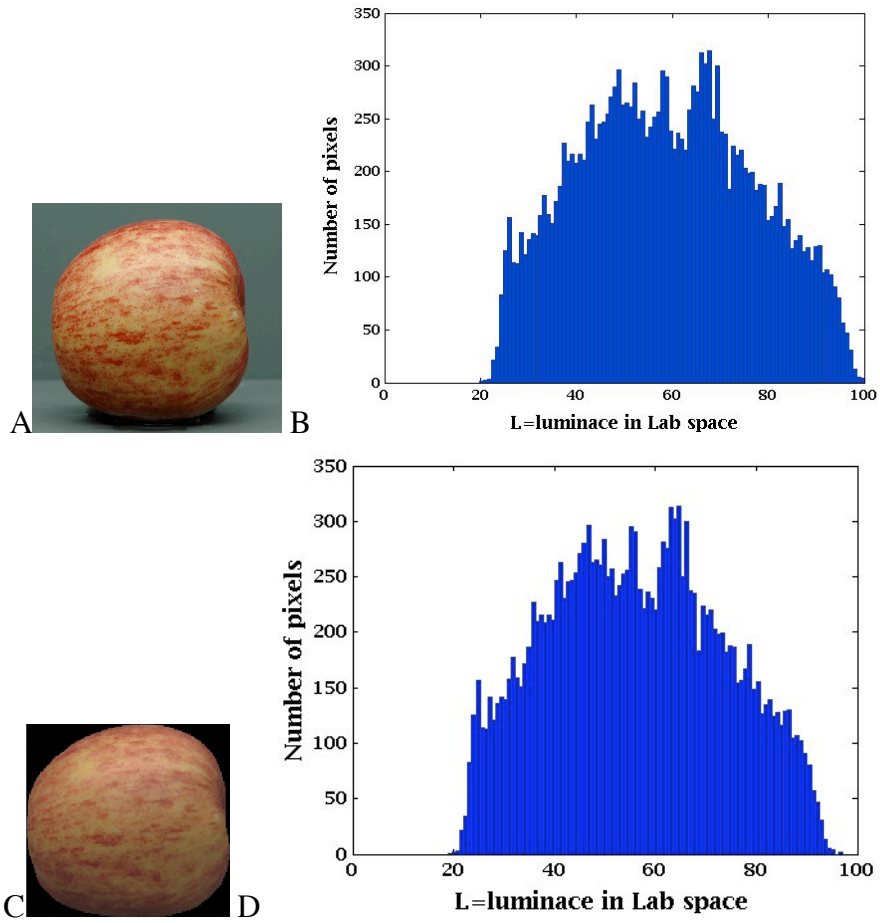

Figure 3.6 – Highlight correction. A) Apple image before correction. B) Histogram of the number of pixels as a function of the luminance value (in Lab) for the apple image before highlight correction (i.e. subfigure A). C) Apple image after highlight correction. B) Histogram of the number of pixels as a function of the luminance value (in Lab) for the apple image after highlight correction (i.e. subfigure C).

Despite the removal of the natural glossiness of the object surface through the application of matt transparent paint (see above), some objects may still have some residual glossiness before the shading removal. In these few cases, we found that less than 0.5% of image presented highlights (i.e. less than the 0.5% of the pixels). Thus, for these singular cases the following method was applied. The image was first transformed from XYZ to Lab using, for each illuminant, as reference white the measure obtained from a Ocean Optics WS-1 Diffuse Reflectance Standard. Then all points found with a luminance L greater than 99 were equalized to the mean of the 8 neighbour values below 99. The reason for the choice of this value (i.e. L=99) is explained in 3.4.1.3. Figure 3.6 shows the image of apple before (A) and after (C) highlight correction and the corresponding histograms of the luminance (before=subfigure Figure 3.6B and after=subfigure Figure 3.6D). Note that only the tail of the histograms is modified.

3.4.1 SHADING CORRECTION

Shading might be caused by non-uniform illumination, non-uniform camera sensitivity, or even dirt and dust on glass (lens) surfaces. A general model of the shading effect (Gonzalez&Woods (2007)) states that for each point in the field-of-view, the illumination E(x,y)

interacts in a multiplicative way with the object surface reflectance R(x,y) to produce the image J(x,y) Equation 3-2. In general, it is assumed that E(x,y) is slowly varying compared to R(x,y).

$$J(x,y) = R(x,y) \bullet E(x,y) \qquad \text{Equation 3-2}$$

In general, the camera response C(x,y) will be proportional to the image irradiance J(x,y) by a multiplicative factor (the gain (*g*)) and an additive constant (the offset (*off*)) (Gonzalez&Woods (2007), Russ (1995)), as in Equation 3-3.

$$C(x,y) = g(x,y) \bullet J(x,y) + off(x,y) \qquad \text{Equation 3-3}$$

In literature, we can distinguish between two cases to determine R(x,y) starting from C(x,y), and leading to the estimation of the terms g(x,y)*E(x,y) and off(x,y). While in the first case it is assumed that only the recorded image C(x,y) is available (a posteriori estimate), in the second case it is possible to record one or more additional calibration images (a priori estimate) (Madisetti et al. (1998)).

3.4.1.1 A POSTERIORI ESTIMATE

This retrospective method attempts to extract the shading estimate exclusively from C(x,y), i.e. from only one final image. There are two main techniques following this approach: lowpass filtering and homomorphing filtering (Russ (1995)). In both, the underlying assumption is that all luminance changes due to the illumination rather than the surface reflectance of the object are slowly varying over space and therefore may be removed by removing the low spatial frequency components of luminance.

In the lowpass filtering method, C(x,y) is first transformed into its Fourier image, then smoothed through a lowpass filter and subtracted from itself ((Gonzalez&Woods (2007)). Last, the zero-frequency component is restored (*const* in Equation 3-4). This process is expressed in the following equation:

$$H(\omega) = C(\omega) - (C(\omega) \otimes LP(\omega)) + const \qquad \text{Equation 3-4}$$

where H(ω) is the estimate of R(ω) in the Fourier domain, LP(ω) is the lowpass filter, and $\otimes$ represents convolution in the Fourier domain.

The homomorphic filtering starts from the assumption that the camera offset is set to zero (Gonzalez&Woods (2007), review in Tomazeric et al. (2002)). Therefore C(x,y) consists solely of multiplicative terms. If we calculate the logarithm of C(x,y) then we can write:

$$\ln[C(x,y)] = \ln[g(x,y) * E(x,y) * R(x,y)] = \ln[g(x,y) * E(x,y)] + \ln(R(x,y)] \qquad \text{Equation 3-5}$$

Furthermore, since E(x,y) is slowly varying comparing to R(x,y) then we can suppress shading by high pass filtering the logarithm of C(x,y). Finally to restore the image we simply need to do the exponent (inverse logarithm) of the output. Thus, in formula:

$$H(x,y) = e^{HP\{\ln(C(x,y))\}} + const \qquad \text{Equation 3-6}$$

Both techniques involve removing low spatial frequencies. But the threshold for the filters must be determined before their application. Clearly this threshold may vary with the type of image, and thus there are limitations to the generality of the techniques.

3.4.1.2 A PRIORI ESTIMATE

If it is possible to record more than one image of the same object or scene through the cameras system, then the most appropriate technique for the removal of shading effects is to record two images: (1) a baseline or 'black' image, that is generated by covering the lens, implying J(x,y) = 0 which leads to $C_b(x,y) = off(x,y)$, and (2) a 'white' image, generated by using R(x,y) = 1 which gives $C_w(x,y) = g(x,y)*E(x,y)+off(x,y)$ (Young (2000)).

The formula is:

$$H(x,y) = k * \frac{C(x,y) - C_b(x,y)}{C_w(x,y) - C_b(x,y)} = \frac{g(x,y) * E(x,y) * R(x,y)}{g(x,y) * E(x,y)} = R(x,y) \qquad \text{Equation 3-7}$$

The constant term *k* is chosen to produce the desired dynamic range. When possible, this technique was found able to best correct the presence of shading on an image (Young (2000))

3.4.1.3 SHADING REMOVAL: THIS THESIS APPROACH

Being able to record more than one image of the object in the same settings, a modified technique of the latter approach is implemented in this thesis. To produce the "white" image, the original object is photographed after being spray-painted in grey in the exact same location as in the original image. Then the object is segmented from the background and, pixel-by-pixel, converted into CIE XYZ tristimulus values. An example of the original object image and of the "white" image is shown in Figure 3.7 (Gala apple). Next the luminance channel of the white object is filtered with a Gaussian bidimensional low pass filter of mean 10 pixels and standard deviation [5,5], to reduce the effect of noise produced by the paint's roughness. The values were determined to be optimal by trial-and-error and tested over two subjects (YL and AH) and the author. For each point of the texture of original image (e.g. Figure 3.7A for the Gala apple), its luminance value is divided by the respective luminance value of the "white" image and scaled by a constant *k* (which accounts for the projector's dynamic range as in Equation 3-7. The black image was obtained recording an image with the lens covered. The final result is shown in Figure 3.7C.

A)  B) 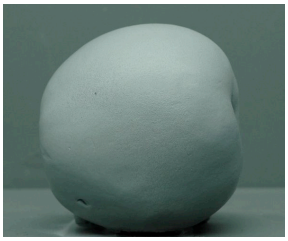 C) 

Figure 3.7 – A) Image of a Gala apple used in the experiments; B) its "white" reference image; C) Extracted texture 2D image without highlights and shading.

It should be point out that the luminance value of the "white" reference object was found everywhere lower than 99 (after the image was converted in Lab space as described above). Since, as explained above, the luminance of the original object was equalized to values below 99 (in Lab space) then the luminance of the original object is lower or equal to the luminance of the white reference. Note that the choice of the equalizing value was therefore determined "a posteriori" considering that the matt grey painted object does not possess highlights. Numerous alternatives were tested, resulting in this the most visually pleasant as tested by two subjects (YL and AH) and the author in a pilot experiment.

A two-piece removable platform was designed to constrain the position of the object to the same location before and after painting. The platform consists of two parts: 1) a PVC rectangular plate (dimension 180mmx80mmx3mm) from which has been cut out a cross shaped piece in the centre; 2) a wood block (dimension 170mmx80mmx100mm) in which centrally has been carved a cross of the same dimension of the plate and inserted a removable PVC cross. The wood block is solidly attached to the Verivide cabinet while the plate can be removed and wedged back in again at the position of the cross. A photograph of the platform is shown in Figure 3.8.

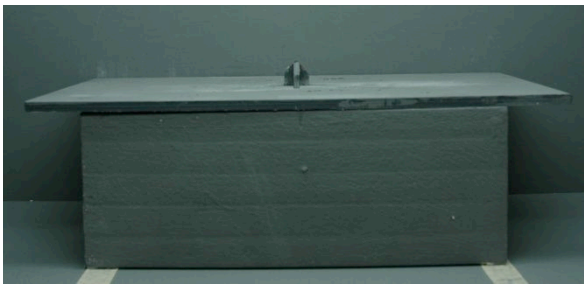

Figure 3.8 – Photograph of the platform used to place the objects in the cabinet.

3.5 GEOMETRY CALIBRATION AND PROFILE DESIGN

As said previously, the contour of the ROI has to overlap exactly the outline of the real object placed inside the box. Therefore a method was developed to estimate the surface geometry of the object as seen from the observer point of view and obtain the stencil of the object. As two experimental apparatuses were used (Chapter 2) we require two separate geometrical calibration algorithms.

For the experimental box described in Chapter 2 section 2.2, the description of the algorithm can be found in Y. Ling's thesis (2005). For the second apparatus (Chapter 2 section 2.3), I developed a similar algorithm taking into account the geometry of the chamber.

The general idea of our geometrical calibration can be briefly described as follows. A Canon PowerShot S45 digital camera takes a picture of the object in the location in which it will be placed during the experiment. Then the object's contour is segmented from the background; a coloured paper is positioned behind the object for better segmentation and precision of the stencil. Next, a programme creates a lookup table which stores each pixel projected by the data projector and its corresponding position within the picture taken by the camera. The two sets are compared to generate the stencil. Ultimately, the stencil is manually refined for a perfect overlap.

3.6 TEXTURE MORPHING AND RESIZING

At this point we have two images: the first contains information about the original object (the shape of its contour and its texture) while the second contains information about the shape of visible surface of the object seen from the observer point, i.e. the stencil of the object. Since the final goal is to obtain an image possessing the original object texture (first image) within the contour of stencil (second image), we had to perform a "morphing" of one image over the other. This is a delicate process consisting of altering the spatial distribution of the points of the original image without corrupting the original texture structure. A piecewise linear spatial transformation was employed (Goshtalsby&Ardeshir (1986)). Initially a number of control point pairs were selected on the two images (examples in Figure 3.9). Each point on the original image and its correspondent on the stencil were selected on the contour or on salient location (relatively to the original texture). The transformation was performed using the control points as anchors for the spatial morphing triangulation (Goshtalsby&Ardeshir (1986)). The number N of control pairs and the technique of transformation (i.e. the piecewise linear spatial transformation) were chosen taking in account: (1) the fidelity in reproducing the same texture, (2) the overlap of the 2D texture structure with the surface of the object 3D structure, and (3) the percentage of point of the morphed texture overlapping the stencil. First I searched for an N value such that less than 1% of the points of the morphed texture were outside the stencil contour (Figure 3.11). Then the appearance of the texture was tested by the author and two subjects (AH and YL) as to satisfy the points (1) and (2) above. Specifically point (2) requires that the texture spatial characteristics overlap the 3D object surface features; for example the real banana texture possess distinctive features where its skin bends (indicated by the red arrow in Figure 3.10A). Therefore, when representing the texture on the 3D replica, the same features have to overlap the replica surface in similar locations (black arrow in Figure 3.10B). Points within the ROI and the stencil of each object were selected in order to assure

such overlap (points A and A' in Figure 3.11 – salient locations). The final number of pairs varied between 47 for the banana to 6 for the apple. Note that, being the transformation linear, the value of each point chromaticity is not altered, but only the position of a point relative to its neighbours will linearly change.

The method above can be applied when the chromatic texture of one object A had to be rendered on: (1) the stencil of same object A, (2) the stencil of a smaller object B or (3) the stencil of an object C 1.1 times bigger than A. Conversely, when the chromatic texture of the object A had to be applied on the stencil of a different object D similar in size to A, the texture of the object A was *first* process using a technique described in Chapter 7 and then morphed to fit the stencil of object D.

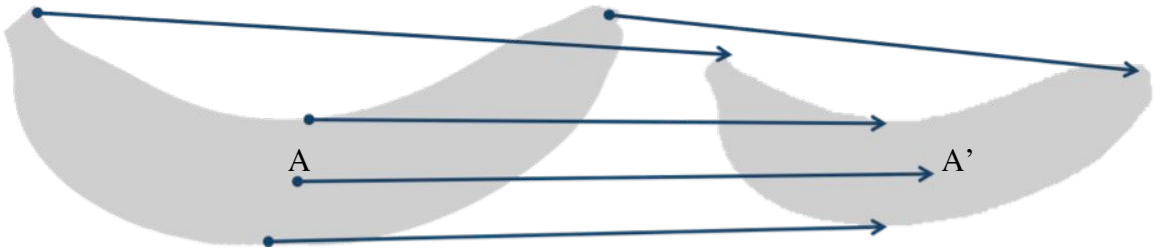

Figure 3.9 – Example of the morphing between the banana original profile and the banana stencil. The point A of the ROI is paired with the point A' inside the 3D object stencil.

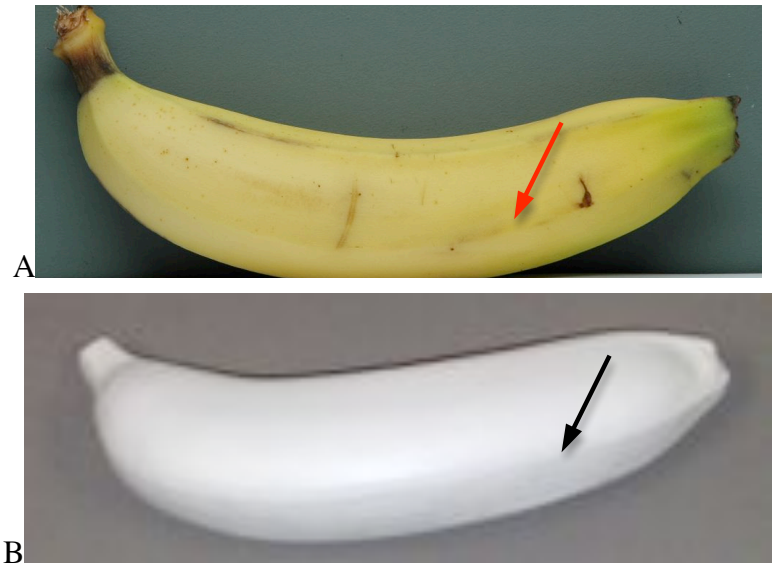

Figure 3.10 – A) Photograph of the original banana. The red arrow points to a location of the banana chromatic surface texture with a specific direction correspondent to a shape cue. B) Photograph of the replica of the banana replica. The black arrow points to the correspondent location of the one pointed by the red arrow in subfigure A.

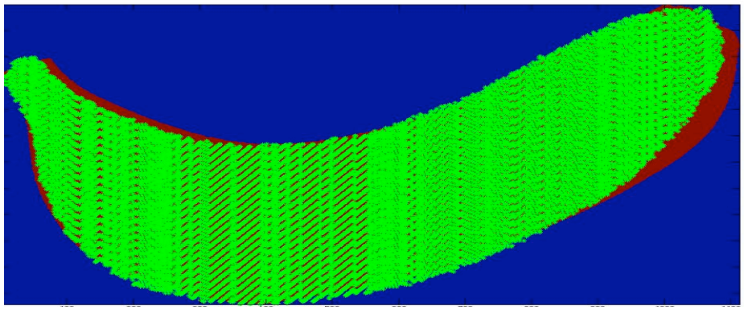

Figure 3.11 – Difference between the banana stencil (green dots) and the banana morphed texture profile (filled red outline) for 42 points.

3.7 IMAGE DISPLAY

After the image has been transformed in XYZ, the texture extracted, manipulated and morphed to the new geometry, then a new image has to be produced. The new image contains the processed image in a selected uniform background (of which colour depends on the experiment). Finally the image has to be transformed from CIE XYZ coordinates in RGB device dependent coordinate of the output device, that is, in our case, the data projector.

Because two different projectors were used in this thesis, two methods were used to characterize the projectors. Details on the first set up projector (DLP data projector Infocus LP530, 1024x768 resolution, 400:1 contrast ratios, and 2000 lumens brightness) and its characterization are presented in Y. Ling thesis (2005). Next section, instead, describes the characterization of the second projector used in the set-up described in Chapter 2 section 2.3.

3.7.1 STIMULUS DISPLAY AND PROJECTOR CHARACTERIZATION

The general paradigm for the characterisation of colour devices requires the derivation of suitable transformations that map between device dependent colour system and device independent colour system, i.e. to convert one colour space to another. Empirically, for output devices, this can be accomplished by showing a set of colour patches with preselected device dependent values and measuring its device-independent colour coordinates. Such data are used to derive a forward device response function that maps device control values to the device independent values. This approach demands collecting a large number of measurements. Finally to validate the model a second set of colour patches are measured and compared with predicted values of the model. Thus, two sets of data are usually required: a training set, to produce the model of the projector, and a test set, to validate the model and measure its accuracy.

However, before characterising the colour space of the projector, the optimal combination of its "contrast" and "brightness" settings had to be determined, to ensure both broad dynamic range and light contrast. The projector used in the second set-up was LCOS reflective display projector (Canon Xeed SX6) with a spatial resolution of 1280X1024 at a refresh rate of 70Hz, and driven by a 256MB ATI® Radeon™ HD 3650 graphics card. The projector image was set at "Standard" and the image was tilted by 2%. Figure 3.12 shows the relationship between digital input and luminance measurements (the tone reproduction curve), under different "brightness" and "contrast" settings. The measurements were taken from the viewing hole of the chamber on the opposite screen using a Photo Research PR-650 spectroradiometer. Figure 3.12 shows that at "brightness" 20 and "contrast" 0 the maximum dynamic range is achieved and the linear relationship between digital input and

luminance is optimized. However at "brightness" 0 and "contrast" 0, with a digital input of RGB=[0,0,0] we obtain the lowest luminance for a wide dynamic range. Hence, depending on the experiment needs, one of the two settings was employed.

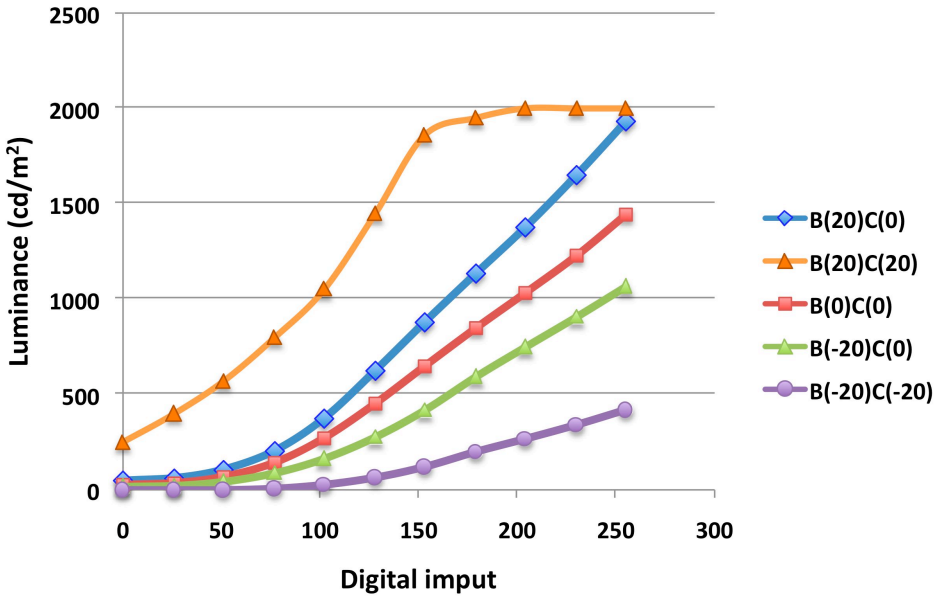

Figure 3.12 – Relationship between digital input and luminance measurement under different "brightness" and "contrast" settings. Green curve 100% brightness-100% contrast. Blue curve: 100% brightness-50% contrast. Red curve: 50% brightness-50% contrast. Black curve: 50% brightness-100% contrast.

In addition, the spatial non-uniformity of the projected image had to be tested. This is because most of the experiments in this thesis were made using several objects at different position of the screen, thus it was extremely important to calculate the variation of colour along the two main axis of the display (vertical and horizontal). A uniform grey image filling the screen was projected and then 35 points at a distance of 100 mm from each other in both cardinal directions were measured with a PR-650 spectroradiometer. The data show that the chromaticity values are only slightly affected by the position of the recording (x=0.3353±0.0035, y=0.3981±0.0042), whilst the luminance varies moderately across locations (1443.4±100.76). To circumvent such problems, for each experimental location and object, I collected a different set of training and testing data.

The training set consisted of colour patches of the red, green and blue primaries at increments of 5 in digital counts from 0 to 255, i.e. RGB=[0 0 0, 5 0 0, …, 255 0 0, 0 5 0, …, 0 255 0, 0 0 5, …, 0 0 255]. Their CIE XYZ coordinates were measured using a Photo Research PR-650 spectroradiometer in the experimental location on central point of the objects. Thus, for *each object*, (N for experiment 1; M for experiment 2; etc) 52 recording were stored for each of the three primaries. Presuming the additivity and linearity of the three colour channels, then the linear combinations of the data could be a possible transformation between these colour spaces. However, it resulted that to preserve the additivity along the three colourimetric dimensions, the data had to be first corrected by subtracting out the tristimulus values for the black measurement (RGB=[0 0 0]). After this correction, the data where used to obtain the transformation matrix of the linear

combination, or forward device-response function. The inverse of this function maps each device-independent colour to the device-dependent colour value of the output device. Finally a set of 100 random RGB values is generated and its XYZ values measured as above. These values are then compared with the predictions obtained by the inverse model of the projector. The mean CIE Luv colour difference between the measured and the predicted test colours were 3±0.4 for the inverse model and 2±0.3 for the forward (average over all points and objects).

The stability of the projector was checked for few days priori each experiment with no noticeable difference among measurements. Note that we have as many model matrices as are object-location combinations (27 different matrices in all). Furthermore the projector was recalibrated at least twice for each experiment.

3.7.2 GAMUT MAPPING ALGORITHM

A common problem when reproducing colour images between dissimilar media (as are a digital camera and a data projector) resides in the diversity of the set of colours that the two devices can reproduce. Therefore, after transforming the image in digital RGB values of the projector using the colour conversion model describe above, we face the problem of overcoming the differences between the set of colours obtainable by the projector and the one required to faithfully reproduce the image perceptual appearance. With the term "colour gamut" we will refer to the range of colours achievable by a given colour reproduction medium under a certain viewing condition and with "gamut boundary" the surface of the volume containing such set of colours. Thus, if the colour gamut of the projector does not include the entire set of colours necessary to reproduce a certain image, then we need a method for assigning the colours that cannot be properly converted (colours "out-of-gamut") to the least perceptually different colours contained in the projector colour gamut. A large number of algorithms have been implemented to map between two gamuts. In the next two sections (3.7.2.1 and 3.7.2.2) is given a concise overview of the two most popular approaches: gamut clipping and gamut compression. A more extensive review can be found in Green&MacDonald (2008) or Reinhard et al. (2008). The last section describes instead the approach pursued in this thesis.

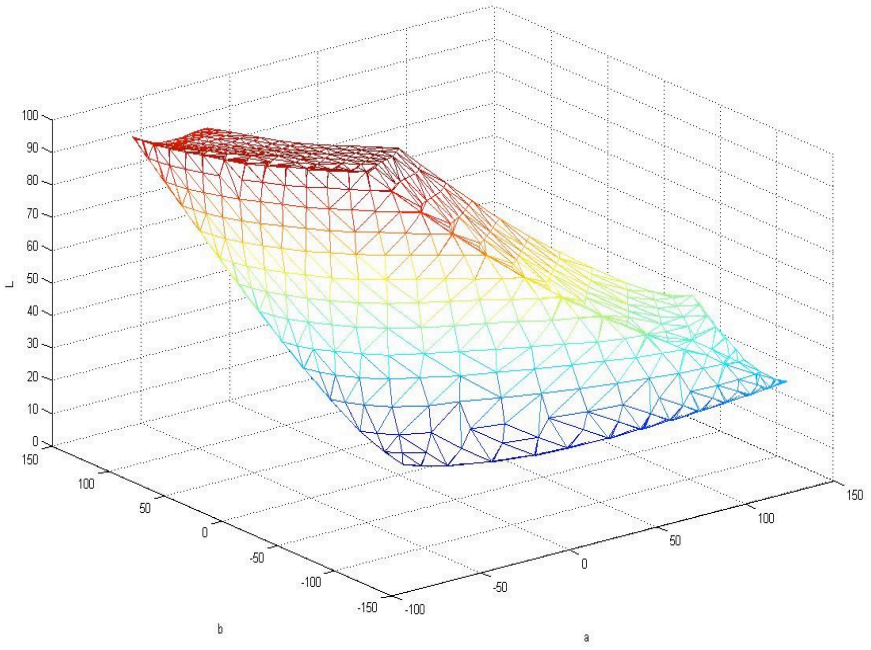

Figure 3.13 – Gamut volume of the projector (Canon Xeed SX6) used in the chamber.

Even though there are diverse mapping techniques, they all have in common an initial step: the need of computing the gamut boundary of the projector (output medium). First, we sample each primary of the RGB device-dependent colour space in N steps and converted in Lab colour space using the projector model described above. Then each sample is connected with its neighbour forming a tessellation that can be approximated by a volume. The surface of this volume is our gamut boundary. For the purpose of this thesis I used N=16 steps, aiming to obtain a tessellation not too coarse or too fine. The projector's gamut volume is shown in Figure 3.13.

3.7.2.1 GAMUT CLIPPING

Gamut clipping algorithms change only the colours that are outside of the destination gamut either from the very beginning or after lightness compression (review in Reinhard et al. (2008)). Although there are a variety of them, all clipping algorithms specify a mapping criterion, which defines how to project the out-of-gamut points of the input medium onto the output reproduction gamut boundary. The input points inside the destination gamut remain untouched. Different clipping algorithms differ mainly in the direction of the projection (Morovic&Luo 2001)). Hue-Preserving Minimum ΔE (HPminΔE) algorithms (CIE 2004) search for the points on the output gamut boundary that have the same hue as of the out-of-gamut colours and are closest to them (i.e. with minimum ΔE colour difference). SCLIP algorithms project the input points toward the colour space centre (the 50% grey point) onto the gamut boundary at constant hue. Finally, CUSP algorithms move out-of-gamut colours toward the point on lightness axis with the lightness of the

cusp on the gamut boundary at constant hue, where the cusp is the point in the output gamut with the maximum chroma at that hue (Morovic&Wang (2003)).

3.7.2.2 COMPRESSION ALGORITHMS

It is important to notice that a colour mapping transformation must modify the image so that all the colours are within the destination gamut, while trying to give a reproduction that is pleasant to look at and as accurate (close to the original) as possible. In other words, the presence of visible artefacts, as inappropriate contours in the image, or lack of contrast in the image would conflict with the above consideration. For such reasons clipping could be a "crude" gamut mapping technique when there is a large difference between the source and the destination gamuts. In such a case, compression methods are more adequate for the task (Morovic&Luo (2001))

Compression gamut algorithms not only change the colour of the input image that are outside the destination gamut, but also changes at least some of image's colours inside the output gamut (Bonnier (2008), Bonnier (2007)). The idea is to use both the input and output gamut to compress the entire set of colours of the first to the range of the second. The compression can be performed uniformly or non–uniformly with different parameters depending on the original gamut. An simple example of uniform compression is the linear compression in which all point are compressed linearly along the luminance axis maintaining chroma and hue unchanged (for review see Fairchild 2005, Braun&Fairchild (2001)); alternatively, luminance and hue can be maintained and the chroma can be compressed toward the centroid of the output gamut (Fairchild 2005, Braun&Fairchild (2001)).

On the other hand, the most common non-uniform approaches are the following (reviews in Morovic&Luo (2001), Bonnier et al. (2007), Bonnier et al. (2008)): (1) non–linear compression; (2) piece–wise linear compression; (3) higher–order polynomial compression; (4) combination of different compression methods (e.g. topographic algorithms).

3.7.2.3 GAMUT MAPPING: THIS THESIS' APPROACH

In this thesis we applied a blend of two different techniques: gamut clipping and linear compression. This non-linear compromise between gamut mapping methods is close to the "soft clipping". This choice was performed considering that the image's and the projector's gamuts are rather similar in terms of hue and chroma and differ mostly in the lightness' range. Furthermore, less than 0.1% of the colours of the points of the original image resulted out-of-gamut after lightness scaling, and even after extreme alteration of the texture's original hue for certain experiments, no more than the 20% of the colour resulted out of the projector's gamut.

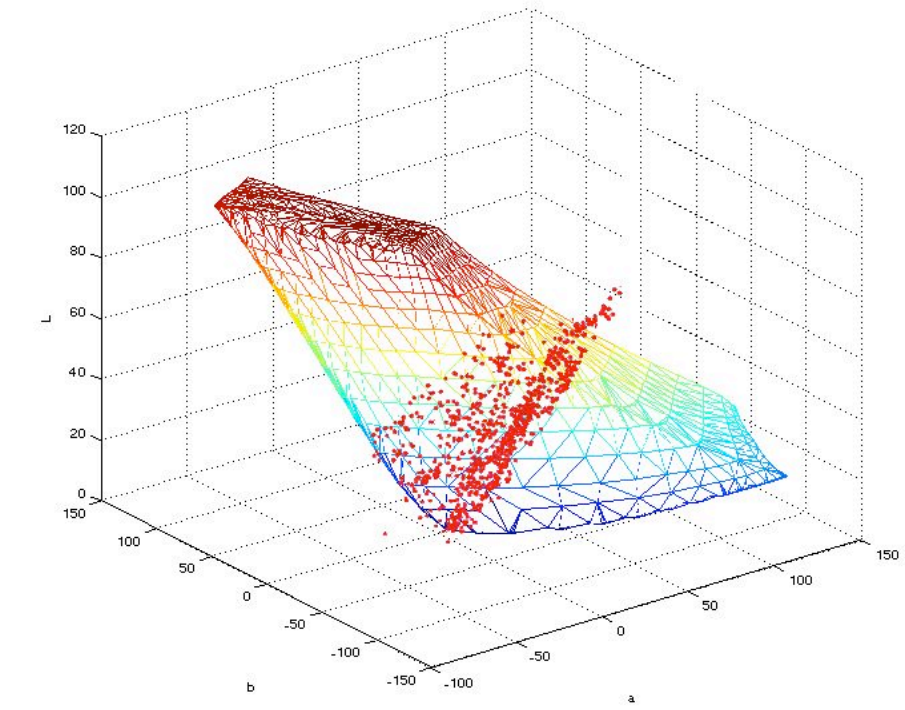

A

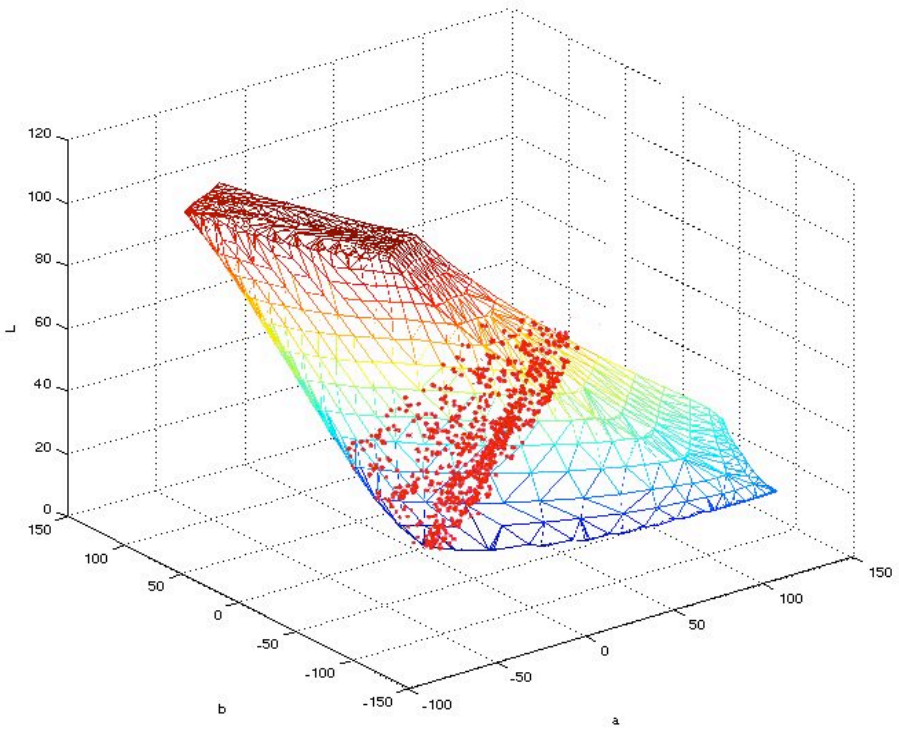

B

Figure 3.14 – Gamut of the projector (cloured tessellation) and distribution of the colours of the banana (red points) before (A) and after (B) gamut mapping for the illuminant D65.

In general the goal of soft clipping is to compress the highest chroma value (or lowest luminance values) of ***all*** points and leave unchanged the lowest chroma values (or highest luminance values). In the presented approach the points compressed are only the ones close to the gamut boundary (using a 5th order regression model). Specifically, this method selects the points with ΔE difference from the boundary equal or lower than 1 and compresses the lowest luminance values to the lowest boundary of the gamut and the highest luminance values to the highest

boundary of the gamut, maintaining hue and chroma unchanged. Then, the remaining out-of-gamut points are clipped using the HPminΔE or the SCLIP algorithm depending on the pleasantness of appearance of the output image (in general 3 figures were generated for each texture/condition in the following experiments, with an average of 200 images for one object). Figure 3.14 gives an example of the distribution of colours of the banana under the illuminant D65 before (Figure 3.14A) and after (Figure 3.14B) gamut mapping, while Figure 3.15 shows the image of the same banana before (A) and after (B) gamut mapping. Note that Figure 3.15B is also a final image of the banana, i.e. the one showed to the observer. Figure 3.15C is the original image of the banana

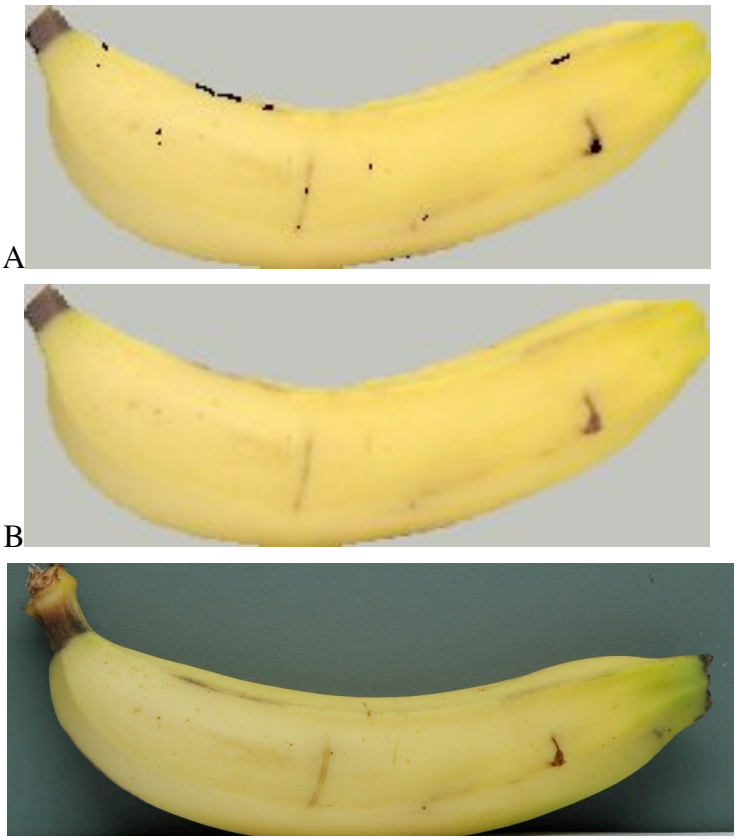

Figure 3.15 – Banana image before (A) and after (B) gamut mapping simulated under the illuminant D65. Note that subfigure B is the final image of the original banana texture on the actual banana stencil. Subfigure C shows the original banana image taken in the Verivide cabinet under D65.

3.8 LIGHT CHARACTERIZATION

As seen in Chapter 2, the sidewalls of the experimental chamber were covered with white paper as well as the lateral sectors of the front wall. Two sets of three Osram fluorescent lamps (red, green, and blue) illuminated these surfaces from the two sides, whereas side panels and a frame block the light reaching the wall at the far front end of the chamber where the image is projected. The lights were characterized using a simple function to linearly interpolate the relationship between voltage input and lamps output. In other words, the chromaticities of the three lamp pairs were measured on four locations one for each reflecting surface (two sidewalls an two lateral

sectors) with the Photo Research PR650 spectroradiometer. The examination of final data indicated that the lights produce practically the same gamut on all four surfaces. Hence, a singular set of look-up tables can be used to correct for nonlinearities in the voltage input/light output relationship. In addition, the additivity of the three lamps was verified holding over the range of intensities used in this thesis.

Finally, it should be noticed that the intensity of the light produced by the projector measured at the far end wall is 2000 cd/m^2, while lamps' light that reaches the same area has an intensity of at the most 3 cd/m^2. Therefore, the contribution of the lighting system is irrelevant. Nevertheless, for completeness' sake, I tested 200 projected colours in an experimental location and measured the colour differences between when the lamps were switched on at their maximum intensity and when they were switched off completely. The maximum difference obtained was 0.5ΔELUV.

3.9 CHANGING THE SURFACE COLOURS OF THE TARGETED OBJECTS BY A JOYSTICK

The observers were asked to perform different task in relation to the experiment. In some they were asked to give simple responses (yes/no, 2AFC, or 3AFC), whilst in other they were allowed to modify the surface appearance of the objects. At the beginning of each experimental session, the subjects were handed an input device, which was either an 8 button USB gamepad or a 3-button Logitech mouse. To operate the joystick within the Matlab environment, the Data Acquisition Toolbox's Adaptor Kit in Microsoft Visual Studio C++ 6.0 environment was used to build an adaptor that enables Matlab to exchange data with the gamepad. To use the mouse instead I employed Matlab 7.4 with the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Then a routine was designed specifically for each experiment.

As a final note, the projected images used in this thesis' experiments are stored either as indexed colour images or as numbered TIFF images. Therefore the access of the specific image is straightforward and the input devices described above are sufficient to allow the subjects complete manipulation of the apparent object surface colours.

Chapter 4

INFLUENCE OF SHAPE CUES ON COLOUR APPEARANCE OF FAMILIAR NATURAL OBJECTS UNDER ILLUMINATION CHANGES

4.1 INTRODUCTION

The main goals of the experiments and analysis presented in this chapter are to (1) demonstrate that humans possess a range of memory colours per familiar object, (2) examine the relationship between memory colour and shape under (a) constant illumination and (b) change of illumination, and (3) explore the relationship between speed of the response, memory colour and shape.

The *incipit* of the first objective is in the words of Hering's (1964): "the memory colour of an object need not be rigorously fixed but can have a certain range of variation depending on its derivation". Although a large number of experiments have studied memory colour, finding "a certain range of variation" across subjects (Bartleson, 1960; Humphrey, Goodale, Jakobson, & Servos, 1994; Newhall, Burnham, & Clark, 1957; Perez-Carpinell, Baldovi, de Fez, & Castro, 1998; Siple & Springer, 1983; Yendrikhovskij, Blommaert, & de Ridder, 1999; see also Chapter 1), the variation within individuals has been overlooked. In fact, in the majority of these experiments the observer had to select ***one*** colour as the colour of a familiar object, provided with only the object name, employing either method of adjustment or selection from a number of alternatives (e.g. Bartleson 1960, Newhall et al 1957, Perez-Carpinell et al 1998). Thus, the observer was restricted in his/her choice to pick only one colour for the object. But if this is the most salient colour associated to such object, we then can ask: "Is this the only one?" In other words, it is possible that instead of ***one and only*** memory colour of the familiar object, we possess a set of alternatives? In the main experiment presented in this chapter, I allow the subject to select all the colours that s/he considers "appropriate" for a typical familiar object (e.g. a standard banana). Here I indicated with "appropriate colour" the colour "oftenest seen" (Hering, 1964) for a specific object, i.e. the memory colour of the object. Using such a novel task we are able to evaluate the observer's memory colour variability for a specific object, if it exists. Moreover, because the subjects are asked to perform the task as quickly as possible we can also explore, for each condition, its difficulty as a function of the speed of the responses.

While the first goal of this study consists in finding out whether a memory colour range exists and to which degree, the second objective is concerned with higher-level implications of memory colour, and if it indeed affects colour appearance and colour constancy. Specifically, one key test is to find robust interactions between familiar object shape and its colour percept. A frequently cited example is that of Duncker's experiment (Duncker, 1939), in which a leaf and a donkey were cut out of green paper and alternately lit with red or neutral illumination (see also Chapter 1). The observers were asked to match the colour of the leaf and donkey using a colour wheel. The results showed that more than half of the observers judged the leaf greener than the

donkey, providing evidence of a memory colour effect. Other later experiments supported these results finding a strong memory colour effect (Delk & Fillenbaum, 1965), while others argued that this effect only exists when the stimulus information was much reduced or the task difficult (Bolles, Hulick, & Hanly, 1959; Bruner, Postman, & Rodrigures, 1951), and the actual memory colour effect may be much smaller than others have suggested (Fisher, Hull, & Holtz, 1956). Alas, in most studies on memory colour and shape, the observers were not allowed to match the colour themselves, or were asked to give a verbal response (Bruner, Postman, & Rodrigures, 1951; Delk & Fillenbaum, 1965) which may have led to biased responses. Although, there are two exceptions to these studies: the works of by Siple and Springer (1983) and Gegenfurtner and colleagues (Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Olkkonen, Hansen, & Gegenfurtner, 2008). Siple and Springer (1983) asked the subjects to adjust the colour of a neutral disk, a familiar object contour with no texture information, or a familiar object contour with texture information, in order to match the memory colour for that particular object. In their experiments, changes in context produced no change in memory colour results, and they therefore conclude that object memory colour is independent of shape and texture information (see also Chapter 1). On the contrary, a similar experiment performed by Gegenfurtner and colleagues (Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Olkkonen, Hansen, & Gegenfurtner, 2008) demonstrated an effect of memory colour on colour appearance, with the size of the effect increasing with the naturalness of the image presented (see also Chapter 1). They asked the observers to adjust the colour of a familiar object photograph with texture information (and "3D" information), a familiar object photograph with scrambled texture information (and "3D" information), a familiar object contour with no texture information (uniformly coloured outlines), or two control disks in order to match the memory colour for that particular object. The results showed the effect of memory colour on the photographs, but no significant effect on the control stimuli. Therefore, while they demonstrate an effect of naturalness of the image presented (i.e. context) on memory colour, their results do not quantify the distinct effects of shape and chromatic variation. In conclusion, we can see that there are contrasting results on the effect of shape on memory colour. In the opinion of the author, the relationship between shape (or context) and memory colour should be examined not only on its mean value, but also on its range. Furthermore, although some studies have evaluated the effect of shape on memory colour under changes in illumination (e.g. Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Olkkonen, Hansen, & Gegenfurtner, 2008), the relationship between memory colour variability, shape and illumination change is still an unexplored field.

Lastly, to the author's knowledge, no study has yet explored the interaction between colour and response time as a function of shape.

To achieve the aims listed above, I carried out a psychophysical experiment (Main Experiment, section 4.2) that probes memory colour of familiar natural objects in three different shape configurations under three illumination conditions. The observer's task was to judge, as quickly as possible, whether a colour was appropriate for a given object. In addition, I performed (1) a control experiment to test the recognisability of the presented shapes (section 4.3) and (2) a control analysis to evaluate the natural variability of the object mean colour and the extent of the polychromaticity (section 4.4).

Results illustrate robust interactions between object shape and memory colour range, as well as response time, and they are consistent under all three illuminants. The memory colour effects on range and speed measures are stronger than on memory colour itself. Our results show a wider implication of memory colour effect than previously known, and demonstrate the complex interaction between colour perception and object shape, as well as object identity.

4.2 EXPERIMENT 1: MEMORY COLOUR AND COLOUR CONSTANCY OF FAMILIAR OBJECTS

This section describes the main experiment of this study and its results, discusses its findings, and introduces the further investigations presented in 4.3 and 4.4.

4.2.1 METHODS

The following experiment was carried out using the set up described in 2.2. The observers were asked to look inside the box through the viewing hole and perform a forced "yes/no" colour match task. They had to judge, as quickly as possible, whether a colour was "appropriate" for a previously named object. The stimuli consisted of uniformly coloured objects in three possible shape configurations (2D disk, 2D filled profile of the object or 3D solid real shape of the object). Since in this experiment were displayed only uniformly coloured surfaces then the image processing described in Chapter 3 was reduced to block 3 and, partially, block 4 only (sections 3.5 and 3.7 respectively). Therefore, the scene geometry calibration methods, the segmentation software, the projector characterization, and presentation software required are explained in Chapter 3. The experiment was run under three different illuminants.

Three familiar objects were employed: a banana, a carrot and a courgette. For each object, a representative (or exemplar) of the category was selected and the spectral power distribution of 10 different points of the same object was taken using a Photo Research PR-650 spectrascan. The

measurements were taken in a Verivide Colour Assessment cabinet under an illuminant metameric to the CIE standard light at 6,500 K correlated colour temperature (D65). Then, for each point, the surface spectral reflectance was calculated using the spectral power distribution of the illuminant measured also in the cabinet. The average of these 10 values is used as the object's final surface spectral reflectance and used as reference in the following experiment.

4.2.1.1 OBSERVERS

Fifteen observers took part in the experiment, all of them students in Newcastle University, 6 male and 9 female, age range between 20 and 26 years old. They all tested as normal on the Farnsworth-Munsell 100 hue test and were naïve to the purpose of the experiment. Before starting the main experiment, they acquired familiarity with the task with a 2 minutes practice session.

4.2.1.2 COLOUR CONDITIONS

As said above, the colour stimuli were uniformly coloured objects, each presented as illuminated by three different lights (i.e. the background illumination). For all objects, the standard daylight illuminant at 6,500 K correlated colour temperature (D65) was chosen as reference. In addition, standard daylight illuminants at 4,000 K and 25,000 K correlated colour temperature (respectively labelled as D40 and D250) were selected for the banana and the courgette. Illuminant D250 and a greenish light (GE) were used for the carrot. However, in the actual experiment, the illuminants' spectrum power distribution (SPD) differed from the CIE standard lights. As a matter of fact, a data projector generated the illuminants, as well as all the other stimuli, and only an approximation of the real SPD of the light could be obtained.

Illuminants	Y(cdm2)	x	y
D40 (yellowish)	7.086	0.3814	0.3828
D65 (neutral)	7.061	0.3136	0.3273
D250 (bluish)	6.887	0.2511	0.2561
Greenish (GE)	7.007	0.2494	0.3576

Table 4.1 - CIE Yxy chromaticity values for the four illuminants measured on extra-white paper.

As a consequence, metamers of these illuminants were employed in the experiment; their CIE Yxy values are shown in the Table 4.1. The real D65, D40, D250, GE spectra, measured in the set-up by a PR-650 spectrascan on matt ultra-white paper, are shown in Figure 4.1.

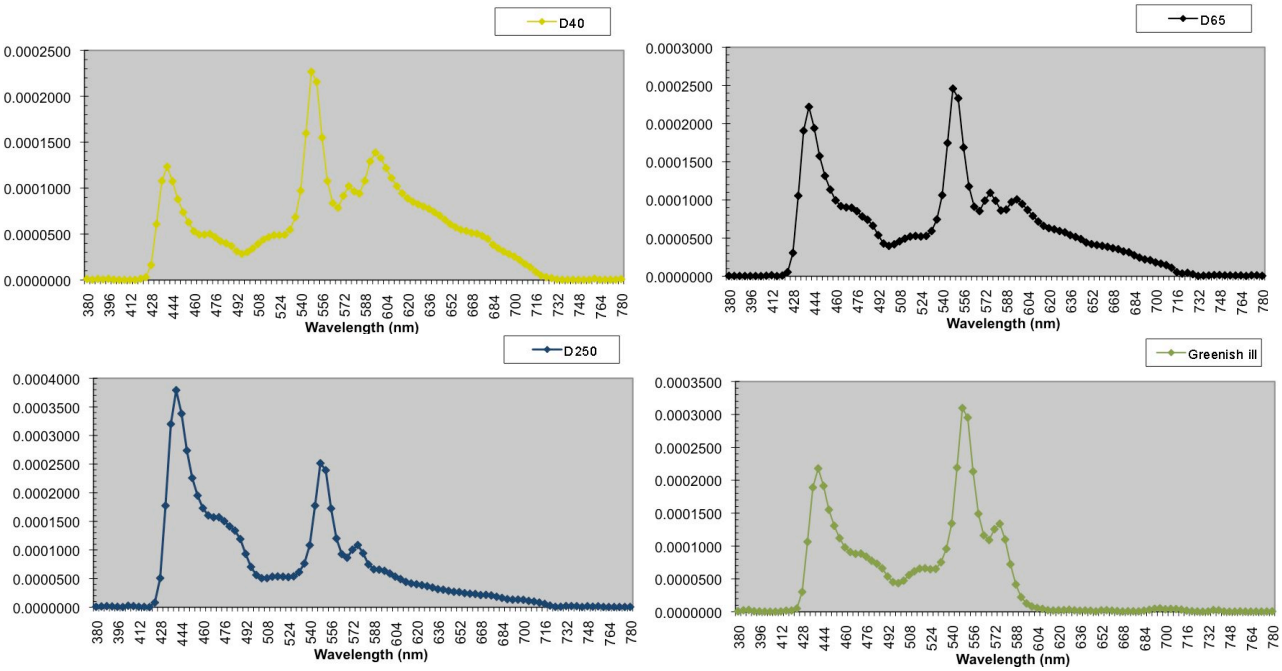

Figure 4.1 – Measured spectral power distribution of the four illuminations. Top left: D40; top right: D65; bottom left: D250; bottom right: greenish illuminant (GE).

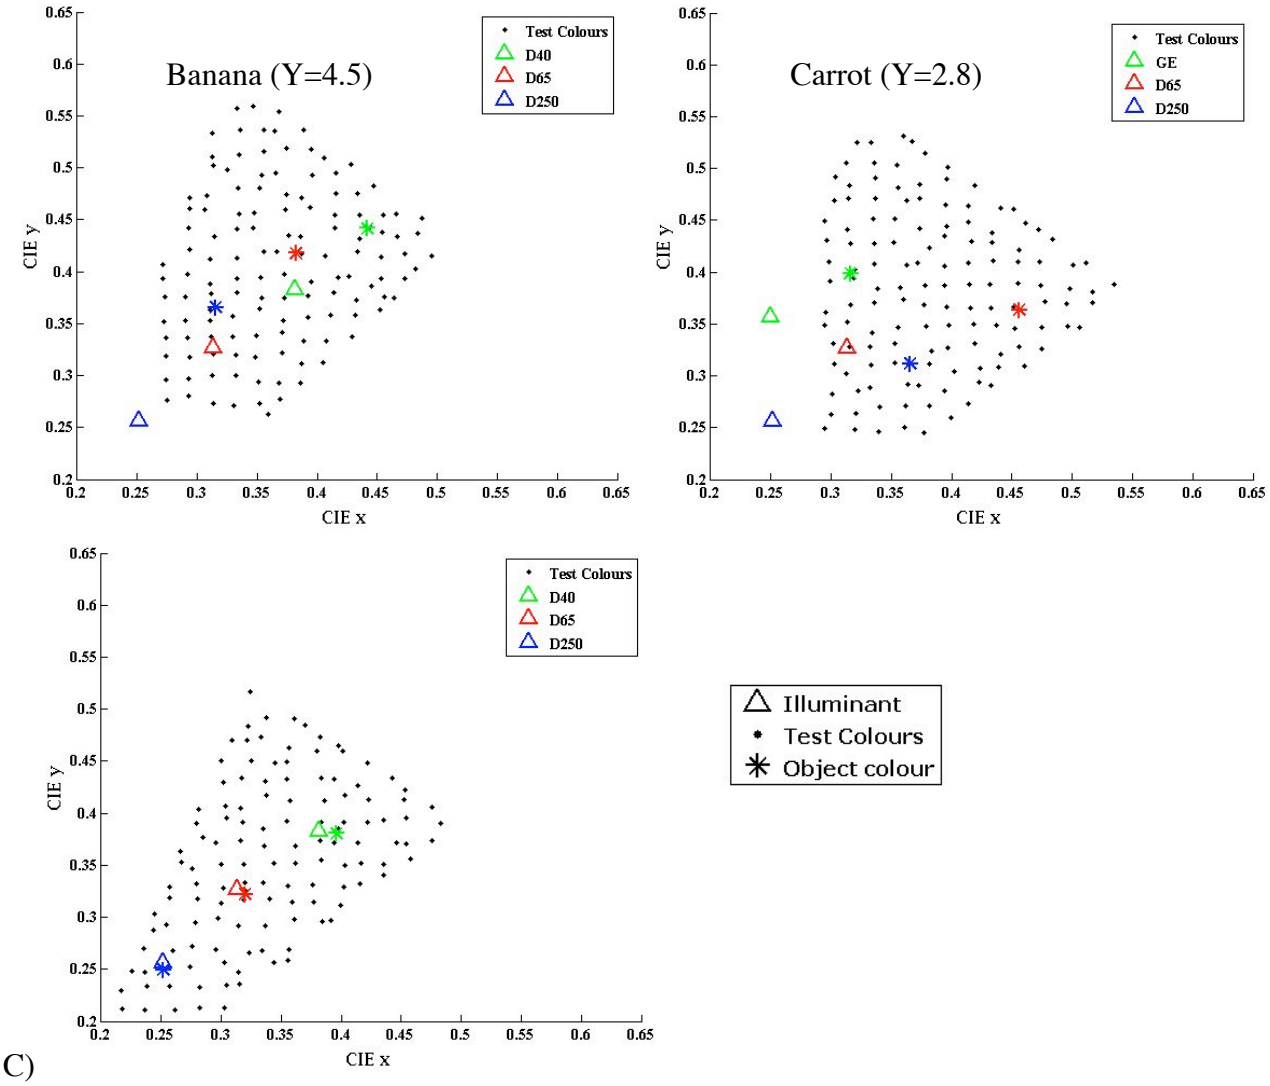

C)

Figure 4.2 – Group of stimuli used for each object (black points), simulated colour of the mean of the real natural object (coloured stars), and illuminant colours (coloured triangles). Blue=D250; Red=D65; Green=D40 or GE for the carrot). A) Banana (N=131); B) Carrot (N=134); C) Courgette (N=127).

Next, the CIE Yxy values of each object were simulated as under these illuminations, using the object's final surface spectral reflectance calculated as above and CIE 1976 colour matching functions. Subsequently a group of isoluminant test colours were selected around the computed Yxy value ("mean object chromaticity coordinates"). These colours were selected to sample as regularly and widely as possible the entire chromatic plane around the objects' xy chromaticities coordinates, under the constrains of the data projector's display gamut for a specific isoluminant plane. Figure 4.2 illustrates the measured CIE xy values for test colours of each object (131 colours for banana, 134 for the carrot and 127 for the courgette) measured under the actual experimental setup, using a PR-650 spectroradiometer. In the same figure are shown for each object its mean colour under the three lights and the neutral points.

4.2.1.3 SHAPE CONDITIONS

For each object, an artificial replica of the object 3D shape was spray painted with white matt acrylic paint. As a result, these objects preserve all the shape cues of the test objects, but can be seen as "blank canvasses" on which we can project different test colours. Figure 4.3 shows the 3D white-painted replicas of the three objects (banana, carrot and courgette).

Figure 4.3 – Photograph of the three objects used in condition c) of Experiment 1. From left to right: banana, carrot, courgette

The object colour is presented in three different geometric configurations: a) two-dimensional disc (generic configuration or *2DG,* covering 3.5x3.5 deg); b) two-dimensional filled "silhouette" of the object (*2DN*, average 4x14 deg); c) solid three-dimensional real shape[4] of the object (*3DN*, average 4x14 degree). The projected image for the 3D condition consists of a filled outline, which exactly aligns with the visible surfaces of individual objects placed in the box, segmented from the background. Figure 4.4 illustrates an example of the three configurations. Figure 4.4c shows an image of the real object taken in the set up. In this chapter, the two-dimensional filled "silhouette" will be alternately called outline configuration or *2DN*.

[4] Dimension: banana= 15.5*6.5*2.8 cm; carrot= 16.4*5.5*3.0 cm; courgette= 22.6*5*5 cm

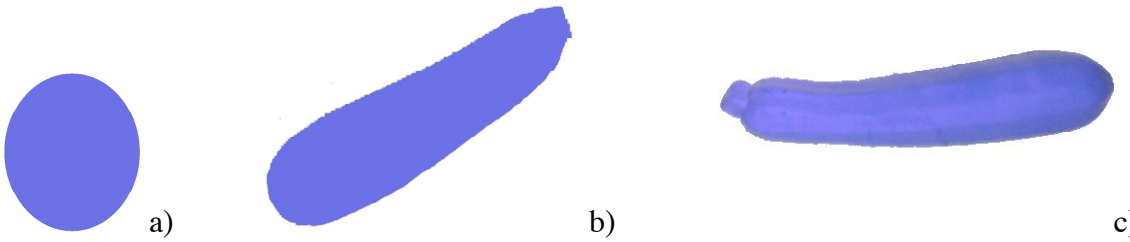

a) b) c)

Figure 4.4 - Configurations presented in the experiment- courgette. a) 2D disc; b) 2D silhouette; c) 3D object (photograph of the object taken in the box).

Note that the filled contour used in condition b) is identical to the image project on top of the white object, and nevertheless the subject's perception in the third condition is that the object's surface is coloured. The subject perceived the object as a 3D solid shape, or a 2D shape, changing colour.

4.2.1.4 PROCEDURE

The subject had to perform a forced "yes/no" colour match task, consisting in judging, as rapidly as possible, whether or not the colour shown was appropriate for a previously named object. The observer had to complete 9 sessions, one for each test object and illuminant combination. Each session was split in 3 blocks, one for each of the 3 shape conditions. Therefore, each of the 27 blocks contained the number of trials correspondent to the number of test colours for the particular named object in one experimental condition. Each block required 10-15 minutes to complete.

The subject was asked to place his/her chin on the chinrest and look inside the experimental box. The general scene displayed a test object centrally placed on a white board. At the beginning of each session, the subject read the name of the object (black word on white background shown in front of the observer), followed by a 60 seconds adaptation phase to the selected illumination. Then the first stimulus was presented for as long as the subject needed to make his/her judgement. This time was recorded as the observer's response time to the stimuli. The observer expressed his/her choice through a joystick: "right" button for "no" - the colour is not appropriate, and "left" button for "yes" - the colour is appropriate. After a decision was made, the response was recorded and the test colour disappeared. Next a mask, composed of random isoluminant colours, was presented for 1 second to exclude afterimages; the mask's mean chromaticity is equal to the illumination's. Then a new stimulus was displayed under the same illuminant and on until all stimuli are presented. The general protocol is shown in Figure 4.5.

For each object, all chosen colour stimuli were presented only once for each illuminant-shape combination (9 times in total). For each observer, under each experimental condition, the presentation sequence of the test colour was randomized, determined at the start of the experiment. The object sequence was permuted so that only one ninth of the observers had the same order.

Conversely, the shape configuration sequence was identical for all the subjects. Configuration 2DG was always presented first followed by condition 2DN and then 3DN. This order ensured a gradual introduction of more cues. Initially only the linguistic information about the familiar object was given, with no shape reference; in the second its silhouette shape was introduced; in the third the full 3D shape representation was displayed.

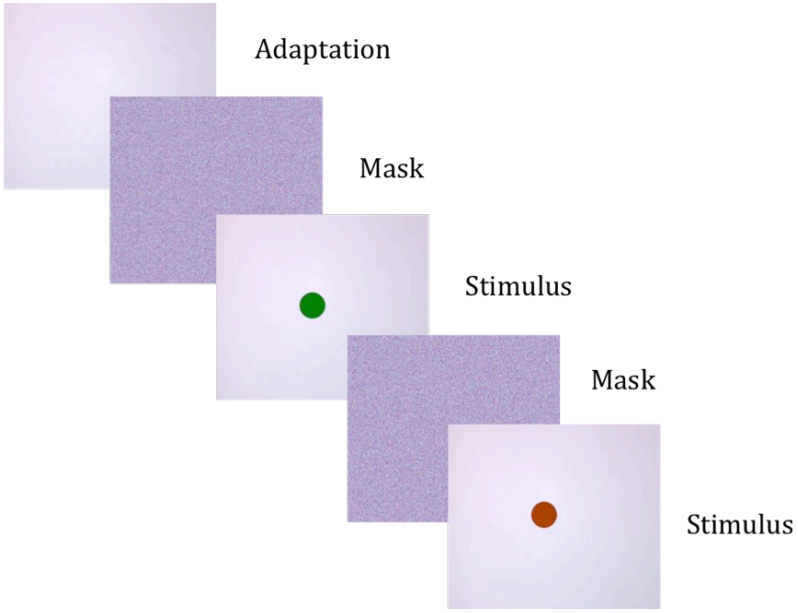

Figure 4.5 – General procedure: example of the experiment images sequence. Note that the showed colours are only illustrative.

4.2.2 RESULTS

As said above, Experiment 1 was designed with the intent of exploring the effect of shape on memory colour and on constancy of memory colour under changing of illumination. Thus, this section has three aims: (1) evaluate the characteristics of the set of colours judged appropriate for the named object, denoted here as "selected colours" or "yes" responses, which are measures of the memory colour of the object; (2) assess the time the subject required to answer the question, denoted as response time, which indicates the facility in accessing memory; (3) investigate the effect of shape on colour constancy.

4.2.2.1 SELECTED COLOURS

4.2.2.1.1 Percentage of "yes" responses

The mean of the number of "yes" responses for each condition and object and the standard error of the mean were, respectively, calculated as follow:

$$\mu\% = \frac{\sum_{i=1}^{N} y_i}{N * K} * 100 \qquad \text{Equation 4-1}$$

$$S_{c,o} = \frac{\sigma_{y_i}}{\sqrt{N}}$$ Equation 4-2

where N is the number of subjects, K is the number of possible choices, y_i is the number of "yes" responses for subject *i*, and σ_{yi} is the standard deviation.

Before computing the mean and standard deviation of the distribution, we discarded the outliers. For each subject, an outlier was defined as any value falling more than 1.5 times the interquantile range below the first or above the third quantile of the distribution. It is worth remembering that the interquantile range and a quantile are calculated with respect to the median of the distribution. Figure 4.6 illustrates an example of the subject choices and his/her discarded outlier.

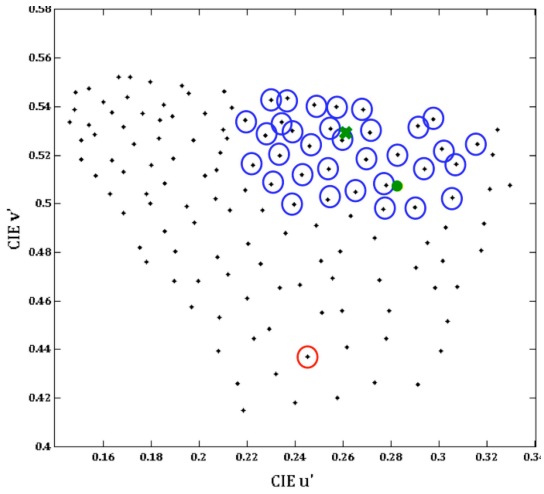

Figure 4.6 – An observer's responses for 2D carrot silhouette, plotted on a CIE u'v' chromaticity diagram. Black dots illustrate all test colours; blue open circles represent all colours the observer responds as appropriate for a banana; the red circle illustrates the outlier; and the green cross represents the computed memory colour of the observer while the green close circle is the measured colour.

Figure 4.7 illustrates the mean percentage of "Yes" as a function of condition, illumination and object. These graphs demonstrate a consistent trend among illuminations; the difference in the number of selected colours between conditions varies similarly for the same object under different light. This finding demonstrates a solid effect of shape on the definition of "appropriateness" of an object colour. A three-way repeated measures ANOVA was performed on the mean percentage of "Yes" responses (factors: object, illuminant, and shape). Mauchly's test indicates that the assumption of sphericity wasn't violated. All effect are reported as significant at $p<0.01$. The ANOVA reported a significant effect of illuminant ($F(1.205,16.87)=55.93$) and shape for all objects, $F(1.908, 26.714)=10.019$. Contrasts reveal a significant overall difference for 2D disk versus 2DN and 3DN configurations ($F_{2DN}(1,14)=18.31$ and $F_{2DG}(1,14)=15.3$), but not between 2DN and 3DN, $F(1,14)=0.001$, $p=0.982$. Yet the direction of the variation and the strength of the effect are both object-dependent and illumination-dependent. To break down of this diversity a two-way ANOVA was performed on each object. Mauchly's test shows failure of the sphericity assumption

of the main effect of dimension for single object analysis. Therefore the ANOVA was corrected using Greenhouse-Geisser estimates.

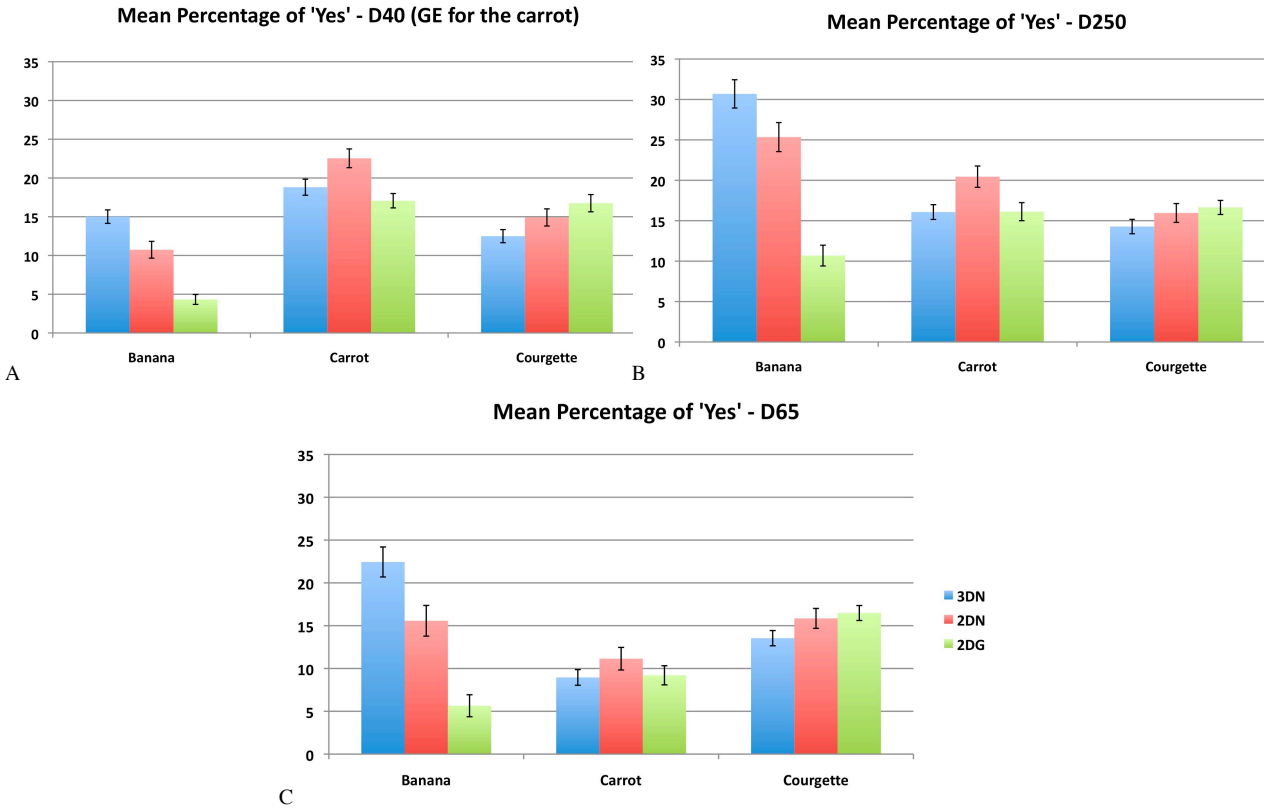

Figure 4.7 – Mean percentage of "Yes" responses over all subjects. Illuminants: A) D40 (note that is illuminant GE for the carrot); B) D250; C) D65. Blue bar: 3DN. Red bar: 2DN. Green bar: 2DG

For the banana, the mean percentage of selected colours for the 3DN configuration is significantly bigger than for the silhouette (contrast $F(1,14) = 6.9$) and the latter bigger than the 2DG configuration ($F(1,14) = 54.3$), as shown in Figure 4.7. This indicates that the number of "appropriate" colours for the banana significantly increases with the introduction of more cues, independently of the light illuminating the scene. In addition, the ANOVA shows a significant effect of illuminatin on percentage of correct responses ($F(1.1,15.2) = 58.23$).

For the carrot, Figure 4.7 illustrates a significantly larger amount of selected colours for the silhouette compared to the disk and the 3DN configuration under all light conditions, $F(1.88,26.37)=4.46$. Thus, the presentation of the object's silhouette increases the number of choices, while the introduction of three-dimensional cues narrows this number down. In other words, we have an increase of percentage of selections from the 2DG to the 2DN condition and a decrease from the 2DN to the 3DN configuration. For the latter the percentage is similar to 2DG condition. This behavior is consistent under all three lights (D65, D250 and GE), although the ANOVA shows an effect of illumination, $F(2,28) = 45.8$.

On the contrary, for the courgette, the mean number of selected colours for the 3DN configuration is smaller than for the silhouette, and this number, in turn, is smaller than for the disc configuration. Even though the same variation is replicated in all light condition, the effect is not

significant ($F(2,28)=3.166$, $p=0.69$). Furthermore there is no significant effect of illuminantion, $p = 0.058$.

Note that the percentage of selected colour varies with illumination ($p<0.001$), with D250 having the highest percentage of "yes" responses, but the overall behaviour as a function of object and shape is similar. Thus, we can study these shape effects under only one illuminant.

4.2.2.1.2 Selected points distribution

A further way to examine the degree of acceptance of the selected colours is to calculate the overall frequency on which the entire tested population (15 observers) selects a specific point. Using D65 as example, and remembering that the same pattern can be drawn for all three illuminants, Figure 4.8 illustrates the distribution of 'Yes' responses, summed over all observers, in the CIE chromaticity plane (u', v'). The number of 'Yes' responses is plotted using a pseudo-colour density map, such that its magnitude determines the colour of the plotted point (the redder the colour, the larger the number of 'Yes' responses). Thus I can compute the mean of the chromaticity coordinates of the points selected by one subject averaged across subjects. For the banana, Figure 4.8 illustrates such mean (black cross), its actual measured chromaticity coordinates (red cross) and the neutral point (D65, red dot). Each subfigure in Figure 4.8 depicts one shape condition.

Although there are individual differences, most observers' colour choices cluster in a small range of the colour space (redder area in each subfigure in Figure 4.8). For each condition and object, the peak of the subject choice distribution is defined as the loci with the maximum number of "yes" responses. The mean of this area and the mean of the entire distribution tend to be more saturated compared with the actual object colour – i.e. the observers remember the banana as more yellow; the carrot as more orange and the courgette as more green (alas, due to gamut limitation, not all natural colours are available for the observer to choose from and therefore we cannot determine the full extent of this shift). These results agree with previous memory colour studies (Bartleson, 1960, Siple and Springer 1983).

The distribution peak and mean do not significantly vary between shape conditions, indicating that the mean memory colour of the object is not affected by the different shape configurations (u': $F(2, 119)=1.74$, $p=0.1802$; v': $F(2,119)=2.02$, $p=0.1368$).

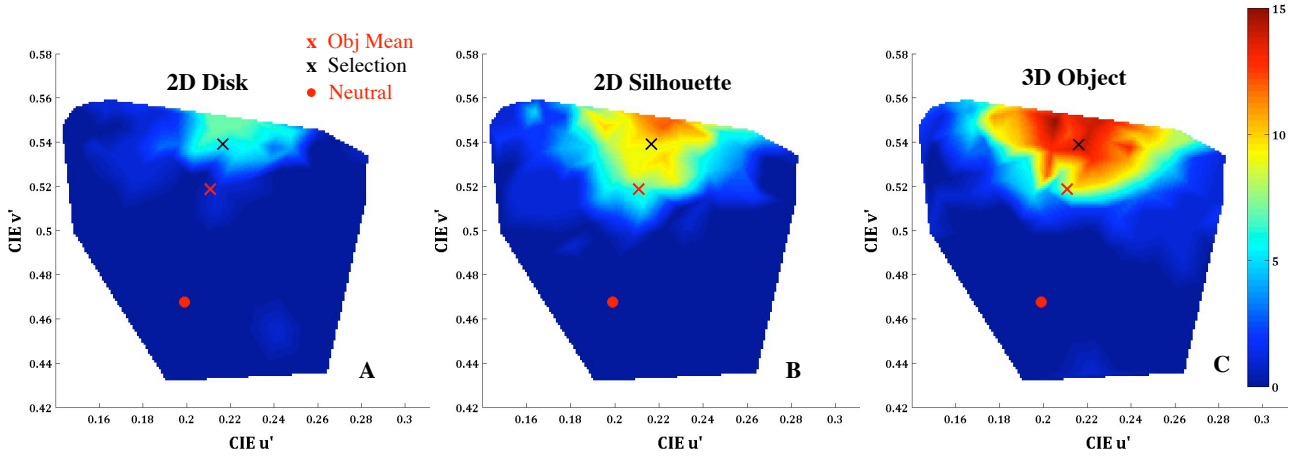

Figure 4.8 - Total number of "Yes" responses for the banana in all three conditions, as a function of CIE x, y chromaticities, for 15 observers under illuminant D65. Red cross: mean colour of the real natural banana. Black cross: mean chromaticities of the observers' selection. Red dot: neutral point. Colour bar on the right defines the number of observers selecting a point; redder is an area more subjects have selected a point in it. A) 2D disc configuration. B) 2D silhouette. C) 3D object.

4.2.2.1.2.1 RANGE

It is worth pointing out that, even if two conditions have the same number of choices and peak of the distribution, they might have a very different range of responses. In other words, all the selection can be located in the peak area, leading to a very small range, or on the contrary can have a central nucleus but also a sparse distribution. In both cases, the measured memory colour can be the same, but the range of colour the observer actually judged as appropriate may differ.

Thus an alternative way to examine the data can be looking at the range of choices made by the total number of subjects. For simplicity, let us explore first the variation under D65 (same results can be drawn for the other lights conditions - see Appendix 1). Figure 4.9 plots the percentage of observers agreeing on the same appropriate colour as a function of its xy coordinates; the black dots in Figure 4.9 represent the test colours. The blue line in each subfigure delimits the area of the chromatic plane containing all the points judged by more than 10% of the observers as appropriate (this area includes 95% of the total number of choices over all subjects). On the other hand, the green line encloses the subgroup of points that was selected by more than 50% of the observers. Lastly the red line delimits the subset of points that was judged as appropriate by more than 90% of the observers. Each panel in Figure 4.9 represents one object-shape configuration combination.

Figure 4.9 visually shows that the effect of shape on range and peak of the response distribution is strongly object-dependent. As a matter of fact, the observer's response for the banana is strongly affected by shape cues, while no effect of shape is shown for the carrot or the courgette. Specifically, for the banana (Figure 4.9, top row), the observers' selection range is small in 2D disc configuration and becomes increasingly wider with the introduction of additional shape cues (i.e. the silhouette range is larger than the 2D disc range and the 3D range is wider than the silhouette's). Furthermore, the number of observers agreeing on the same "appropriate" colour for the banana

increases with the introduction of more cues to the object shape. In fact, for example for the 2DG case, less than 50% of the subjects selected the same colour as appropriate for the banana, while more than 90% judged the points inside the red contour as appropriate in 3DN.

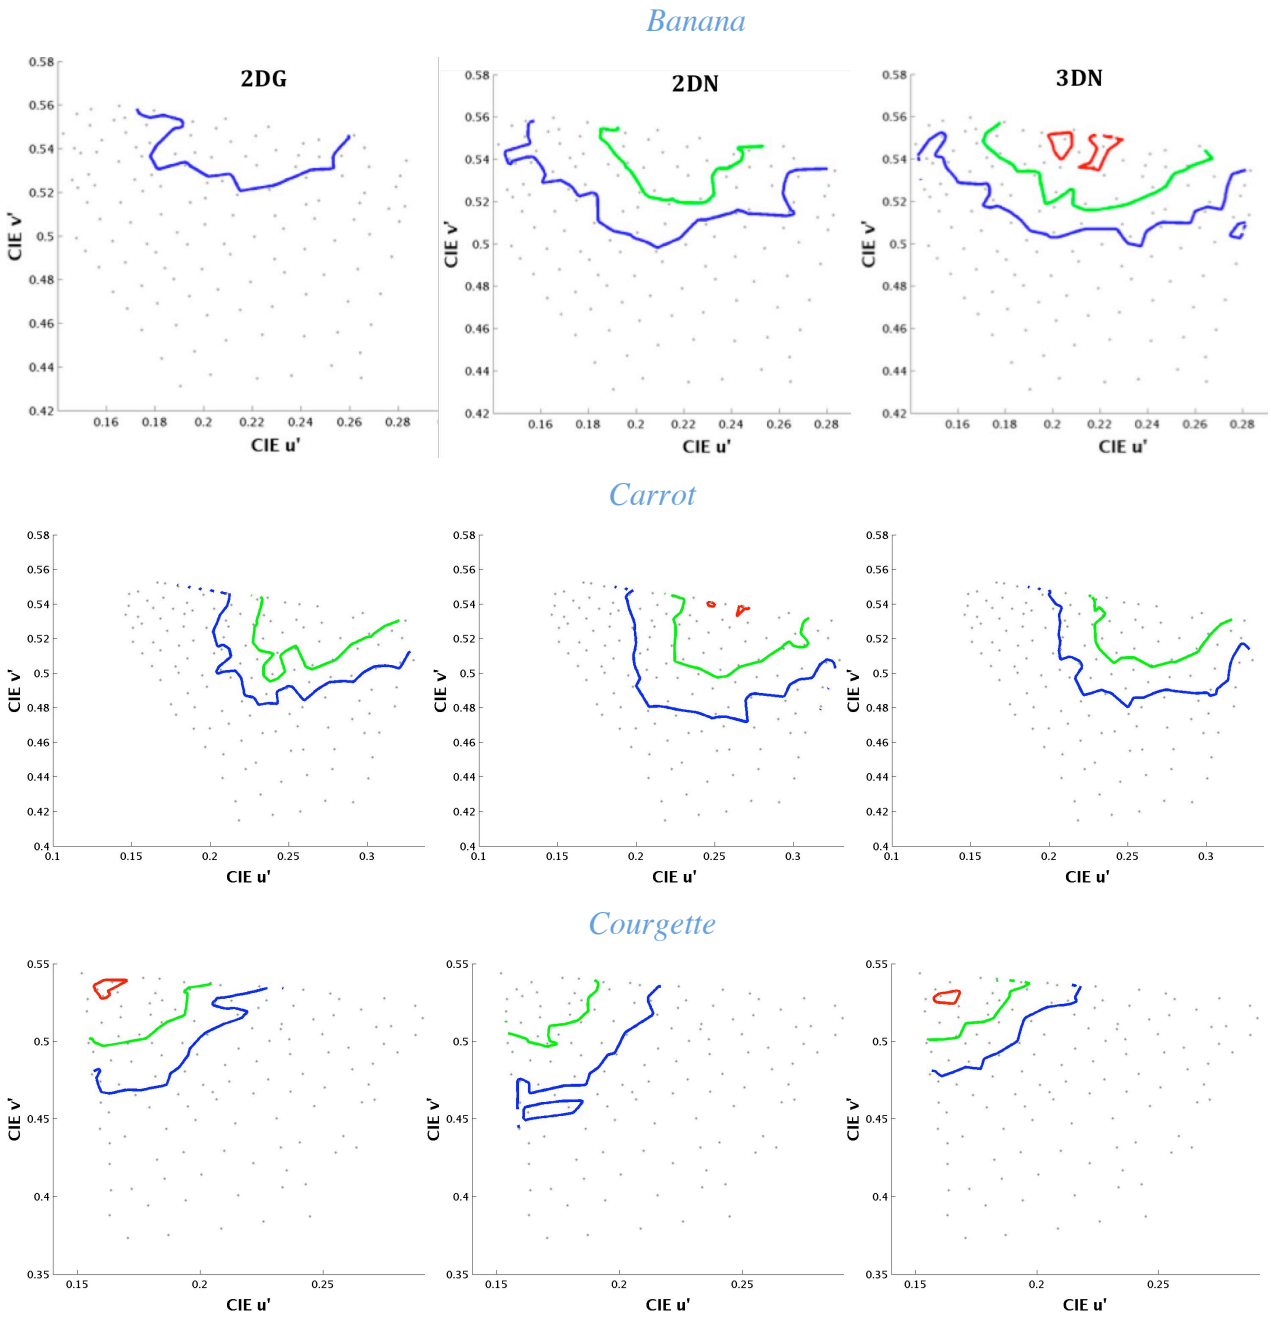

Figure 4.9 – Total number of 'Yes' responses for all 3 objects, as a function of CIE u'v' chromaticities, for all 15 observers, for D65. Where present, the red line: contour of the area wherein more than the 85% of the observers judged the colours to be appropriate for the named object. Green line: contour of the area wherein > 50% of the observers judged the colours to be appropriate. Blue lines contour of the area wherein > 10% of the observers judged the colours to be appropriate. Top row: Banana. Middle Row: Carrot. Bottom row: Courgette. Columns: Left - 2D disk; Centre – 2D silhouette; Right – 3D solid object.

In conclusion, in the case of the banana, the addition of shape cues increases the acceptance rate of a colour by the tested population (i.e. the peak of the distribution stay the same, but becomes more pronounced with the introduction of more shape cues), and enlarges the spread of choices from a central core (i.e. the distribution has the same mean but wider range). Conversely, the carrot

and the courgette do not behave similarly; the range and peak are only slightly affected by shape configuration.

4.2.2.2 *RESPONSE TIME*

As described above the subject had not a fix time to respond, however he/she was requested to answer as quickly as possible. The time of response, defined as the time lapse between the presentation of an experimental colour and the observer's manual response, was recorded on a trial-to-trial base during the experiment. Therefore, for each experimental condition and object, the recorded response time was averaged across all the 'Yes' trials and then again across all 15 observers, after discarding the outliers (defined as 4.2.2.1.1).

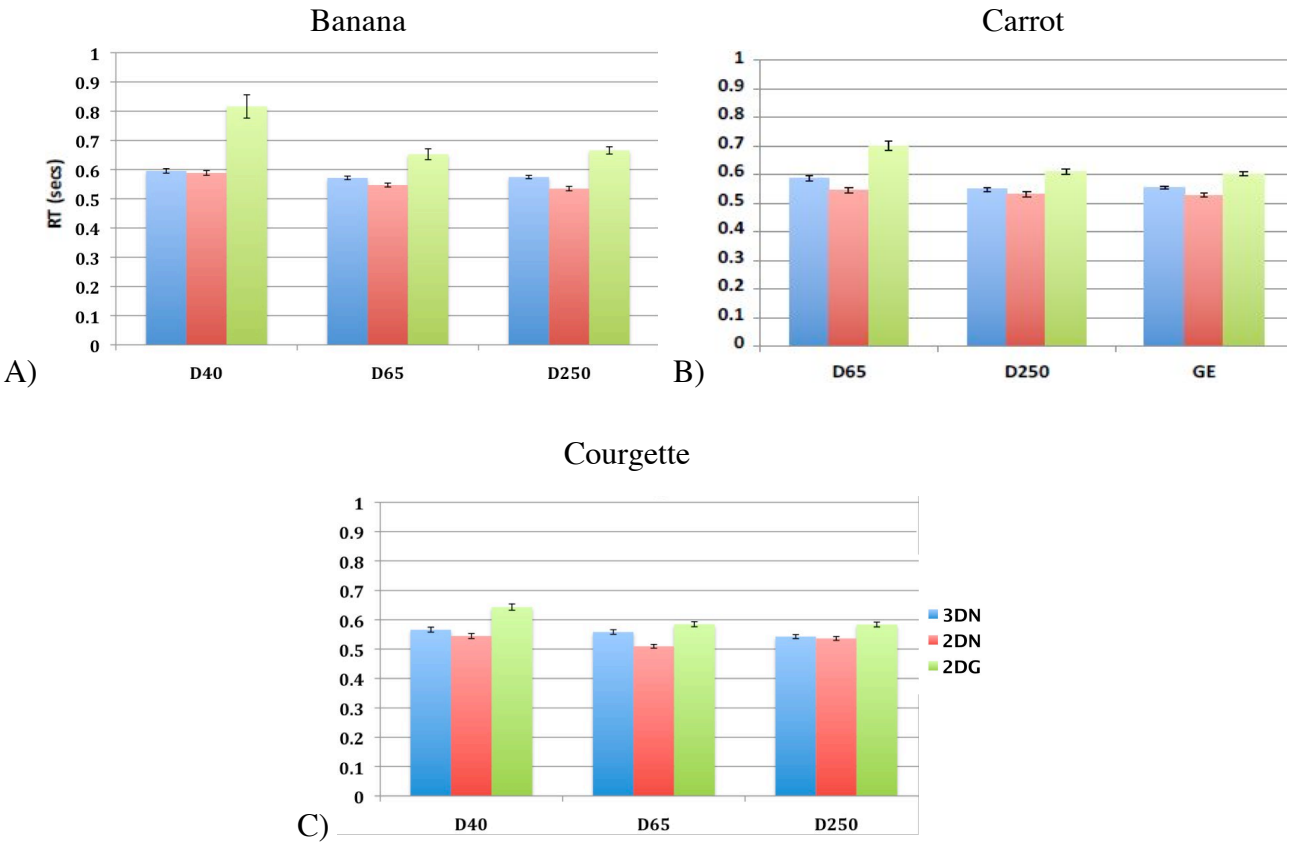

Figure 4.10 – Mean response time for the "Yes" responses (in seconds) after discarding outliers – A) Banana; B) Carrot; C) Courgette. Blue bar: 3DN. Red bar: 2DN. Green bar: 2DG. Error bars represent the standard error of the mean.

Figure 4.10 shows the response time (RT) for all objects in any condition. This illustrates a straightforward pattern, identical and consistent for all objects; the observers are on average significantly slower to react to the task in the 2D disc configuration compared to the 2D silhouette or the 3D object configuration ($F(1.941, 27.17)=7.41$, $p<0.005$). Response time is to some extent smaller for the silhouette configuration compared to its 3D shape, this difference being significant as revealed by the contrast analysis ($p<0.05$). It is worth remembering that the longer the response time, the harder the task. Thus the results prove that the observers found difficult to confirm the object colour appropriateness in absence of its natural shape cue. Conversely, the presence of the

object's shape facilitates the task, but the three-dimensional cues required a slightly longer time of evaluation than the 2D. Furthermore, the RT results show a significant main interaction between the object and the shape, $F(3.3, 46.42)=6.0$, $p<0.001$, and illumination and shape ($F(3.12, 43.72)=13.73$, $p<0.001$).

In addition, I computed the response times for all the 'No' responses, using the same method as for the 'Yes' responses; the results are illustrated in Figure 4.11. The overall pattern of the response time for the 'No' responses shows that the subjects are quicker in discarding a stimuli when presented in a familiar shape than in a disc ($F(2, 126)=3.22$, $p<0.05$). Even if, for the banana, there is a longer response time responding to 3DN banana stimuli, compared with 2D disk or 2DN banana, this difference is not significant ($p=0.2477$). Note that in general RT for "No" responses is lower than for "Yes" responses, showing more confidence in rejecting a colour than accepting it.

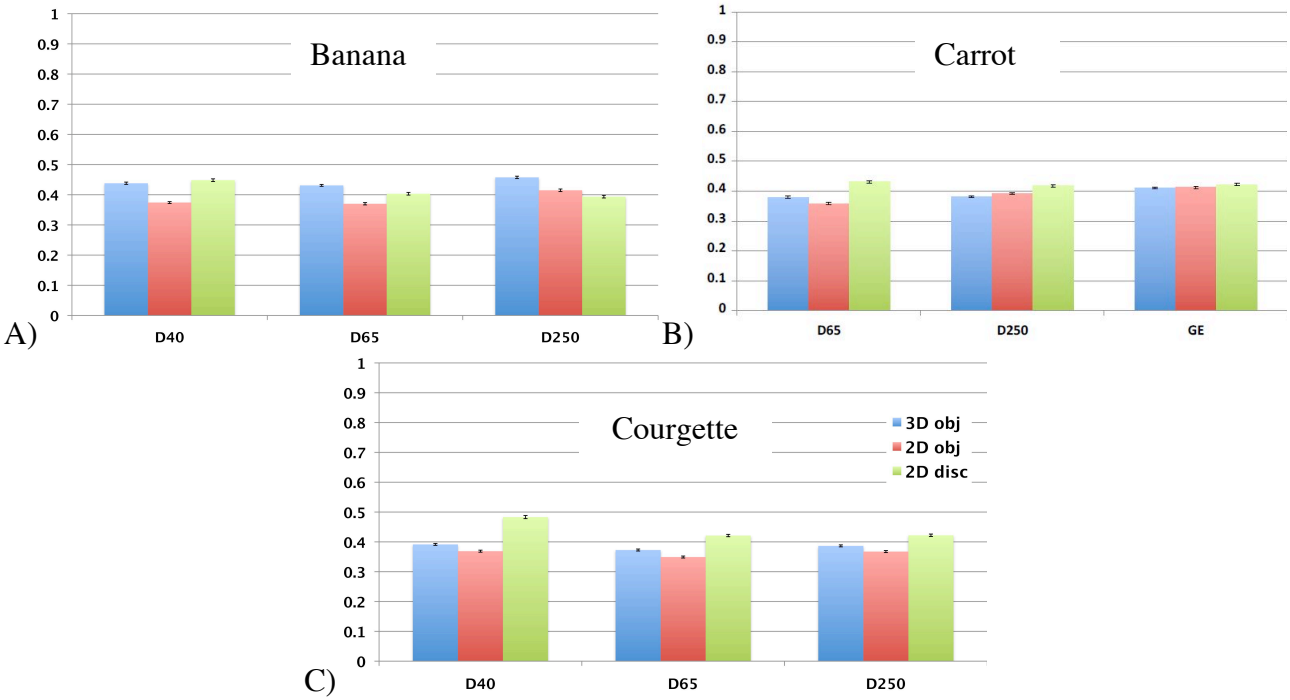

Figure 4.11 - Mean response time for the "No" responses (in seconds) after discarding outliers – A) Banana; B) Carrot; C) Courgette. Blue bar: 3DN. Red bar: 2DN. Green bar: 2DG. Error bars represent the standard error of the mean.

4.2.2.3 COLOUR CONSTANCY

As seen before, classically the colour constancy index for this experiment would be calculated as follows:

$$CCI = 1 - \frac{S'_{ill} - S''_{ill}}{S'_{D65} - S''_{ill}} \qquad \text{Equation 4-3}$$

where S'_{D65} is the measured mean of the real object's colour under D65 (in a certain colour space), S'_{ill} is the measured mean of real object's colour under illuminant *ill* and S''_{ill} is the perceived object mean colour under illuminat *ill*. Although, Ling&Hurlbert (2008) study

demonstrate the need of including a memory shift into the colour constancy index (CCI) in successive colour constancy task. Thus, they modified the above formula as follow:

$$CCI = 1 - \frac{(\vec{S}_{ill} - \vec{S}_{D65}) \bullet (-\vec{S}_M)}{\left\| \vec{S}_M \right\|} \qquad \text{Equation 4-4}$$

where the vector $\vec{S}_{D65}$ is the subjects' matching error with respect to the measured value obtained both under D65 (the "pure memory shift" using Ling&Hurlbert terminology), the vector $\vec{S}_{ill}$ is the subjects' matching error with respect to the measured value obtained both under the chosen illuminant *ill* (the "constancy shift"), and the vector $\vec{S}_M$ is the physical chromaticity shift between the original measured mean object colour distribution under D65 and the one under the chosen illumination (the "physical shift").

In this experiment, the subject has to recall a long-term, well-established memory of the familiar object and match it under a reference light and a test light. Thus, also here, it is necessary to introduce a memory factor in calculating the CC index. However, an altered version of the Ling&Hurbert (2008) index was applied for this experiment.

Let $P^r_{D65} = (u'_{D65}, v'_{D65})$, $P^r_{ill} = (u'_{ill}, v'_{ill})$ and $\overrightarrow{\Delta S_r}$ be the real measured chromaticity coordinates in CIE L'u'v' space of an object under D65 and an illuminant *ill*, and the vector connecting them, respectively. Thus, the vector $\overline{\Delta S_{ill}}$ is the physical "illuminant" shift. Additionally, let $P^r_{D65}\big|_p = (u'_{D65}, v'_{D65})_p$, $P^r_{ill}\big|_p = (u'_{ill}, v'_{ill})$ and $\overline{\Delta S_p}$ be the means of the subjects' perceived chromaticity coordinates in CIE L'u'v' space of the same object under D65 and *ill*, and the vector connecting them, respectively. Then $\overline{\Delta S_p}$ will indicate the perceptual illuminant shift. Therefore, we can calculate the following polar values:

$$\alpha_{rj} = \arctan\left(\frac{v'_{ill} - v'_{D65}}{u'_{ill} - u'_{D65}}\right)_{r\,j} \qquad \text{Equation 4-5}$$

$$\alpha_{pj} = \arctan\left(\frac{v'_{ill} - v'_{D65}}{u'_{ill} - u'_{D65}}\right)_{p\,j} \qquad \text{Equation 4-6}$$

$$R_{r\,j} = \sqrt{\left(u'_{ill} - u'_{D65}\right)^2_{r\,j} + \left(v'_{ill} - v'_{D65}\right)^2_{r\,j}} \qquad \text{Equation 4-7}$$

$$R_{\mathrm{p}\,j} = \sqrt{\left(u'_{ill} - u'_{D65}\right)^2_{\mathrm{p}\,j} + \left(v'_{ill} - v'_{D65}\right)^2_{\mathrm{p}\,j}} \qquad \text{Equation 4-8}$$

Note that α_x is the angle formed by the *x* vector and the u' axis and R_x is the length of the vector *x* for the object *j*. Figure 4.12 illustrated hypothetical configurations of the variables' relations in CIE u' v' isoilluminant plane.

If $\alpha_r=\alpha_p$ and $R_r=R_p$ then the vectors $\overrightarrow{\Delta S_r}$ and $\overline{\Delta S_p}$ in Figure 4.12 are parallel and have the same length. Consequently, the subject will have perfect colour constancy, but memory shift equal to $\overline{\Delta S_m}$ (red arrow in Figure 4.12) On the contrary, if any pair is dissimilar then the colour constancy will worsen in proportion of the combination of angle and radius differences. In Figure 4.12 the dot-dashed black vector has the same angle as the physical vector change ($\overrightarrow{\Delta S_r}$, black arrow), but different modulus, while the green arrow has the same modulus but different angle.

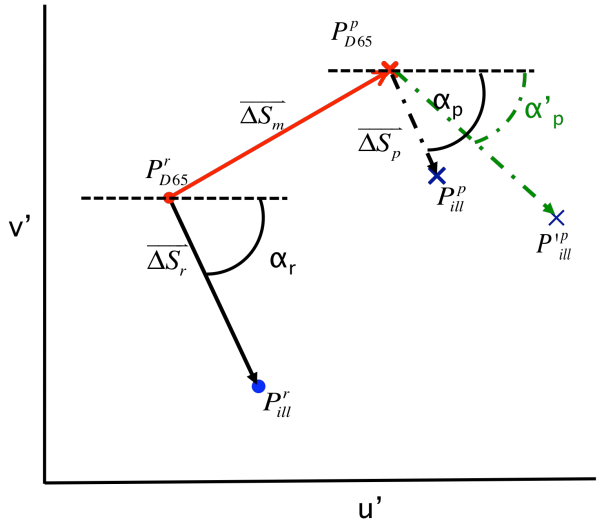

Figure 4.12 – Hypothetical configuration of colour constancy settings. Variables defined in the text. In specific, α_x are the angles that vector x makes with the u' axis.

Therefore, a new CCI index can be defined as follows:

$$CCI = CCI_\alpha * CCI_R \qquad \text{Equation 4-9}$$

where:

$$CCI_\alpha = 1 - \left\| \frac{\alpha_r - \alpha_p}{180} \right\| \qquad \text{Equation 4-10}$$

$$CCI_R = 1 - \left\| \frac{R_r - R_p}{180} \right\| \qquad \text{Equation 4-11}$$

Thus, CCI is equal to 1 if both CCI_α and CCI_R are 1, to 0 if any of the two is 0 and a value between 0 and 1 in any other case. Note that ΔS_m could be equal to any value, even zero, but the colour constancy index would not be affected. The results for all the cases are plotted in Figure 4.13.

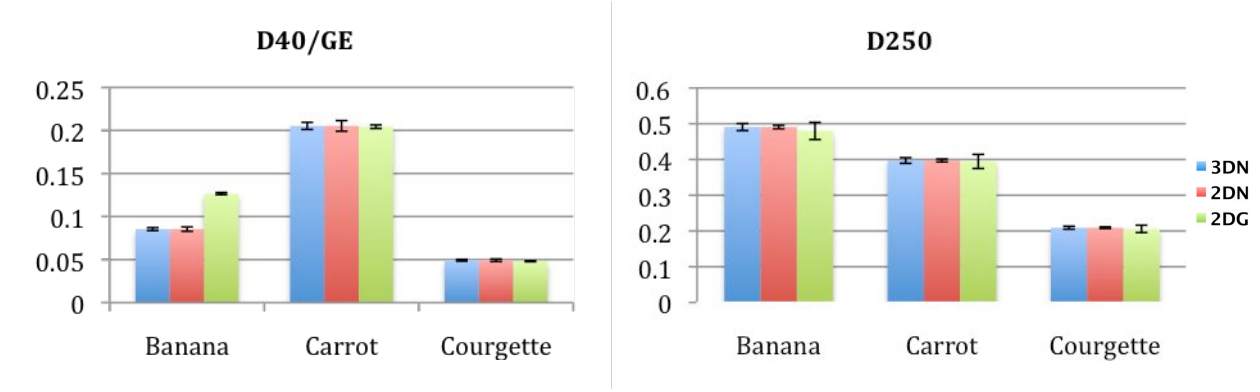

Figure 4.13 – Colour constancy indices for all objects and cases. Blue bar: 2D disk. Red bar: 2D silhouette. Green bar: 3D object. Left panel=D40 (GE for the carrot). Right panel=D250.

Colour constancy is poor (between 0.08 and 0.5), although much better under D250, because of the tendency to oversaturate in memory. Clearly, there is no significant difference in CCI between conditions ($p=.67$), while there is between objects ($p<0.05$). The only peculiar case is the 3D banana, for which colour constancy is slightly better under D40 (this difference being significant; paired T-test, $p<0.05$).

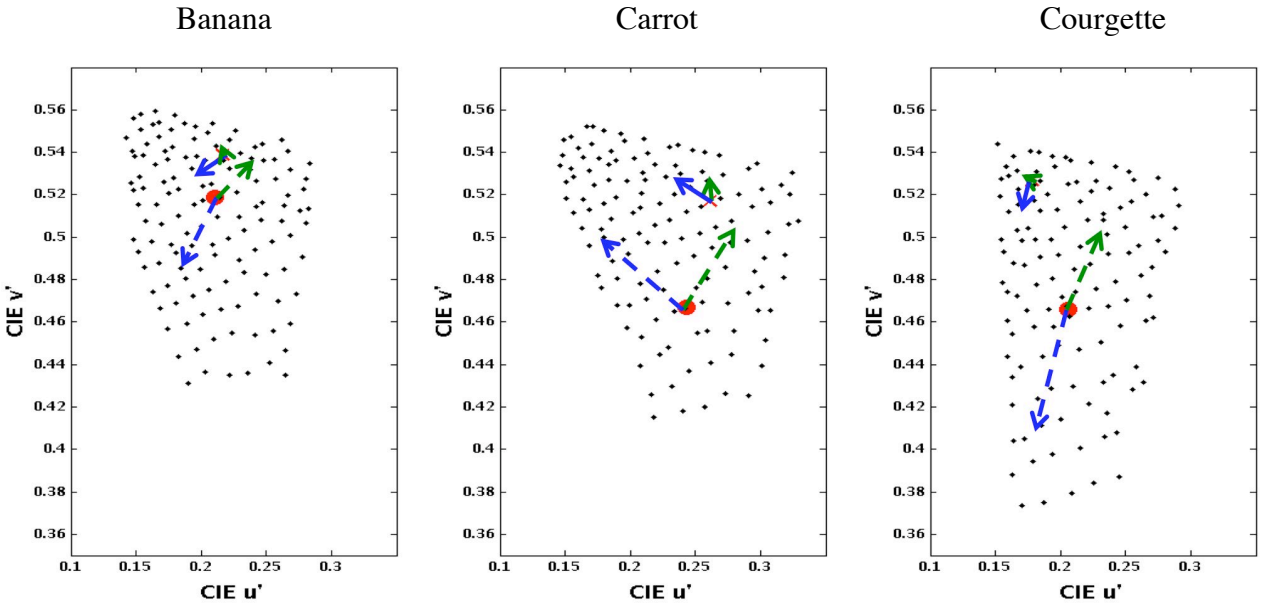

Figure 4.14 - Mean subjects' perceptual chromaticity shifts under change of illumination, from neutral D65 (red symbols), for the 2D disk. Arrows indicate mean colour shifts from the D65 to D250 (blue arrow) or to D40 (green arrow; note that is GE for the carrot). The heads of the solid arrows denote subjects' choices mean under chromatic illuminants D40 (green) and D250 (blue). Dashed arrows' heads denote measured object colour under chromatic illuminants D40 (blue) and D250 (green). Black dots: test colours. Left: banana. Centre: carrot. Right: courgette.

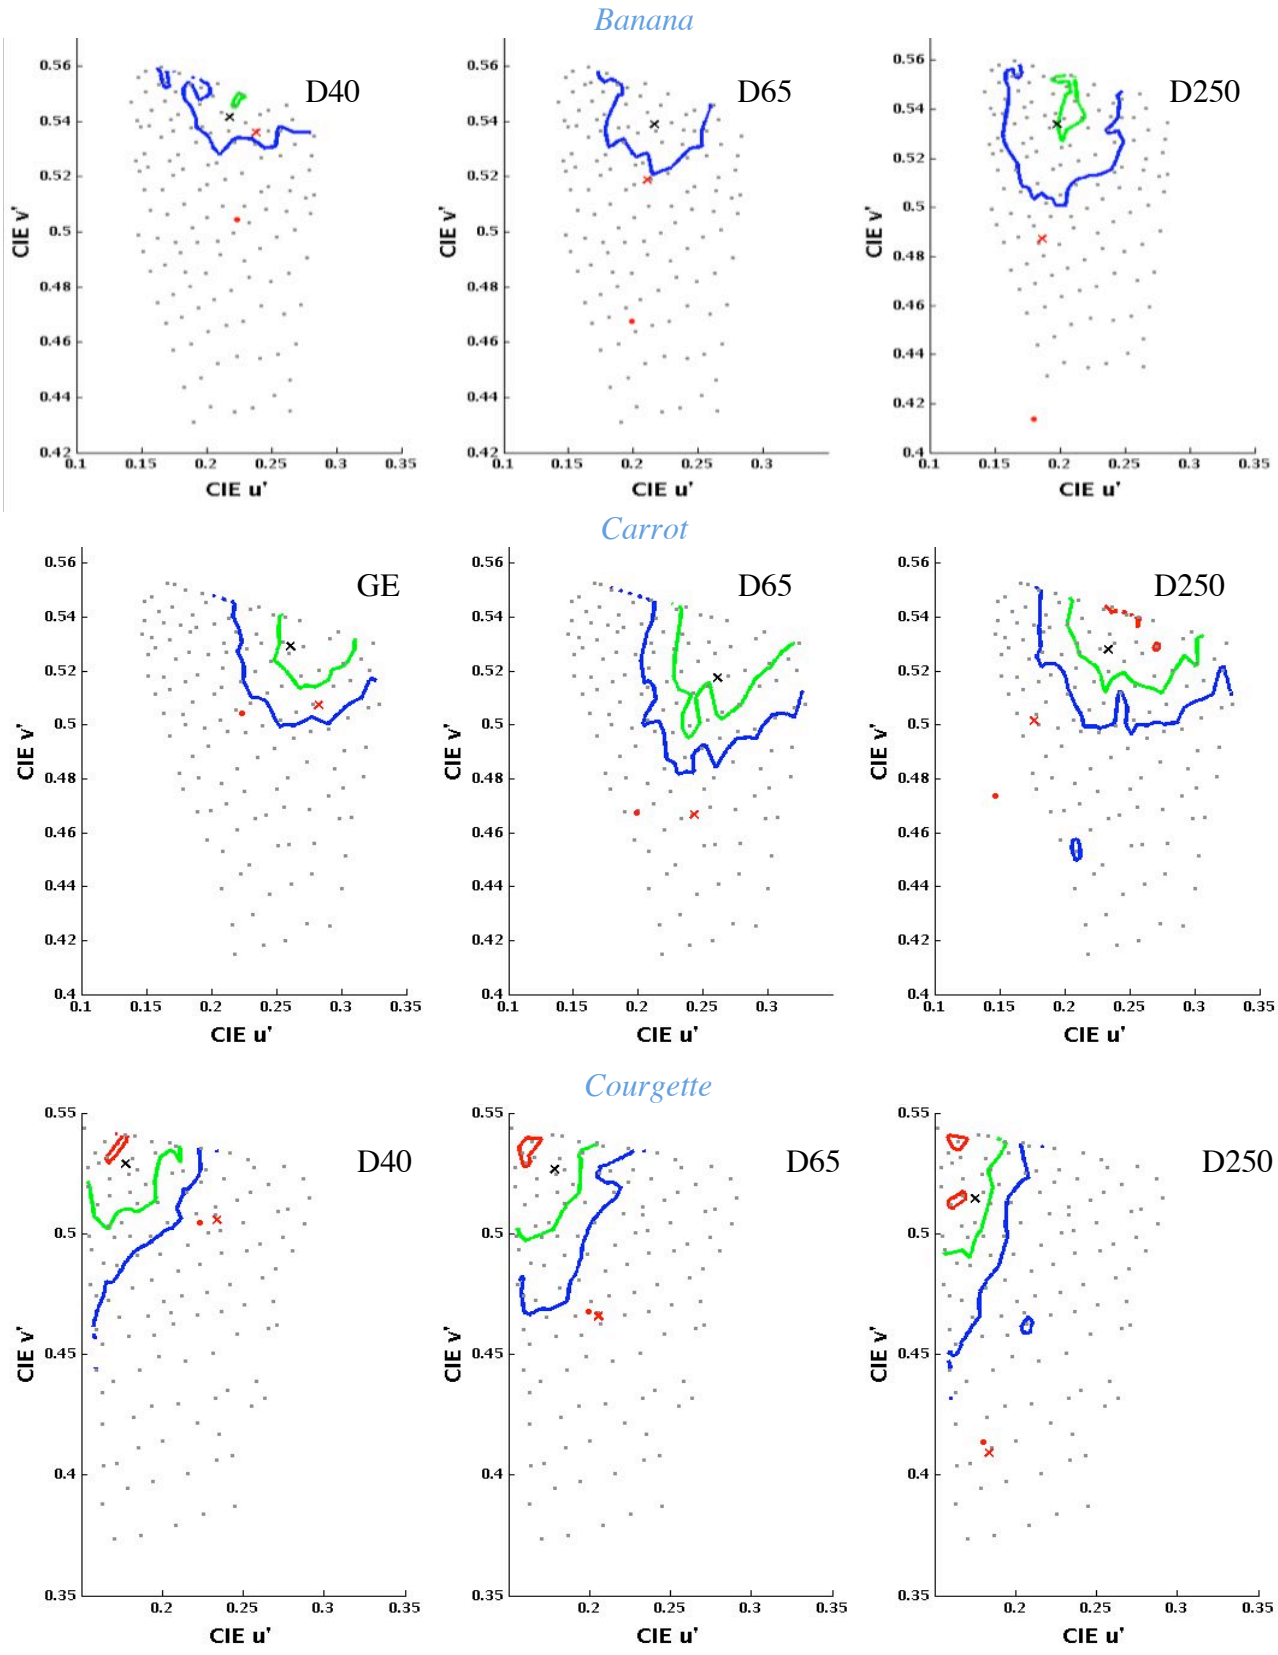

Figure 4.15 - Total number of "Yes" responses in condition 2DN, as a function of CIE x, y chromaticities, for 15 observers under all illuminants. Top row: Banana. Middle row: Carrot. Bottom row: Courgette. Black points: test colours. Red close circle: neutral point. Red cross: mean colour of the natural object. Black cross: mean of the selected colours. Illuminants, from right to left: D40 (GE), D65, D250. Where present, the red line: contour of the area wherein more than the 90% of the observers judged the colours to be appropriate for the named object. Green line: contour of the area wherein > 50% of the observers judged the colours to be appropriate. Blue lines contour of the area wherein > 10% of the observers judged the colours to be appropriate.

Only for illustration, Figure 4.14 plots physical illuminant shift and perceptual shift vectors in u'v' isoilluminant plane, for the 2DG case for the three objects and two light (D40 and D250). Clearly the subjects' choices are restricted in a small area of the chromatic space.

A further way to examine the effect of light change is to observe how the range of the distribution changes under different light conditions. Figure 4.15 shows, for 2DG, that the loci of maximum number of "yes" responses (redder area) moves with the illumination change. Moreover, the range of selected colours varies between illuminants within conditions, implying that differences in illumination affect the range of suitable colours (Figure 4.15). This difference is significant for banana and carrot but not the courgette ($p<0.05$).

4.2.3 DISCUSSION

Results of this experiment can be divided into the effect of shape on (1) colour constancy and (2) memory colour.

In general colour constancy appeared poor for all objects and conditions. Even if there is an effect of illumination, which tends in the direction of constancy, other factors preclude complete colour constancy. Moreover, there is no significant difference between shape conditions, although there is between objects. An incomplete light adaptation of the observers and/or the gamut limitations of the projector may explain these findings and will be discussed further in section 4.5.4. The new apparatus design copes with these problems as proven in Chapter 6 of this thesis.

On the contrary, the effect of shape on memory colour is not as simply described. Its complex behaviour depends on the type of the analysis: (1) the mean number of selected colours, (2) the mean memory colour; (3) the range of memory colours, and (4) the response times.

Firstly, mean memory colour and response time are object-independent. As a matter of fact, *for all objects*, mean memory colour is consistently more saturated than the natural mean colour and not affected by shape, and the response times are significantly smaller for selection in the 3D condition. On the contrary, the mean percentage of number of selections and the memory range have an *object-specific pattern*, although significantly dependent on shape. Thus, for all objects, the mean chosen colour is not affected by shape, and response time difference is sharpened by it, showing a faster decision making in 3D shape, while the object-dependent effect on number and range of memory colour is puzzling. The reason for these different trends may be manifold.

A first explanation might be that these different directions are a symptom of the same behaviour: the subjects remember better the true range of colours of the real object when the colour is displayed onto the object's appropriate shape. In other words, let us suppose that without any cue to the object's shape we can recall only the same fixed amount of the information about the object colour independently of the number of colours its surface may actually possess. In this case a

subject would recall only a smaller percentage of the surface chomaticities of a greatly chromatically variegated object with respect to a less chromatically variegated object. Thus, if the presence of more cues allows better access to the information stored in memory then such a percentage would increase greatly for a very polychromatic object and in a much less degree or remain unchanged for an almost uniform object. Based on this assumption, the introduction of shape cues would largely expand the range of acceptable colours of objects possessing a wide range of colours and to a lesser degree objects with smaller distributions. Then it is logical to ask: what is the difference between the chromaticity distributions of the three objects? Is the banana's distribution much wider that the courgette's?

A further reason can reside in the recognition of the familiar object's shape. As said, memory could be enhanced if the colour is shown in a familiar shape appropriate for the object. Then, what if the shape presented is not easily recognizable? Is the shape of the carrot used in the experiment more recognizable in 2D than in 3D? Or, for example, is the banana shape more identifiable than the courgette?

While the colour constancy effect needed no further investigation, all the above questions prompt the following control tests and analyses to explain the memory colour results.

4.3 CONTROL EXPERIMENT: ESTIMATION OF RECOGNIZABILITY OF THE OBJECTS FROM THE DEPICTED SHAPE

The experiment was divided in two parts, and in total 87 observers, aged between 20-29, participated in the experiment. All the observers were undergraduate or postgraduate students in the Psychology course at Newcastle University, and they were all naïve to the purpose of the experiment. None of them had participated in Experiment 1.

4.3.1 METHODS

A first group of 45 subjects, 15 male and 30 female, was placed in a university classroom. A commonly used projector presented the images of 5 objects to the students, as in Figure 4.16.

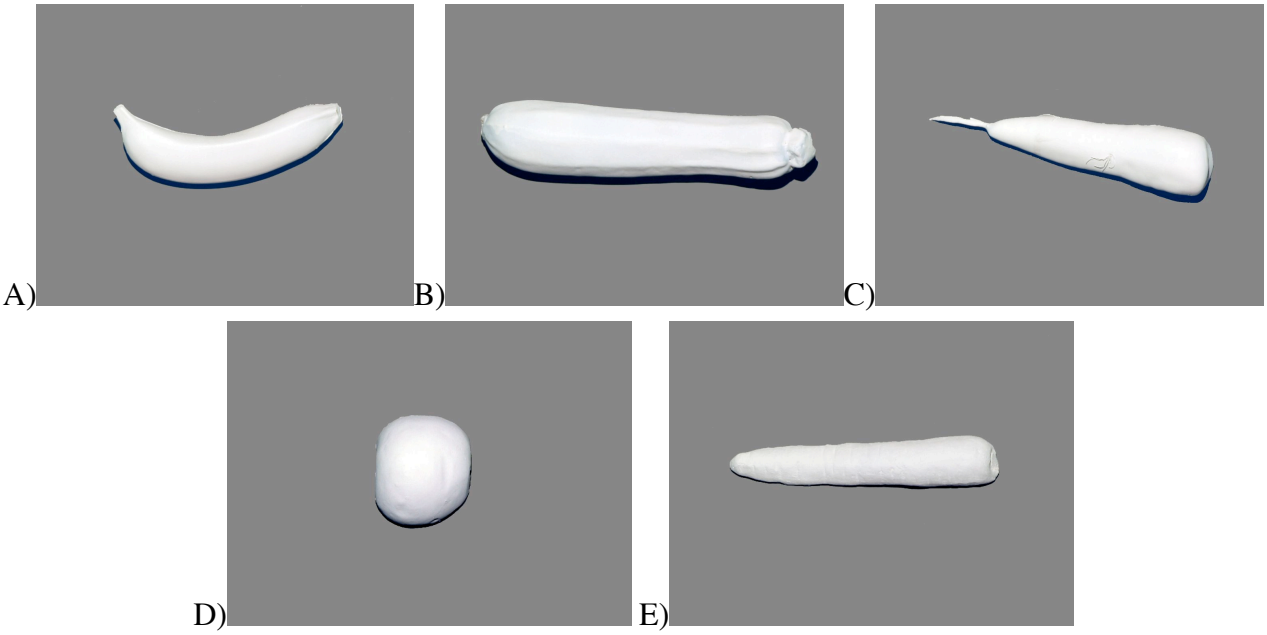

Figure 4.16 – Images of the 3D replicas presented to the first group of observers. A) Banana. B) Courgette. C) First carrot. D) Apple (see Chapter 6). E) Second carrot (see Chapter 6).

These objects were photographs of the original 3D replicas of familiar objects used in Experiment 1 (banana, carrot and courgette) and two supplemental objects (apple and a second type of carrot) that were used in a subsequent experiment (described in Chapter 6). The room was normally illuminated and the image was presented on the entire screen and totally visible to all the participants. A questionnaire was distributed at the beginning of the experiment listing the following questions for each object (example of the answer sheet in Appendix 3): a) what is the object? b) how certain are you in a scale of 1 to 7 (with 1 as very uncertain and 7 as 100% certain)?

The subjects had no time restriction in performing their choices. It was requested that the subjects did not share their opinion with their neighbours and they were seated in such a way to make this interaction difficult. The experiment lasted about 15 minutes.

A second group of 40 participants, 25 female and 15 male, performed the same task and procedure, but a different set of images was shown, as in Figure 4.17. In this case, the objects are the filled outlines of the images presented to the first group.

A difference in the object recognition performance between objects, and between 2D and 3D configurations will help establish how much information the contours and surface shading cues of the objects in Experiment 1 provide.

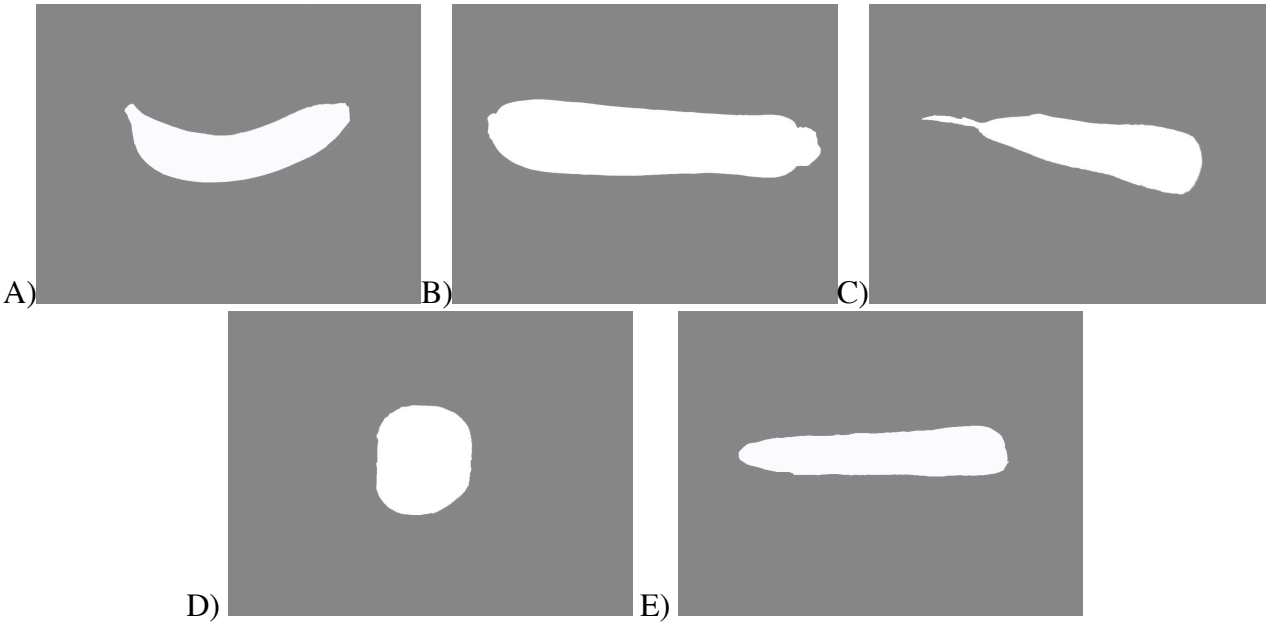

Figure 4.17 - Images of the 3D replicas' outlines presented to the second group. A) Banana. B) Courgette. C) First carrot. D) Apple (see Chapter 6). E) Second carrot (see Chapter 6).

4.3.2 RESULTS

Here are reported only the results of the first three objects (banana, first carrot, and courgette); the results of the apple and the second carrot will be discussed in (Chapter 6 section 6.4.3).

Let us examine first the results from the first group, i.e. the 3D case. The blue bar in Figure 4.18 shows the percent of correct response for the 3D presentations, over all 45 subjects. Only 1 observer recognizes the carrot and with low certainty (4 out of 7). Two observers identified the carrot as parsnip, which is similar in shape but not in colour, yet they are not certain with their responses (4 and 5 in certainty). The courgette is fairly recognized, with 29 out of 45 correct answers, but with a certainty ranging from 2 to 7 (mean 5.75). For the banana, 40 out of 45 observers are 100% certain with their responses, with a correct response certainty mean of 6.82.

Note that the maximum certainty is 7 and the minimum is 1. Thus, I set to zero the subject's "certainty" for a specific object when it is wrongly recognized. Then, the mean of the object's certainty over all observers (certainty index – C) can be computed. Subsequently, it is possible to calculate the overall capacity in recognizing a specific object, as expressed in the following formula:

$$IS_{3D}^{j} = C_j * \frac{A_j}{N} \qquad \text{Equation 4-12}$$

where A is the number of correct answers over all 45 subjects for object j, N is the number of subjects, C_j the certainty index of object j and IS_{3D}^{j} indicates the "identification-by-shape" score for 3D presentation, i.e. the weighted certainty for the specific object j. These indexes for the three

objects are shown in Figure 4.19. The difference between objects' recognisability is visibly great (paired T-test, in both cases $p<0.001$).

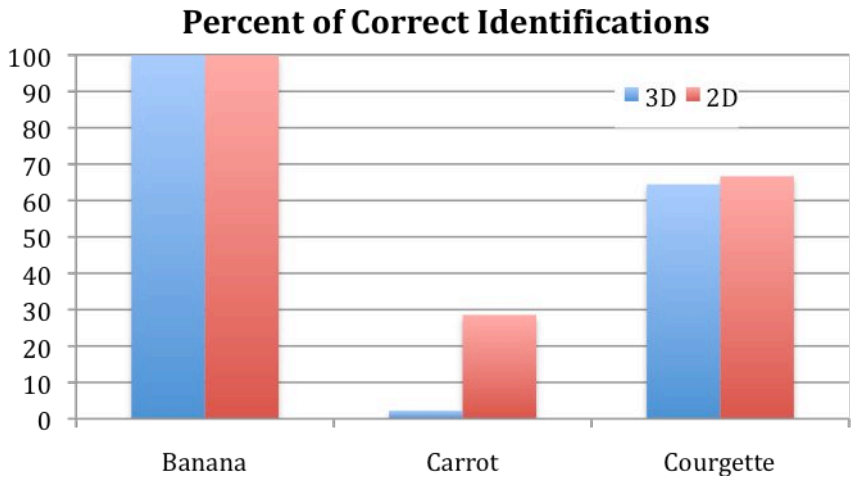

Figure 4.18 – Percent of correct identifications of the three objects, over all participants. Blue bar: 3D presentation (N=45). Red bar: 2D presentation (N=42).

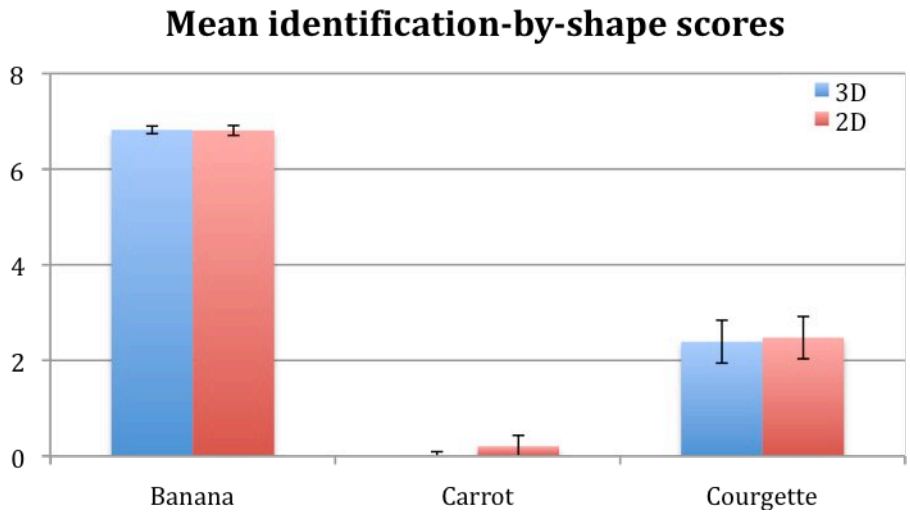

Figure 4.19 – Mean identification-by-shape scores, over all participants. Blue bar: 3D (N=45). Red: 2D (N=42). Error bars: standard error of the mean.

For the 2D case, as for the 3D, I compute the percent of correct responses and the recognition-by-shape index for all three objects over all 42 subjects (as before). Also in this case, all subjects correctly identified the banana shape, with a mean certainty of 6.81. Then, similarly to 3D, the courgette shape is moderately less recognizable that the banana, with 28 out of 42 correct answers, and a certainty ranging from 3 to 7 (mean 5.57). Thus, banana and courgette presented as flat silhouette or 3D shape do not alter the observer responses, as the percent of correct response is approximately the same as shown in Figure 4.18. Instead, the percentage of correct responses for the 2D carrot is higher than in the 3D case, with 12 observers rightly identifying its 2D shape with a certainty ranging from 1 to 7 (mean 2.5). Thus, the carrot 2D shape is more recognizable than its corresponding 3D version, although it is still less identifiable compared to the other two objects. A

paired T-test showed the effect of shape on the carrot certainty to be significant ($p<0.05$). Figure 4.19 plots identification-by-shape scores.

The banana was so uniquely recognizable, with a 100% score, that the possibility of a roof effect was explored with a further test. Ten participants of the first group, were asked to also name the banana profile and then grade, in a scale between -5 and 5, how identifiable was a 3D shape of a banana compared to the 2D shape, with -5 greatly worse, 5 exceptionally better and 0 not difference. The results show that the 3D shape scored a mean value of +1.5, indicating to be much more detectable. The adjusted identification-by-shape score of the banana, calculated adding 2.1 (1.5/5*7) to the 3D shape respect the 2D, is show in Figure 4.20A, together with the other two objects scores unchanged. Therefore, because actually the maximum score achievable is 7+5/5*7=14, we can divide the scores by 14 and obtain indices ranging from 0 to 1, termed "*recognition-by-shape*" indices (RS).

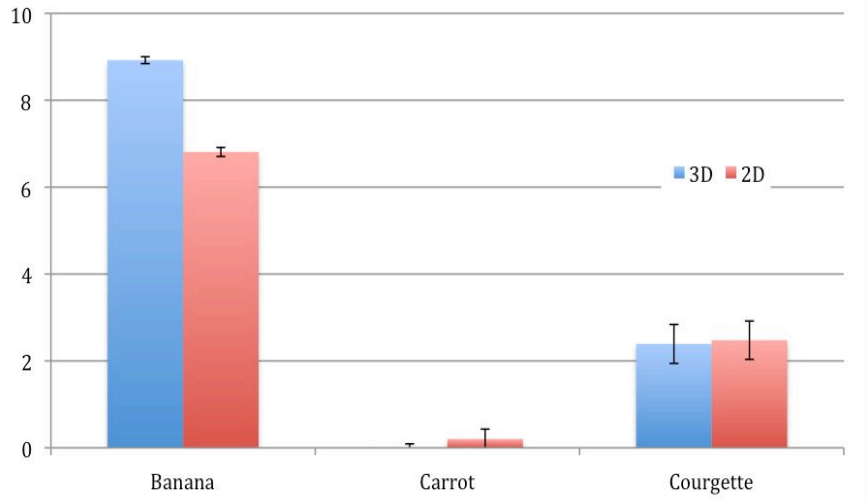

Figure 4.20 – Recognition-by-shape indexes, over all participants. Blue bar: 3D (N=45). Red: 2D (N=42). Error bars: s.e.m.

4.3.3 DISCUSSION

The results clearly demonstrate that among all 3 objects, only the banana is fully identifiable by its profile and shading cues or its profile alone. The courgette is equally recognizable in 2D and 3D, while the carrot is more difficult to recognize in 3D than 2D shape. The latest trend indicates that shading may play a role in object recognition, and may explain to certain extend the perplexing result in Experiment 1. In fact, the painted carrot in the experiments originates from a decorative carrot replica, which it may not be the best example of a carrot seen in a typical British supermarket, although, with its original colour and texture, it was perfectly identifiable. Thus, the presence of surface shading can induce a bigger deviation from the usually seen carrot, introducing additional contours, and lead to lower degree of identification. In the next experiments presented in this thesis a real carrot is used to cast its "perfect" reproduction, hence improve the subject recognition.

In conclusion, even though the fruit and vegetable object used in this experiment are familiar, their contours and shading cues or profiles alone encompass different levels of information. Therefore, this finding offers a first element in explaining the particular pattern in the carrot results in Experiment 1, in that, while the 2D carrot is more recognizable and leads to a bigger range of selection, the 3D carrot shape is as unrecognizable as the 2D disk, from which the same acceptance rate.

4.4 ANALYSIS OF SURFACE CHROMATICITY VARIABILITY OF FAMILIAR OBJECTS

In the previous section it has been proven that the peculiar trend in the carrot may depend on the difference in recognizing its shape when presented in 2D or 3D. This section, instead, we ask if there is a difference between the chromatic distributions of the three studied objects and if it can explain the results obtained for the courgette.

As said in 4.2.3, a possible reason for object-dependent trends in the observer's responses may be that, in the presence of the 3D object, the subject tends to remember better the range of colour of the object. In other words, let us hypothesize that for each object we have more than one memory colour, thus a range of memory colours associated to a specific object. In addition, let us assume that when asked to recall the colour of an object only through linguistic information, i.e. the name of the familiar object, we can remember always the same small amount of colours independently of the object's natural colours distribution. Let us call this range, i.e. the range of memory colours recalled in absence of any object shape cues, the "*pure memory colour range*" of the object. Therefore, if the size of object's natural colours distribution is small, the two quantities (distribution size and pure memory range) might almost coincide. Conversely, if the size of object's natural colours distribution is large, the pure memory colours range would be a fraction of the full distribution. Hence, if we postulate that the observers have a memory colour range wider than the pure memory colour range that can be better accessed if appropriate cues are given, than we can propose shape cues as a possible candidates. Base on these assumptions, the subject's range of choices would increases for objects with large chromatic variation and remain constant for objects with small distributions. The following experiment shows the difference between natural bananas, carrots, and courgettes distributions, proving our theory to be valid.

4.4.1 METHODS

Chapter 3 describes a method, developed in this thesis, to record accurately and analyze the chromaticity distributions of natural objects under various illuminations with the purpose of

investigating chromatic texture. Here we employ such method to examine the colour information possessed by the natural objects used in Experiment 1.

A)  B) 

C) 

Figure 4.21 – Objects analysed in this experiment. A) banana; B) carrot; C) courgette.

A banana, a carrot and a courgette bought from a high street supermarket were selected for this analysis. Each object was photographed using the Nikon D70 S camera and the procedure described in 3.3.1. To reiterate, the object is placed in a Verivide viewing cabinet on the fixing platform, photographed, segmented from the background and pixel-by-pixel converted from RGB dependent colour space in CIE XYZ colour space. The light used for analysis is CIE standard light at colour temperature 6,500 K (focusing on the memory colour task). Figure 4.21 shows the pictures of the three objects analysed in this experiment.

4.4.2 RESULTS

As described in our previous works (Hurlbert, Vurro, and Ling (2008)), when the original chromaticity distributions of a natural object are transformed into EMG cone contrast space (Eskew, McLellan, & Giulianini, 1999), they form distinct signatures under D65. A more in-dept analysis of these signatures and their variation under different light source will be discussed in Chapter 5, while in this chapter I will only examine the over all spread of the distributions.

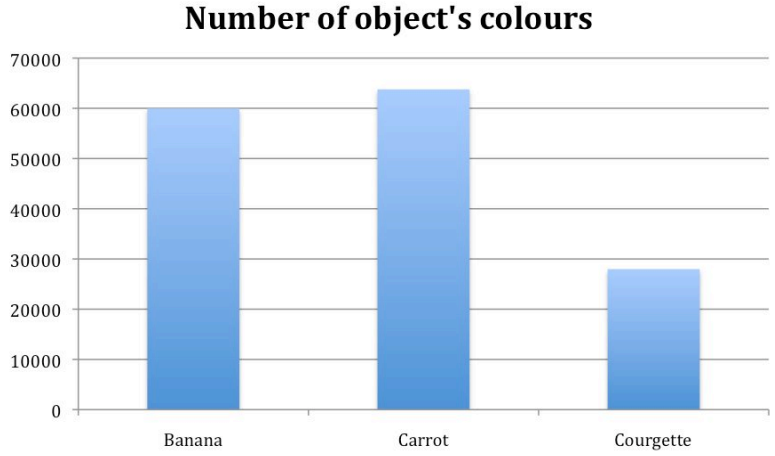

Figure 4.22 – Number of surface colours for the banana, carrot and courgette used in this experiment.

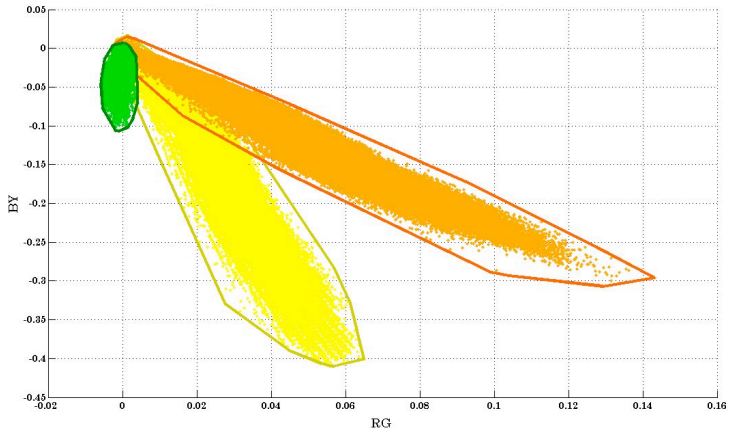

Figure 4.23 – Distribution of the banana (yellow dots), carrot (orange dots), and courgette (green dots) in EMG cone contrast space. Objects' distribution contours: dark yellow line=banana; dark orange line=carrot; dark green line=courgette. Note that [0,0] coordinate is the neutral point

Figure 4.22 shows the number of distinct colours possessed by each object recorded at a resolution of 3008x2000, i.e. the number of pixels with different XYZ values. All objects demonstrate clear polychromaticity, in that each of them possesses a large number of distinct surface chromaticities. Although, the courgette possesses half of the number of colour compared to the banana and carrot. Moreover, the size of their chromatic spread varies greatly. A way to determine the extension of the object distribution is to calculate the area containing its chromaticities, discarding the points repeated only once.

Figure 4.23 illustrated the objects' distributions in EMG cone contrast space[5] and the contour of such area for all three objects. The courgette has the smallest variation in EMG plane, while the carrot distribution covers a bigger area. Finally, the banana possesses the largest spread in colour, both in terms of hue and saturation. For each object, the value of the area enclosed by the convex hull drawn in Figure 4.23, i.e. the size of the colour variation, was calculated to quantify the difference between spreads and called *object colour variation* index (OV). The OV indices are: 0.0106 for the banana, 0.0108 for the carrot and 0.00086 for the courgette (Figure 4.24). Clearly, banana and carrot possess a much wider spread than the courgette.

[5] For details on the RG-BY colour space used in this chapter and how is scaled relative to the discrimination threshold see section 1.1.1 and Eskew et al. (1999).

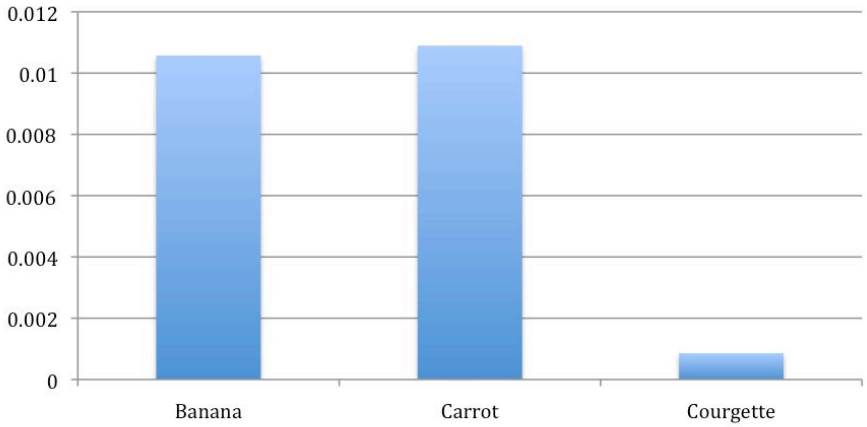

Figure 4.24 – Object colour variation indices.

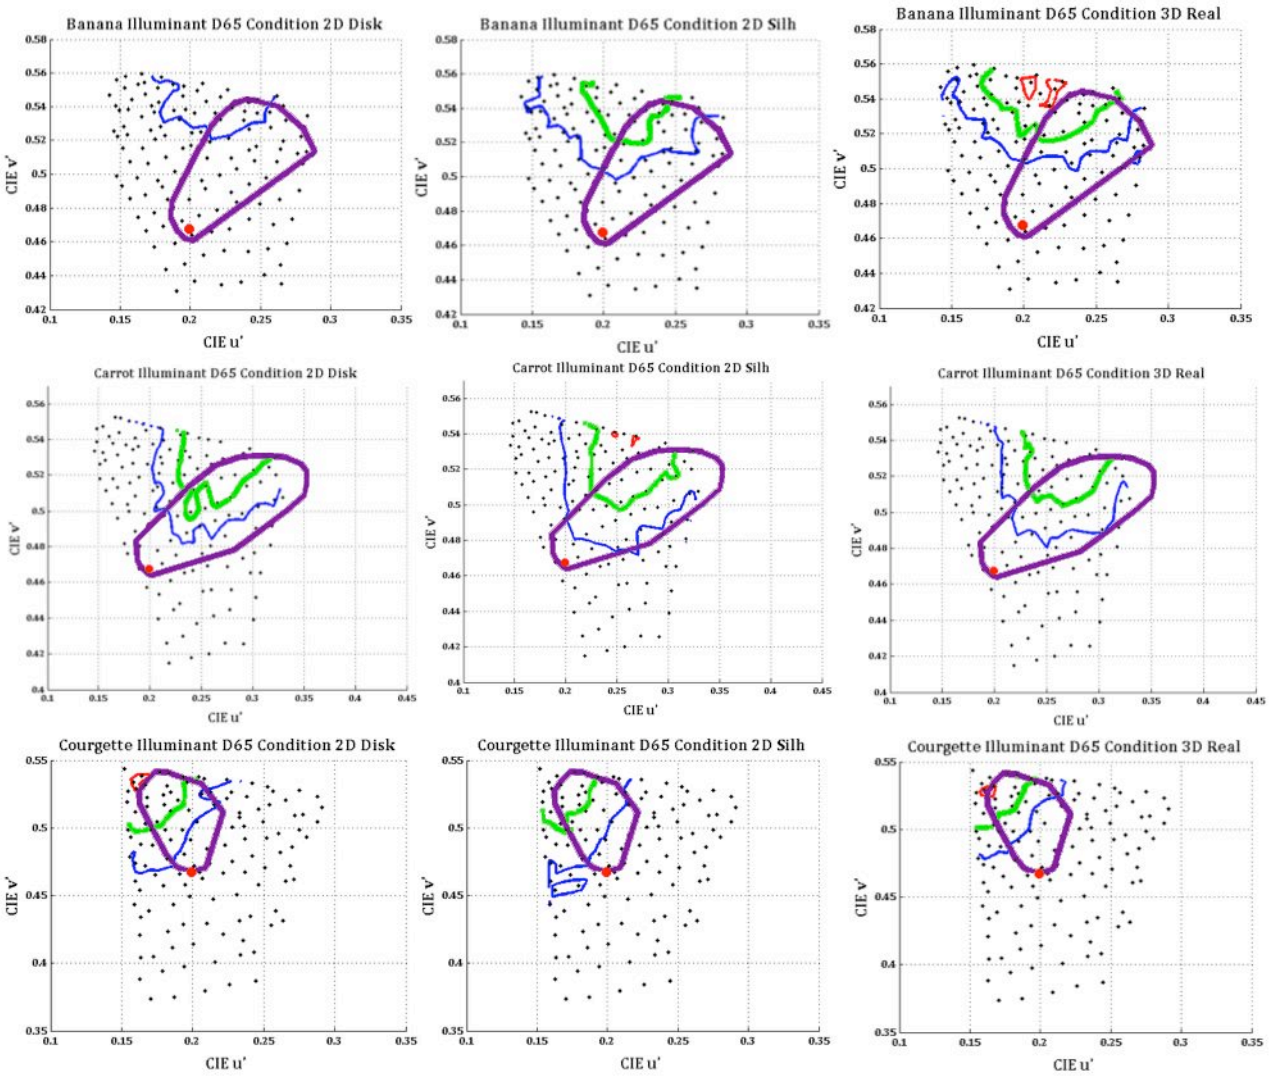

Figure 4.25 - Distributions for banana, carrot and courgette, compared with the total number of 'Yes' responses of Experiment 1 in u'v' isoilluminat space, for all 15 observers. Black dots: test colours. Violet closed line: object distribution contour. Blue, green and red lines: contours of the areas where, respectively, more than 10%, 50% and 90% of the subjects judged the values enclosed by the line as appropriate as described in section 4.2.2.1.2 and in Figure 4.9. Red dot: neutral point. First row: banana. Second row: carrot. Third row: courgette.

Figure 4.25 compares the objects' distributions (violet closed line) with the responses of the main experiment, for all three objects under D65. The figure shows a consistent overlap of the distributions and the 'Yes' responses (in the figure the red, green and blue lines delimit the

chromaticities selected by 10%, 50% or 90% of the subjects as appropriate for the object respectively, as described in section 4.2.2.1.2 and in Figure 4.9). Interestingly, the overlap increases for the banana and carrot with the introduction of a recognizable congruent shape (note that for the carrot when the shape becomes less identifiable, i.e. introduction of 3D shape, the overlap slightly decreases). Conversely, the natural spread of the courgette's distribution coincide, almost totally, with the range of colours accepted as appropriate for the object in the 2D disc condition and the range that does not significantly change among different conditions. Keep in mind that the range in the 2D condition of any object is the previously defined *pure memory colour range* of that specific object.

4.4.3 DISCUSSION

These results prove the predictions expressed in the introduction of this section. If the pure memory colour range of the object is a fixed quantity and the object's surface colour distribution is limited in extent to a similar range, then we would expect no difference in acceptance between different shape conditions. The courgette results demonstrate such a point. On the contrary when the pure memory colour range is only a fraction of the large spread of the object surface colour distribution, as for the banana and carrot, the introduction of more appropriate cues leads to an increase in the number of choices.

4.5 GENERAL DISCUSSION

This chapter presents one main experiment, one control experiment and a final object colour distribution's analysis. The main experiment explores memory colour and colour constancy of familiar objects when presented in diverse shape and shading conditions. While the observers show very low colour constancy, not significantly different for all objects and conditions, more interesting results were obtained in the analysis of their memory colour. Observers' memory colour responses of natural fruit and vegetables were measured as a function of shape cues to object identity. Then the effect was examined on distinct levels: mean percentage of selected colours, mean memory colours, range of memory colours, and the response times. The results suggested the existence of a range or a distribution of memory colour associated with each familiar object. Furthermore, whereas no significant interaction is found between shape and mean memory colour itself, there is a robust, object-independent interaction between object shape and response times. Conversely, puzzling results were obtained for the interaction between shape and mean percentage of selected colours, as well as, shape and range of memory colours. A control experiment and an object's colour distribution analysis were designed to explain such findings.

The control experiment measured the recognisability of the objects' shapes used in the main experiment. The shape of the banana was the easiest to be recognized by just its shape and shading or shape alone, while the carrot was the hardest, being particularly difficult in the 3D configuration.

The analyses of the surface colour distributions of the objects were designed to measure the chromatic spread of the object colour distribution used in the main experiment. Results indicate that natural bananas and carrots possess a much larger variation in their chromatic spread compared to the natural courgettes. Below I discuss the final conclusions we drew from these findings.

4.5.1 INTERACTION BETWEEN FAMILIAR OBJECT SHAPE AND MEMORY COLOUR MEAN

First, the main experiment proves that familiar objects have more than one memory colour, i.e. a range of memory colours. Next, a consistent shift on the memory of the object's colour occurs along the saturation dimension. This is in line with other studies (Bartleson, 1960; Newhall, Burnham, & Clark, 1957; Perez-Carpinell, Baldovi, de Fez, & Castro, 1998; Siple & Springer, 1983) showing that usually a remembered colour is more saturated than the original stimulus, regardless of the time lapse. Lastly, no significant effect of shape is found on the mean memory colour, i.e. the core of the memory colour doesn't change. This is accordance with numerous studies (Bolles, Hulick, & Hanly, 1959; Bruner, Postman, & Rodrigures, 1951; Siple & Springer, 1983; Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Olkkonen, Hansen, & Gegenfurtner, 2008). Specifically, as seen in the introduction, Gegenfurtner and colleagues (2006 and 2008) demonstrate that the mean memory colour bias for uniformly coloured outlines shapes were not significantly different from the neutral control stimulus. Furthermore, they showed that, in the most natural presentation of the object (photograph of the original object), the memory colour bias was on average 9.5% deviated from the neutral control stimulus (supposedly less that 5 ΔEuv - see section 4.1). Therefore, because in this chapter experiment the observers in this experiment can choose between a limited number of colours, with a perceptual colour difference between two adjacent points vary in a range between 5-20 CIE ΔEuv colour difference[6], then, if any effect of shape exists, such effect is too small to be detected in this particular experiment. Hence I can conclude that the mean memory colour is not influenced by shape or only marginally.

[6] Note that the main concern in selecting the stimuli for the main experiment was to sample the entire chromaticity plane while keeping the number of trials in the experiment to a minimum.

4.5.2 *INTERACTION BETWEEN FAMILIAR OBJECT SHAPE AND THE RANGE OF MEMORY COLOURS*

To date, this study is the first that demonstrate an effect of shape on the range of memory colours, both in percentage of selections and extension of the area selected in colour space. More importantly, this effect is object specific, with diverse trends for different objects, and it is larger than the effect produced on the mean memory colour. Here I will demonstrate that this effect is a combination of mainly two cues in the test object representation: its original chromatic distribution and shape recognizability.

First, if the shape of the tested object carries congruent recognizable cues about that object, it may help the observers to better access their visual memory. On the contrary, if is not recognizable or incongruent with the memory of the object shape, this makes the retrieval of information from memory less efficient and incomplete. Consequently: (1) presenting in their congruent form (2DN or 3DN) the banana and courgette, which are both clearly identifiable by shape alone as shown in the results of Experiment 2, will enhance memory retrieval; (2) presenting the flat silhouette of the carrot will improve it slightly, being much less identifiable; (3) presenting the 3D carrot will not improve it at all, being almost totally unrecognizable.

Secondly, there are also clear differences in the chromatic distribution patterns of the natural objects, as illustrated in the analysis in section 4.4. The natural banana and carrot textures contain a large range of chromaticities, which can be better accessed in visual memory by adding more shape cues. In contrast, the natural courgette's surfaces do not possess such an extent of variation, and even with better accessibility through shape cues, will not impact as much the memory colour range. In addition, as shown in Figure 4.25, the more cues to the object identity are added the more the range of the observers' selection overlaps the object distribution in CIE u*v* isoilluminant space.

Combining the conclusions drawn above it is possible to convey the contribution of both factors in a model, as in the following formula:

$$\sigma_{MC_j} = A_j * RS_j * OV_j + \sigma_{PM_j} \qquad \text{Equation 4-13}$$

where, for object j, σ_{MC} is the memory colour range, RS_j is the recognition-by-shape index as calculated in section 4.3 , OV_j is the object colour variation index calculated in section 4.4, σ_{PM} is the "pure" memory range of the object, and A_j is an parameter the same for all conditions, but dependent on the object.

For example, using this formula to simulate the observer behaviour for the three objects *under D65*, and assuming σ_{PM} equal to ratio of the subjects' colours selection by the total of

possible selections for object j (it is logical to assume this as the memory baseline), and RS_j zero for the disk (there is no shape information), and A_j=1 for all object, we obtain the results in Figure 4.26. This graph reproduces well the observed effects in the main experiment (graph in Figure 4.7C also re-showed in Figure 4.26B), concluding that this simple model does in fact explain the results. Obviously these are not the only factors influencing memory colour range and other object features, such as surface material properties or surface texture, may also play a role. Nevertheless, in the matter of this experiment, this model can explain the majority of the results relative to the memory colour range.

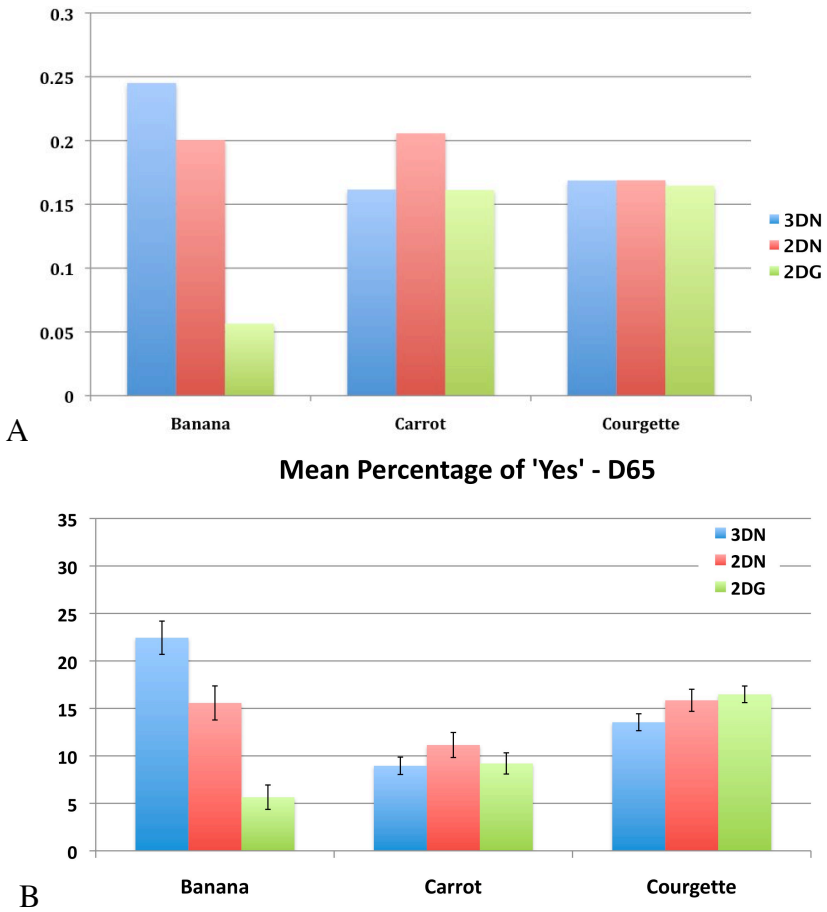

Figure 4.26 – A) Model: prediction of the observers' behaviours for the memory colour range for D65. Green bar: 2D disk. Red bar: 2D profile. Blue bar: 3D solid object. B) Same graph as Figure 4.7C; for comparison.

4.5.3 INTERACTION BETWEEN SHAPE AND COLOUR RECOGNITION DELAY

Supporting evidences that 'a priori' knowledge of the object colour could help the retrieval process and improve object recognition and classification comes from studies mainly using single objects and scenes (Hansen, Olkkonen, Walter, & Gegenfurtner, 2006; Oliva & Schyns, 2000; Ostergaard & Davidoff, 1985; Tanaka & Presnell, 1999; Biederman & Ju, 1988; Davidoff & Ostergaard, 1988; Humphrey, Goodale, Jakobson, & Servos, 1994; Oliva & Schyns, 2000; Price & Humphreys, 1989, Liebe, Fisher, Logothetis, & Rainer, 2009). These works employ mostly objects

with diagnostic colours, such as a yellow banana or a green forest, and reported an advantage in object or scene recognition with the presence of colour. For example, Tanaka and Presnell (1999) found that recognition of objects with high colour diagnosticity (i.e. a taxi or a fire engine) benefited more from colour than object with a lower indicative colour (i.e. a lamp); additionally, Liebe et al. (2009) demonstrate a strong effect of colour in achieving higher performance in visual memory and recognition tasks both in human and primates.

On the other hand, very few studies have analyzed the temporal factor on information access (Spence, Wong, Rusan, & Rastegar, 2006; Gegenfurtner & Rieger, 2000; Wichmann, Sharpe, & Gegenfurtner, 2002). For example, Gegenfurtner and Rieger (2000) showed that natural image presented in colour required less time to be recognized, while Spence et al. (2006) demonstrated, in successive delayed matching-to-sample paradigm, that at equal stimuli exposure times, colour images had a better recognition performance.

Finally, to date, the author is not aware of any study exploring the interaction between colour and recognition time as a function of shape, this experiment being the first of its kind. For all object and illumination, the results clearly, strongly, and consistently prove that a sensible colour–shape association accelerates the response time, presumably as the result of an enhanced image representation in memory due to the additional attributes.

Thus, the interaction between colour and shape cues implies a conceptual knowledge about the object, and indicates that shape influences the recognition of an object by its colour in respect of temporal factors at the time of retrieval.

4.5.4 CHANGE IN ILLUMINATION EFFECT

Results show very poor colour constancy and no effect of shape across illumination change. The reason for these results may be two-fold: border effect due to the projector gamut limitations or failure of constancy of the observers group.

The latest assumes that the *entire* population studied in this experiment is incapable of normally performing in a colour constancy task. Because the subjects had normal colour vision and four of the 15 observers performed normally in the following experiment (Chapter 6), this hypothesis can be safely discarded. However, the incomplete adaptation of the subject during the experiment cannot be totally discarded. Hence, the need of designing a new and better set-up for future experiments as described in previous chapters.

The second explanation presumes that the projector gamut is too small to allow reaching chromaticities that the subjects would consider colour constant relative to their selection under D65. Consequently their colour constancy would be very poor. For example, Figure 4.27 illustrates the subjects' mean selection under D65 (i.e. their memory shift, red cross in the figure) and the

hypothetical situation in which they would achieve perfect colour constancy under D40/GE[7] (blue) and D250 (green) for the 3D condition. The heads of the continuous arrows (i.e. subjects' illuminant shift vector) indicates points of perfect constancy for D40/GE (blue arrow) or D250 (green arrow). Each subfigure in Figure 4.27 represents one object. Keep in mind that to achieve perfect constancy the subjects' illuminant shift vector must start from his/her memory colour of the object and be parallel and have the same length that the illuminant shift vector (dashed arrows in the figure; D40/GE=blue dashed line and D250=green dashed line). Therefore, in Figure 4.27 the green (blue) dashed arrows are parallel and have the same length that the continuous green (blue) arrows, and the arrowheads indicate. Then, if, for each object, we look at the test points (black points in the figure) used for a specific object in the experiment, it is clear that the set of colours given to the observer to choose from his/her colour constant value was insufficient. Remember that the choice of test points depends on the projector gamut at the specific isoilluminant plane (dark red contour in Figure 4.27). Therefore the size of projector gamut precludes achieving constancy. This is also confirmed by the difference in range selection and peak of selection under different lights, indicating an incomplete effect of illuminant. Lastly, if the selection is squeezed on the border of the projector gamut (Figure 4.14 and Figure 4.27) any effect of shape cannot be seen in this experiment. This unravels the main limit of the set up used in this experiment and the reasons for design and develop the second apparatus (section 2.3).

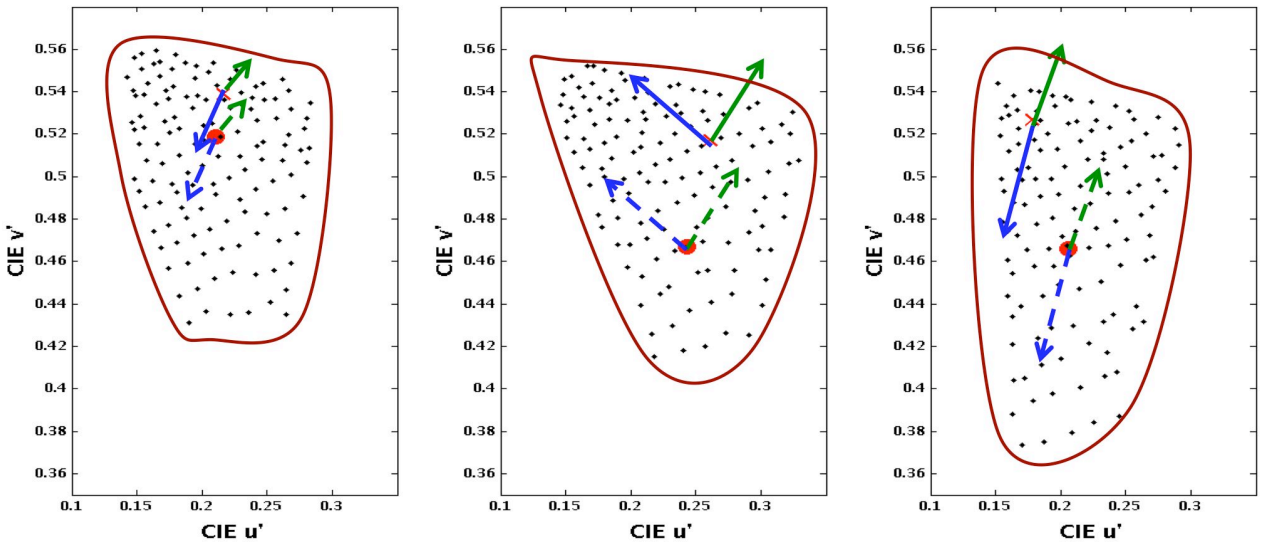

Figure 4.27 - Hypothetical mean subjects' perceptual chromaticity shifts in case of perfect colour constancy, for the 3D shape for the three objects. Dashed arrows indicate mean subjects' perceptual shifts in case of perfect constancy, from their memory colour setting under D65 (red cross) to D250 (green) or to D40/GE (blue). Solid arrows: physical shift of the object colour under the three lights; this is equal to zero memory shift and perfect constancy. Arrowheads denote settings under chromatic illuminants D40/GE (blue) and D250 (yellow). Filled circle: measured object colour under D65 (red), D250 (yellow) and D40/GE (blue). Black dots: test colours. Bordeaux coloured contour: projector gamut for the chosen Y value (object-dependent). Left: Banana. Centre: Carrot. Right: Courgette.

Finally I can speculate that, considering the conclusions drawn for memory colours, if the total extent of the effect could be measured a difference between shape cues would be found.

[7] GE illuminant replaces the D40 illuminant for the carrot.

4.6 CONCLUSIONS

The effect of familiar object shape was examined on colour constancy and memory colour at multiple levels. The results demonstrated the existence of a range of memory colours associated with a familiar object, the size of which was strongly object-shape-dependent. I proposed a relationship between memory colour range and (1) shape cues to object identity and (2) natural chromatic distribution of the object. Furthermore, I suggested the influence of shape cues to the object identity affect directly the range of memory colour proportionally to the original object chromatic distribution. In particular, object's shape interacts with memory colour range facilitating the cognitive task of recognition. Results also showed that 3D shapes, if more indicative of the object identity, increased the subject performance. These results were consistent across all illuminants. A simple model of memory colour range was introduced which explains the majority of the result relative to the memory colour range and the subjects' behaviour in different conditions. Consistently for all objects and illuminations, memory retrieval is significantly faster when the object shape is introduced, indicating an advantage of coherent shape-colour interaction in recognition. Shape effect on the memory colour itself is not significant or presumably too small to be detected in this experiment. Colour constancy was poor due to the limitation of the current set-up.

In conclusion, this study proves that coherent object-shape association facilitates memory colour in time performance and information retrieval, and aids object recognition. These results are consistent across all test illuminants.

Chapter 5

CHROMATIC TEXTURE ANALYSIS OF NATURAL FAMILIAR OBJECTS

5.1 INTRODUCTION

"How does the visual system achieve constancy across all possible scenes?" This is a problem that has puzzled computational vision scientists for years. I have previously described several algorithms that try to solve this intriguing problem, but all fail to be applicable in any situation. Yet, if observers possess the memory colour of one or more objects in the scene, it is the opinion of this author that it could be "easily" solved. In fact, for example, a typical strawberry's colour is "red" and this information can be utilized as an "anchor" to determine the illuminant colour. In the previous chapter, I have analysed the chromatic variegation of natural familiar objects' postulating a possible interaction with the memory colour of the object. Furthermore Tominga&Wandell (2002) have shown that it is possible to estimate from an image of a scene the unknown light illuminating it, if we know the image range and distribution.

Based on these considerations, and aiming to create the ground for developing an algorithm for the manipulation of natural chromatic distribution for successive experiments, I have analysed certain features of a number of natural objects. The objects could belong (1) to different object categories (e.g. banana, potato etc.) and were examined under diverse illuminations or (2) to the same category, i.e. exemplars of the same object category (e.g. three carrots), and were examined under same or different illuminants. This chapter concisely presents the results of such analyses, additionally examining findings on a possible relationship between the range and distribution of familiar objects' colours and (1) memory colour and (2) colour constancy. Furthermore, this "computational" chapter act as a preamble to the experiments presented in the following chapters.

5.2 METHODS

For the first part of the experiment, a database of 8 natural objects was collected. The objects are common fruits and vegetables found in a popular UK supermarket: Melrouge apple, banana, carrot, clementine, lime, plum, potato, and strawberry. These were chosen in order to cover a large range of object colours. Figure 5.1 represents thumbnails of the complete set of 8 objects, taken under CIE standard daylight at correlated colour temperature 6,500 K in the Verivide Colour Assessment cabinet. A second group of 9 objects was considered in the second part of experiment: three bananas at different stages of ripeness, three apples of different variety and three carrots of different types (Figure 5.2). Objects belonging to the same category (or "family") were photographed in different orientations and examined under three illuminants listed below (for example Figure 5.3 shows the three apples under all illuminants).

Figure 5.1 – Image cutout of the 8 objects used in this experiment photographed under standard daylight D65. From top left corner to bottom right: banana, Melrouge apple, clementine, lime, potato, plum, carrot, and strawberry.

Figure 5.2 –Photographs of the objects used in the second part of the experiment. A) Three apples of different quality (left to right: Melrouge, Gala, and Pink Lady). B) Three bananas in different stage of ripeness. C) Three carrots of different type (left to right: Organic, typical British, and Imperator carrot (typical Italian)).

The images were obtained with the Nikon D70 SRL camera calibrated as in Chapter 3at a spatial resolution of 3008 by 2000 pixels and mounted with an 18-70mm kit lens. The pictures were taken in the Verivide Colour Assessment cabinet under the three illuminants described in section 3.3: D65, CWF, and F. The spectra of the three lights, taken by a PR-650 spectroradiometer in the same Verivide Colour Assessment cabinet, are illustrated in Figure 5.4. Images were converted into tristimulus values (XYZ) for the CIE 1931 standard colorimetric observer using the model described in 3.3. Next they were converted into EMG cone contrast space as described in 1.1.1 (see also Eskew et al, 1999). This space was chosen because (1) it separates chromaticity and luminance,

thus they could be adjusted independently, and (2) it is directly based on the physiological properties of the visual system, simulating its early stage adaptation and encoding.

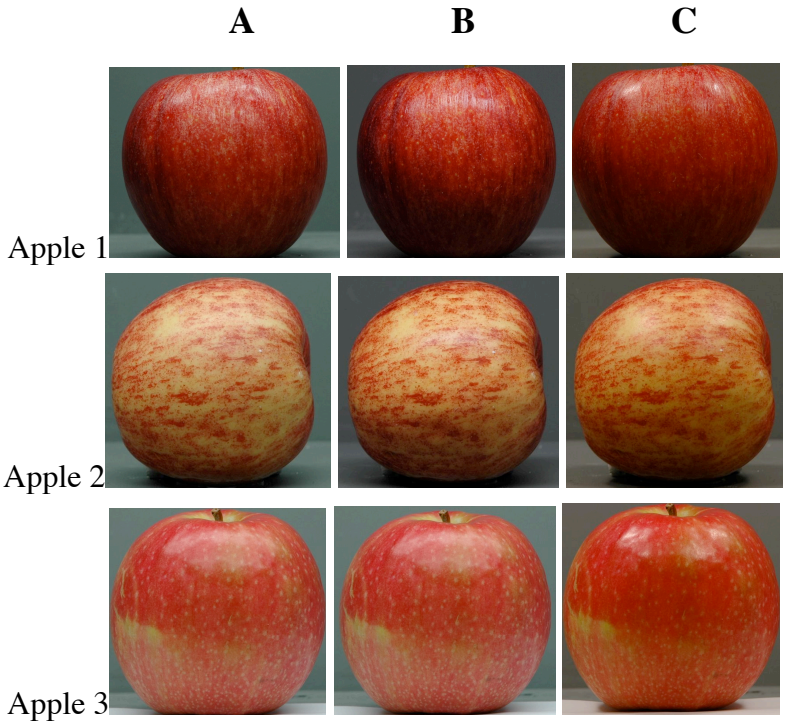

Figure 5.3 – Photographs of the three apples used in this part of the experiment under the 3 lights. A) D65. B) CWF. C) F. Apple 1: Melrouge variety. Apple2: Gala variety; Apple3: Pink lady variety.

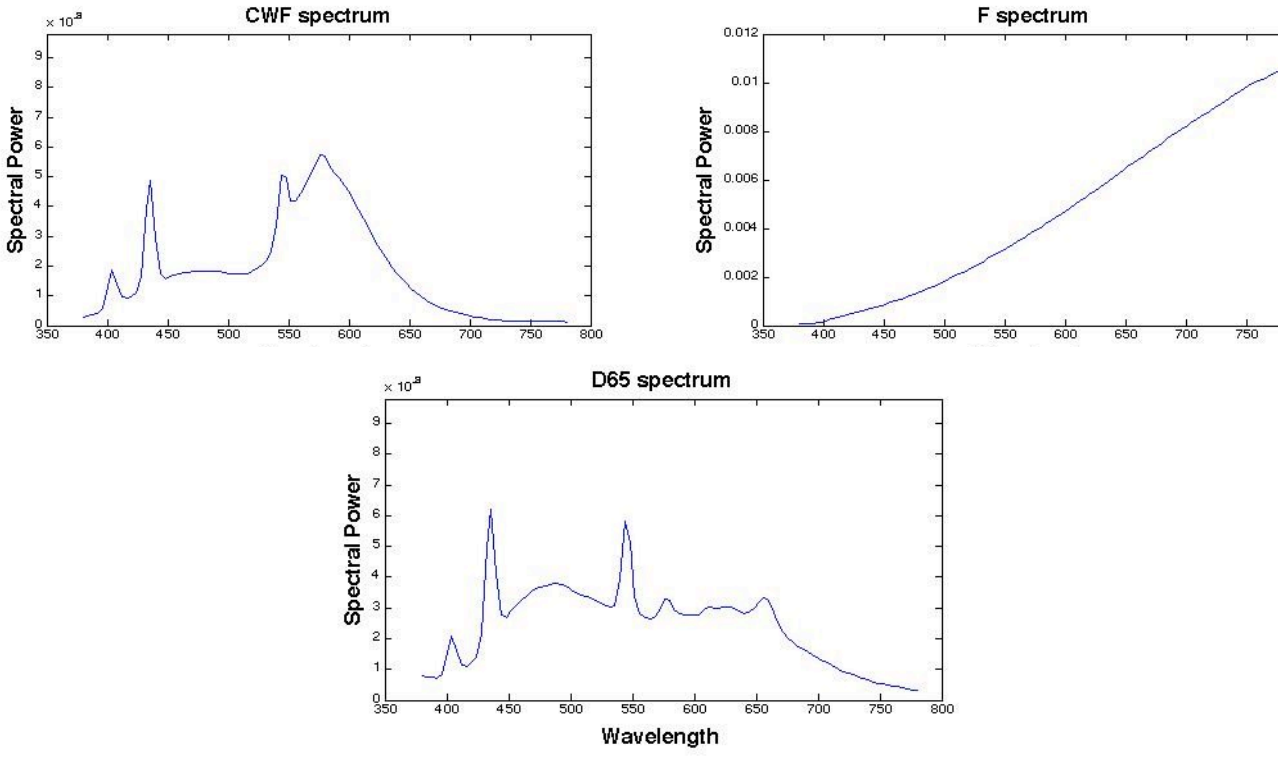

Figure 5.4 – Spectra of the three illuminations as taken by a PR-650 spectroradiometer in the Verivide Colour Assessment cabinet used in this thesis. Top left: CWF; top right: F; bottom: D65.

As described in section 1.1.1, the EGM cone contrast space is an orthogonal space with an achromatic luminance axis, the L + M axis, and two chromatic axes, the L - M axis, and S - (L + M)

axis, that form an isoluminant plane. These three axes intersect at the white point of the scene. Equation 1.2 shows the EMG convertion formula. To reiterate what it was stated above and in Chapter 1, the origin of this colour space represents the adaptation point of the scene (the value of the cone excitations for a white sample in the scene), the axes represent the change in cone excitation from the adapting point, and the coefficients in the formula in Equation 1.2 were obtained from discrimination experiments (see Eskew et al. 1999).

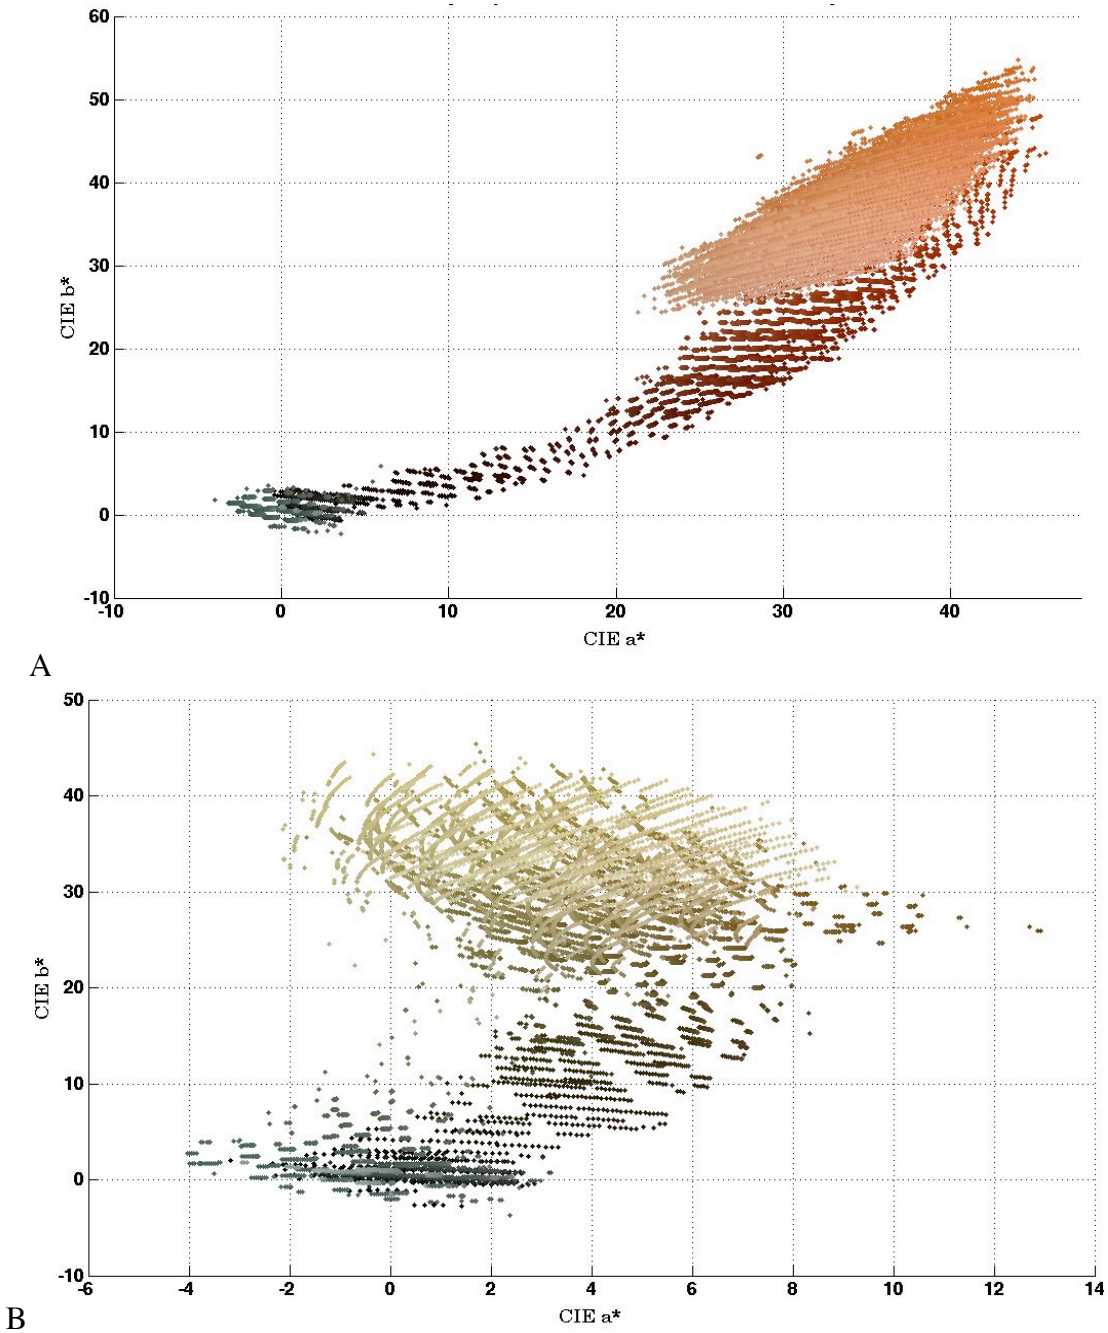

Figure 5.5 – Chromaticity distribution of two analysed objects in L*a*b* space; a*b* isoluminant plane under D65. A) Carrot; B) Potato. Displaying only colours present more than 5 times in the object's image.

In this system, a point can be defined using cylindrical coordinates: radius, azimuth and height (r, ϑ, h). The radius, i.e. the distance between the projection of the point on the isoluminant plane and the origin, can be related to saturation changes. The azimuth, i.e. the counter-clockwise

angle formed from the RG axis, defines the chromatic direction and can be assimilated to a hue variation (in angles between 0 and 2π). The height defines the luminance value.

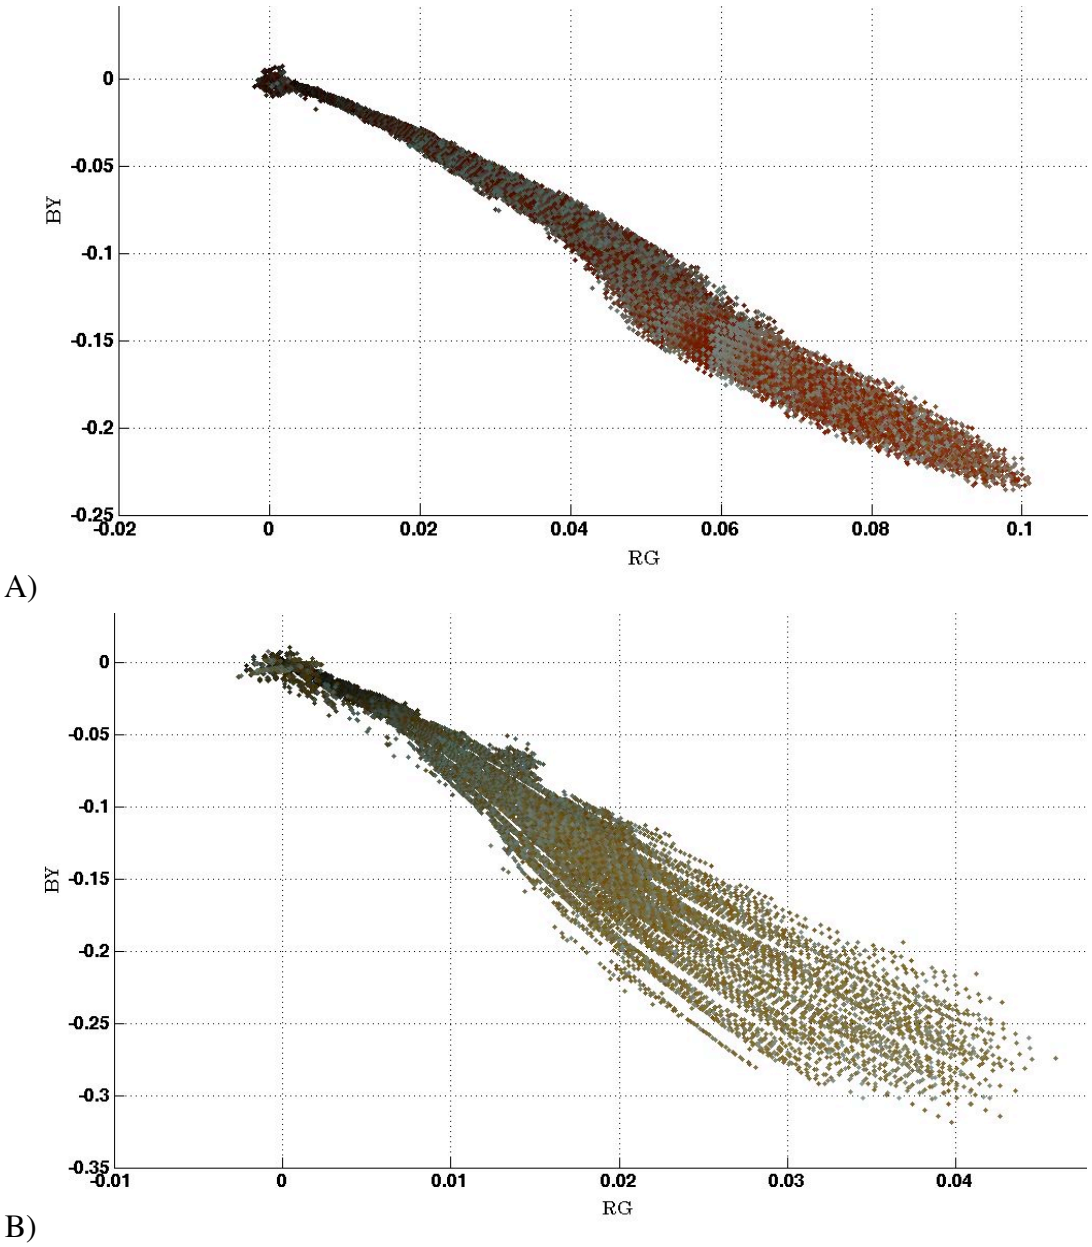

Figure 5.6 - Chromaticity distribution of two analysed objects in EMG cone contrast space isoluminant plane under D65. A) Carrot; B) Potato. Displaying only colours present more than 5 times in the object's image

As an example, Figure 5.5 and Figure 5.6 show, for D65, the colour distributions of two objects in the database plotted in CIELAB space and in EMG cone contrast space[8] respectively. For clarity, only the data points present more than 5 times in the image are represented in the graphs, but all points were used in the following analyses. The depicted colour of each point represents an approximation of its actual colour. In contrast to Linhares et al (2008), all points of the image were considered not only the discernible colour, since the final scope is to determine an overall index (or indices) for the entire distribution. It is worth pointing out that a scaling unit of the RG axis (horizontal axis) is equal to one scaling unit of the BY axis (vertical axis) because (1) the EMG

[8] It should be mention that the axis scales in the figures are not in units of equal threshold discriminability.

cone contrast coordinates are obtained as a linear combination of fraction of cone differences, and (2) they define the contrasts between the adaptation point and the cone response to a colour stimulus in the scene.

The quantity of highlights was considerably reduced by using vanish matt spray paint as described in 3.4. Alas, for some objects the removal was incomplete. If some points show values of luminance larger than 99 in Lab space, they were considered to be corresponding to highlights; these points were less than 1% of all data and were excluded from further analysis.

It is possible to divide the analysis in two classes: (1) "between object category" effect and (2) "within object category" effect. In the first case, I studied a group of objects belonging to different vegetable or fruit categories (i.e. one banana, one carrot, one cucumber and so on) under three different light sources. The aim here is to find common parameters for the illumination-dependent change of the chromaticity distributions. In the second case, all the objects belong to the same category (i.e. N bananas or N carrots etc) taken under one light or different lights. The goal in this case is to examine the overall effect of illumination on the diverse distributions of one category. Differences between effects were tested with repeated measures ANOVA.

5.3 RESULTS

The results are organized in two sections. In the first I compare objects of different fruit or vegetable category; in the second, I present results for objects that are different exemplars of the same fruit or vegetable category.

5.3.1 BETWEEN OBJECT'S CATEGORY EFFECT

Figure 5.7 shows, as examples, the chromatic distributions of 5 of the 8 objects under D65 plotted in RG-BY cone contrast space (for clarity only points repeated more than 5 times are plotted): apple (in red), banana (yellow), clementine (orange), lime (green) and plum (violet).

Interestingly, for all objects, the distribution of surface colours clusters in a certain area of the cone-contrast space, forming a characteristic signature, distinct for each object. Clearly, this object-dependant colour segregation provides, to some extent, a way to segment the image into individual objects. In addition, their shape is quite peculiar: they form an elongated cluster that follows a radiant direction from the neutral point. The regularity of these chromatic signatures suggests that they may provide robust information about the surface colour.

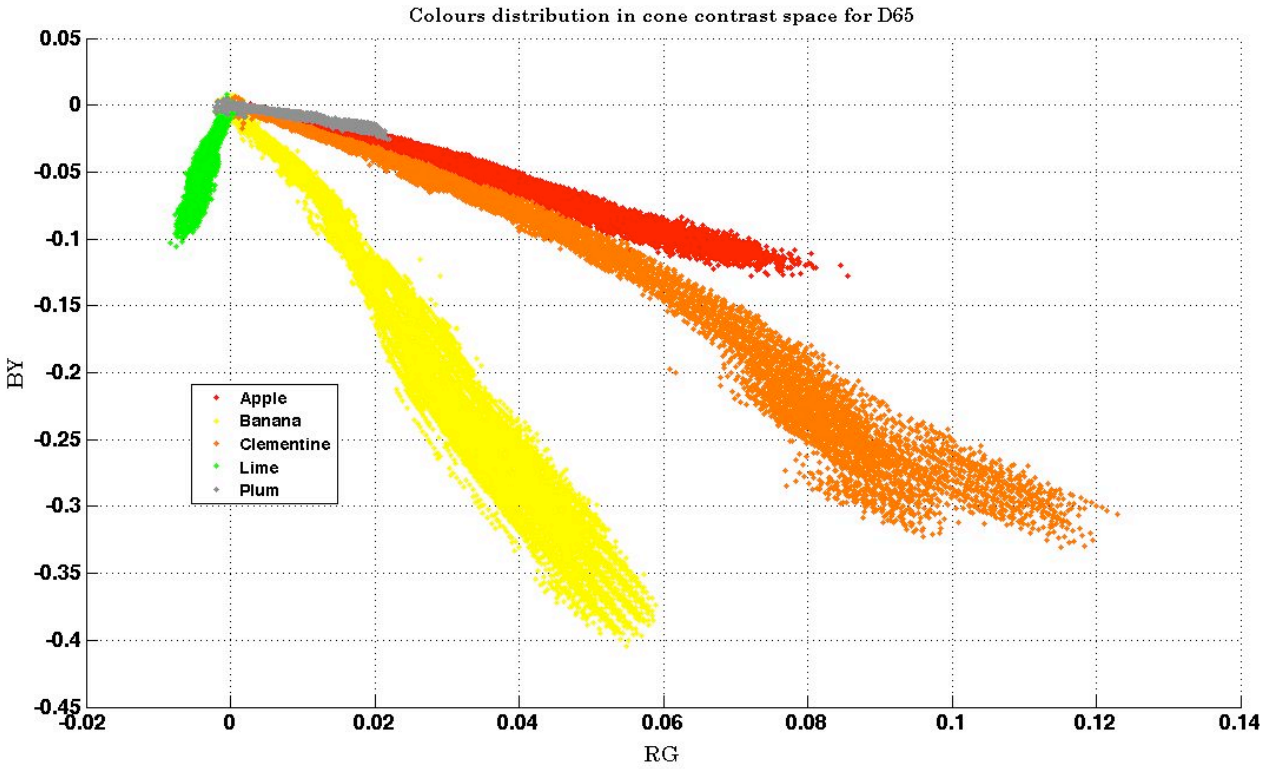

Figure 5.7 – Colour distributions in cone contrast space under D65 for apple (red dots), banana (yellow), clementine (orange), lime (green) and plum (violet). Note that for clarity only points repeated more than 5 times are depicted.

Therefore, let us examine the variation of these distributions under different light conditions evaluating three features of the distribution: (1) the mean; (2) the "hue" angle, as described in the following paragraph; and (3) the spread of the distribution.

5.3.1.1 MEAN AND HUE ANGLE OF THE DISTRIBUTION FOR DIFFERENT OBJECTS AND ILLUMINANTS

Figure 5.8 illustrates the colour distributions of the clementine for D65, CWF and F. The blue star in this figure represents the mean of the distribution, while the blue line represents the "hue" vector of the distribution defined as follows. Let v_{iJ} be the vector that has its origin in the neutral point and end at a point *i* of the clementine's distribution, for a certain illumination J. Lets then calculate the vectorial mean for all the v_{iJ} in the entire distribution under J, obtaining the mean vector V_J. This vector defines the mean "hue" of the distribution and it is illustrated as a black line in Figure 5.8. It is now possible to compute the angle ϑ, or "hue angle", that V_J forms with the RG axis, as follows:

$$\vartheta = \arctan\left(\frac{BY}{RG}\right)$$

Equation 5-1

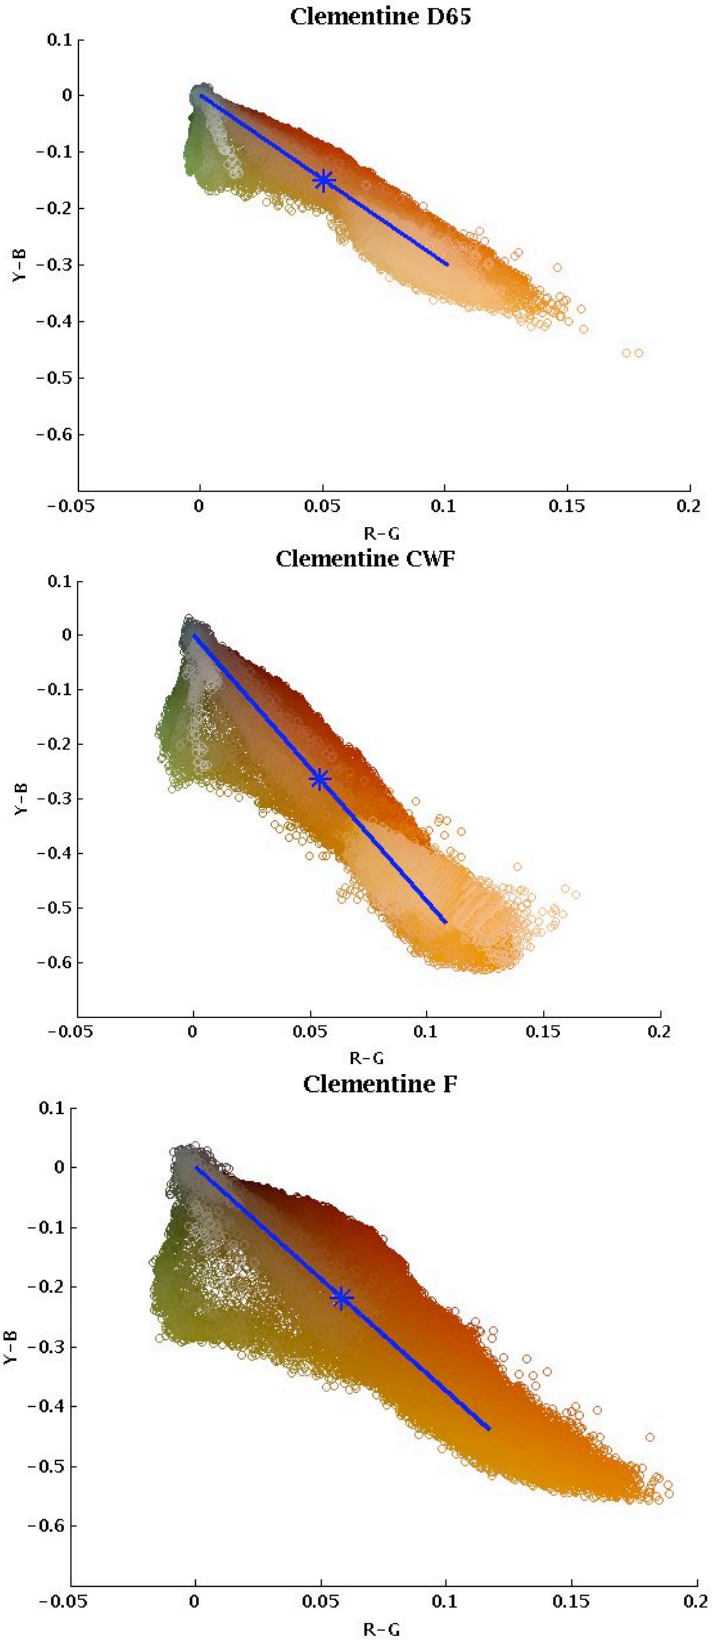

Figure 5.8 – Chromaticity distributions of the clementine under the three illuminants in EMG cone contrast space. A) D65; B) CWF; C) F. Blue line: mean vector of the distribution (hue vector). Blue star: mean of the distributionThe same colour is displayed only once; all recorded colours are displayed.

To illustrate, Figure 5.9 compares on the same graph (using the origin of the coordinates as, simultaneously, the D65 neutral point and the CWF neutral point) the mean and the hue vector for the Melrouge apple under D65 and CWF. Visually the hue vectors of the distributions appear very similar between distributions under different illuminants, while the means seem to vary more. Note

that the hue vectors are expressed in radian. A two-way repeated measures ANOVA was performed on the RG and BY coordinates of the mean of the distribution for all objects under the three illuminants. Since the Mauchly sphericity was violated, the ANOVA was corrected using Greenhouse-Geisser estimates. The results reported as significant the main effect of illuminant on the mean of the distribution within object ($F(1.029,7.02)= 27.074$, $p<0.001$), with a mean difference between the object's distribution mean under different illuminants of 17%, and a significant difference between objects ($F(1,7)=22.16$, $p<0.005$).

Moreover, there is a significant interaction between illuminant and coordinates factor on the mean of the distribution of the object ($F(1.39, 9.77) = 33.12$, $p<0.001$). Therefore, for all objects, there is a significant difference between the means of the object's distributions under different illuminants.

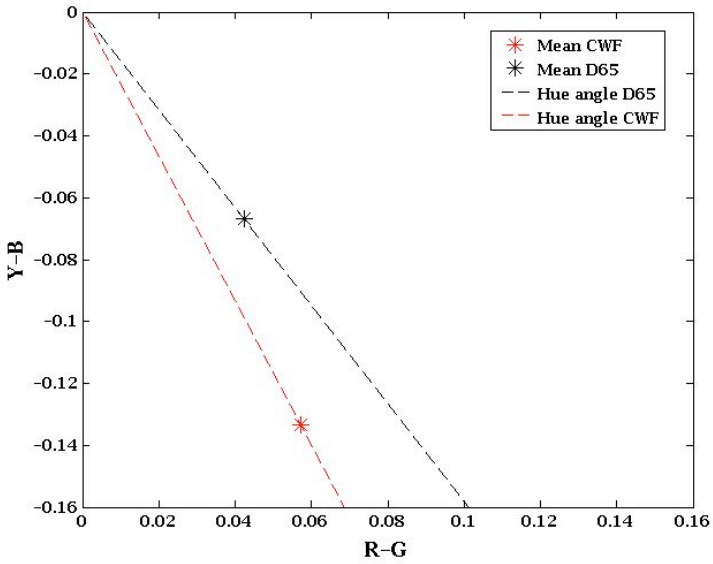

Figure 5.9 – Mean and hue vector of the Melrouge apple distribution under D65 (in black) and CWF (in red). Means: star points. Hue vectors: dashed lines.

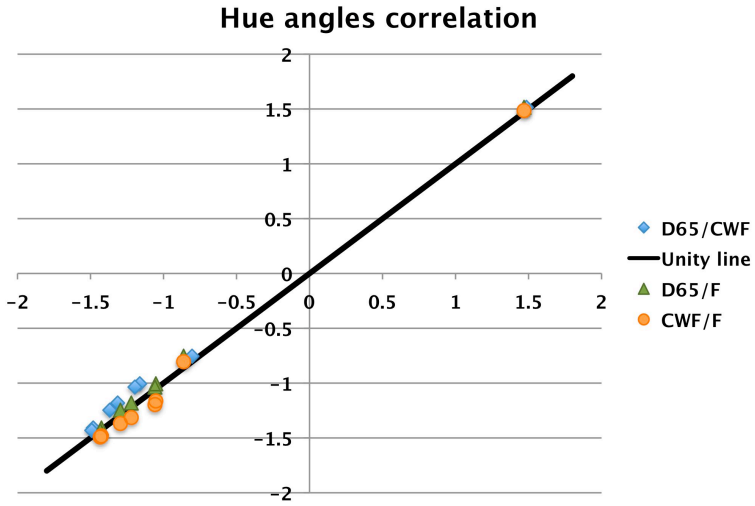

Figure 5.10 – Comparison betweene hue angles of one object under one illuminant and a second illuminant, for all 8 objects. Each point represents one object in one illuminant combination. Blue diamond: D65/CWF. Green triangle: D65/F. Orange close circle: CWF/F. Black line: unity line. The angles are expressed in radian.

Conversely, the hue angles of the object's distributions under two illuminants appear more similar. Figure 5.10 compares the hue angle of one object under one illuminant against the hue angle of the same object under a different illuminant, for all objects and illuminant. Each point represents one object in one illuminant combination. It is possible to observe that all points lie on or about the unity line. The ANOVA reported no significant difference between the hue angles under different illuminations for all objects (corrected with Greenhouse-Geisser estimates, $F(1.065, 7.45)=1.61$, $p=0.235$; mean standard deviation=0.052). Yet, Figure 5.11 shows that the hue angle varies significantly across objects (ANOVA corrected with Greenhouse-Geisser estimates, $F(1, 7)= 18.461$, $p<0.004$; mean standard deviation=0.98). Specifically, all negative hue angles are reported as being significantly different ($p<0.03$). It is worth pointing out that although the mean of the distribution lie on the hue vector, only the hue angle proves to be an invariant parameter of the distribution.

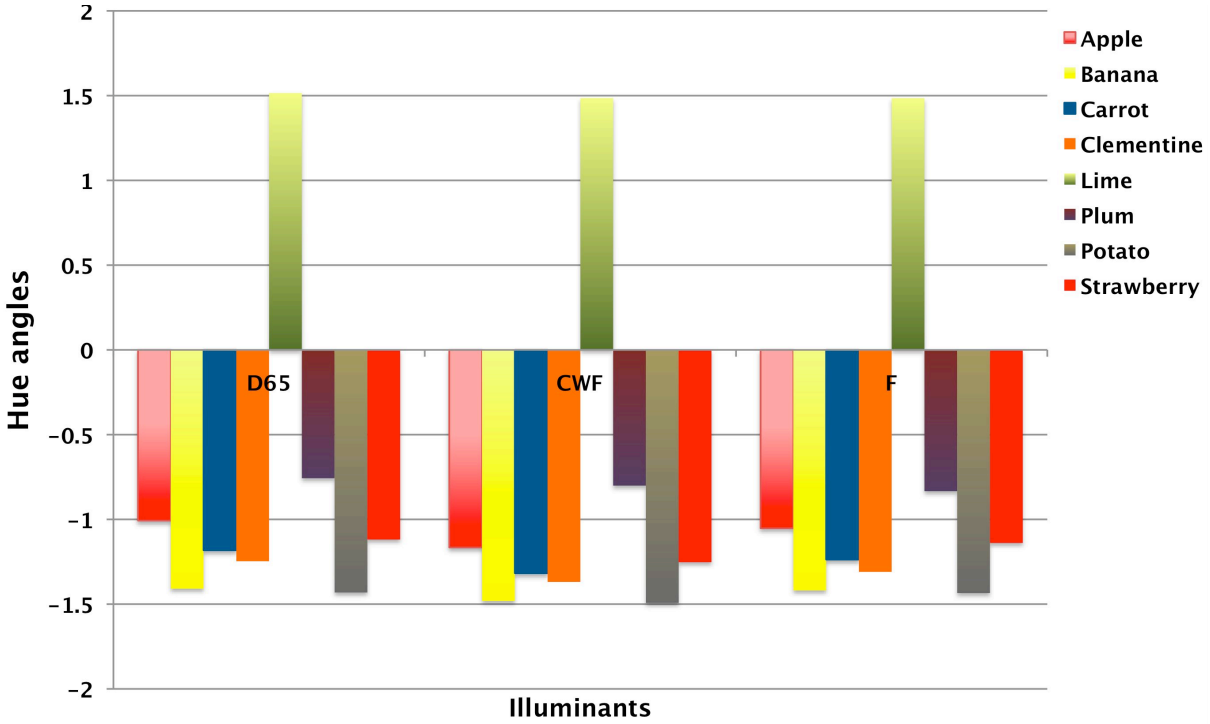

Figure 5.11 – Hue angle difference across the 8 objects. Angles expressed in radians.

5.3.1.2 SPREAD OF THE DISTRIBUTION FOR DIFFERENT OBJECTS AND ILLUMINANTS

A further way to examine the distributions is measure the difference between their spreads. The spread of a distribution will be computed as the area containing 99% of the object's chromaticities (for the purpose of this analysis named the "*gamut*" of the object). Thus, an easy way to look at the variation in the object's colour distribution under distinct illuminants is to compare the change of the perimeter of this area (or gamut). Figure 5.12 shows the strawberry's colour distribution in the EMG isoluminant space and its gamut's contour.

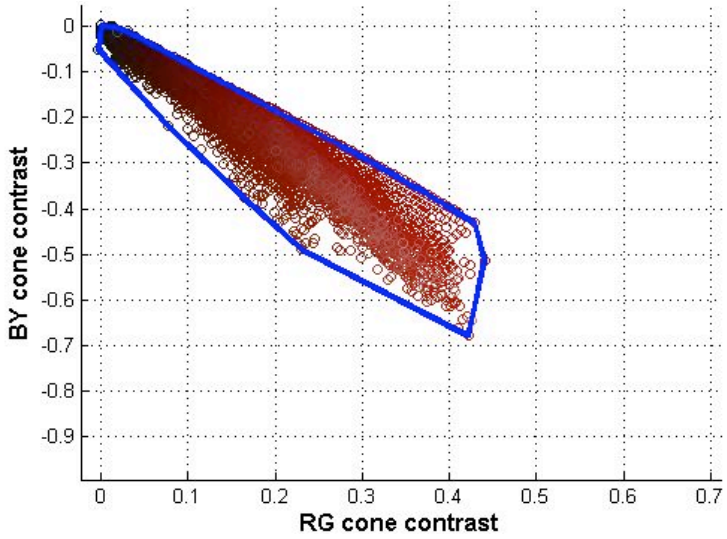

Figure 5.12 - The gamut of the EMG cone contrast space chromaticities of the Strawberry under D65 in EMG cone contrast isoluminant plane. The blue contour inscribes the convex hull that contains 95% of the distribution's points. The origin of the plane [0,0] represents the chromaticity of the illuminant under which the observer adapts.

Figure 5.13 plots the object's gamut contours of the object distribution under the all three illuminations; it is worth pointing out that the origin of the EMG cone contrast isoluminant plane [0, 0] represents the chromaticity of the illuminant under which the subject adapts, thus in this graph the three adaptation points are overlapped. By eye, it is evident that the object's gamut tends to expand or contract depending on the illuminant. In addition, the contours shift to some extent in the chromaticity plane in relation to this change.

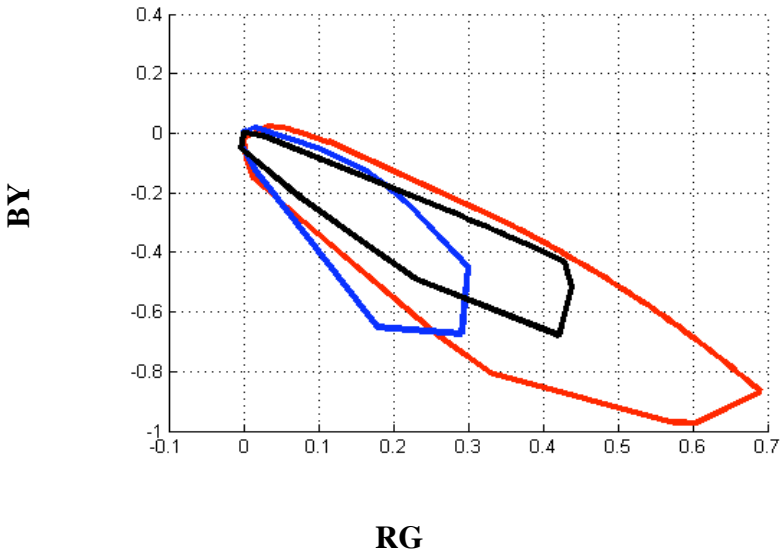

Figure 5.13 - Strawberry: gamuts' shift between lights in EMG cone contrast space. Black line: D65; Blue: CWF; Red: F. Note that the origin of the EMG cone contrast isoluminant plane [0, 0] represents the chromaticity of the illuminant under which the subject adapts, thus in this graph the three adaptation points are overlapped

To quantify this change, i.e. the distribution difference among illuminations, it is necessary to define a parameter, which measures the difference between gamuts' contours. For this purpose, I

modify a technique used in image processing for illuminant classification in a camera, described in Tominaga and Wandell (2002).

A suitable measure of the similarity between distributions with relation to their size can be the correlation between them, and computed as in the following formula:

$$Cor = \frac{A_{12}}{\sqrt{A_1 * A_2}}$$

Equation 5-2

where, in our experiment, *Cor* is the "gamut correlation index", A_1 is the area containing the 99% of the first object chromatic distribution, A_2 is the area covered by the second, and A_{12} is area shared by the two when plotted on the same plane. Unlike Tominaga and Wandell (2002), all areas are calculated not in the camera sensor-dependent RGB space, but the EMG cone contrast isoluminant plane; thus, both normalization of the image and image scaling required by their algorithm to compensate for intensity differences between images, are not necessary. For similar reasons, their image pre-processing was replaced with the one described in Chapter 3.

	D65/F	D65/CWF	CWF/F	Mean
Apple	0.2623	0.6099	0.3521	**0.4081**
Banana	0.4482	0.4468	0.6304	**0.5085**
Carrot	0.3319	0.2048	0.5101	**0.3489**
Clementine	0.2718	0.2671	0.3717	**0.3035**
Lime	0.1957	0.2965	0.2713	**0.2545**
Plum	0.1345	0.7323	0.1837	**0.3502**
Potato	0.224	0.3503	0.4299	**0.3347**
Strawberry	0.2945	0.7161	0.4054	**0.4720**
MEAN	**0.2704**	**0.4529**	**0.3943**	
STDEV	**0.0882**	**0.1946**	**0.1286**	

Table 5.1 – Gamut correlation indices between the distributions of same object under different illuminants, for all objects and illuminant combinations. Column on the far right: means of each object correlations across illuminations.

Figure 5.14 illustrates the gamut correlation indices between object's gamuts under D65 versus CWF or F, for all 8 objects. Table 5.1 lists such correlations for all objects illuminant combinations (D65/CWF, D65/F, and CWF/F) and the means per object and per illuminant combination.

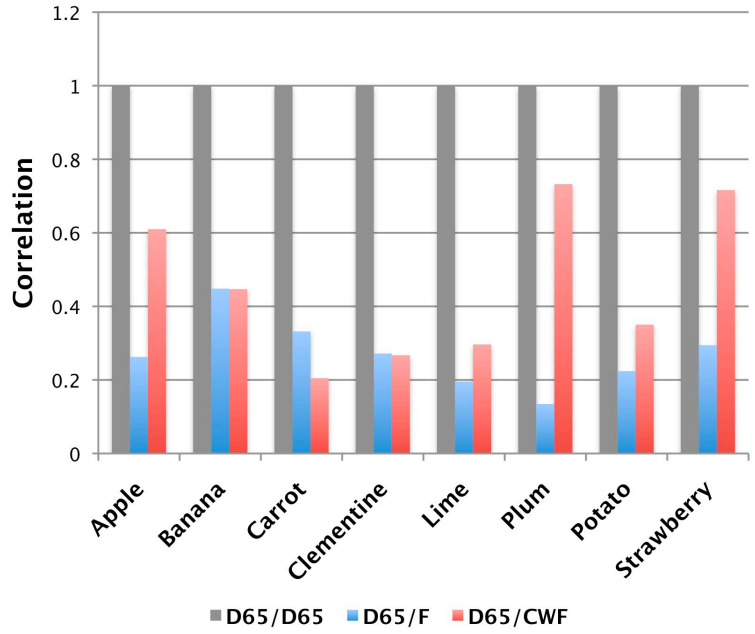

Figure 5.14 – Correlation between object distribution under D65 and D65 (grey bars), CWF (red bars) or F (blue bars).

Note that the same illuminant combination (D65/D65, CWF/CWF, and F/F) have, obviously, gamut correlation indices equal to 1 being the same distributions, and are not reported in Table 5.1.

First interesting finding is that the gamut correlation indices vary significantly between illuminants, i.e. there is a significant difference between the spreads of the distributions of the same object under the three different illuminants ($F(1.25, 8.76)=45.67$[9], $p<0.0001$). Specifically, a significant difference was found between: (1) D65/D65 and any other combination (Bonferroni pairwise comparison: p-values are less than 0.001; e.g. contrast D65/D65 against D65/F, $F(1,7)=478.96$, $p<0.0001$); (2) D65/F and CWF/F (within-object contrast $F(1,7)=37.59$, $p<0.001$); D65/F and D65/CWF (within-object contrast $F(1,27)=5.38$, $p<0.05$). Yet, contrasts revealed no significant difference between D65/CWF and CWF/F ($F(1,27)=0.33$, $p=0.58$). Hence, the gamut of one object belonging to a certain category (an exemplar of the object category) under D65 is weakly correlated with the gamut of the same object for another illuminant. Note that the same comment can be stated for CWF or F. As a consequence of this result, a question that naturally arises is "If we analyse more exemplars of to the ***same*** object category (apples, bananas, cucumbers etc.), what is the correlation between their distributions under the same illuminant or under different illuminants?" The next part of this analysis explores this question, and will be discussed in section 5.3.2.

A second finding is that *different* objects result in different gamut correlation indices with respect to the same illuminant combination (D65/CWF or D65/F). In fact, lets compute the gamut correlation between one object under one illuminant and a different object under the same

[9] Mauchly's sphericity was violated and the ANOVA was corrected using Greenhouse-Geisser estimates

illuminant, and perform the same calculation for all object combinations. Using the carrot as an example, Figure 5.15 shows that the carrot correlates to a very small degree with the other objects except the clementine (*cor*=0.547). On average the correlation between the distribution of different objects of diverse category under the same illuminant ("between-category correlation") is: 0.21 (±0.19) for D65, 0.18 (±0.16) for CWF, and 0.19 (±0.18). Moreover, the correlation between the distribution of different objects of diverse category under different illuminants ("between-category-between-illuminant correlation") is, on average: 0.16 (±0.12) for D65/F, 0.23 (±0.18) for D65/CWF, and 0.20 (±0.17) for CWF/F. Figure 5.16 shows between-category-between-illuminant correlations for the apple. Finally, the correlation between the same object of one category under different illuminants ("within-category-between-illuminant correlation") is on average: 0.27 (±0.09) for D65/F, 0.45 (±0.19) for D65/CWF, and 0.39 (±0.13) for CWF/F. Hence, on average, the difference between the distribution spreads of different objects under the same illuminant is bigger than the difference between the distribution spreads of the same object under different illuminants[10] ($t = -23.849$, $p<0.0001$). However, single objects can have a better correlation with others under the same light than themselves under different lights, as shown in Figure 5.17. Nevertheless, the 75% of the between-category correlation indices are equal or lower than 0.25 (black line in Figure 5.17) while only 25% of the within-object-between-illuminant correlation indices are equal or lower than 0.25. Table 5.2 lists the between-category correlations across all illuminants. Note that within-object-within-illuminant correlations are higher than all these correlations (as they are equal to 1).

Objects	Apple	Banana	Carrot	Clementine	Lime	Plum	Potato	Strawberry
Mean	0.3323	0.2836	0.3041	0.3403	0.2135	0.2527	0.288	0.3369

Table 5.2 – Means of the between-category correlations for all category combinations across all illuminants.

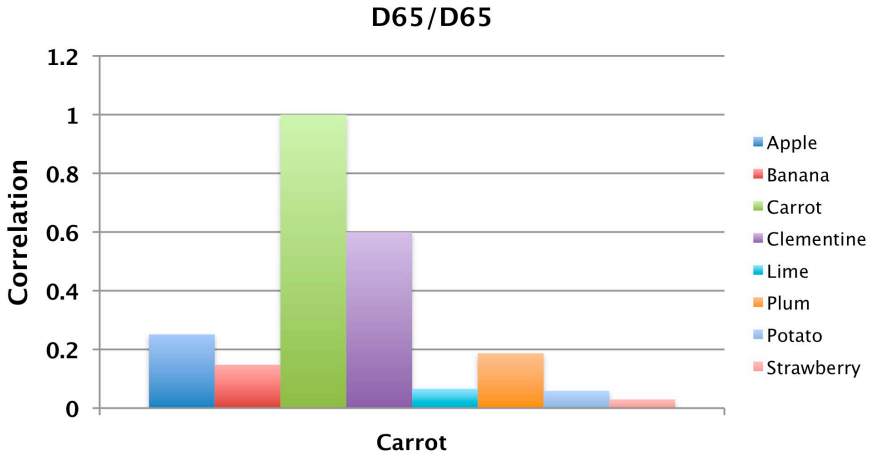

Figure 5.15 – Gamut correlation indices between the carrot and all 8 objects (the carrot itself included) under D65 illuminant (D65/D65). Each bar is the correlation with one object.

[10] One-way ANOVA between groups $F=18.255$, $p<0.001$.

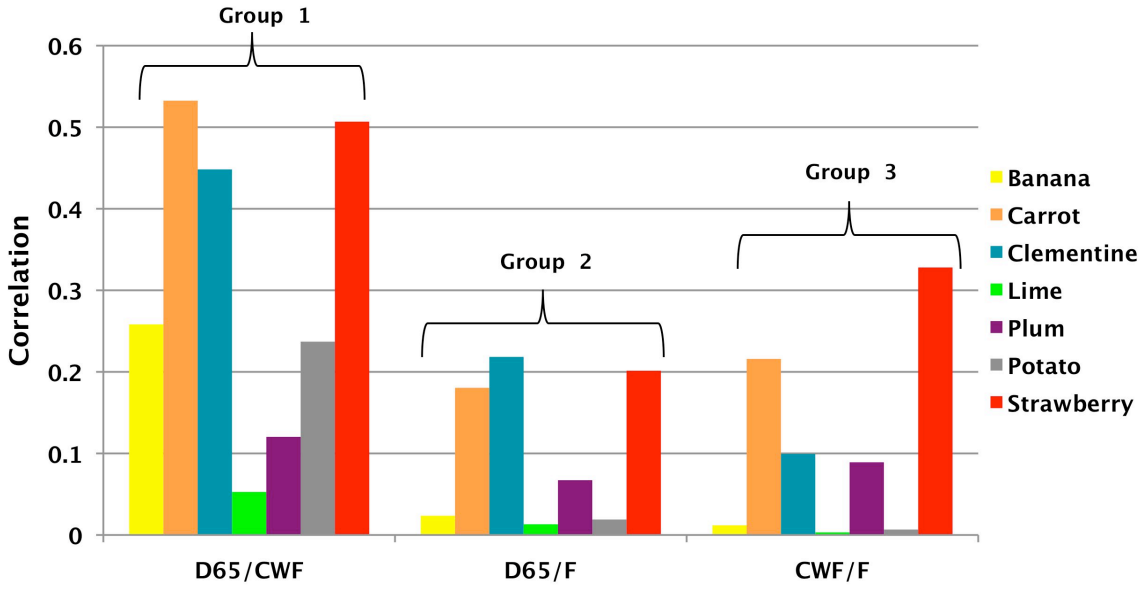

Figure 5.16 - Gamut correlation indices between the apple and all 7 objects (apple itself excluded) under different illuminant. Group 1: carrot under D65 versus other objects under CWF. Group 2: carrot under D65 versus other objects under F. Group 3: carrot under CWF versus other objects under F.

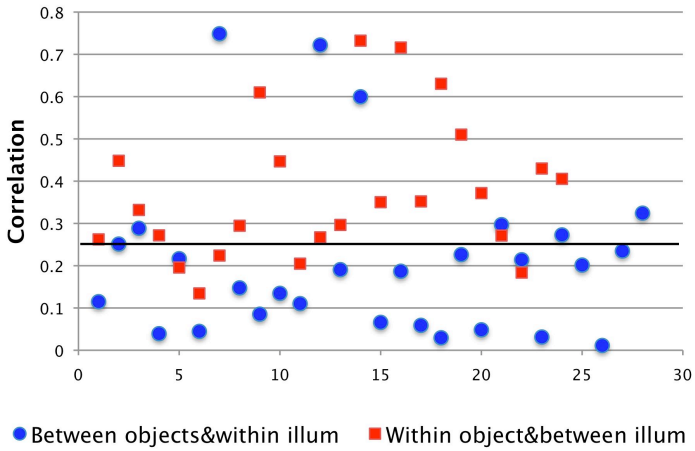

Figure 5.17 – Gamut correlation indices between objects' distributions. Blue close circles: correlation between the distributions of different objects under illuminant D65 (between-category-within-illuminant); each point is one object combination. Red squares: correlation between the distributions of the same object under different illuminants (within-object-between-illuminant). Each point is one object in one illuminant combination; 75% of the blue points and 25% of the red points lie below the black line.

In addition, let us compare the between-category correlations under one illuminant with respect to the ones under a second illuminant for all objects; Figure 5.18 compares D65 against CWF between-category correlations (each point is one object combination). Note that all point cluster around the unity line (black line in the figure). A repeated measures ANOVA was performed on the between-objects correlations under all three illluminats, and reported[11] no significant main effect of illuminant on the between-objects correlation ($F(1.427,9.991)=1.767$, $p=0.219$), whereas the main effect of object was found to be significant ($F(1,7)=348.689$, $p<0.0001$). In other words, the correlation between one object distribution and the other objects is illuminant-independent.

[11] Mauchly's sphericity was violated and the ANOVA was corrected using Greenhouse-Geisser estimates.

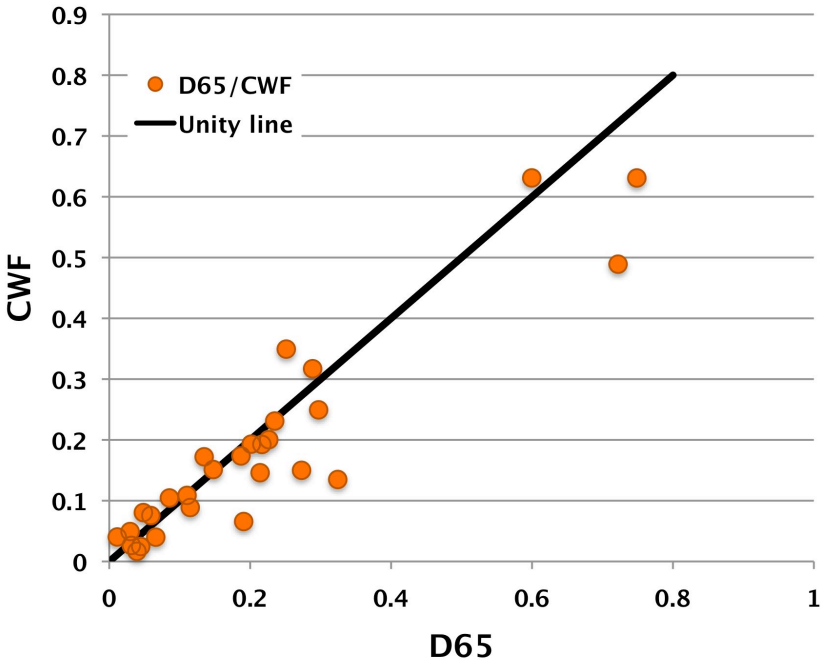

Figure 5.18 – Scatter plot comparing the correlation between any two objects under D65 against the correlation between the same two objects under CWF. Each point is one object combination.

The combination of hue angle and gamut correlation may help to discriminate between objects (and eventually recognize them) through only their surface chromatic variegation.

5.3.2 WITHIN OBJECT CATEGORY EFFECT

As a consequence of the above findings we can ask: "If we have several exemplars of the same object category, how would the features of their distributions vary under the same illumination?" Furthermore, how would they vary with respect to illuminant changes? As in the previous section, I analysed three features of the distribution of each exemplar of the object category: (1) the mean, (2) the "hue" angle, and (3) the spread of the distribution. Three categories of natural objects were examined: apple, banana, and carrot.

A 3-way repeated measures ANOVA was performed on the RG and BY coordinates of the mean of the distribution for all object under the three illuminants. Since the Mauchly's sphericity was violated, the ANOVA was corrected using Greenhouse-Geisser estimates. The results reported as significant the main effect of illuminant on the mean of the distribution within exemplar ($F(1.865, 3.73)= 146.874$, $p<0.001$), as well as between object category ($F(1.28, 2.56)=13.34$, $p<0.05$). Moreover, there is a significant interaction between illuminant and coordinates factor on the mean of the distribution of the exemplar ($F(1.395, 2.79) = 347.126$, $p<0.001$). Therefore, for all exemplars, there is a significant difference between the means of the object's distributions under different illuminants.

Instead, Figure 5.19 compares the hue angle of the three exemplars under the test illuminant D65, for all objects. Each point represents, for each category, the hue angles of one exemplar with

the other two. Note that the hue angles appear not differ between exemplars of the same category (the points practically overlap), while they differ between categories as they cluster distinctly from each other. This result was expected and indicates that objects belonging to the same category have the same hue angles, distinct from other objects. Quantitatively, two-way repeated measures ANOVA reported a significant effect of object category on hue angle ($F(1.048, 2.089)= 359.725$, $p<0.002$). Moreover, contrasts show no significant difference on hue angle for exemplars of the same category between D65 and CWF ($F(1,2)= 2.026$, $p= 0.291$) and D65 and F ($F(1,2)= 4.448$, $p= 0.169$). Hence the hue angles are similar across illuminants.

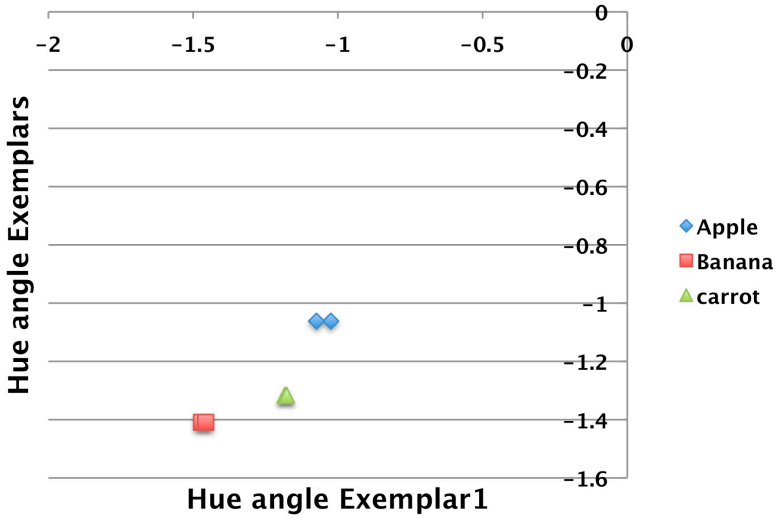

Figure 5.19 – Comparison between the hue angles of the three exemplars under D65, for all objects. Angles in radiants.

On the other hand, I have previously demonstrated that the object's distribution range varies across illuminants, thus it is logical to ask: Does the object distribution also differ across exemplars? For example, lets examine three bananas in three different stages of ripeness and positions (Figure 5.2B). Figure 5.20 shows the convex hulls containing the distributions of the three bananas under D65 in EMG cone contrast space. Visually, note that the gamuts of the bananas under D65 are rather similar, although the unripe bananas have an expanded range in the greener cone-contrasts area. Quantitatively, it is possible to calculate the gamut correlation between different exemplars of the same category under the same illuminant. Figure 5.21 shows the results for the banana (Figure 5.3B). Interestingly the distributions correlate best under the same illuminant. For example, the gamut of 'Banana 1' under D65 correlates best with a different banana also under D65, while the correlation with a different banana under a different illuminant (CWF or F) is lower. In other words, the gamut of ***any object*** belonging to the "banana" category brings consistent information about the specific light illuminating it. Similar results are obtained for the two other object categories: apple (Figure 5.2A) and carrots (Figure 5.2C).

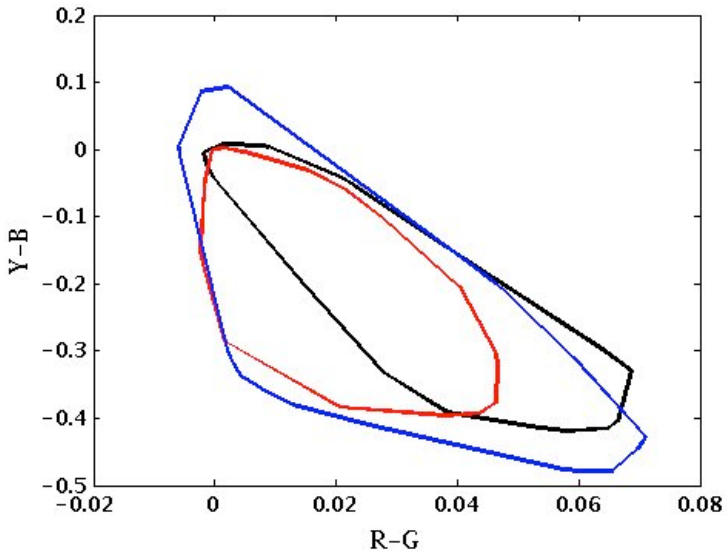

Figure 5.20 – Contours of the gamuts of the three bananas under D65 in EMG cone contrast space. Black close line: Banana 1. Red closed line: Banana 2. Blue closed line: Banana 3.

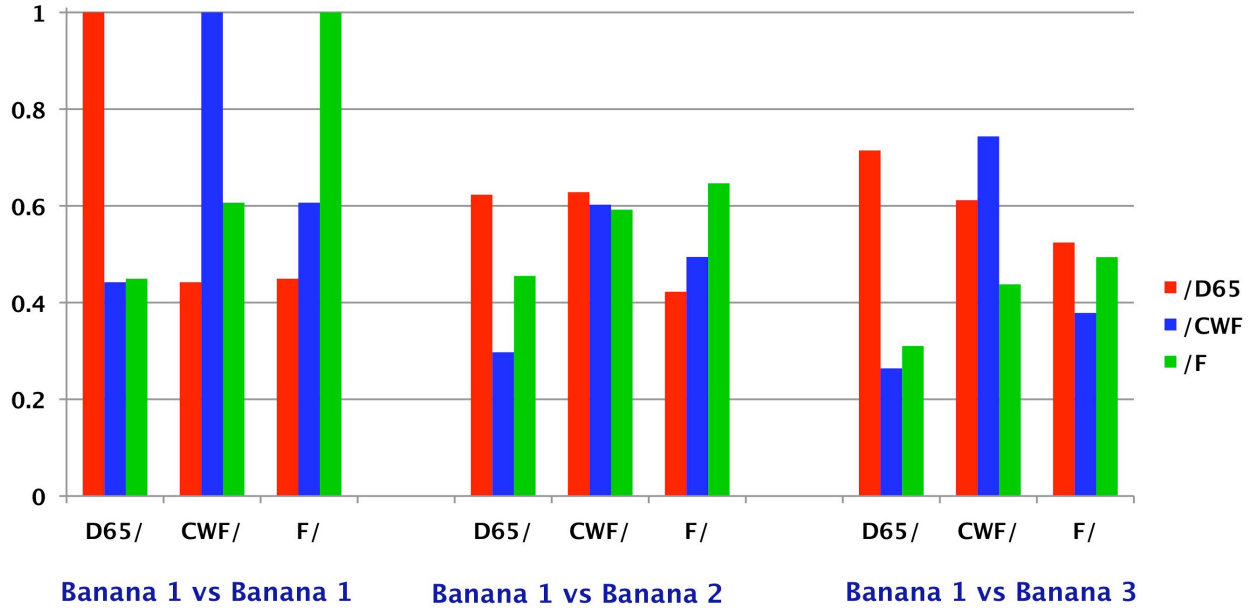

Figure 5.21 – Correlation index between three bananas. Correlation index for bananas under the same lights are always higher than under different lights. Y-axis: Illuminants under which the *first* object is taken. Bars: correlation with the second object taken under D65 (red), CWF (blue), or F (green). Note that obviously the correlation between the 'Banana 1' and itself under the same light is 1.

Figure 5.22 compares the correlation between the gamuts of two exemplars of the same category both illuminated by D65 ("within-category-within-illuminant correlation") and the correlation between the gamuts of the same two exemplars one illuminated by D65 and the second by another illuminant ("within-category-between-illuminant correlation"). For example, in subfigure C) the point indicated by the black arrow represents the correlation between Carrot1 and Carrot2 both under D65 compared to the correlation between Carrot1 under D65 and Carrot2 under CWF. It is important to notice that all points lie above the unity line for all three object category, denoting that objects belonging to the same category correlate best when illuminated by the *same* illuminant. For completeness' sake, Table 5.3 lists the correlations for all three categories of objects in three illuminant combinations not depicted in Figure 5.22. The mean of the correlations between different exemplars of the same category under the same illuminant is 0.606, while the mean of the

correlations between different exemplars of the same category under different illuminants is 0.3071. Hence, exemplars of the same category under the same illuminant correlate ***twice*** as strongly as those under different illuminants, this difference being significant. In fact, a two-way repeated measures ANOVA reported as significant the main effect of illuminant combination ($F(1.759,5.27) = 12.341$, $p<0.005$). In addition, contrasts revealed a significant difference between the D65/D65 and any other combination illX/D65 or D65/illX (where illX is CWF or F), as show in Table 5.4. Furthrmore, there is a significant main effect of object category on gamut correlation ($F(2,6)=5.4$, $p<0.05$), but no significant the interaction between illuminant and object category effect (i.e. the effect of illuminant is independent of the object category, $F(16,48)=1.115$, $p=0.369$).

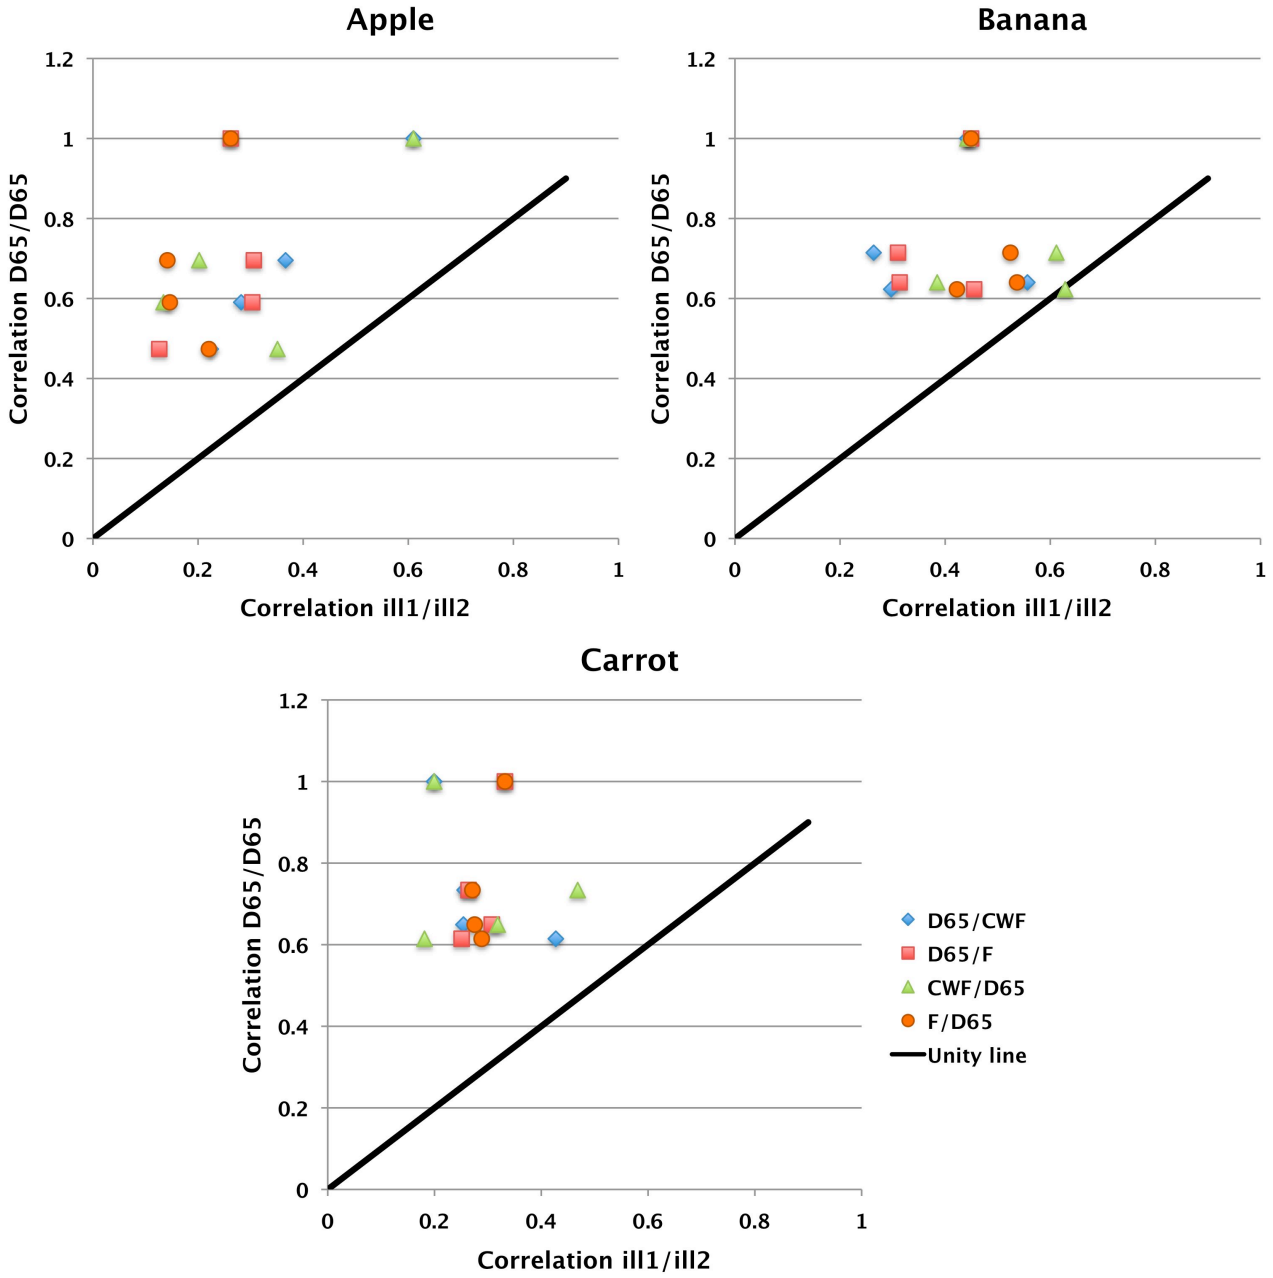

Figure 5.22 – Comparison between correlations of different combination of illuminants. Each subfigure plots one object category; from left to right top to bottom: Apple, Banana, and Carrot. Each dot is one object exemplar pair under D65/D65 compared with the same pair under the combination ill1/ill2 (see label in the figure).

APPLE	D65/D65	CWF/CWF	CWF/F	F/F
1 vs 1	1	1	0.3521	1
1 vs 2	0.8028	0.7679	0.2985	0.659
2 vs 3	0.8368	0.8471	0.3103	0.8574
1 vs 3	0.817	0.8672	0.2557	0.7336
BANANA	D65/D65	CWF/CWF	CWF/F	F/F
1 vs 1	1	1	0.606	1
1 vs 2	0.7088	0.6272	0.5084	0.7329
2 vs 3	0.9509	0.9432	0.566	0.9649
1 vs 3	0.7239	0.6632	0.2701	0.7597
CARROT	D65/D65	CWF/CWF	CWF/F	F/F
1 vs 1	1	1	0.508	1
1 vs 2	0.8663	0.8057	0.3695	0.8752
2 vs 3	0.8358	0.8569	0.3037	0.7556
1 vs 3	0.85	0.7797	0.3047	0.8156

Table 5.3 - Correlation indices between objects of the same category. Apples (Figure 5.2A). Bananas (Figure 5.2B). Carrots (Figure 5.2C). Each row contains the correlation of the object J versus object K under the same or different lilluminant. Each column represents one combination of light (for example CWF/D65 is the correlation of one object under CWF versus another object of the same category under D65).

	F value	P-value
D65/CWF	23.231	0.017
D65/F	29.757	0.012
CWF/D65	13.574	0.035
F/D65	24.151	0.016

Table 5.4 – Contrasts p-values for the comparison D65/D65 combination against illX/D65 or D65/illX combinations

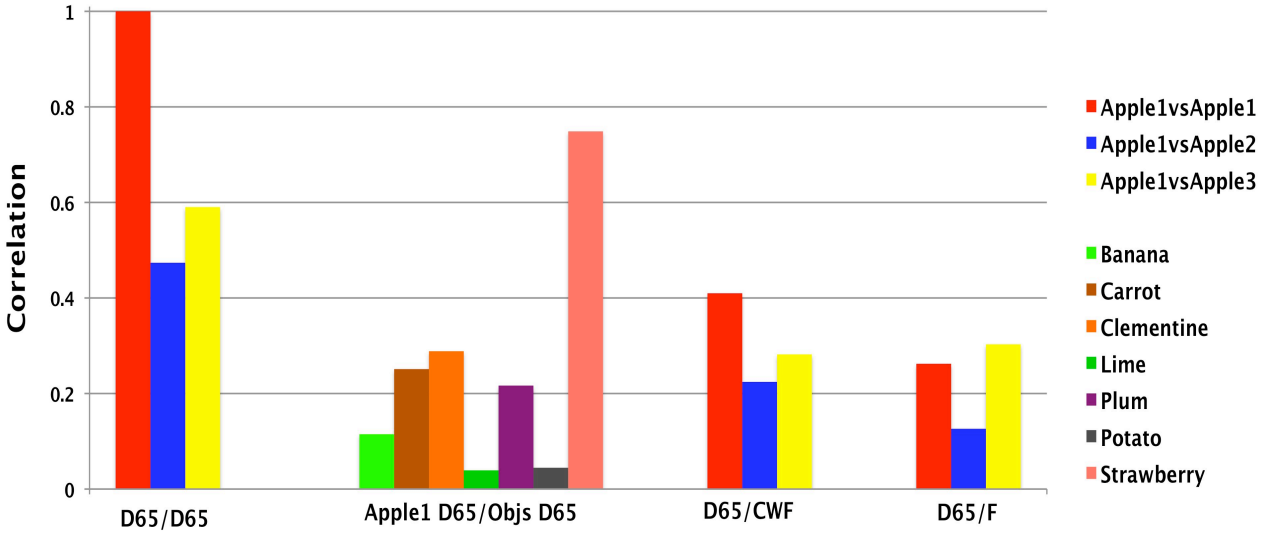

Figure 5.23 – Correlation comparison between Apple1 and other two apples under the same illuminant or different illuminant. In the same figure comparison between Apple1 and all the other objects.

5.3.3 ADAPTATION

The above analyses effectively assume that the visual system is *fully* adapted to the illuminant chromaticity (as said the origin of the EMG cone contrast space is the illuminant chromaticity itself). Hence under imperfect adaptation these analyses would not hold. Besides, would the adaptation to another point influence these findings? For example, let us hypothesize that the visual system has adapted to the object mean chromaticity under that light and not to the illuminant chromaticity. Selected a set of three exemplars of one category (for example the carrot) lets compute the correlation between them under the same illuminant and under different illuminant. If the results are similar as the one presented in the previous section then we can conclude that this propriety of the object distribution (i.e. the gamut correlation) is independent from the origin of the EMG colour space (i.e. the adaptation point). Note that this does not signify that it is independent from the illuminant as this information is contained in the mean of the distribution (as seen in section 5.3.1.1 the mean of the distribution significantly changes with the illuminant colour). Figure 5.24 compares the correlation between different exemplar of the apple, the banana and the carrot category calculated using the object mean chromaticity under the specific illuminant as adaptation point. Results show a similar trend as in section 5.3.2 (see Figure 5.22), although constrast reported no significant the comparison D65/D65 against illX/D65 or D65/illX (where illX is CWF or F; for all combination $p>0.1$). Therefore the adaptation to the illuminant is necessary.

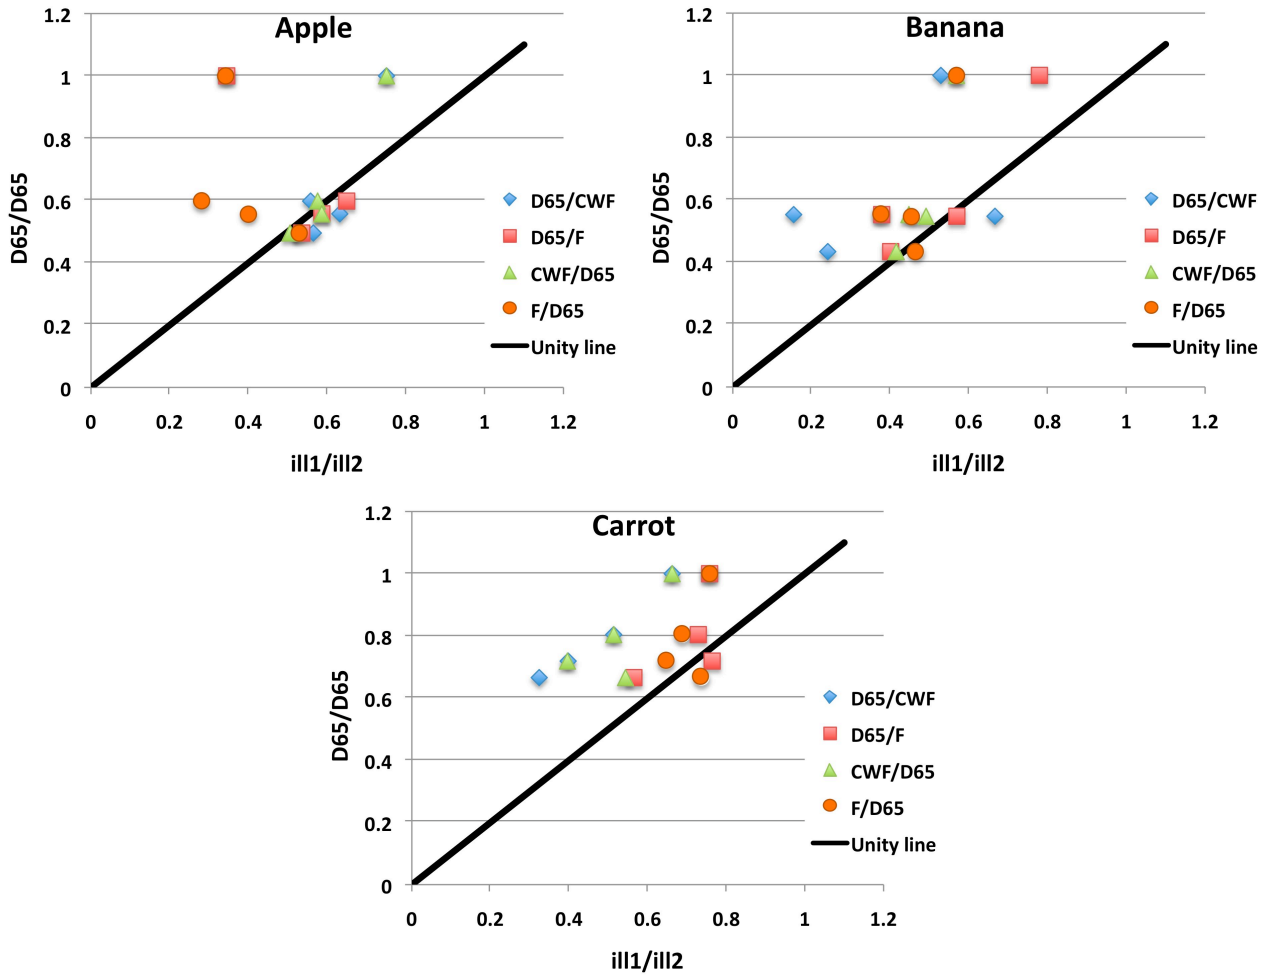

Figure 5.24 – Comparison between correlations of different combination of illuminants considering the adaptation to the object mean. Each subfigure plots one object category. From left to right, from top to bottom: Apple, Banana, and Carrot. Each dot is one object exemplar pair under D65/D65 compared with the same pair under the combination ill1/ill2 (see label in the figure).

The difference between correlations of the distributions plotted relative to different adaptation points is significant ($p<0.00001$ for between-category correlations and $p<0.00001$ for within-object-between-illuminant correlations).

5.4 DISCUSSION

As previously mentioned, to achieve colour constancy the observer has to solve a computationally underdetermined inverse problem. Therefore, the visual system requires constraints on the illuminant spectral power distributions and the surface reflectance functions to develop appropriate estimations of the actual colours in the scene. Here the results demonstrate that this problem might be easily solved in presence of "memory colour".

This chapter demonstrates that natural objects possess physical features that, if known, can be used as anchors, i.e. constrains, when comparing its colour under different illuminations.

5.4.1 THE OBJECT'S SIGNATURE, MEAN AND HUE ANGLE

We have seen that the distribution of each object belonging to diverse fruit or vegetable category clusters in a specific area of the EMG cone contrast space (or signature of the object category under a certain illuminant). Moreover, due to the elongated form of this signature, it is possible to describe the dominant hue of the distribution with a vector. Results showed that the direction of this vector is stable under changes in illumination and across objects belonging to the same category, while it varies between objects belonging to the different category. On the other hand, the mean of the distribution in this colour space is significantly different under diverse illuminants, and varies significantly between object category and object exemplars. These interesting results can be interpreted as follows: *the knowledge of the hue angle of* ***any*** *exemplar of an object's category under* ***any*** *illuminant may provide a constraint of the object colour distribution when we perform a colour constancy task*. Conversely, the mean of the distribution, although correlated to the object category, varies across illumination. *Hence, the analysis of the hue angle in a colour constancy task could be the most appropriate parameter for the evaluation of an observer's performance*.

5.4.2 BETWEEN-CATEGORY EFFECT

Results showed that objects of different categories could be distinguished by their chromatic distribution through their mean, hue angle and gamut correlation. In fact, each parameter of the distribution is significantly affected by the object category. In other words, we can use the combination of these features of the distribution to distinguish one object from the other based *only* on its surface chromatic appearance and eventually identifying it in *absence* of shape. Furthermore, since the correlation between the distribution of an exemplar of an object category and the distribution of an exemplar of other object categories is *illuminant-independent*, then we can distinguish between different categories under *any* light condition and in any change in illumination. If a combination of all the parameter could identify the object based only on its chromatic appearance then we can pose the question of which of this factor or factor pair could be the most informative. Further studies on this direction are needed, possibly using methods such as the multiple regressions.

5.4.3 THE WITHIN-CATEGORY CORRELATION

One of the most *intriguing* results of this experiment is *the demonstration that the chromatic distribution of **any** sample of an object category correlates best with a **different** sample of the same category when presented under the same illumintant*. Specifically, exemplars of the same category under the same illuminant correlate twice more than under different illuminant. This correlation weakens in dependence of the difference between the illuminants under which the objects are presented. Hence, if we know the gamut of any carrot under a reference light, it can be compared with any other carrot placed under an unknown light condition. Therefore, we can establish if the latter illuminant is the same as the reference or different. Furthermore, if we know the object gamut's correlation between the reference and the test illuminant, we can determine the latter. However, this presumes that the memory bank holds 'exemplars' of objects under M different illuminants (M being a natural finite number) and does the correlation with each one. Hence we do not possess a continuum space in which we choose the illuminant chromaticity. In spite of that, we can obtain a more precise estimate of the illumination using the following process.

Lets suppose that an observer has memorized the distributions of ***one*** familiar object O_1, belonging to the category A, under N different illuminants (N_1, N_2, …, N_N), and calculated the Q correlation indexes[12], C_{NiNj}. Then, let C_{HNj} be the correlation between the distribution D_H of ***another*** object O_2 belonging to the category A under an ***unknown*** illuminant H and the distribution D_{Nj} of the object O_1 under a ***known*** illuminant N_j. Thus we can search for the N_j such that C_{Hkj} is the greatest. Let be $K=N_K$ such illuminant. Therefore, we can compare C_{HK} with the correlation indexes C_{KNj}, for all j=1,..,N, and find the C_{KNj}th with the smallest difference. Finally, if we hypothesize the linearity of the correlation as a function of the illuminant (Tominaga and Wandell (2002) proved that correlation steps are approximately constant in mired steps, although not colour-temperature steps), we can speculate that the unknown illuminant is somewhere in between N_j and N_i. Using this technique we can obtain an estimate of the light illuminating the scene possibly sufficient to perform effectively a colour constancy task. The unsolved question is then: "Is this an estimate of the illuminant or a function that can be incorporated in a Bayesian model of human colour constancy?" (Forsyth (1990); Barnard, Cardei, Funt; Tominaga and Wandell (2002), Brainard et al. (2006), Gijsenij&Gevers (2007)).

One final interesting point is that, if we observe a scene with 2 (or more) objects of which the observer knows the memory colours, or better, the memory colours' distributions, then the ratio between the C_{Hkj} and C_{Hki} of the first object would be different from the ratio between the C'_{Hkj} and

[12] Note that there are $\binom{K}{2}$ combinations of light and thus Q=$\binom{K}{2}$ correlations

C'_{Hki} of the second object. The comparison between these ratios could reinforce the estimation of the illuminant.

5.5 CONCLUSIONS

In conclusion, I have presented two features of the chromatic surface of familiar object categories that encompass information helpful in an object recognition task and colour constancy task. Firstly I have demonstrated that the hue angle of the distribution (represented in EMG cone contrast space) of any exemplar of an object category is more stable than the mean of its distribution. Provided with these results I have postulated that the constancy of the hue angle under different light suggests that it can be an element in a colour constancy task.

Secondly I have demonstrated that different exemplars of the same category correlate best when illuminated by the same illuminant. Therefore I have suggested that we may possess, for one exemplar of each familiar object's category, a set of illuminant-dependent chromatic distribution and use them to find the best-matching illuminant by comparing the observed gamut with stored (memorized) distributions. Employing this method we could estimate the colour of the illuminant when performing colour constancy task.

Hence this chapter has highlighted two factors in the study of colour appearance and specifically colour constancy: (1) the hue angle is an appropriate variable to evaluate in assessing the performance in a colour constancy task (i.e. the more the subject chooses similar angles under different illuminations, the more s/he is colour constant); (2) in presenting a chromatic surface under two different illuminants a simple *translation* of the entire distribution gamut to a different adaptation point is insufficient (e.g. Hansen et al. (2008)), as *the distribution gamut changes form and this is correlated with the illuminant colour*. In view of these findings, a series of algorithms were designed to simulate the change of illumination and the change of hue angle of natural object. These were applied to test the effect of shape and chromatic proprieties of the stimulus surface in a set of experiments (memory colour, constancy and simultaneous and delayed discrimination experiments).

Chapter 6

PERCEIVED NATURALNESS OF SURFACE FAMILIAR TEXTURE OF 3D OBJECTS

6.1 INTRODUCTION

Several studies have shown that colour may be used to identify objects (Brainard, 2004; Oliva & Schyns, 2000; Tanaka, Weiskopf, & Williams, 2001) and enhance performance across a broad range of tasks. These include perception of shape from shading (Kingdom, 2003), scene recognition (Gegenfurtner & Rieger, 2000; Wichman, Sharpe, & Gegenfurtner, 2002), and object recognition (Tanaka & Presnell, 1999). On the other hand, observers' expectations of an object's prototypical colour influence the perception of the scene in viewing. For example, an ambiguous hue, such as a yellow-orange, is more likely to be categorized as yellow on a banana than on a orange. This categorization bias also affects the perception of other colour-generic objects, e.g. socks, cars or crayons (Mitterer & de Ruiter, 2008). Furthermore, Hansen et al. (2006) demonstrate an effect of memory on colour appearance. I have previously discussed the main points of this study, thus I will here only briefly summarized them. When observers were asked to adjust the colour of photographs of natural fruits or vegetables until they appear achromatic, they actually adjust the photograph colour to a setting shifted away from the observers' actual neutral point. Specifically, the adjustment was in the direction opposite to the typical colour of the object (Hansen et al 2006). Hence, the knowledge of the characteristic colour of a familiar object prompts a bias on the perception of its actual colour. It remains an open question as to how, to what degree and exactly which factors influence the access to this colour knowledge or memory colour.

As shown by the main experiment in Chapter 4, two factors seem be involved in a simple model of colour memory: the object shape recognition and the chromatic distribution of the original object texture. In the same chapter I have proposed that linguistic instructions allow the observers to access a small portion of the actual memory of the object's colours, and that the introduction of identity cues (e.g. shape) may help the observers to *better* access their visual memory. Several studies support this statement, showing indeed that colour is not processed in isolation from other types of visual information. In fact, colour appearance is influenced by form (review, Gegenfurtner, 2003), spatial frequency and orientation (Johnson, Hawken, & Shapley, 2001, 2004), and several other features (e.g. colour contrast, Wolf&Hurlbert, 2004), interacting at different levels of the visual pathway from even the earliest processing stages for form to higher cortical areas for more complex features (Edwards, Xiao, Keysers, Foldiak, & Perrett, 2003). Furthermore, scene geometry and three-dimensional shape strongly influence colour appearance (Bloj, Kersten, & Hurlbert, 1999, Ling&Hurbert, 2004). Recently, Olkkonen and colleagues (2008) demonstrated the effect of object texture and, to some degree, the effect of shape on memory colour. Their findings show that congruent object-shape configurations have stronger colour memory than uniform disc shapes. However, this work evaluated only the *mean* memory colour difference between conditions, as

opposed to the difference in the range of the memory colour. Secondly, they presented flat images combining information about the texture and 3D shading of the object. Thirdly, no direct comparison between shapes was performed. In fact, the identity of the texture was not maintained in the different test shapes (disc, 2D profile, 3D object rendering). For such reasons the effective weight of the object's cues (chromatic texture, shape and three-dimensionality) cannot be evaluated.

On the other hand, in a world in continuous light change, this object-dependent bias can stabilize perception in a way that observers might achieve better colour constancy. As a matter of fact, observers need to compensate for the influence of illumination to see the object's colour as invariant over different illumination conditions. Bottom-up information, such as global and local means of the light (Kraft & Brainard, 1999) and brightness-hue correlations in the image (Golz & MacLeod, 2002), can help in such compensation. However, top-down or other mechanism may also play a role (Beeckmans, 2009). In fact, when faced with ambiguity, the observer looking at a banana will be correct more often than wrong when assuming that it is actually yellow. Based on these assumptions, the more stable and accurate the colour memory of a specific object is, the better its colour constancy should be.

In light of the above considerations, and as a logical continuation of the experiment in Chapter 4 and Chapter 5, the experiment presented in this chapter aims to answer the following questions: "How, and to what degree, do the variegation of the object's chromaticity, and 2D and 3D cues to shape affect colour appearance? Are these effects independent of illumination colour? Do these factors influence colour memory and colour constancy equally?"

Similarly to the paradigm described by Olkkonen et al. (2008) this experiment's task measures observers' colour perception by allowing them to freely adjust the stimulus colour. Note that this study does not require any verbal report during the experiment. The observer can freely select a colour setting that s/he considers typical for a familiar object. Homogeneous and chromatically variegated (texture) surfaces were tested together with four different object-shape representations: 2D generic shape (rhomboid), 2D silhouette of the object, 3D generic shape (3D rhomboid), and object-congruent 3D shape. Twenty-eight participants took part in the experiment. Significant differences were found between conditions and the sizes of these main effects were evaluated for each condition. Results show that the presence of the object's contour influenced memory colour the most, while 3D cues the least. Chromatic conditions had an overall consistent but lower effect relative to the object's contour effect size.

In the second part of the experiment we evaluated the effect of change in illumination on memory colour and calculated the constancy of the perceived object colour in relation to the variegation of the chromaticity of the stimulus, and the 2D and 3D cues to shape. Comparing the results of the two lights, they showed the same pattern. However, while there was a significant

difference between the magnitudes of the effects due to shape cues under the two illuminations, the size of the effect due to texture variegation was the same. Lastly, colour constancy indices were calculated for all experimental factors (chromatic and shape cues). Results showed an effect of the variegation of the chromaticity of the stimulus (texture), but not of shape.

6.2 METHODS

The following experiment was carried out using the set up described in section 2.3; using this apparatus I can digitally manipulate the surface appearance of a solid object in real time and point-by-point. The illumination of the chamber is produced by two sets of three fluorescent lights (i.e. the side lights or lighting system, see section 2.3.1.1). The projector image (see section 3.7.1) was set at 'Standard', with contrast 0 and brightness at -20 or 0 when simulating, respectively, illuminant D65 or CWF (see section 6.2.2). All software packages were written in Matlab 7.5 (The Mathworks Inc.).

The stimuli consisted of two alternative chromatic configurations: (1) chromatically textured objects or (2) uniformly coloured objects, each on a uniform background whose colour defined the adaptation point. Each chromatic configuration could be presented in four possible shape configurations: (1) 2D generic shape (rhomboid); (2) 2D natural silhouette of the object; (3) 3D generic shape (3D rhomboid); or (4) 3D natural shape of the object. Therefore, I evaluated three factors of the stimulus per fruit or vegetable object: (1) variegation of its surface chromaticities (texture/uniform; called chromatic factor), (2) presence/absence of object contour shape (silhouette/rhomboid; called contour factor), and (3) 3D cues to the object shape (2D/3D; called dimensionality factor). Three natural, familiar fruit and vegetable objects, bought at an ordinary supermarket, were employed for this experiment, namely: (1) a Gala apple, (2) a ripe banana, and (3) a carrot. The experiment was carried out using two simulated illuminations (see below). Table 6.1 lists all the combinations of the elements (or conditions) of the factors of this experiment and illuminations for one object (for example carrot). In the following, the term "shape factors" will refer to contour factor and dimensionality factor together. Figure 6.1 illustrates the comparison we can make between elements of different factors: texture versus uniform, object natural contour versus rhomboid generic contour, and 2D versus 3D configuration. The first comparison tests the effect of the representation of the chromatic properties of the surface on colour appearance (first factor). The second comparison tests the effect of shape diagnosticity on colour appearance (second factor), while the third comparison evaluates the effect of shape dimensionality on colour appearance (third factor).

Objects	*Illuminations*	*Chromatic factor*	*Contour factor*	*Dimensionality factor*
Carrot	D65	**Texture condition**	Natural condition	2D (silhouette)
				3D
			Generic condition (rhomboid)	2D (silhouette)
				3D
		Uniform condition (mean, most saturated)	Natural condition	2D (silhouette)
				3D
			Generic condition (rhomboid)	2D (silhouette)
				3D
	CWF	**Texture condition**	Natural condition	2D (silhouette)
				3D
			Generic condition (rhomboid)	2D (silhouette)
				3D
		Uniform condition (mean, most saturated)	Natural condition	2D (silhouette)
				3D
			Generic condition (rhomboid)	2D (silhouette)
				3D

Table 6.1 – Example of the combinations of the elements (or conditions) of the all factors and the two illuminations used in this experiment. Object: carrot.

The observers were asked to look inside the experimental chamber through the viewing hole and perform a colour adjustment task. After an object was named, the subjects had to adjust, as quickly as possible, the colour appearance of such object to its "natural" colour appearance using a gamepad (section 6.2.5). Two background illuminations were simulated: (1) standard daylight at 6,500 K correlated colour temperature (D65), chosen as reference illuminant, and (2) Cool Flourescent White light (CWF), selected as test. During the experiment the background illumination and the chamber illumination (i.e. the light reflecting from the chamber walls) had the same chromaticity.

The three objects were photographed using the method explained in section 3.3 and then processed as described in Chapter 3 (Figure 3.1) to obtain the basic object chromatic texture for the experiment. The objects' colours distributions and proprieties are analysed in section 6.2.3.1 under both lights, while section 6.2.3.1.1 explains in detail the texture manipulation technique and the adjustment software. Shape configurations characteristics are described in 6.2.4.

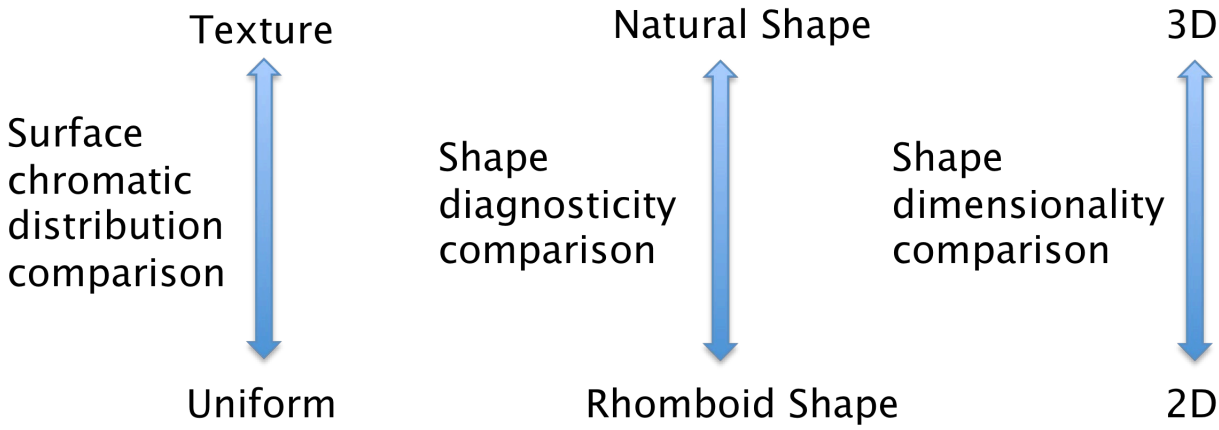

Figure 6.1 - Comparison between elements (conditions) of different factors: texture versus uniform, object natural contour versus rhomboid generic contour, and 2D versus 3D configuration

6.2.1 OBSERVERS

Twenty-eight observers, 20 female and 8 male, all in the age range of 19-28 years and all students at Newcastle University, took part in all conditions of the experiment. They all tested normal on the Farnsworth-Munsell 100 hue test (mean of the total error score 23.9 (Kinnear&Sahaire (2002))) and were naïve to the purpose of the experiment. An ethical agreement was obtained before performing the experiment. Before starting the main experiment, they acquired familiarity with the task during a 3 minutes practice session.

6.2.2 ILLUMINANTS

In the experiment the background colour simulated the CIE standard D65 and CWF illumination recorded in the Verivive cabinet in which the original photographs of the objects were taken (see Chapter 3). Similarly to Chapter 4 (see section 4.2.1), the spectrum power distributions of the backgrounds are the best approximations to the original lights that the projector can generate. In other words, metamers of these illuminants were employed in the experiment. The D65 and CWF spectra measured in the set-up by a PR-650 spectrascan are shown in Figure 4.1. Keep in mind that the intensities of the side lights were adjusted so that the chromaticity of the light reflected from the chamber walls and the background chromaticity were the same. Therefore, I will use the term illumination to indicate in general the simulated background illumination or the chamber walls illumination. Both CIE Yxy values are shown in Table 4.1.

	Background			Walls		
Illuminants	Y(cd/m	X	Y	Y(cd/m	X	Y

D65 (neutral)	7.061	0.3136	0.3273	3.991	0.314	0.332
CWF	6.887	0.2511	0.2561	4.033	0.261	0.249

Table 6.2 - CIE Yxy chromaticity values for the two illuminations. On the right: chromaticity coordinates of the background. On the right: mean chromaticity coordinates of the light reflected by the chamber walls.

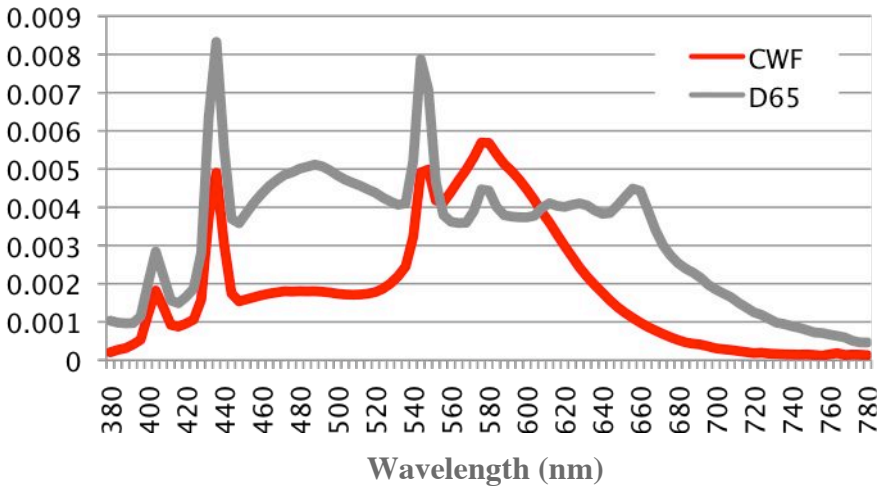

Figure 6.2 – Spectral power distribution of D65 (grey line) and CWF (red line) as measured in the set up.

6.2.3 CHROMATIC CONDITIONS

During the experiment I varied the representation of the chromatic proprieties of the surface of the displayed object between: (1) full chromatic texture (texture or TEX condition), (2) mean colour only (MM condition), and (3) most saturated colour of the object (MS condition). The MM and MS conditions will be also grouped under the name of "uniform conditions", as they represent cases in which the colour stimuli are uniformly coloured surfaces. The surface colours were simulated under the two different lights (i.e. the background illumination). Contrary to studies by Olkkonen et al. (2008), the effect of the illuminant change was not simulated as a rigid shift in the direction and amount of the light change, but was obtained through the procedure described in the next sections.

6.2.3.1 TEXTURED OBJECTS

Photographs of three familiar objects, a Gala apple, a ripe banana and a carrot, commonly found in any supermarket, were recorded under standard daylight illuminant at 6,500 K correlated colour temperature (D65), chosen as reference, and Cool Fluorescent White illuminant (CWF), selected as test light. The images were subsequently converted to CIE XYZ colour space using the method described in Chapter 3.

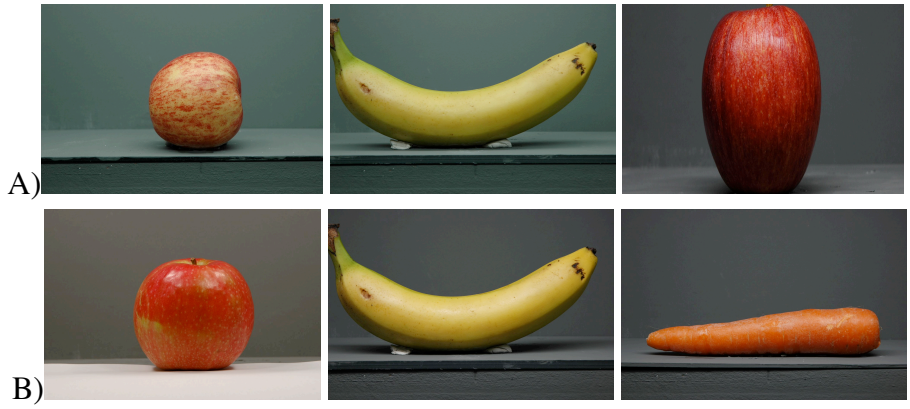

Figure 6.3 – Thumbnails of the three objects' photographs. A) D65; B) CWF.

The six photographs, three for each light, are shown in Figure 6.3. Figure 6.4 illustrates the chromaticity distributions for the three objects under both illuminants in the CIE chromaticity plane (u', v') obtained after using the camera model transformation, while Figure 6.5 plots them in EMG cone contrast space.

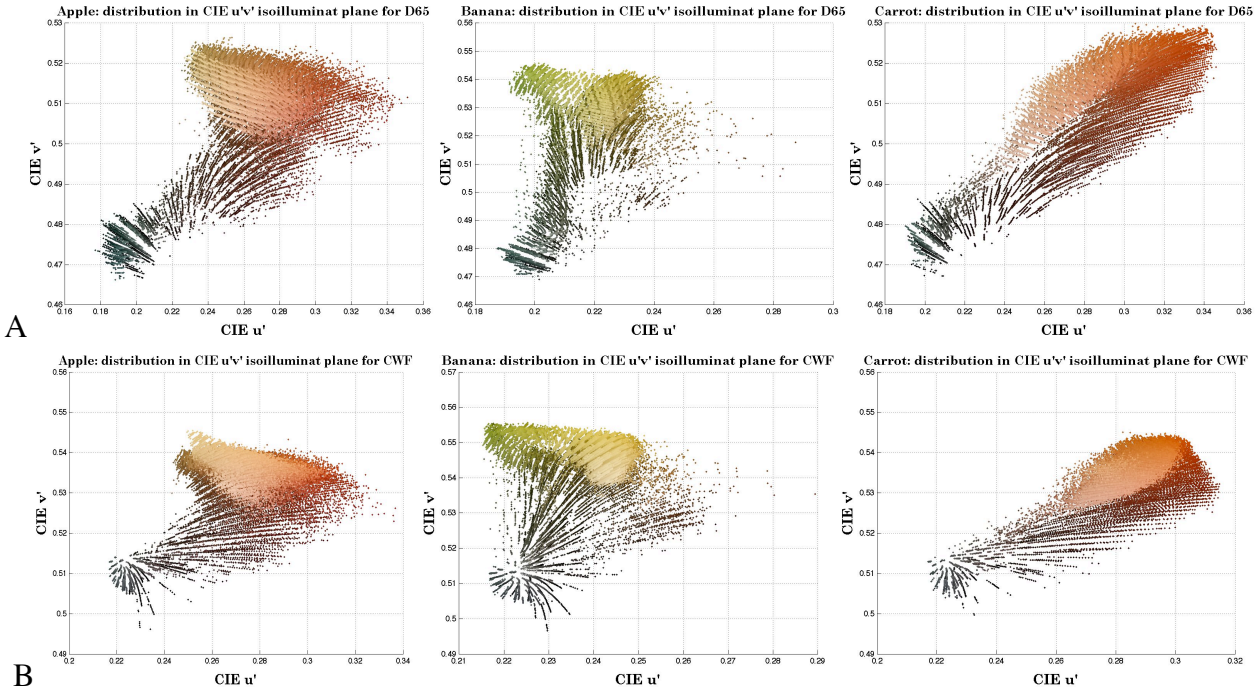

Figure 6.4 – Chromaticity distribution of the three objects for the two illuminations in CIE L'u'v' space. Columns from left to right: Gala apple, banana, and carrot. A) D65. B) CWF. Note that the colour of each point of the distributions in the figure is only indicative of its real colour. Each point is represented only once even if occurring more frequently in the distribution.

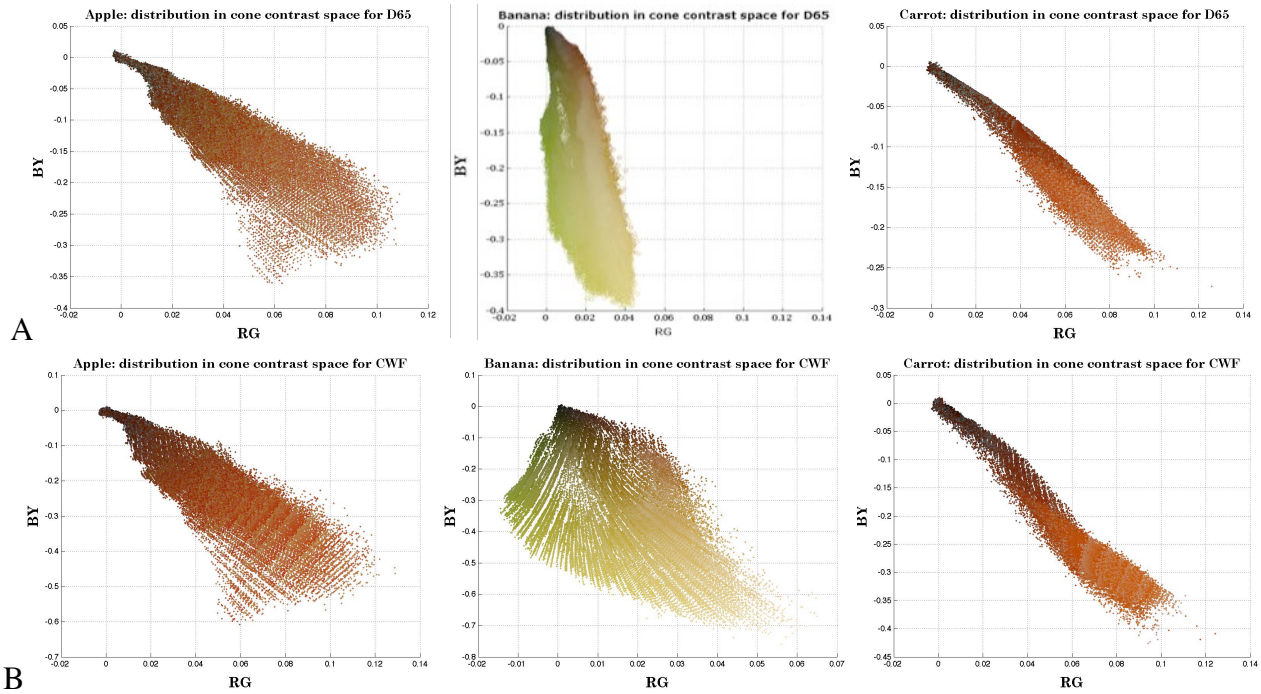

Figure 6.5 – Chromaticity distribution of the three objects for the two illuminations in EMG cone contrast space. Columns from left to right: Gala apple, banana, and carrot. A) D65. B) CWF. Note that the colour of each point of the distributions in the figure is only indicative of its real colour. Each point is represented only once even if occurring more frequently in the distribution.

As in Chapter 4 and Chapter 5, it is possible to calculate the area covered by the chromatic distribution of an object in EMG cone contrast space, i.e. the object colour variation indices or OV (see sections 5.3.1.2 and 4.4), and evaluate the difference between the extent of the objects' spreads (Table 6.3 for D65). Then, the distributions were manipulated as described in the general framework in Chapter 3. Specifically, block 2 will be described in detail in section 6.2.3.1.1, while block 1, block 3, block 4, and block 5 were described in Chapter 3 section 3.3 to 3.7 (see Figure 3.1).

Object	OV*100
Apple	2.315
Banana	1.2845
Carrot	0.9634

Table 6.3 – List of the object-variation indices (OV) for D65.

6.2.3.1.1 Textured stimuli generation

As shown in Chapter 5, a natural object's chromaticity distribution, when plotted in EMG cone contrast isoillumant plane, forms a distinct signature. In Figure 6.5 we can observe the signatures for the tested objects. As described in Chapter 5 section 5.3.1.1 we can compute, for each object under a specific light, the mean hue vector of its surface colours distribution. Table 6.4 lists

the object hue angles, calculated as in Equation 5-1 in section 5.3.1.1, for both the D65 and CWF simulated illuminants.

	Apple	Banana	Carrot
D65	70.13	84.13	67.73
CWF	73.62	87.10	65.83

Table 6.4 - Angles of the three objects under the two illuminants (D65 and CWF) (angles calculated from the RG axis in clockwise direction).

Therefore, a convenient method to change the overall colour of the object, and maintain the relative distance between colour points, is to rotate the entire distribution solidly with its hue vector around the adaptation point, i.e. the cone contrast space origin. The luminance value of each point was kept constant and only the RG and BY coordinates changed in the rotation. In order to fit the projector's widest gamut for a specific object, the luminance was scaled by a fixed factor dependent on the object. Figure 6.6 shows an example of the gamut of the projector and of the banana colours' distribution with the luminance (L) scaled by 2.5.

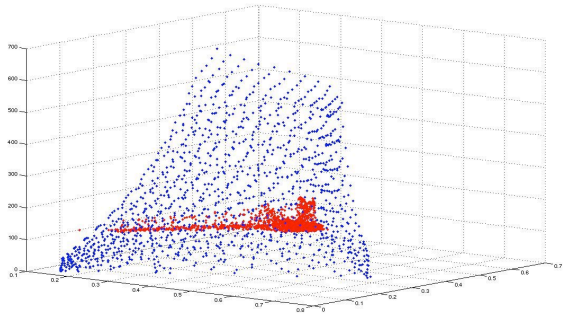

Figure 6.6 – Projector gamut in the setting of this experiment and banana distribution in La*b* colour space. The L value of the banana distribution is scaled of 2.5. Blue dots: sampled projector gamut. Red dots: banana distribution, each colour point shown once.

The hue disk of rotation was sampled at 1 degree of angle in EMG cone contrast space; thus, 360 images were generated per object/configuration combination (24 combinations) for each illuminant (2x24 final combinations=17280 images). As an example, Figure 6.7 illustrates the chromaticity distribution of the original banana (0° rotation) and the same distribution rotated 180° in EMG cone contrast space. Note that, due to uncalibrated reproduction of the colour on paper, the colour of each point in the figure is only indicative of the actual colour of displayed in the chamber.

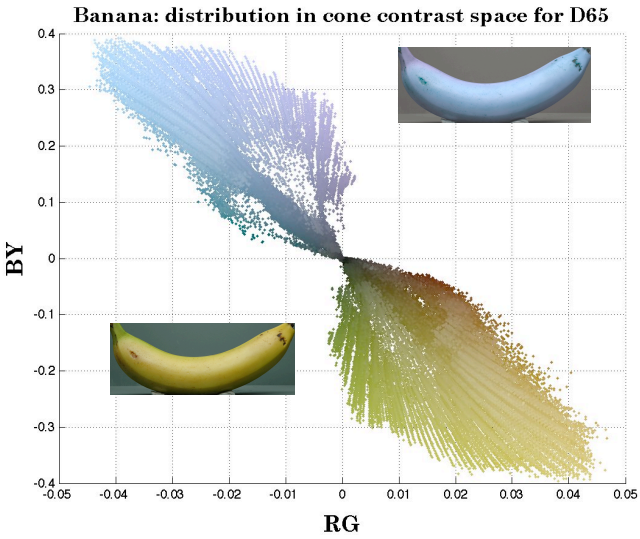

Figure 6.7 – Banana's distribution in its original configuration (yellow-brown) and in its 180° rotated configuration (blue-violet) in EMG cone contrast space. Note that, due to the uncalibrated reproduction of the colour, the colour of each point of the distributions is only indicative of its actual colour in the experiment. Each colour is represented only once even if it occurs more frequently in the distribution. Bottom left insert = original banana image. Top right insert = altered banana image.

For each point in the chromatic distribution, the position after rotation was determined as follow. Let be $c_i = (r_i, \Theta_i)$ a point of the initial distribution in the EMG isoluminant plane, given in polar coordinates, Φ_N the angle of the rotation and $c_i' = (r_i', \Theta_i')$ the position of the new adjusted point after rotation. Then, the new position for each point in the distribution can be computed, in polar coordinates, as:

$$r_i' = r \qquad \text{Equation 6-1}$$

$$\Theta'_i = \Theta_i + \Phi_N \qquad \text{Equation 6-2}$$

While, in cone contrast space, rotating the distribution around the neutral point implies that its barycentre draws a circle centred in the origin, in the CIE L'u'v' system, the same rotation corresponds to a pseudo-elliptical rotation also around the origin (Figure 6.8).

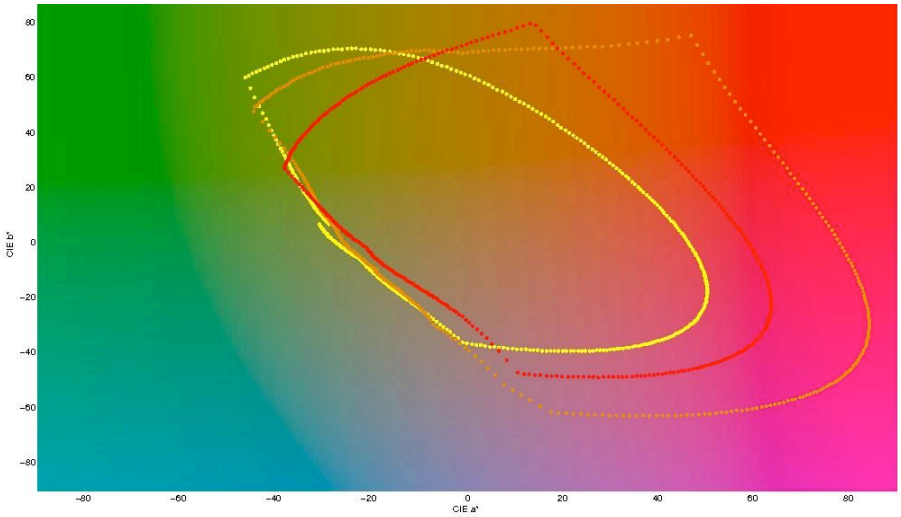

Figure 6.8 - Rotation of the barycentre of the objects' distributions under D65 in CIE Lab space. Each dot represents the location of the barycentre of the distribution after 1 degree of rotation compared to the neighbours.

Discriminability threshold between colours depends on the location of the colour in colour space. Alas, defining the perceptual difference between successive angles for the entire distribution is neither possible, nor convenient. However, we can calculate the difference between the mean of the distribution at different angles of rotation.

One degree of rotation in cone contrast space corresponds to a max ΔEuv difference between consecutive angles of 1.38, 1.58 and 2.77, respectively for apple, banana and carrot. Table 6.5 lists the ΔE difference between the mean hue angle of original distribution (indicated as the angle at 0 degree of rotation) and 24 adjacent angles at 2 degrees of rotation in clockwise direction (positive rotation) or counter-clockwise (negative rotation). Note that the ΔEuv difference between steps is of the same range order for all objects, but the carrot has a slightly bigger difference between steps. Nevertheless, the carrot has the smallest distribution variation, thus the overall difference of the distribution between angle steps can be considered similar.

Positive	Apple	Banana	Carrot	**Negative**	Apple	Banana	Carrot
24	47.755	36.948	67.206	-24	44.778	37.419	60.311
22	44.761	33.994	62.413	-22	41.317	34.389	55.761
20	41.948	31.01	57.351	-20	37.803	31.336	51.128
18	37.903	27.997	51.991	-18	34.238	28.261	46.41
16	33.411	24.958	46.313	-16	30.622	25.167	41.608
14	29.017	21.896	40.328	-14	26.956	22.055	36.721
12	24.721	18.812	34.144	-12	23.242	18.929	31.747
10	20.494	15.709	28.162	-10	19.48	15.79	26.686
8	16.316	12.59	23.622	-8	15.671	12.641	21.536
6	12.179	9.4564	17.369	-6	11.818	9.4855	16.296
4	8.0807	6.312	11.422	-4	7.9206	6.3249	10.963
2	4.0209	3.1589	5.6456	-2	3.9809	3.1622	5.5328
0	0	0	0	0	0	0	0

Table 6.5 – Perceptual difference between the original mean of the objects' distribution and the 24 angle steps around it for all three objects. Examples of perceptual difference between 2 degree of scaled angle in EMG cone contrast space under the illuminant D65.

6.2.3.2 UNIFORMLY COLOURED OBJECTS

In addition to chromatic texture (TEX), two types of uniform colours were employed: the average of the distribution and the most saturated point of the distribution. For each angle of rotation of the distribution described above, a pair of homogeneously coloured stimuli was generated. One stimulus chromaticity was equal to the mean chromaticity of the rotated distribution (MM) and the other to the most saturated (MS) value. The luminance of the uniform stimuli was

equal to the mean luminance of the object's distribution. Thus, the perceptual difference between consecutive points for the MM case, i.e. one degree of rotation in cone contrast space, corresponds to a max ΔE differences of 1.35, 1.56 and 2.32, respectively for apple, banana and carrot. Instead, for the most saturated case the max ΔE differences are 2.15, 2.5 and 2.9, respectively. Again, in general, the ΔE differences between rotation steps are of the same range order for all objects. Henceforward when referring to uniform mean colour condition together with the uniform most saturated condition I will use the term "uniform conditions" for simplicity and brevity. Finally, it is worth pointing out that chromaticity coordinates of all stimuli were simulated with respect to the illuminant D65 or CWF generated by the projector.

6.2.4 SHAPE CONDITIONS

Four solid objects were used in this experiment, of which three were 3D artificial replicas of the fruit and vegetable objects listed before and one was an unfamiliar handmade object. Specifically, apple and banana were copies in polystyrene of the real object shape, while the carrot was hand-made as described next in 6.2.4.1. The generic object was an irregular rhomboid-based pyramid made using air-drying modelling clay (DAS white by GIOTTO/FILA). For the purpose of this thesis it will be called "3D rhomboid" or 3D generic shape, alternatively. All objects were finished off with matt white acrylic-based paint. As a result, these objects preserve all the shape cues of the test objects, but can be seen as a "blank canvas" on which different test colours can be projected. Figure 4.3 shows the 3D white-painted objects (apple, banana, carrot and "3D rhomboid").

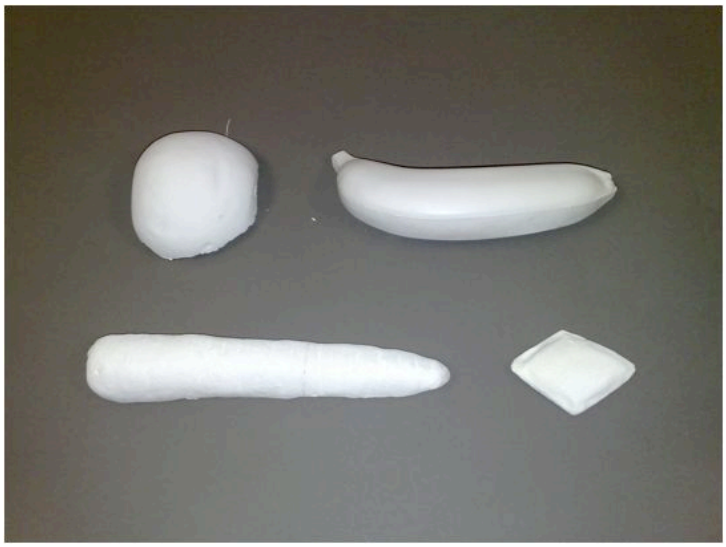

Figure 6.9 – Photograph of the three objects used in condition d).

Thus, for each object, the stimulus can be presented in four different geometric configurations: (a) 2D rhomboid (2DG, covering 3.00x2.15 deg); (b) 2D "silhouette" of the object (2DN, see Table 6.6); (c) 3D rhomboid (3DG, covering 3.00x2.15 deg); (d) solid 3D natural shape

of the object (3DN, see Table 6.6). Henceforward, 2D and 3D rhomboid shapes together will be referred with the term "generic shape condition", while "natural shape condition" will refer to 2D "silhouette" of the object and 3D natural shape of the object.

Figure 4.4 illustrates pictures of the four configurations for the apple (texture case). Specifically, subfigure Figure 4.4B) and D) are photographs of the 3D configurations (generic and natural contour, respectively) taken in the experimental chamber using a Canon digital camera. Keep in mind that colours in Figure 4.4 are only an approximation of the ones displayed in the experimental chamber.

Objects	Degrees
Apple	3.814x3.148
Banana	7.360x3.101
Carrot	8.718x1.670

Table 6.6 – Visual degrees of angle covered by the objects used in the experiment.

A) 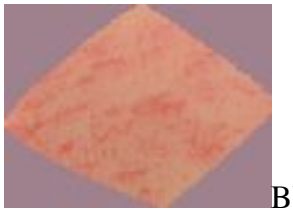 B)  C)  D) 

Figure 6.10 - Configurations presented in the experiment for the texture case: Apple, D65 illuminant. A) 2D rhomboid; B) 3D rhomboid (photograph); C) 2D silhouette of the object's natural shape; D) the 3D natural object shape condition (photograph). Note that the subfigure B) and D) are photographs of the condition take it in the actual experimental chamber. Keep in mind that image colours are only an approximation of the ones displayed in the experimental chamber.

As for the experiment in Chapter 4, the object cut out in the projected image exactly aligns with the visible surfaces of individual objects placed in the box, segmented from the background. Furthermore, the images used in condition (a) and (c) are identical to the image projected on top of the white object (case (b) and (d)). The subject's perception of the object is either as a 3D solid shape changing colour, or alternatively as a 2D shape. It is important to say that the texture condition of the generic object (or rhomboid) consists of a cut out of the original object chromatic texture that *subsequently* is manipulated as described in 6.2.3.1.1. In the following, each combination of the condition of the factors (or "factors combination") will be indicated as "chromatic label-shape label", e.g. the texture condition combined with the 2D generic shape will be referred as TEX-2DG.

6.2.4.1 OBJECT CASTING

The sequence of object casting is shown in Figure 6.11. First a soft mould of the object was made using algae gelatine. Before the gelatine settled the object was pressed inside the mixture. After solidification, the object was slowly removed. Lightweight white plaster of Paris was poured into the form. Before the mould was sprayed with synthetic oil to allow easier extraction of the cast after the compost had hardened. The same die was used to obtain 3 identical copies of the object.

Figure 6.11 – Casting process sequence. From left to right: from mould making to final cast.

6.2.5 EXPERIMENTAL PROCEDURE

The subjects were asked to adjust the colour (uniformly colour or texture) of an object until it appeared natural with respect to a previously named fruit or vegetable (Gala apple, banana, carrot), i.e. to set it to its typical colour. They had to complete two blocks, one for each illuminant. Each block consists of 12 sessions, 4 for each object (2DG, 3DG, 2DN, and 3DN). In each session, texture and uniform conditions were alternated, counterbalancing mean and most saturated case, such that in total 10 trials presented the MM condition, 10 trials the MS and 20 trials the TEX condition. The duration of each session was 15-25 minutes.

Initially, the subject was asked to stabilize his/her head on the chin rest and look inside the experimental chamber; then s/he was handed a joystick. The lights inside the chamber were switched on and set at the illuminant chromaticity values; the rest of experimental room was in complete darkness. The general scene displayed a test object centrally placed on a uniform background.

At the beginning of each session, the subject read the name of one of the test objects (black word on white background shown in front the observer), followed by a 60 seconds adaptation phase to the selected illumination. Then the first stimulus was presented, chosen randomly between the three chromatic conditions (texture, uniform mean and uniform most saturated). Each session consisted of one object/shape factors combination. The starting point for the adjustment was randomly picked among the set of images generated for the specific combination of factors (object/contour/dimensionality/chromatic factors combination). Two buttons of the gamepad allowed the subject to change the colour appearance of the object until he/she judged it to be 'natural'. One button induced a clockwise rotation of the distribution while the second a

counterclockwise. The buttons' functions were counterbalanced between subjects. No limitation on the rotation was given; in other words, the participant could use the clockwise (or counterclockwise) button to complete an entire 359° rotation and smoothly continue in the same direction returning to 0° and so on. Also the observers were not given a limit of how many times they could repeat the same rotation in one trial and pass the same point. However, they were asked to complete the task as quickly as possible. The stimuli were displayed for as long as the subject needed to make his/her judgement. Once the decision was made, the observer pressed a third button and the selection was recorded, as well the time of adjustment and the initial setting. Next, a fixation cross was presented for 0.5 sec and the next trail initiated. The first stimulus chromatic condition was randomly chosen between texture, mean and most saturated condition, and counterbalanced over all subjects/sessions. The general protocol is presented in Figure 6.12.

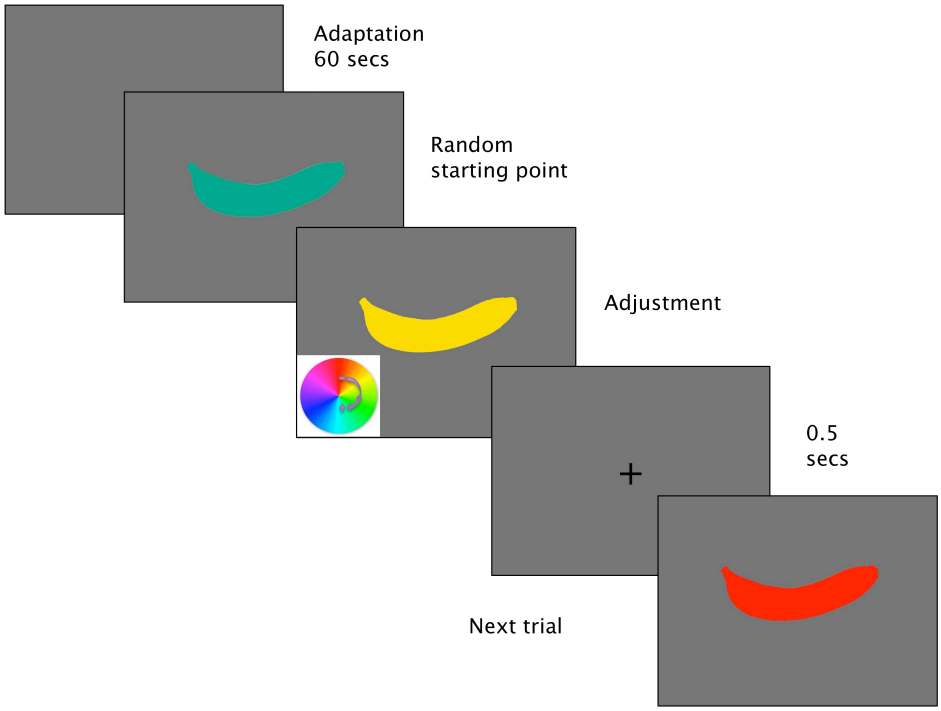

Figure 6.12 – General procedure: example of the experiment images sequence. Note that the colours in the figure are only for illustration and are not the ones displayed in the experiment.

There were 9 possible object sequences; each observer was assigned to one of them (i.e. one ninth of the observers had the same object sequence). Conversely the shape configuration sequence is identical for all the subjects. The generic configuration is always presented first, counter balancing the 2D and 3D order of presentation. The last two conditions are pseudorandomised, with 50% of the participants shown the 2D object's shape first and then the 3D shape, and 50% the reverse. This order ensures a gradual introduction of more identity cues. Initially only the linguistic information about the familiar object is given, with no shape reference; next the familiar shape is introduced, adding richness to the scene.

6.3 RESULTS

A convenient approach to analyze the data is to divide the results in: (1) effect of shape and familiarity under a constant light source (memory colour); (2) the effects under varying light (colour constancy).

6.3.1 MEMORY COLOUR

Memory colour was analysed under the ***reference light D65***. Figure 6.13 shows the observers' selections for the apple in 3D rhomboid shape for texture condition under D65 (or apple-TEX-3DG-D65 combination). Visibly, the range of choices is wide and, although the selections cluster around the original distribution hue angle (black line), they are spread both clockwise (positive angles) and counterclockwise (negative angles) in respect to the original angle. Note that the red line in the figure is obviously an outlier; therefore any analysis was computed after discarding the outliers for each subject, which were defined as the points falling more than 1.5 times the interquartile range above the third quartile. Also, for each shape configuration, only 10 of the 20 texture trials were randomly selected to have the same number of trial as the uniform mean and uniform most saturated cases. Henceforth, all effects are reported as significant at $p<0.05$, unless otherwise indicated.

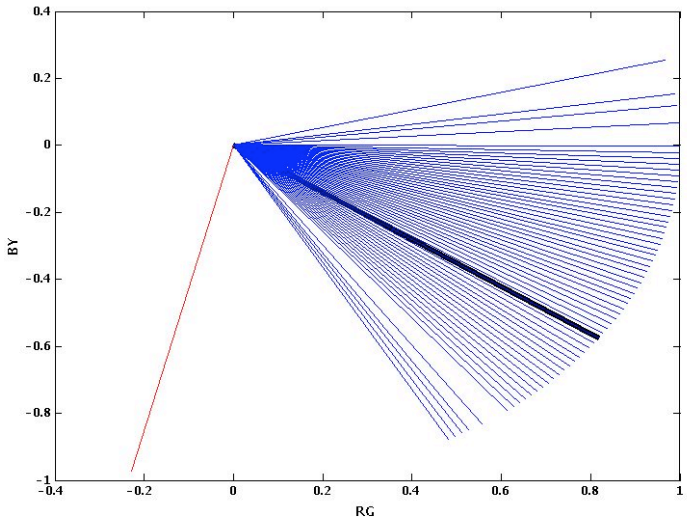

Figure 6.13 – Subjects' selection for the apple in 3D rhomboid shape for texture condition, D65. Black line: original hue angle. Red line: outlier.

6.3.1.1 THE SUBJECTS' SELECTIONS

Figure 6.14 plots the number of times in which the observers selected a specific angle as a function of the angle value itself, for the apple-MM-3DN-D65 combination (i.e. the apple in the mean uniform condition in 3DN shape configuration for the D65 illuminant). For example, if a certain angle α is selected by only 5 observers with the two choosing it in 3 trial, two in 1 and one in 2 trails, then the number of times that this angle was chosen would be equal to 10 (thus the maximum time per condition that a point can be choose is 10 trials*28 subjects=280). Each grey dot

in Figure 6.14 represents the number of times the angle on the x-axis was selected in the specific combination of conditions by *all* observers in *all* trials. This distribution of the settings chosen by all observers/trial represents the global behaviour for this specific condition (apple-MM-3DN-D65). It is important to note that the zero point (grey vertical line) in the figure represents the original hue angle of the specific object under D65, and that the angles selected by the subjects are reported relatively to such value. In other words, if the object's original angle is 25 degree (0 degree in the graph), but the subject selected an angle of 5 clockwise-degrees bigger, and then the value plotted in the graph would be 5 (not 30). In Figure 6.14 are also indicated the mean of the distribution (red dashed line), the spread of the distribution (horizontal red arrow), and the mode of the distribution, i.e. the value of the distribution that occurs with the greatest frequency (red arrow). The mode represents the angle that is most selected by all observers in all trials belonging to the same condition.

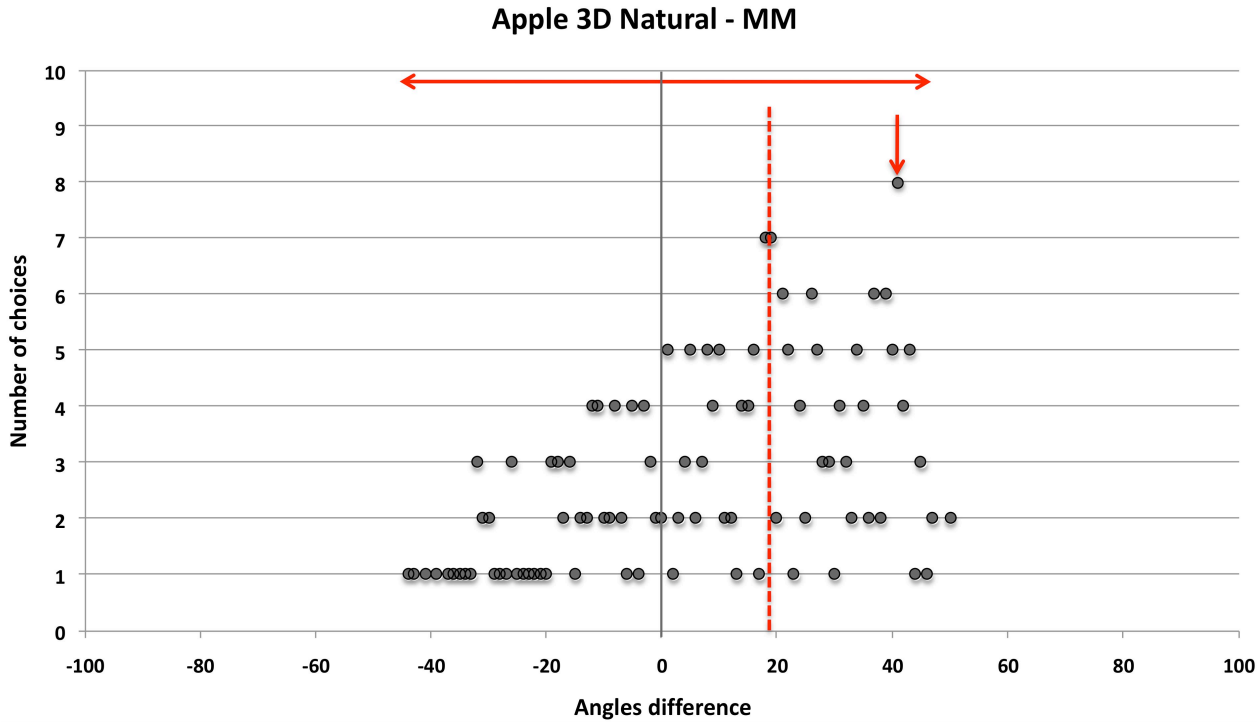

Figure 6.14 - Plot of the number of trials in which a specific angle was selected across all trials of all 28 subjects. Condition combination: apple-MM-2DN-D65. Each grey dot represents the number of times the angle on the x-axis was selected in the specific combination of conditions by *all* observers. The red arrow indicates the mode (i.e. the angle selected the most over all trials. Red dashed line: mean of the distribution. Horizontal red arrow indicates the range of the selection.

Let us then visually compare the selection distribution for the apple showed above with the ones for the TEX and MS condition also in the 2DN shape and under D65 (Figure 6.15; TEX=blue dots, MM=red dots, and MS=green dots). In Figure 6.15 the modes of the three distribution, indicated by the vertical arrows, are different, in particular of the mode of the MM condition distribution (red arrow) is largely shifted to the right of both of them (more negative). The mean of the TEX condition distribution (dashed blue line) and the MS condition distribuition mean (dashed green line) are rather close to each other similar, while the MM condition mean (dashed red line) is greatly shifted right from them. Lastly, while the spreads of the TEX condition and the MM

conditions distribution are fairly similar, TEX and MS condition distribution spreads of choices differ (horizontal arrows – TEX=blue, MM=red, and MS=green).

Therefore, the memory colour difference between conditions can be analysed as the difference between these distributions on at least four levels: (1) the shift between the mode of the selections' distribution and the typical object hue angle (defined "dominant setting" or "dominant angle"); (2) number of times (or trials) in which the dominant angle was selected by all observers; (3) mean of the distribution (μ); (4) and spread of the distribution or range of the selected angles. In section 6.3.1.4, I will introduce an additional way in which the data on memory colour can be analysed.

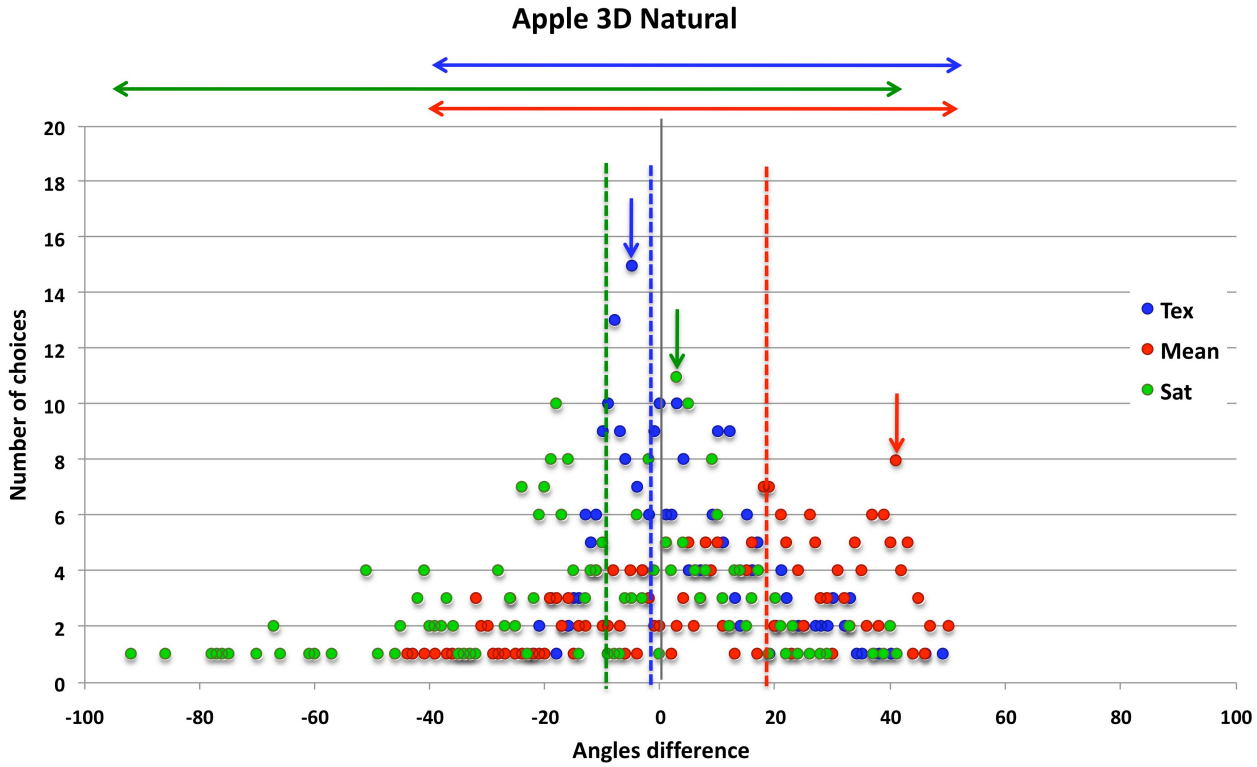

Figure 6.15 – Comparison between chromatic conditions: number of selections per angle in all trials (M=280) performed by all observers (N=28). Apple in 3DN condition (illuminant D65). Vertical axis: number of selections. Horizontal axis: difference between the original hue angle and the subject angle setting. Each dot represents data from one chromatic condition: texture (TEX, blue), mean (MM, red), or most saturated (MS, green). Dashed vertical lines: mean of the distributions – blue=TEX, red=MM, green=MS. Arrowheads indicate the mode of the distribution – blue arrow=TEX, red arrow=MM, green arrow=MS. Horizontal arrows: maximum spread of the distribution – blue=TEX, red=MM, green=MS.

6.3.1.2 DOMINANT ANGLES

As said before, I calculate the mode of the selections' distribution of all trials belonging to the specific condition (10 trials x 28 subjects = 280 observation per condition; see Figure 6.15). This angle represents the most prominent memory of object colour over all subjects for the specific condition. The absolute angular difference between mode and the typical object hue angle is the dominant angle of the distribution (in Figure 6.15 the angle position of the peak of the distribution). Figure 6.16 compares the dominant settings of all objects with respect to the surface chromatic representation (TEX vs MM or MS) for natural shape (Figure 6.16A) and generic shape (Figure

6.16B) condition. In each subfigure, each point represents one object in one chromatic/dimensionality condition combination. By eye, the majority of the points lie below the unity line (in black) indicating that in general the dominant angle for the TEX condition is smaller than for the uniform conditions. Quantitatively, we can compute the mean (μ) of the dominant angles per condition and the mean difference (ν) between the dominant angles of two conditions. The mean of dominant angles selected by the all subjects for texture condition is μ_{TEX} =6.66, for the mean uniform condition is μ_{MM} =12 and in the most saturated trials is μ_{MS} =8.66. Additionally, we can compute the for the natural shape condition, the mean difference between the dominant angles of texture and uniform conditions are $\nu_{TEX\text{-}MM}$(Nat) = -7.33 (TEX<MM deviation, $\sigma_{TEX\text{-}MM}$(Nat)=10.5) and $\nu_{TEX\text{-}MS}$(Nat) = -1.23 (TEX<MM deviation, $\sigma_{TEX\text{-}MS}$(Nat)=2.90), respectively. Similarly, for the generic shape condition the mean difference between the dominant angles of texture and uniform conditions are $\nu_{TEX\text{-}MM}$(Gen) = -3.333 (TEX<MM, $\sigma_{TEX\text{-}MM}$(Gen)=9.22) and $\nu_{TEX\text{-}MS}$(Gen) = -1.07 (TEX<MM deviation, $\sigma_{TEX\text{-}MM}$(Gen)=3.21), respectively (Figure 6.16B). Thus, the presence of the natural texture of the object shifts the subject settings *consistently* towards the original object hue angle. This effect is stronger for the natural condition, $\nu_{TEX\text{-}MM}$(Nat) >$\nu_{TEX\text{-}MM}$(Gen).

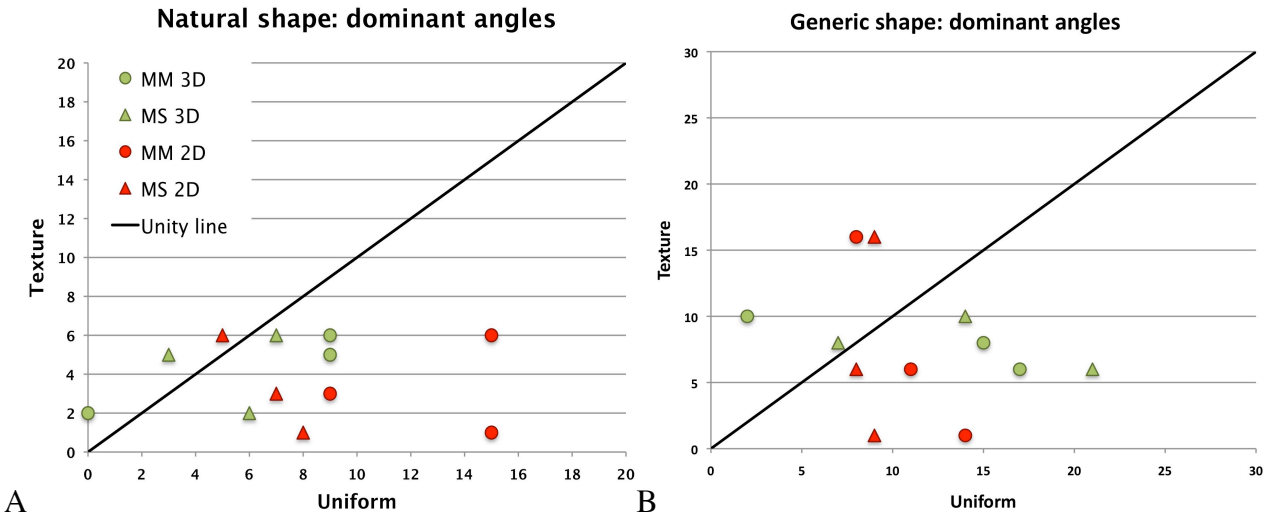

Figure 6.16 – Scatter plot of the dominant angles as a function of chromatic conditions. A) Natural shape condition. B) generic shape condition. Each point represents one object in one chromatic/dimensionality conditions combination. Closed circles: relationship between texture and mean uniform (MM) conditions. Triangles: relationship between texture and most saturated uniform (MS) conditions. Green coloured points: 3D shape. Red coloured points: 2D shape. Black line: unity line.

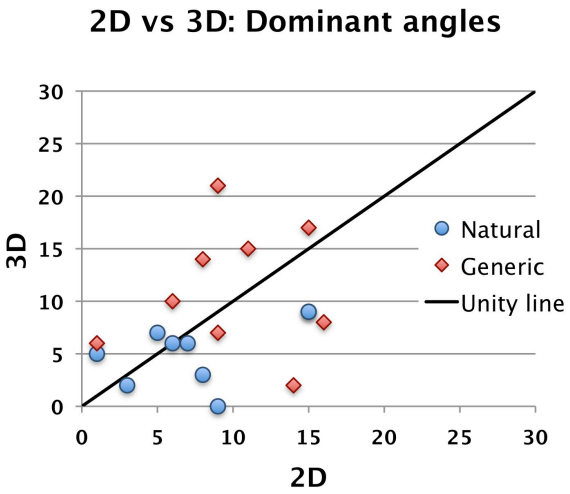

Figure 6.17 – Scatter plot of the dominant angles as a function of the shape dimensionality. Each point represents one object in one chromatic/contour factor combination. Blue closed circles: Natural shape. Red diamond: Generic shape. Black line: unity line.

Instead, Figure 6.17 compares the dominant settings with respect to the shape dimensionality (2D vs 3D). Each point in each subfigure represents one object in one chromatic/contour condition combination. The means are μ_{2D}= 8.55 and μ_{3D}=9.67. Although the plot shows consistent shifts from the object's original angle for 3D *natural* contour condition relative to the 2D *natural* contour condition, the mean difference between the two is small (ν = -1.22±1.27, 3D<2D deviation; Figure 6.17; blue circles). Similarly there is no significant difference between dominant angles for 3D *generic* contour condition relative to the 2D *generic* contour condition, ν = 2±7 (2D<3D, p=0.22; Figure 6.17; red diamonds). Note that the variability between dimensionality conditions relative to the generic condition is more than 3 times higher than the natural condition.

Finally compares the dominant angles as a function of the shape diagnosticity (natural vs generic). Each point in each subfigure represents one object in one chromatic/dimensionality condition combination. The means are μ_{Gen}= 10.11 and μ_{Nat}=8.11, this difference being significant (t-test, p<0.05). Clearly, the dominant angles selected by the subject for the natural contour presentations are closer to the object original hue angle than for the generic condition.

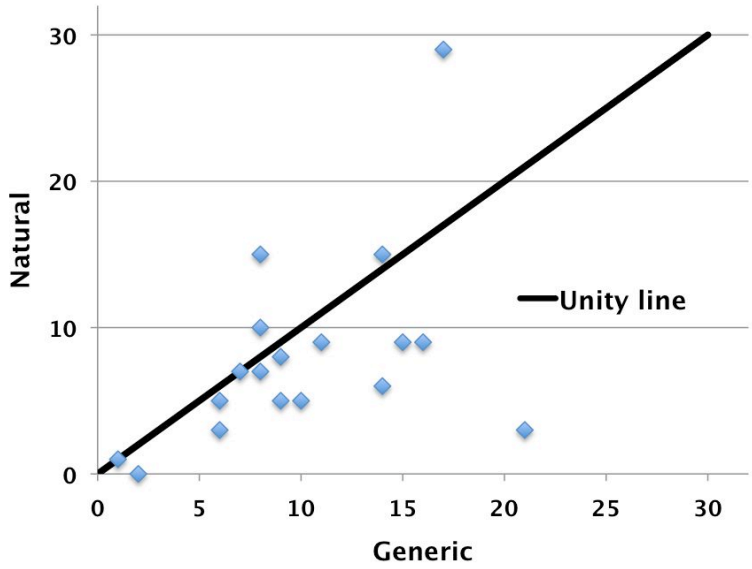

Figure 6.18 - Scatter plot of the dominant angles as a function of the shape diagnosticity. Each point represents one object in one contour/chromatic factor combination. Black line: unity line.

An additional way to evaluate the difference between within-factor conditions is to calculate the percentage difference between them as follows:

$$d_{AB}\% = \frac{\mu_A - \mu_B}{\mu_A}$$ Equation 6-3

where μ_A and μ_B are the means of dominant angles for two generic conditions (A and B) of the same factor. Therefore, it is possible to say that on average the dominant angle is: (1) 80.0% (16.67%) closer to the object's typical hue angle for the texture condition than for the mean uniform condition (most saturated condition), $p<0.001$; (2) 24.66% closer to the object's typical hue angle for the natural condition than for the generic condition ($p<0.05$); and (3) 11.58% closer to the object's typical hue angle for the 3D condition than for the 2D condition, only significant for the natural condition ($p<0.05$).

6.3.1.3 FREQUENCY OF THE SELECTIONS AT THE DOMINANT ANGLE

As seen in 6.3.1.1, while two chromatic/shape factors combinations of an object can have the same dominant angle, the frequency of selecting that angle may differ, leading to diverse amplitudes of the fitting curves in Figure 6.15. Note that the higher this value, the stronger the global memory of the object colour for the particular condition.

Figure 6.19 shows the number of times, over all the observers' selections, that the dominant angle was picked as a function of chromatic condition. The occurrence of the dominant angle for the polychromatic stimulus varies along the vertical axis, while the occurrence of the dominant angle for the homogeneous conditions along the horizontal. Therefore, each point represents the

amplitude of the distribution curve of the chromatically variegated stimulus with respect to the uniform mean condition (indicated with circles – MM) or the uniform most saturated (squares – MS). In the same figure the 2D and 3D conditions are plotted respectively in red and green colour. Since all points, except one, are found above or on the unity line, the dominant angles for the texture condition are more frequently chosen that in the homogeneous cases (μ(TEX)=16.01,μ(MM)=14.33, μ(MS)=13.08). As before we can calculate the percentage difference as in Equation 6-3 where instead in this case μ_A and μ_B are the means of the occurrence of the dominant angles for two generic conditions (A and B) of the same factor. Hence we can compute that, on average, the dominant angle is chosen 9.95% times more for the TEX condition that for the mean uniform conditions and 17.81% times more than for the most saturated uniform condition. A paired T-test reports the effect of texture condition on dominant angle with respect to uniform condition significant over all objects (versus both MM and MS $p<0.01$). Therefore, textured stimuli encourage consistency in the subjects' colour choices, fortifying the collective memory colour of the object.

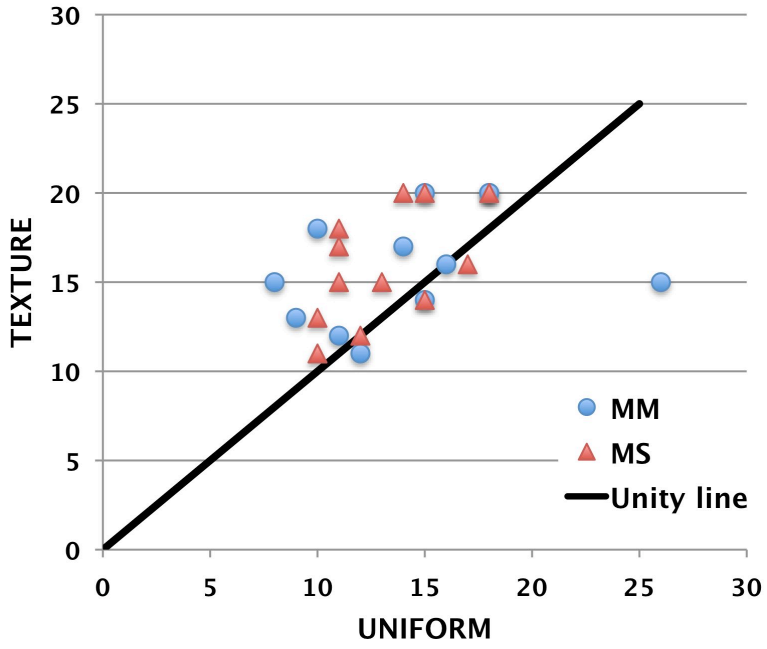

Figure 6.19 - Scatter plot of the occurrence of subjects' dominant angle as a function of the chromatic condition for all objects. Blue closed circles: each point represents one mean uniform (MM) condition and object/shape condition combination. Red triangles: each point represents one most saturated uniform (MM) condition and object/shape condition combination. Black line: unity line.

Figure 6.20 illustrates the relationship between occurrence of dominant angle and shape diagnosticity (y-axis is the natural contour and x-axis is the generic contour), for all objects. In view of the fact that the greater part of the points lies above or on the unity line, by eye we can assert that the natural contour may strengthen the global memory of the object colour. In fact, on average the dominant angle is chosen 12.27% times more for the natural contour condition that for the generic conditions, this difference being significant ($p<0.05$). Therefore, natural contours encourage consistency in the subjects' colour choices.

Conversely, weaker links can be seen between the numbers of choices at the dominant angle in dependence on shape dimensionality (2D or 3D), as shown in Figure 6.20. As a matter of fact, on average the dominant angle is chosen 3.03% times more for the 3D condition that for the 2D conditions, this difference not being significant (p=0.621).

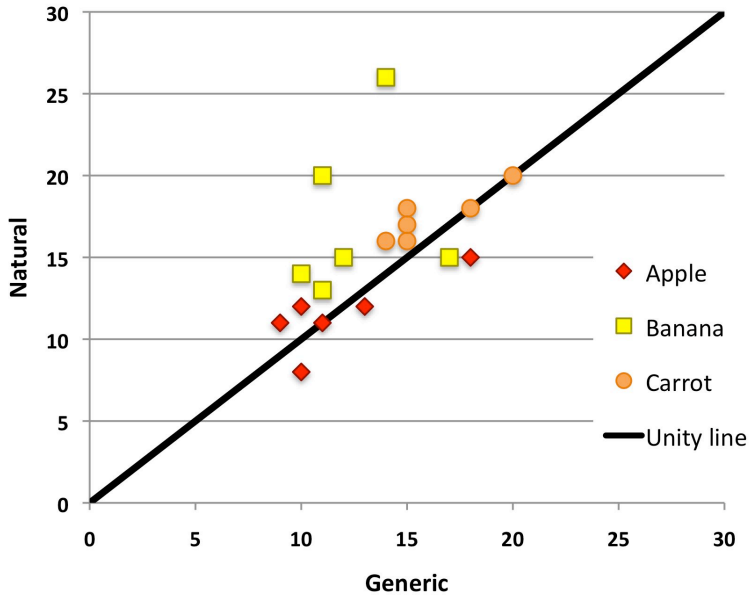

Figure 6.20 – Scatter plot of the occurrence of subjects' dominant angle as a function of the shape diagnosticity for all objects. Each point represents one chromatic/dimensionality condition combination. Red diamonds: Apple. Yellow squares: Banana. Orange closed circles: Carrot. Black line: unity line.

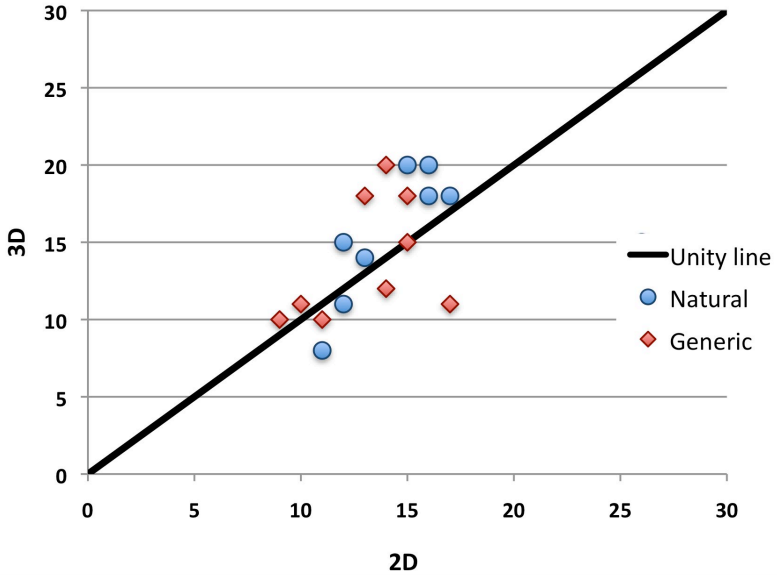

Figure 6.21 - Scatter plot of the occurrence of subjects' dominant angle as a function of the shape dimensionality. Each point represents one object/contour condition combination. Blue closed circles: Natural contour. Red diamond: Generic contour. Black line: unity line.

6.3.1.4 MEAN MEMORY COLOUR AND ABSOLUTE MEMORY COLOUR DEVIATION (OR DEGREE OF ACCURACY)

A further way to examine the distribution is to calculate its mean. For each object/condition combination, I averaged the mean of each subject settings (mean over 10 trials) across all subjects

(N=28) discarding the outliers for each subject as described before. The mean memory colour (MMC) is defined as the angular difference between this average and the original angle of the object. Figure 6.22 contrasts mean memory colour angle of the uniform conditions with texture condition. The ANOVA reported not significant the main effect of chromatic factor on mean memory colour angle ($F(1.23,33.4)=3.5$, $p=0.62$[13], pairwise comparisons, $p_{TEX-MM}=0.077$, $p_{TEX-MS}=0.36$ and $p_{MM-MS}=0.84$). Consequently, the global mean of the distribution is not affected by the surface chromatic representations. The means across all subject for all texture, mean uniform, or most saturated uniform trials were $\mu_{TEX}= -2.72$, $\mu_{MM}= -4.82$, or $\mu_{MS}= -5.26$, respectively.

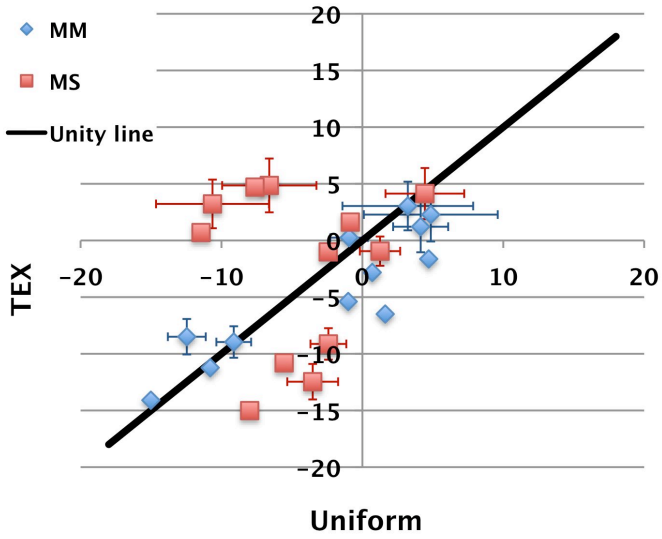

Figure 6.22 – Scatter plot of the mean memory colour angles of all objects/shape combinations as a function of chromatic conditions. Blue diamond: each point represents, for the *mean uniform condition*, one of the object/shape condition combinations. Red squares: each point represents, for the *most saturated uniform condition*, one of the object/shape condition combinations. The capped coloured error bars indicate the standard error of the mean (SEM). Black line: unity line.

Similarly, there is no significant effect of shape diagnosticity on mean memory colour ($F(1,27)=0.091$, $p=0.76$, $\mu_{Nat}= -4.14$, $\mu_{Gen}= -4.38$). Conversely, there is a significant effect of shape dimensionality on mean memory colour ($F(1,27)=23.7$, $p<0.05$), as illustrated in Figure 6.24B. On average the mean memory colour angles for the 3D condition are closer to the object typical hue angle than the 2D shape selections ($\mu_{3D}= -2.42$, $\mu_{2D}= -6.11$). In addition, the ANOVA shows that there is a significant interaction between surface chromatic representation and shape dimensionality ($F(1.72,46.45)=11.22$, $p<0.05$), but not between surface chromatic representation and shape diagnosticity ($F(1.56,42.34)=2.39$, $p=0.14$). This result indicates a relationship between surface chromatic representation and dimensionality cues, whereas the independence of surface chromatic representation from shape diagnosticity.

[13] Mauchly's sphericity was violated and the ANOVA was corrected using Greenhouse-Geisser estimates.

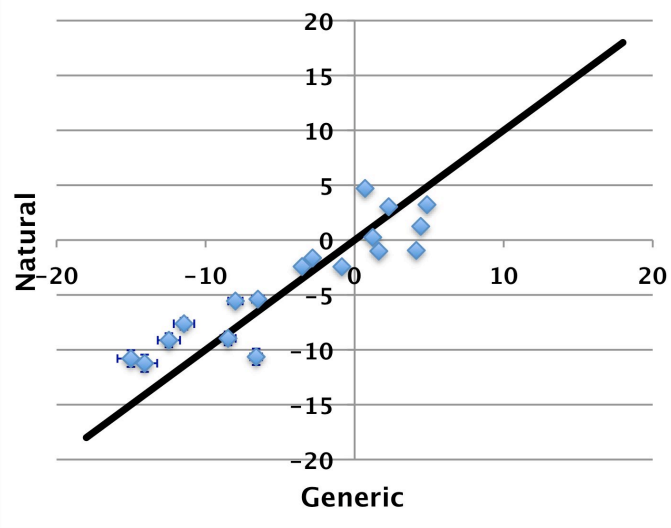

Figure 6.23 - Scatter plot of the mean memory colour angles as a function of shape diagnosticity (natural vs generic contour). Each point represents one of the object/shape/chromatic condition combinations. Error bars denote one standard error of the mean (SEM). Black line: unity line.

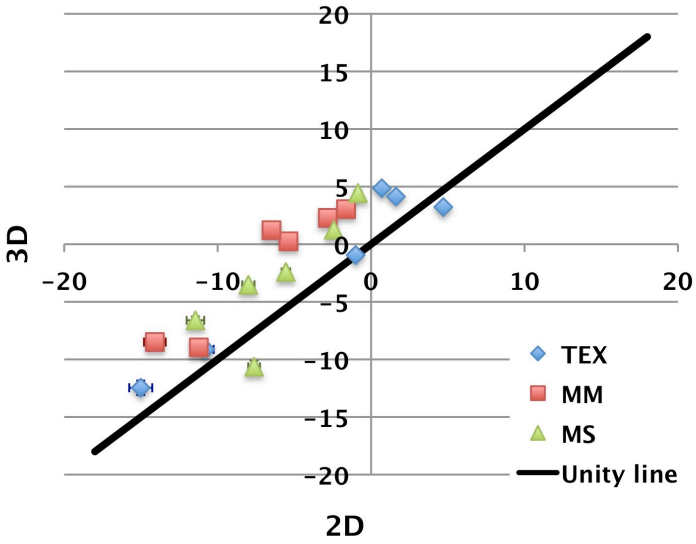

Figure 6.24 - Scatter plot of the mean memory colour angles as a function of shape dimensionality (2D vs 3D). Blue diamond: each point represents, for the *texture condition*, one of the object/contour condition combinations. Red squares: each point represents, for the *mean uniform condition*, one of the object/contour condition combinations. Green triangle: each point represents, for the *most saturated uniform condition*, one of the object/contour condition combinations. Coloured error bars denote one standard error of the mean (SEM). Black line: unity line.

However, the previous mean averages the signed deviations form the object original hue angle. In fact, previously I have defined as negative deviations the clockwise rotations and as positive deviation the counter clockwise rotations. Thus, if the same number of observers chooses two opposite angles of same amplitude, e.g. ±5 degree, then the global mean of the distribution would be zero, implying a perfect overlap of the object's typical colour and the subjects' memory colour. Actually, though, the subjects' memory colour is 5 degrees off from the original angle. Hence, the true angular deviation, or error, of the observers' choices from the typical object setting is the average of the ***absolute*** values of the subjects' means. Such value measures the degree of accuracy of the mean of the subjects' choices from the object typical hue angle (termed "absolute memory colour angular deviation", MCD). Therefore, first we calculate the mean of the subject's

selections, then its absolute value and finally we average over all observers, as express in the following formula:

$$MCD_i = \frac{\sum_{j=1}^{N} \left| \sum_{k=1}^{K} \Delta\alpha_{jk} \right|_i}{N * K}$$ Equation 6-4

where N and K are respectively the number of subjects and the number of trials per condition (N=28 and K=10), while $\Delta\alpha_{jkli}$ is the difference between the angle selected in the *k*th trial by the subject *j* in condition *i* and the object's typical hue angle. In addition, we can compute the percentage of deviation as in Equation 6-3 where in this case μ_A and μ_B are instead the means of absolute memory colour angular deviation over all trials belonging to two generic conditions A and B of the same factor.

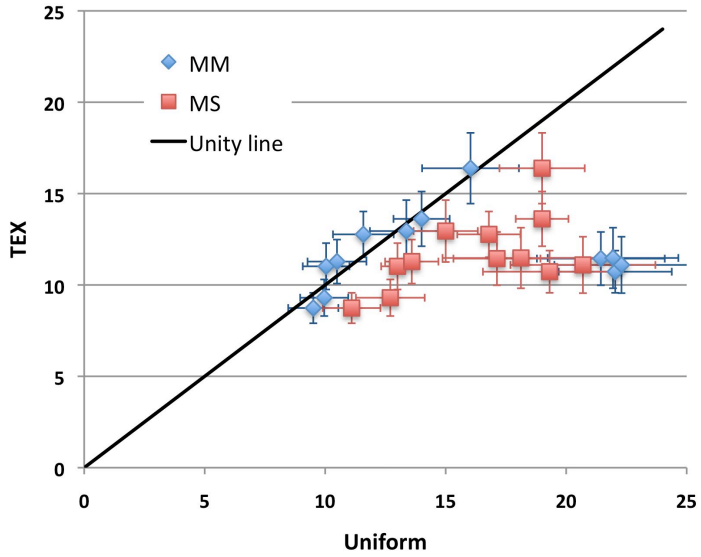

Figure 6.25 - Scatter plot of the absolute memory colour angular deviation as a function of chromatic conditions. Blue diamond: each point represents, for the *mean uniform condition*, one of the object/shape condition combinations. Red squares: each point represents, for the *most saturated uniform condition*, one of the object/shape condition combinations. The capped coloured error bars indicate the standard error of the mean (SEM). Black line: unity line.

Figure 6.25 shows that generally the absolute memory colour angular deviation is 31.70% (15.17%) smaller for the texture condition than for the mean (most saturated) uniform condition, this difference being statistically significant[14] ($F(1.71,46.18)=6.99$, $p<0.05$; $\mu_{TEX}=10.22$, $\mu_{MM}=13.46$, and $\mu_{MS}=11.77$). In other words, the subjects choose angles closer to the specific object's typical hue angle in texture trials than in uniform trials.

[14] Mauchly's sphericity was violated and the ANOVA was corrected using Greenhouse-Geisser estimates.

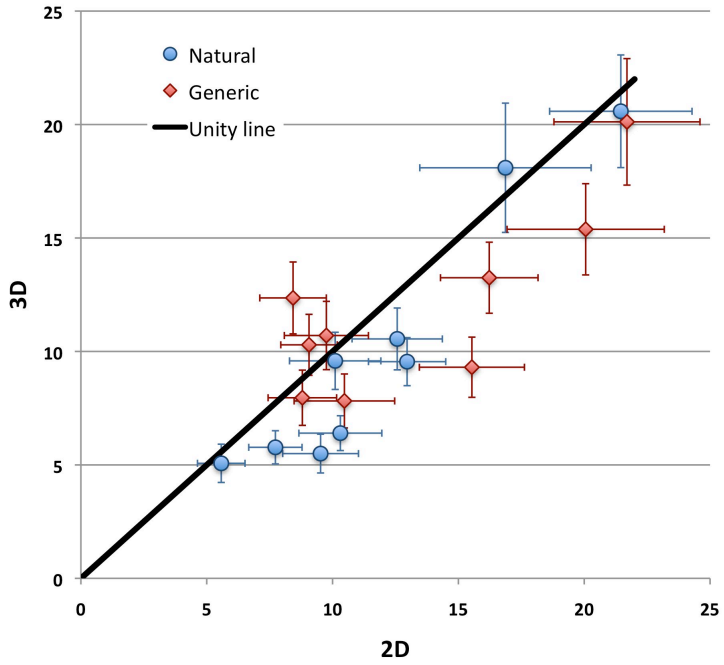

Figure 6.26 - Scatter plot of the absolute memory colour angular deviation as a function of dimensionality over all objects. Each point represents one object/condition combination. Blue closed circles: Natural contour. Red diamond: Generic contour. The capped coloured error bars indicate the standard error of the mean (SEM). Black line: unity line.

Figure 6.26 compares 2D to 3D configuration and almost all points lie below or on the unity line. Results show that on average the subjects select angles 14.01% closer to the typical object hue angle in case of 3D presentations than 2D (μ_{2D} =12.61, μ_{3D} =11.06). This effect is significant for all objects, $F(1,27)=6.05$ ($p<0.05$). It should be noticed that this effect is slightly stronger for natural contour than generic (μ_{2D}(Gen) =13.33, μ_{2D}(Nat)=11.89, μ_{3D}(Gen)=11.90, μ_{3D}(Nat)=10.12, p[15]=0.78).

Additionally, the ANOVA shows a significant effect of the shape diagnosticity ($F(1,27)=6.68$, $p<0.05$). Figure 6.27 compares the absolute memory colour angular deviation as a function of shape diagnosticity, for all objects. The graph shows that globally the MCD is 14.63% larger for generic contour than for natural contour condition, i.e. the subjects select angles closer to the typical object hue angle when the stimuli is presented in the original object's natural contour.

[15] The significance was calculated using a paired T-test between the difference of 2D Generic and 3D Generic absolute memory colour angular deviations and the difference of 2D Natural and 3D Natural absolute memory colour angular deviations.

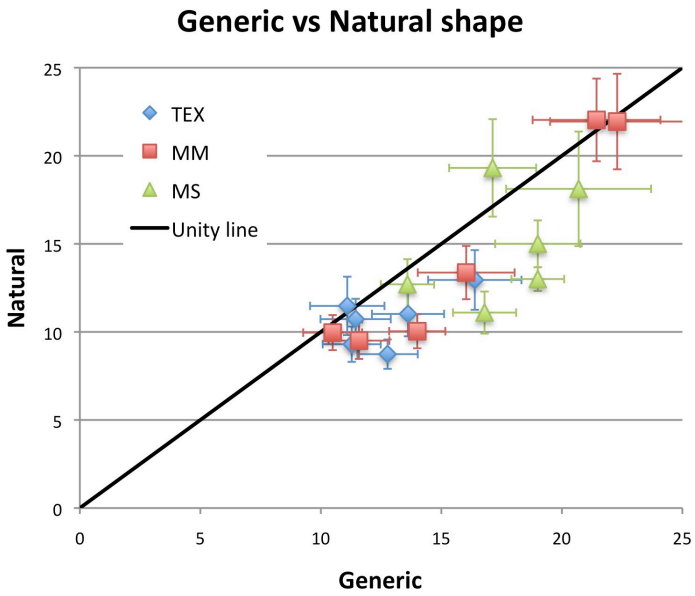

Figure 6.27 - Scatter plot of the absolute memory colour angular deviation as a function of shape diagnosticity. Blue diamond: each point represents, for the *texture condition*, one of the object/dimensionality condition combinations. Red squares: each point represents, for the *mean uniform condition*, one of the object/dimensionality condition combinations. Green triangle: each point represents, for the *most saturated uniform condition*, one of the object/dimensionality condition combinations. Coloured error bars denote one standard error of the mean (SEM). Black line: unity line.

6.3.1.5 MEMORY COLOUR RANGE OR PRECISION

Figure 6.28 demonstrates that although the overall mean value for one subject for two different conditions can be the same, the range of angles selected can be very dissimilar. In fact, the figure plots two sets of 10 angles selected by one subject for two generic conditions (red squares= condition 1; blue diamonds= condition 2). The continuous lines draw the mean of the 10 angles for each condition (red line=condition 1; blue line=condition 2), while the dashed lines the minimum and maximum of variation of the choices (red dashed lines=condition 1; blue dashed lines=condition 2). The area between the two dashed lines of the same colour is the range of the condition selections. Note that the two mean lines (continuous lines) practically completely overlap, whilst the range of choices is largely different. Keep in mind that the range denotes the span of the subject memory of the object colour angle (Figure 6.15), hence the smaller the range, the higher the subject's *precision* in selecting points around his/her memory mean. Thus, a reasonable way to explore the precision of the subjects' choices is to compute for *each* subject the range of choices and average it across all observers Equation 6-5. Note that the analysis of the observer's range of choices is similar to the one performed in Chapter 4 for the main experiment. However, while in Chapter 4 we were exploring the subject memory colour variability with respect to the colours within the *object*, in this experiment we explore the subject memory colour variability respect to the distributions within the object category.

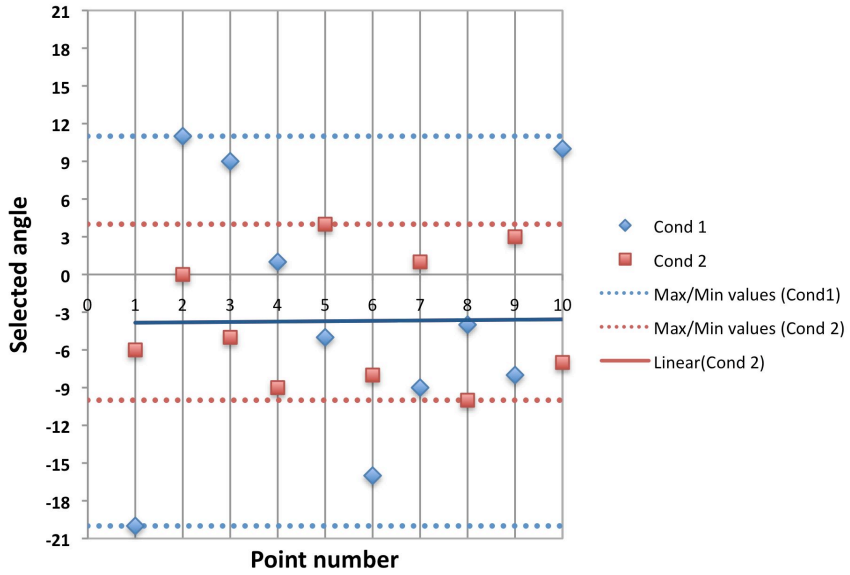

Figure 6.28 - Two sets of 10 angles selected for two generic conditions by one subject. Blue diamonds: condition 1. Red square: condition 2. Blue full line: mean of the choices for condition 1. Blue dashed line: max and min of the choices for condition 1. Red continuous line: mean of the choices for condition 2. Red pointed line: max and min of the choices for condition 2.

Let's define mean memory colour range (MCR) the value calculated as in the following formula:

$$MCR = \frac{\sum_{j=1}^{N} abs(\max(\Delta\alpha_{kj}) - \min(\Delta\alpha_{kj}))}{N} \qquad \text{Equation 6-5}$$

where $\Delta\alpha_{ij}$ is the angular difference between the angle selected in the trial *k* by the subject *j* and the original distribution's angle, $\max(\Delta\alpha_{ij})$ and $\min(\Delta\alpha_{ij})$ are the maximum and minimum values of all deviations, excluding the outliers, and N is the number of subjects (N=28). Similarly to previous sections, we can compute the percentage of mean memory colour range difference between two conditions as in Equation 6-3 whereas in this case μ_A and μ_B are instead the means of mean memory colour range over all trials belonging to two generic conditions A and B of the same factor. This is the "percentage range difference" between two conditions.

Figure 6.29A compares the MCR between TEX and MM condition in all shape cues combination. Each point represents one object in one contour/dimensionality condition combination. By eye, we can observe that the majority of the points in the graph lie above the unity line (in blue) with two exceptions. These latter two points belong both to the banana object in the generic shape condition (and only 2DN is significantly different from the unity (T-test, $p<0.05$)). Similarly, visually comparing the texture with most saturated uniform conditions in Figure 6.29B, the majority of the points is found below the bisect line, and again the only exception belongs to the banana when present on the generic shape (this time in 3D).

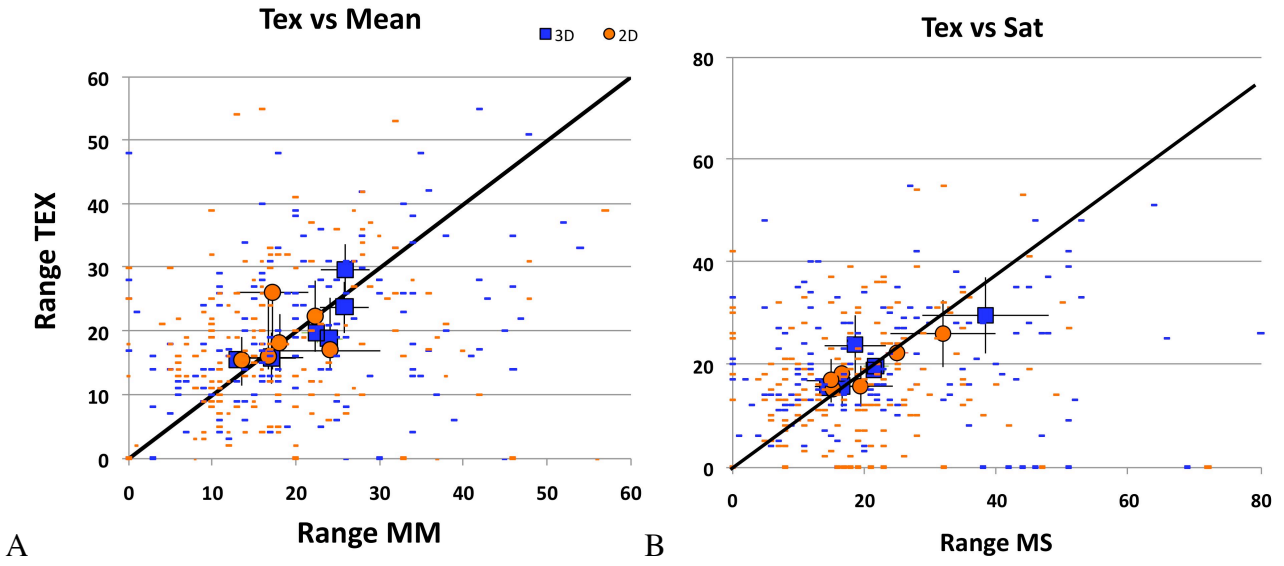

Figure 6.29 – Comparison between the mean memory colour range for the texture chromatic condition (TEX range) and A) the mean uniform condition (MM) or B) the most saturated uniform condition (MS). Blue squares: 3D. Orange filled circles: 2D. Each small point represents one subject's selection for one object for 3D shapes (blue point) and 2D shapes (orange point); 3*28=84 points. Each large point represents the mean memory colour of one object over all subjects for 3D shapes (blue squares) and 2D shapes (orange filled circles). Black line: unity line. Error bars denote one standard deviation.

While qualitatively there is a trend in the effect of texture compared with the MM or MS, is this quantitatively significant? The mean memory colour range of TEX condition (μ_{TEX}=19.86) is 6.45% smaller than MS condition (μ_{MS}=21.14) and 0.55% smaller than the MM condition (μ_{MM}=19.97). Then, a 4-way factorial repeated-measures ANOVA was performed on all the condition under one light (D65). Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using Greenhouse-Geisser estimates. The ANOVA reported no significant main effect of chromatic factor on mean memory colour range ($F(1.49, 40.33)=2.93$, $p=0.079$). Nevertheless contrasts revealed that the mean memory colour range of TEX condition was significantly smaller than MS condition ($F(1,27)=8.441$, $p<0.05$), but not between MM and MS ($F(1,27)= 3.207$, $p=0.085$).

To clarify these effects, let's further examine this case calculating the global memory colour range (GMCR), i.e. we first compute the mean of the each subject's selection and then calculate the differnce between the maximum and minimum mean across all subjects. The global memory colour range can be expressed by the following equation:

$$GMCR_j = \max\left(\sum_{k=1}^{K}\Delta\alpha_{jk}\right) - \min\left(\sum_{k=1}^{K}\Delta\alpha_{jk}\right) \qquad \text{Equation 6-6}$$

where $\Delta\alpha_{kj}$ is the angular difference between the angle selected in the trial *k* by the subject *j* and the original distribution's angle, K is the number of trial (K=10), and max and min are the maximum and minimum values of subjects' means. Results in Figure 6.30 show that the global memory colour range is significantly smaller for texture condition than uniform conditions (paired T-test, p_{MM}=0.0348 and p_{MS}=0.02). Thus, globally the range of choices for chromatically variegate stimuli is smaller than uniform stimuli, but this effect is weak with-in subject.

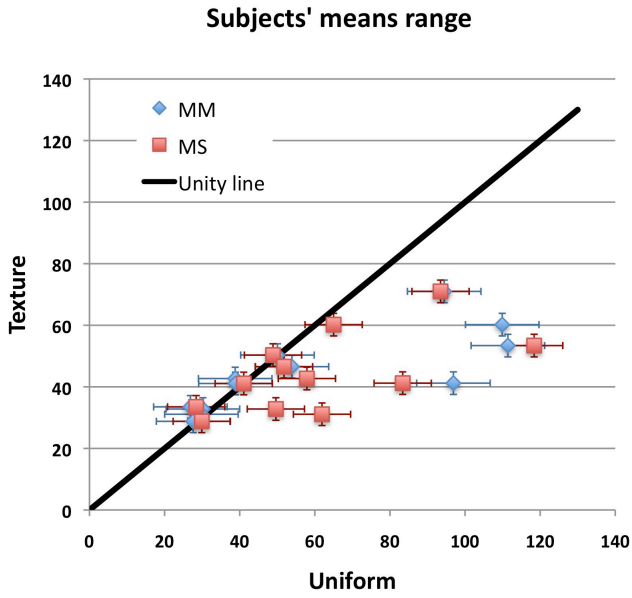

Figure 6.30 – Global memory colour range (GMCR) as a function of chromatic conditions. Blue diamond: texture-related angles plotted respect the MM angles. Red square: texture-related angles plotted respect the MS angles. Coloured error bars denote one standard error of the mean (SEM). Black line: unity line.

Figure 6.31 shows the mean memory colour ranges for all objects as a function of shape dimensionality, split by naturalness of the shape. Each point represents one object in one contour/chromatic condition combination. All objects follow the same behaviour: when presented in a generic shape, mean memory colour ranges are larger for 3D condition (average over all combination μ_{3D}=23.47) than for the 2D (μ_{2D}=20.32), this difference being significant (3-way ANOVA $F(1,27)=16.76$, $p<0.05$).

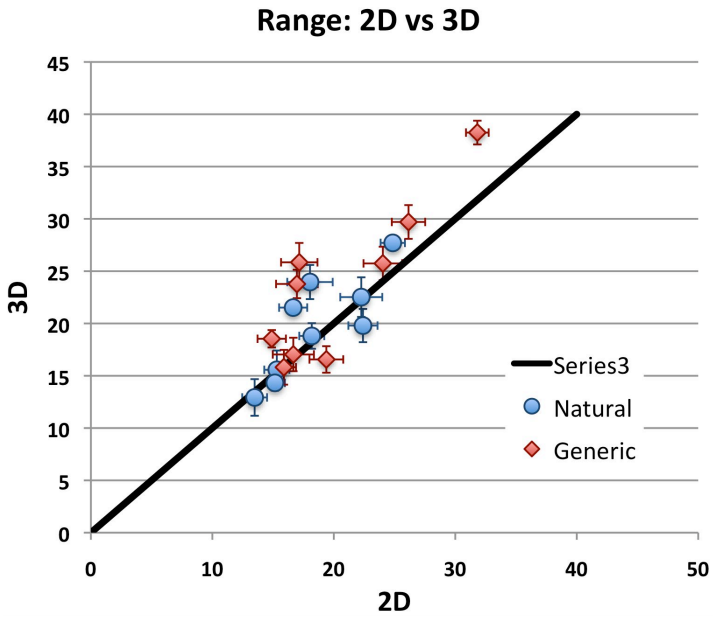

Figure 6.31 – Effect of shape dimensionality: mean memory colour range comparison between 2D and 3D conditions, over all combinations. Each point represents one object/colour/shape-familiarity combination. Blue closed circle: natural shape. Red diamonds: generic shape. Black line: unity line. Error bars denote one standard error of the mean.

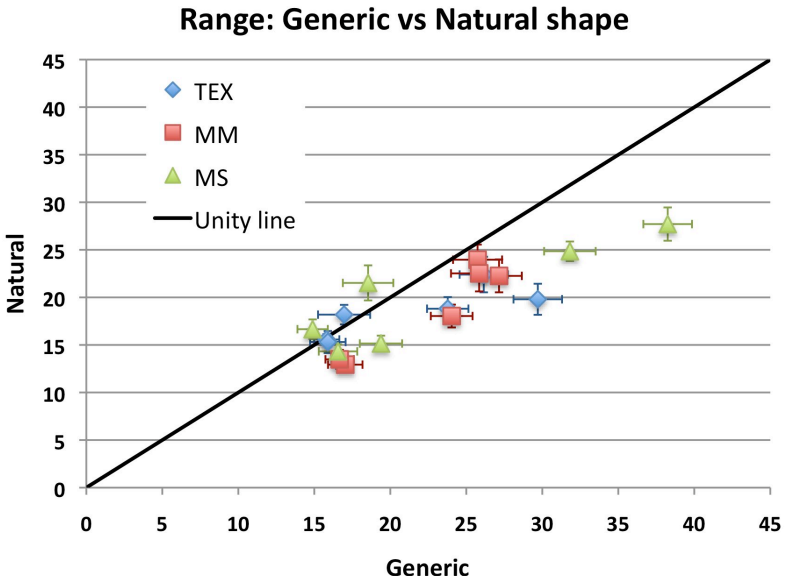

Figure 6.32 – Effect of shape diagnosticity: mean memory range of choices comparison between generic and natural configuration, over all combinations. Each point represents one object/colour/dimensionality combination. Blue diamonds: Texture, in 2D and 3D. Red squares: Uniform mean colour, in 2D and 3D. Green triangles: Uniform most saturated colour, in 2D and 3D. Black line: unity line. Error bars denote one standard error of the mean.

Conversely, there was no significant main effect of shape dimensionality on mean memory colour for the natural shape ($F(1,27)=1.174$, $p=0.288$). This comparison explores the effect of the dimensionality factor (shape dimensionality comparison).

Finally, the mean memory colour range is smaller for natural conditions compared to the generic condition, this being statistically significant for all three objects, $F(1,27)=23.96$, $p<0.05$. Each point represents one object in one contour/chromatic condition combination. Figure 6.32 illustrates such results; the average of mean memory colour ranges for the natural condition across all object/contour/chromatic condition combinations is $\mu_{Nat}=18.74$ and for the generic condition is $\mu_{Gen}=21.90$. This comparison explores the effect of naturalness and diagnosticity of the object shape.

6.3.2 EFFECT OF CHANGE IN ILLUMINATION

In this section I examine the effect of the three factors (chromatic, contour, and dimensionality) on memory colour difference under the second illuminant, CWF. The same considerations stated in section 6.3.1.1 can be repeated for CWF. Paired T-test on dominant angle shows no significance difference with D65 ($p=0.45$), and the frequency of selections of the dominant angle for the CWF was compared with respect to D65, showing the same trend for all factor (the effects of factors on frequency of the selection were found significant; $p<0.05$); nevertheless, a lower number of selections was found under CWF respect D65 (T-test, $p<0.01$).

Hence, this section only focuses on the degree of memory colour accuracy and precision (i.e. absolute memory colour angular deviation and mean memory colour range) and then on the

magnitude of the shift as function of the illuminant change (i.e. the colour constancy). The results demonstrate that the effect of the three factors on memory colour accuracy and precision under the second illuminant (CWF) trace the same pattern as under the reference illuminant, D65 (section 6.3.1). However the size of the effect vary between illuminant. Finally I will discuss the effects of the three factors on colour constancy.

6.3.2.1.1 Memory colour under illuminant changes

Figure 6.33 shows, for CWF, the effect of chromatic conditions on: (a) absolute memory colour angular deviation (MCD, in Figure 6.33A) and (b) mean memory colour range (MCR, in Figure 6.33B). The pattern of the chromatic main effect on MCD and MCR is similar between D65 (see respectively Figure 6.25 and Figure 6.29) and CWF (Figure 6.33A and B), this effect being significant (4-way ANOVA, Greenhouse-Geisser correction; $F_{MCD}(1.56, 42.37)= 31.648$, $p_{MCD}<0.001$; $F_{MCR}(1.95, 52.82)=7.28$, $p_{MCR}<0.005$). Subsequently, a 5-way repeated measures ANOVA was performed on MCD and MCR considering "illumination" as 5th factor. Results showed an effect of illumination on absolute memory colour angular deviation ($(F(1,27)=11.59$, $p<0.05$) and mean memory colour range ($F(1,27)=64.483$, $p<0.05$). Moreover, the size of the effect is on average different. In fact, the mean difference between absolute memory colour deviations of texture condition and uniform condition is 2.41 times larger for CWF than D65 (the means of absolute memory colour deviations for CWF per chromatic condition are: $\mu_{TEX}=10.83$, $\mu_{MM}=15.57$, and $\mu_{MS}=17.62$; while for D65 are, see 6.3.1.4, $\mu_{TEX}=10.22$, $\mu_{MM}=13.46$, and $\mu_{MS}=11.77$). The ANOVA reported a significant interaction on absolute memory colour deviation (MCD) between chromatic factor and illuminant ($F(1, 27)=20.23$, $p<0.01$)

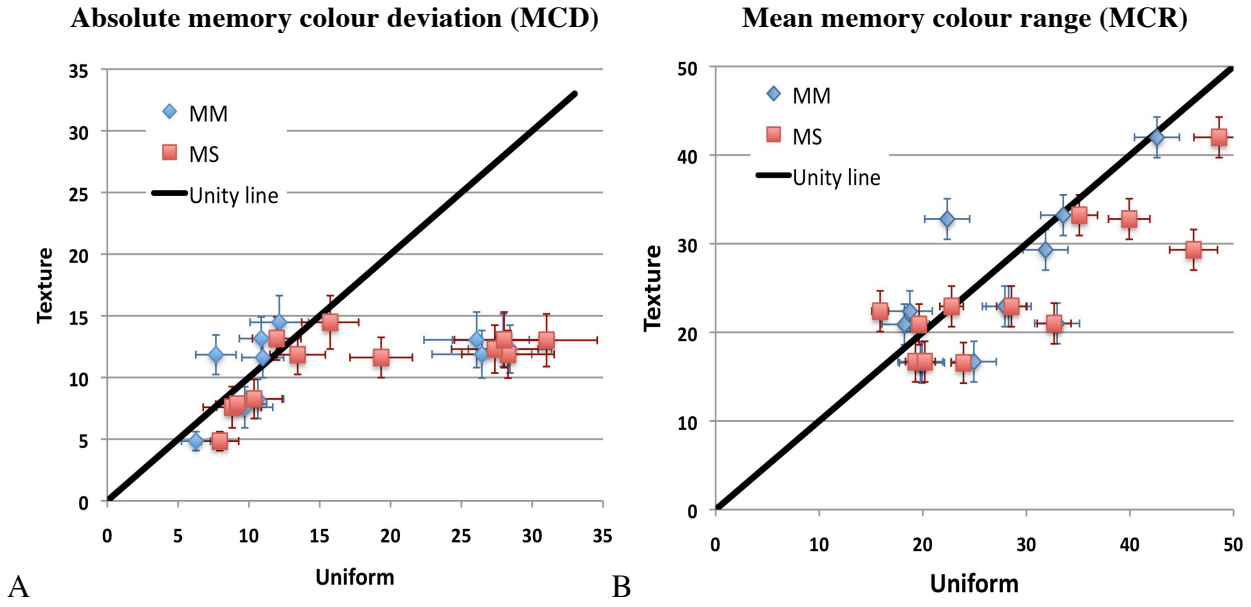

Figure 6.33 – Effect of the chromatic factor on memory colour under CWF. A) Absolute memory colour deviation (MCD). B) Mean memory colour range (MCR). Blue diamond: each point represents one object/shape factor combination in the mean uniform condition. Red square: each point represents one object/shape factor combination in the most saturated uniform condition. Error bars: s.e.m. Black line: unity line.

Conversely ANOVA reported no significant interaction on mean memory colour range (MCR) between chromatic factor and illuminant, ($F(1.43, 38.67)=0.809$, $p=0.416$); therefore there is no significant difference between size of chromatic effects for the two illuminants on mean memory colour range (i.e. they have similar means). Figure 6.34A and B illustrate the last statement comparing the mean memory colour range of texture versus mean uniform and most saturated uniform condition respectively; the dashed lines represent the linear interpolations of the data from the two illuminations. Clearly the two lines overlap indicating the same effect size for the two lights.

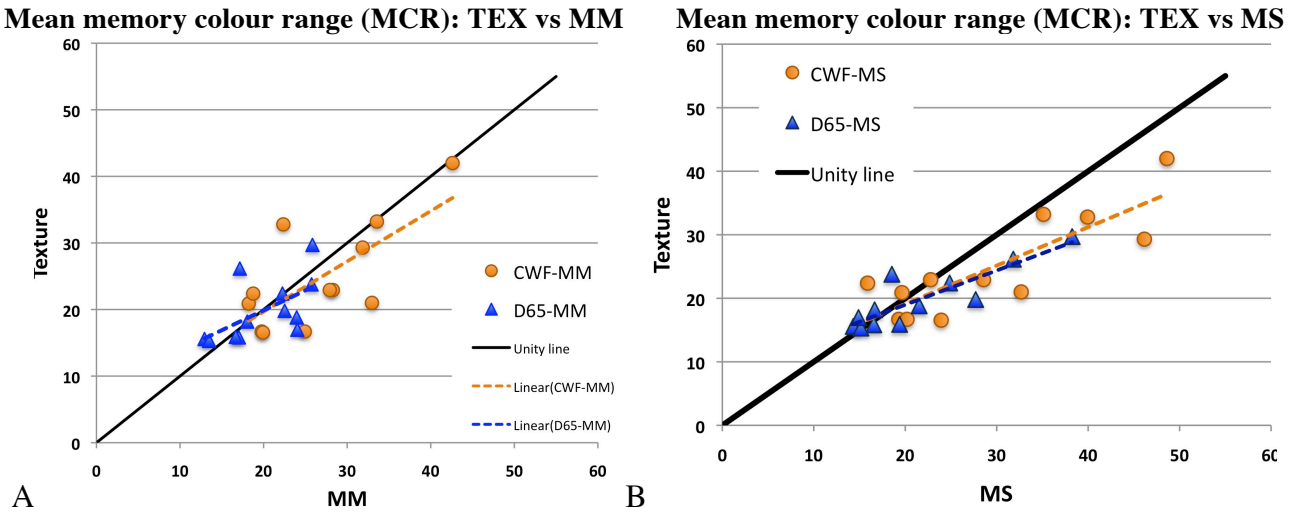

Figure 6.34 – Comparison between the effects of the chromatic factor on mean memory colour range under D65 and CWF. The graph plots the texture condition versus: A) each point represents one object/shape factor combination in mean uniform colour condition; B) each point represents one object/shape factor combination in most saturated uniform colour condition. Orange closed circles: CWF. Blue triangles: D65. Dashed lines: fitted line of the D65 (blue) or CWF (orange) data. Black line: unity line.

Figure 6.35A and B examine the results for the contour factor (shape diagnosticity) on MCD and MCR for CWF; again we can find consistency with the results obtained for D65 (see Figure 6.27 and Figure 6.32 respectively). The ANOVA reported this effect as significant (respectively $F_{MCD}(1,27)=6.18$, $p_{MCD}<0.05$; $F_{MCR}(1,27)=26.13$, $p_{MCR}<0.0001$). However there is no significant interaction between contour factor (shape diagnosticity) and illuminant on mean memory colour range ($F(1,27)=0.236$, $p=0.631$). Nevertheless, Figure 6.35C and D illustrate that there is a difference between the sizes of the contour effect on MCA and MCR for the two illuminations. Figure 6.35C and D plot, on the same graphs, the results for D65 and CWF as a function of contour conditions and the two lines fitting the data of the two illumination (orange=CWF, blue=D65). Note that they deviate from each other, indicating a difference in the magnitude of the effect on MCD (and on MCR) between the two illuminants; specifically the contour effects are lower for CWF and the the difference in the size of the effects on MCD is significant ($p=0.046$).

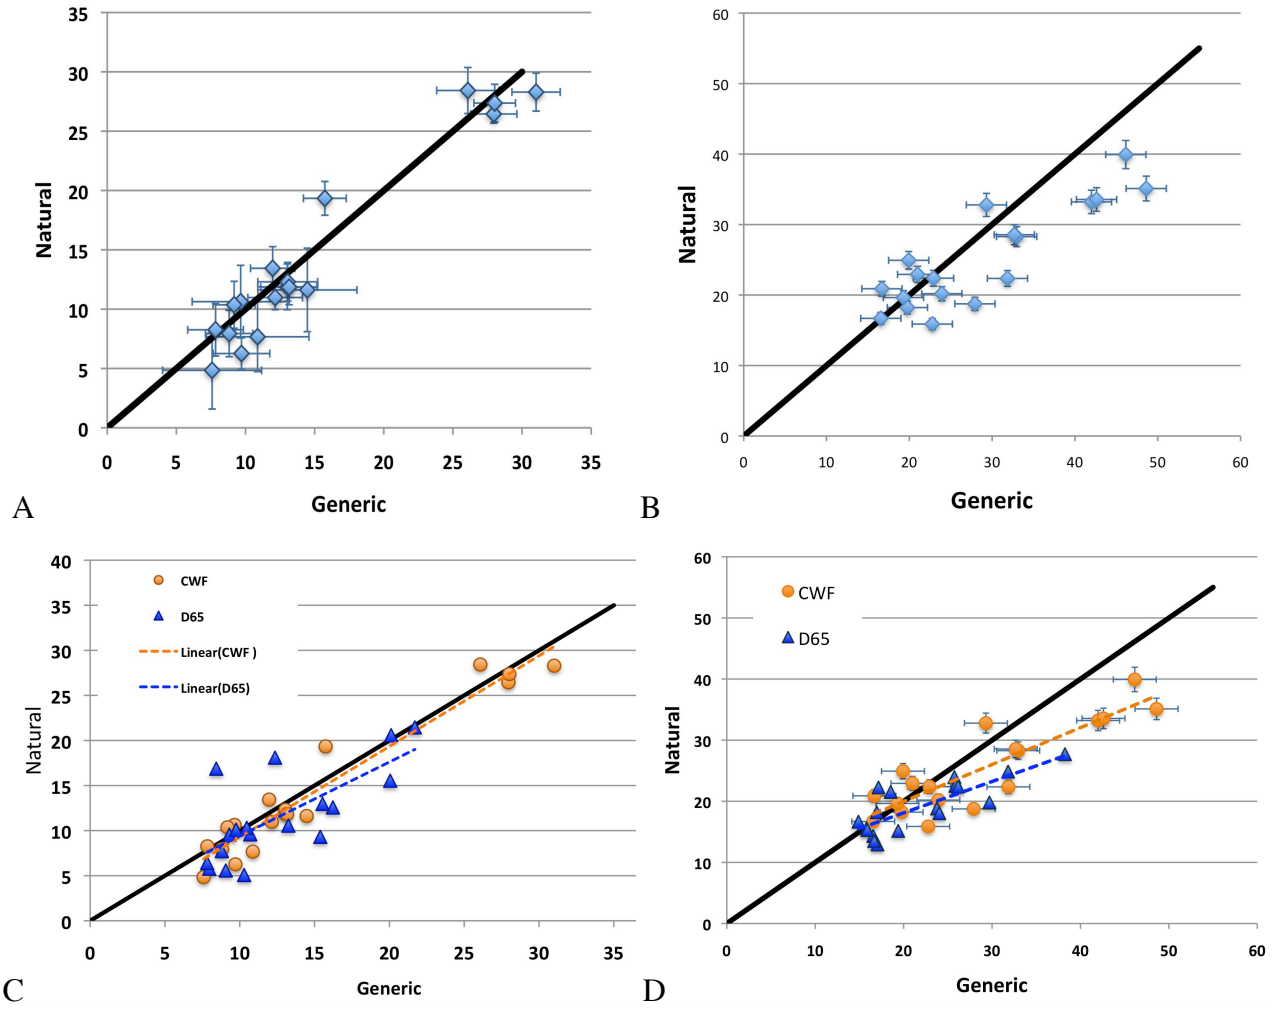

Figure 6.35 – A) Effect of the contour factor (shape diagnosticity) on absolute memory colour deviation (MCD) under CWF. B) Effect of the contour factor on mean memory colour range (MCR) under CWF. C) Comparison of the effects of the contour factor on absolute memory colour deviation (MCD) under D65 and CWF. D) Comparison of the effects of the contour factor on mean memory colour range (MCR) under D65 and CWF. Each point represents one object/condition combination. Orange closed circles: CWF. Blue triangles: D65. Dashed lines: linear fitting of the data under D65 (blue) or CWF (orange). Error bars: s.e.m. Black line: unity line.

Similarly, Figure 6.36A and B plots the absolute memory colour deviation (A) and mean memory colour range (B) as a function of dimensionality factor for CWF. Yet again the results for CWF trace the same trend of D65 (see respectively Figure 6.26 Figure 6.31).However the size of the effect is less strong in CWF (Figure 6.36C and D) than D65, this difference being significant (interaction illumination and dimensionality, $F_{MCD}(1,27)=9.521$, $p_{MCD}<0.01$; $F_{MCR}(1,27)=18.64$, $p_{MCR}<0.001$). Therefore, it is possible to assert that dimension induce a different bias (i.e. size of the effect) on absolute memory colour deviation and mean memory colour range depending on the illumination.

Absolute memory colour deviation (MCD) **Mean memory colour range (MCR)**

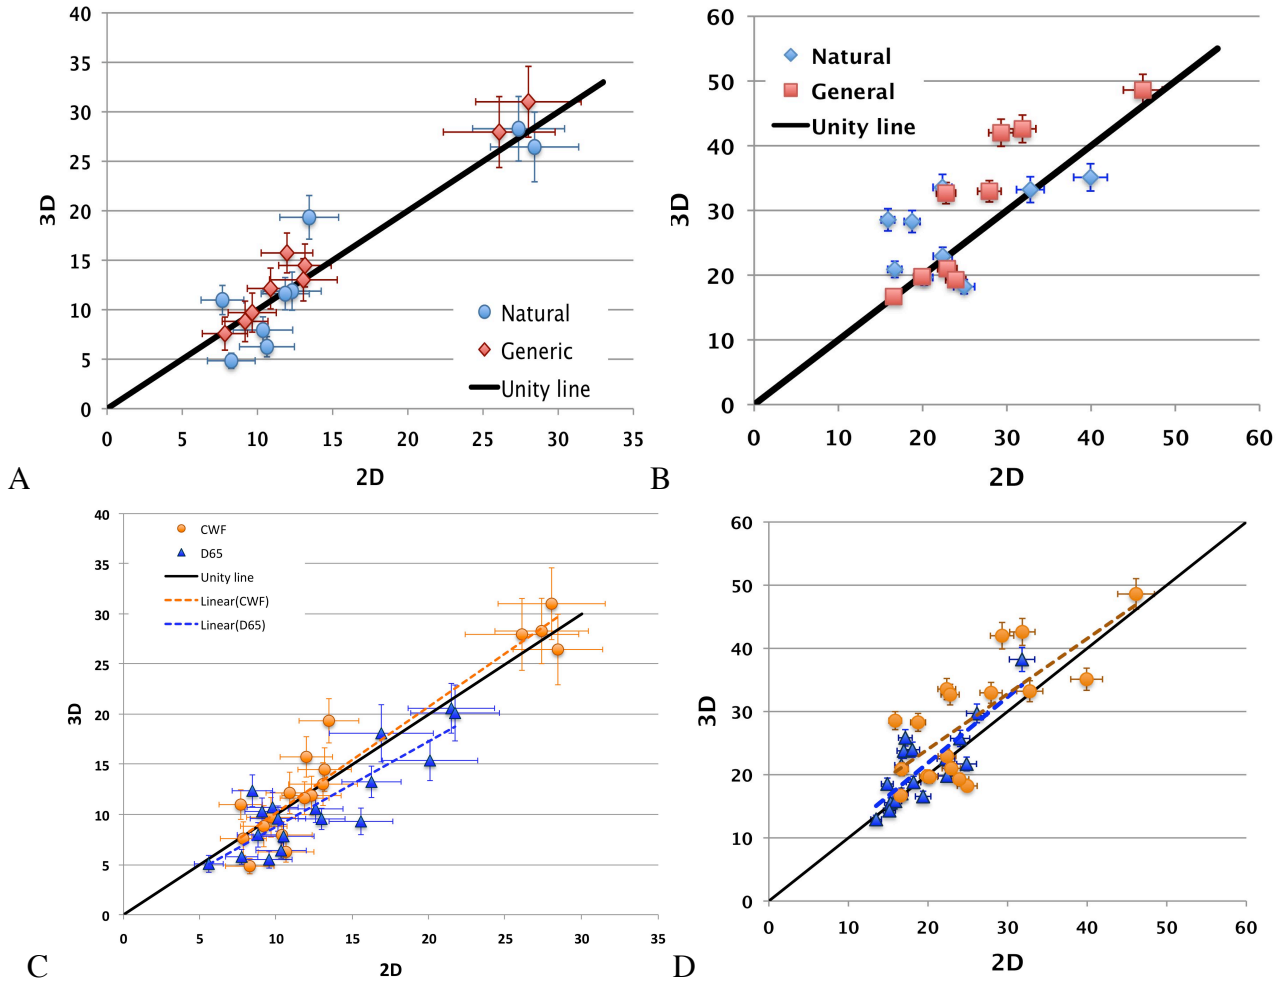

Figure 6.36 – A) Effect of the dimensionality factor on absolute memory colour deviation for CWF. B) Effect of the dimensionality factor on mean memory colour range for CWF. C) Comparison of the effects of the dimensionality factor on absolute memory colour deviation for D65 and CWF. D) Comparison of the effects of the dimensionality factor on mean memory colour range for D65 and CWF. Orange closed circles: CWF. Blue triangles: D65. Dashed lines: linear fitting of the data for D65 (blue) or CWF (orange). Error bars: s.e.m. Black line: unity line.

6.3.2.2 COLOUR CONSTANCY

Figure 6.37 illustrates, as an example, the angles selected for the apple in its natural 3D shape by all subjects under the two illuminants and shows, for each of them, the mean vector of the original apple's chromatic distribution. Because the observers had one-dimension of freedom, that is, the angle of rotation, they would have been perfectly colour constant if they chose exactly the original angle of the distribution under CWF as correspondent to the shift of the distribution they saw under D65. Any deviation from this angle is the error that the subject made under the second light, that is, the deviation from the perfect colour constancy. Thus, the constancy of a subject perception can be computed as the difference between the original angle and the mean of the subjective choices under CWF. In addition, as described in previous chapters (Chapter 4), when memory (of a colour or of an object's colour) is involved in a colour constancy task, then a memory factor must be introduced. In such a case, perfect constancy would be achieved if the subject's memory shift from the typical angle of the distribution under CWF were equal to the shift from the typical angle of the distribution under D65 (same bias under different lights). Therefore, the colour

constancy index in this experiment can be calculated from the mean of difference between the subject's mean selections under the two lights.

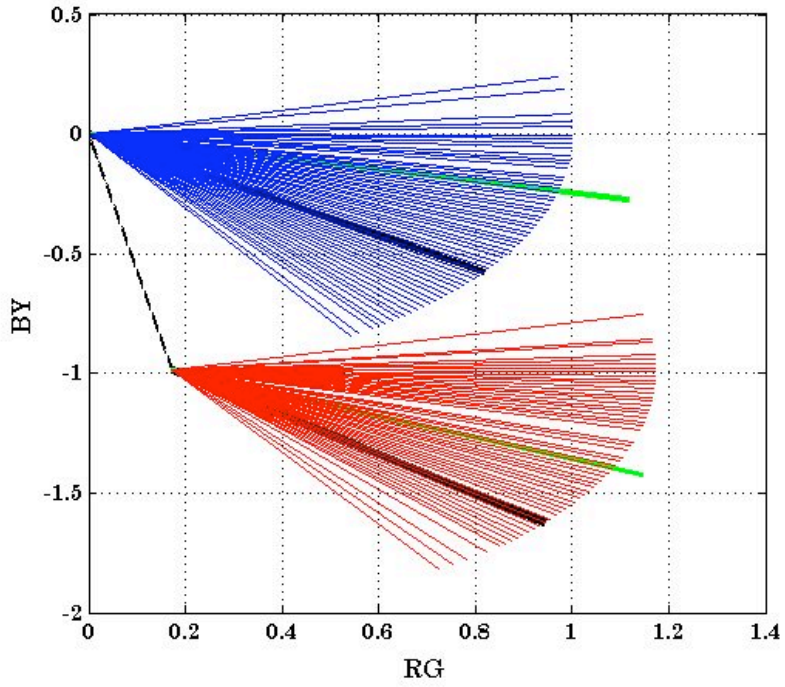

Figure 6.37 – Plot of 10 subjects' selections for Apple in its natural 3D shape for the MS case under standard D65 (blue lines) and CWF (red lines) in RG-BY colour space, with D65 as neutral point. For each illuminant, the black line is the mean vector of the original object picture distribution, while the green line is the mean of the subjects' choices. Grey dashed line: distance between the illuminants neutral point. Note that (0,0) is the neutral point for D65 and the red dot is the chromaticity of the CWF illuminant, origin of the rotation of the distribution in CWF.

Let's define as relative angle α_K the difference between the mean of the subject's selected angles in a generic condition under the K illumination and the original object's angle also under K. Then, the Colour Constancy Error (CCE), for each subject *j*, can be described by the following formula:

$$CCE = 1 - \frac{\sum_{i=1}^{n} abs\left(\alpha_i|_{CWF} - \alpha_i|_{D65}\right)}{\frac{N}{180}} \qquad \text{Equation 6-7}$$

where N is the number of observers (N=28), and α_{ilD65} and α_{ilCWF} are the means of the relative angles selected by the subject i in a certain condition (green line in Figure 6.37), respectively for illuminant D65 and CWF.

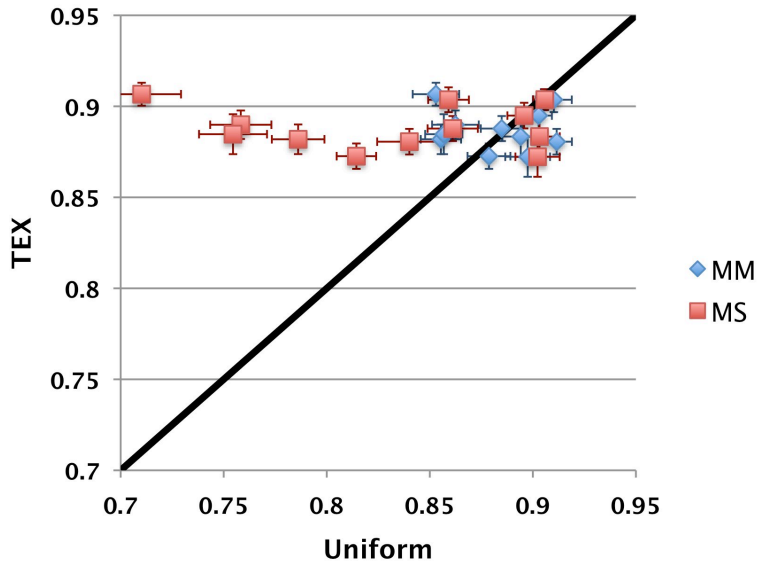

Figure 6.38 - Scatter plot of the mean colour constancy error (N=28) for the three objects as a function of their chromatic conditions for all shape cue combinations. Blue diamond: texture versus mean uniform colour. Red squares: texture versus most saturated uniform colour. Error bars denote the standard error of the mean. Black line: unity line.

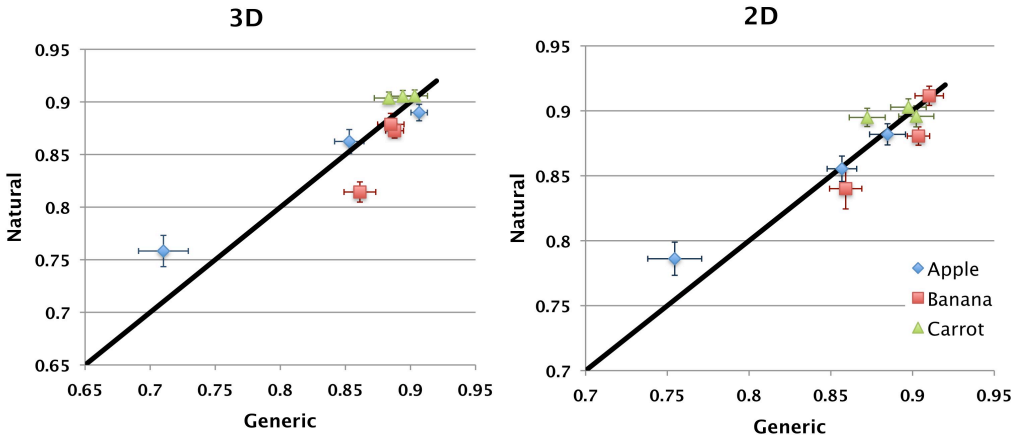

Figure 6.39 – Scatter plot of the mean colour constancy error (N=28) for the three objects as a function of their shape, natural versus generic, in all chromatic conditions for 3D (A) and 2D (B) combinations. Coloured error bars denote the standard error of the mean. Black line: unity line.

Figure 6.38 shows that the chromatic texture condition has in general better constancy compared to the uniform cases, the difference being significant ($F(1.625, 43.881)= 35.418$, $p<0.001$). Instead, there is no statistically significant difference between natural and generic shape ($F(1,27)=0.008$, $p=0.928$) or between 2D and 3D conditions ($F(1,27)=0.533$, $p=.472$), as shown respectively in Figure 6.39 and Figure 6.40.

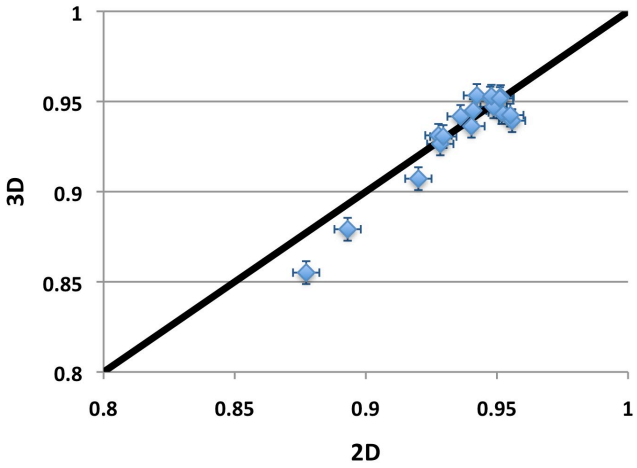

Figure 6.40 - Scatter plot of the mean colour constancy error (N=28) for the three objects as a function of dimensionality, 2D versus 3D, in all chromatic/shape familiarity combinations. Coloured error bars denote the standard error of the mean. Black line: unity line.

For completeness' sake, the colour constancy index was also computed using the dominant angle instead of the mean angle; no difference in the results was found (p=0.824).

6.4 DISCUSSION

To date and to my knowledge, the study presented in this chapter is the first that: (1) examines the effect of presenting a solid 3D chromatic texture on memory colour and colour constancy; (2) allows free adjustment of the stimulus colour by the observer; (3) completely excludes the shading information from the 2D stimuli and controls separately for the effect of chromatic texture diagnosticity and shape diagnosticity. In fact, as described in the introduction (section 6.1) several studies have: (1) examined the effect of memory on colour appearance using simple flat homogeneously coloured patches (e.g. Jin&Shevell (1996)); (2) examined subjects ability to recall the colour of an object by selecting it from a set of samples (Pérez-Carpinell et al. (1998)); or (3), even if more complex stimuli and more freedom were given, as in Olkkonen et al. (2008) (see introduction 6.1), presented the stimulus on a display and the polychromatic control stimulus was not directly comparable with the original chromatic texture.

To investigate whether or not the observers' memory colour is affected by the different visual cues to the object identity, separate analysis of features of the memory of an object's colour were examined: the value and frequency of the dominant angle, the mean memory colour, the absolute memory colour angular deviation (from the object's typical hue angle), and the memory colour range (of the selection distribution). These analyses were performed for both illuminants. In addition, a colour constancy index was computed comparing the mean of the subjects' selection (or the dominant angle, showing equivalent results).

Each distinct factor of object identity (surface chromatic texture, shape diagnosticity and dimensionality) and its effect on colour appearance is described in the following sections.

6.4.1 EFFECT OF CHROMATIC INFORMATION ON MEMORY COLOUR

Here the aim is to explore the effect of different representation of the chromatic proprieties of the familiar object's surface, namely: (1) full chromatic texture (TEX); (2) mean colour only (mean uniform colour, MM); and (3) most saturated colour only (most saturated uniform, MS).

The results showed no significant effect of chromatic conditions on mean memory colour. The main explanation consisted in that the negative and positive deviations from the original angle were averaged, with loss of the real difference from the typical setting. Thus the best way to interpret the results was to evaluate the absolute deviation of the subject mean setting from the object's original hue angle, i.e the absolute memory colour angular deviation (MCD, section 6.3.1.4). Keep in mind that the smaller such deviation, the more accurate is the subjects' selection. Since perfect accuracy implies a MCD=0 (see 6.3.1.4) and complete inaccuracy a MCD=180, then it is possible to define a memory colour accuracy (MCA) index ranging from 0 to 1, as follows:

$$MCA = \frac{180 - MCD}{180}$$ Equation 6-8

The pattern drawn by the MCD (consequently MCA) showed that the subjects were significantly more accurate for polychromatic surfaces than uniform surfaces. In support of this finding are the results on the dominant angle (6.3.1.2). As a matter of fact, the dominant angle for the texture condition is closer to the object's typical colour setting than for the homogeneous conditions, also implying higher accuracy. It should be noticed that the effect of the chromatic factor on dominant angle is stronger for natural shapes that generic, implying an interaction between the congruency of shape and texture. A reason of this difference in size of the effects could be that the object natural contour stabilizes the chromatic factor shift, i.e. increase the confidence in the subjects' choices. Conversely, when no information on the object shape is given (generic contour), the subjects' confidence diminishes.

On the other hand, the observers chose more frequently the dominant angle in the texture condition than in the uniform cases, denoting better global precision in the texture condition (section 6.3.1.3). This finding is supported by the results on global memory colour range (GMCR); in fact it is smaller for the texture condition than for uniform conditions (section 6.3.1.5). In addition, the within subject precision, measured by the mean memory colour range (section 6.3.1.5), is better for texture trials than uniform trails, but this difference is only significant in the most saturated uniform case. Therefore we can say the texture improves global precision compared to uniform stimuli. This behaviour is reasonable if we think that the colours at the edge of the distribution move further than its mean or most saturated chromaticity for each degree of rotation,

so the variations are more visible. However, this could be one of the possible explanations, and an additional hypothesis might be that having more information on the object's colour distribution stabilizes memory colour in such that globally the subjects are more precise in their choices, therefore the smaller range. This would be in line with the previous arguments concerning the better dominant angle consistency of selections for texture with respect to uniform cases, and smaller absolute angle mean deviation. Nevertheless, it is plausible that the possible effects of texture on mean memory colour range are too small compared to other factors to be fully evaluated with this procedure.

In conclusion, memory colour is stronger and more consistent for textured than homogenous stimuli. These results are also robust for the second illuminant, indicating that the effect of chromatic texture is illuminant-independent (see 6.4.5).

Based on these conclusions further study on memory colour should take into account the effect of the complexity of the chromatic stimuli. As a matter of fact, the natural world is not made of homogeneous coloured surfaces, but, on the contrary, is quite variegated in colour, both in chromaticity and luminance. Thus, studies on memory of a familiar natural object, e.g. a banana, cannot disregard the texture of its surface. This study proves that a simple match to a uniform surface is not sufficient to evaluate the real value of memory colour.

Finally the results prompt a question: does this effect depend on the multiple colour surfaces forming the texture or on the specific texture that identifies the object? In other words, is it due to the familiarity/recognisability of the texture or only to the presence of many colours?

6.4.2 EFFECT OF SHAPE DIAGNOSTICITY ON MEMORY COLOUR

As mentioned above, an original difference of this study with respect to previous works consisted in controlling shape independently from the other parameters of the study. For example, previous studies (Olkkonen et al. (2008), Hansen et al. (2008) or Hansen&Gegenfurtner (2006)), varied the shape of the chromatic stimuli but also altered the texture of the control shape relative to the original texture. In other words, instead of using the same base texture for the congruent and incongruent shape, they used, depending on the study, random noise, pink noise or brown noise texture on the incongruent shape, and natural texture on the congruent shape. Despite the fact that this is an adequate way to control for texture diagnosticity, which was one of the aims of those experiments, it does not give the same chromatic information while varying the shape. In addition, the textured objects were presented on a display keeping the shading from the original image intact so to give a 3D effect. The intent here, instead, is to evaluate the overall difference on memory due solely to different shape cues, i.e generic/natural contour shape and 2D/3D geometry, maintaining

the same chromatic conditions. In other words, because it is possible to pair a natural or generic contour presentation with all chromatic conditions and, as well as, with both 2D and 3D configuration, then shape diagnosticity can be treated as an independent factor.

Observers showed better accuracy, measured by the absolute memory colour angular deviation (see Equation 6-8), when the colour stimuli were presented in their natural contour than in the generic. Moreover, the dominant angle is closer to the object typical hue angle for object-congruent shape (natural contour) than rhomboid (generic contour), while the frequency of selection at the dominant angle (section 6.3.1.3) is higher. So it is possible to assert that, when the colour stimuli is presented in the appropriate object-contour combination, the entire population of subjects recalled more consistently a selected set of angles which were closer to the typical object's hue angle. In conclusion, object-shape congruency consistently increases memory colour accuracy and global precision.

Furthermore, the mean memory colour range of natural contours is smaller than the generic. The explanation for this result could be that: (1) the range of colour narrowed because the subjects are more certain of the object's typical hue angle (i.e. its colour or colour distribution); or (2) the natural shape interferes with the memory of the object colour and impairs the observers' ability to access the wide range of memory colours of the object, leading to a smaller range than the one actually available. Clearly, the second reason seems in contradiction with the results in Chapter 4, in which we stated that congruent shape improves access to memory colour range and also that the *direction* of the effect in this experiment and the main experiment in Chapter 4 are opposite. Keep in mind that we should consider prudently the relation between the two experiments. Firstly for the type of task performed, here we asked the subjects to *adjust* the colour of the object to its *natural* colour appearance varying a one-dimensional parameter (the hue angle), while in the previous experiment the observer had to *select* if *any* colour presented was *appropriate* for the specific object among a range. Thus, in this experiment we give the subject control of his choice to determine the most typical colour in his/her memory of the object, while the 2AFC in Chapter 4 they are constrained by the author's selected stimuli to select the appropriate colours not the most typical. Secondly, the experiment in Chapter 4 shows to the subject a two dimensional range of choice equispaced in Luv colour space, while in this experiment they have a one-dimensional setting parameter equispaced in EMG cone contrast space. Thirdly, only one object (the banana) is the same as used before.

Evidently, in the experiment described in this chapter, the observer searched for narrow range of typical hues, while in the previous experiment s/he picked all colours s/he judged appropriate for the object. Based on this assumption, in the first case, improved access to memory implies smaller range, while in the second case a larger range. This explanation not only is in line

with our previous conclusions but also of other studies (e.g. Olkkonen et al. (2008)) that assert that congruent shape helps memory colour. In the end, it is possible to propose that the range of colours narrowed because the subjects were more certain of the hue angle of the typical colour (or colour distribution) of the specific object (point (1) above).

In conclusion, shape diagnosticity (or object shape congruency) consistently increases the subject's memory colour accuracy and precision relative to the object's typical hue angle. These results are also robust for the second illuminant, indicating that the effect of shape diagnosticity is illuminant-independent (see 6.4.5).

6.4.3 EFFECT OF DIMENSIONALITY ON MEMORY COLOUR

To the author knowledge, the experiment reported here is the sole one investigating the effect on colour appearance of chromatically variegated surfaces of natural objects in a physically solid three-dimensional environment, that is, the stimuli is not a computer simulation, rendering the object's three-dimensional cues. This means that the observers are experiencing a more natural perception of the object than in the common colour experiments.

In a complex environment the results of the data analysis are more complicated, but the factorial ANOVA analysis allows the evaluation of the main effect of dimensionality as well as any interaction with other factors.

The principal conclusions from the results is that 3D presentation lead to a more accurate and consistent memory of the colour of the object. This can be proven initially considering the degree of memory colour accuracy for the 3D versus the 2D configuration. In fact, since the absolute memory colour angular deviation (section 6.3.1.4) from the object's typical hue angle is smaller for 3D that 2D configuration (Figure 6.26), then it is logical to conclude that the 3D conditions increase the observers' accuracy in the selection. In addition, this effect is stronger for natural shapes compared to generic, although the T-test found this difference no significant. Consequently it is possible to suggest an interaction between the shape diagnosticity and dimensionality. In fact, no difference in object identity exists between the stimulus displayed as a 3D rhomboid (or better a rhombic based pyramid) and its 2D form. On the contrary, in Chapter 4, Experiment 2 showed that solid shapes of an object might be more or less recognizable than their 2D contour. Thus, it is logical to assume that, since the size of the effect of dimensionality on absolute memory colour **angular deviation** is greater for natural objects than the generic, a 3D object encompasses more information about the object identity than its 2D match. *Therefore, I postulate that the introduction of three-dimensional cues enhances memory colour **accuracy**, and*

that the strength of the effect is relative to the naturalness and recognisability of the displayed stimulus.

To test this hypothesis, I use the results of the control experiment in Chapter 4 (section 4.3) to evaluate the difference in recognisability of 3D and 2D presentation of the three objects: apple, banana and "new" or "second" carrot. Table 6.7 lists the recognition-by-shape (RS) indices as defined in 4.3.2, showing that 3D apple, banana, and carrot were 1.48, 1.31 and 25.9 times, respectively, more identifiable than their 2D counterparts. It should be noticed that RS varies from 0 for the rhomboid in 2D or 3D (they equally do not give any information about the object identity through their shape), to 1 when the subject is 100% certain of identify the object from its shape configuration alone.

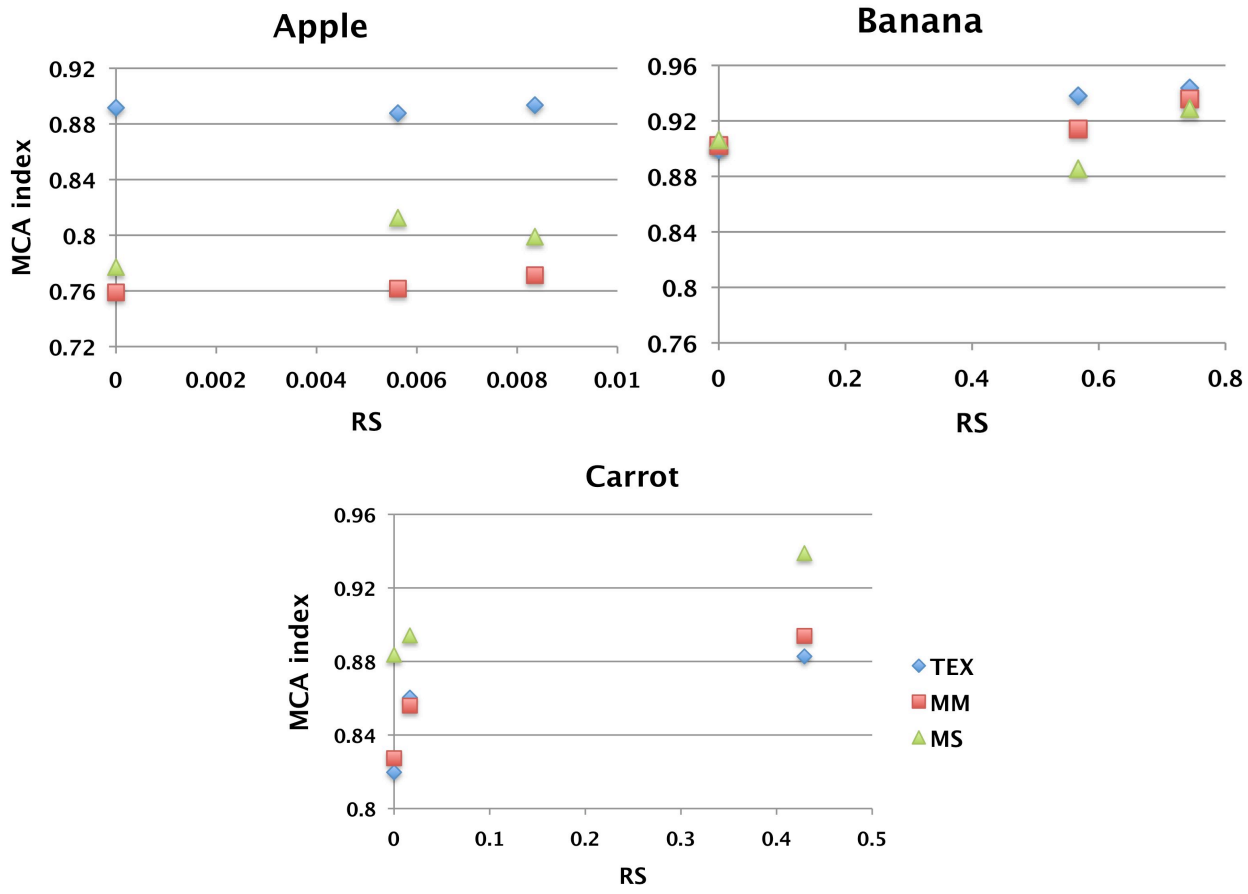

Figure 6.41 – Correlation between recognition-by-shape indices and memory colour accuracy (MCA) index. Each point represents one chromatic/shape recognisability (3D natural contour, 2D natural contour and 2D rhomboid) combination. Blue diamonds: Textured stimuli. Red square: mean uniform stimuli. Green triangle: most saturated uniform stimuli.

Consequently, if the above assumption is correct, the absolute memory colour angular deviation (MCD) should diminish with (i.e. be directly proportional to) the recognisability of the shape of the stimulus. Lets then plot the MCA index as a function of recognisability of the 2D and 3D shape and the 2D rhomboid (RS=0), for all three objects and chromatic conditions.

Objects	RS-3D	RS-2D
Apple	0.0084	0.0056

Banana	0.7435	0.5675
Carrot	0.4291	0.0165

Table 6.7 – List of the recognition-by-shape indices for 2D and 3D shapes.

The plots in Figure 6.41 show that generally the hypothesis is confirmed, although the apple in the most saturated condition does not follow the trend. These results reinforce the above theory, i.e. better object identification through shape implies better subject's accuracy in recalling and selecting an angle closer to the object's actual colour.

Finally, the results for mean memory colour range depict an interesting pattern. Since, the 3D shape configuration has a significant bigger mean memory colour range with respect to the 2D when presented in a generic shape (3-way ANOVA, $F(1,27)=16.76$), the observer is less precise in the 3D condition relative to the 2D condition. This implies that the introduction of three-dimensional cues confounds the perception of the observers, if there is no benefit in the identification of the object. On the other hand, no significant difference in range is detected for the natural shape, indicating that as long as the shape of the object is displayed the observer can reach the same precision and access equally the colour memory range for the specific object.

In conclusion, the introduction of three-dimensional cues: (1) improves the accuracy of memory colour for both for natural and generic shaped object; (2) increase the size of the accuracy effect in natural shaped objects; and (3) reduces the precision of memory colour if no reference to the object actual shape is given (generic contour), otherwise it has no effect. These results are also robust for the second illuminant, indicating that the effect of shape diagnosticity is illuminant-independent (see 6.4.5).

6.4.4 ANALYSIS OF FACTOR COMBINATIONS INFLUENCING MEMORY COLOUR ACCURACY

In the above subsections I have discussed separately the effect of the different variables on memory colour accuracy, precision and stability. Here I discuss the combination of these variables on memory colour accuracy and attempt to reconcile the results of Chapter 4 with the current chapter.

First, let us define as the discernible colours of an object *j*, the colours with a colour difference from the mean colour of the object *j* equal to or bigger of 1.2 CIEDE2000 units (Linhares et al. 2008). Then, any given sample *k* of the object chromatic surface distribution will contain a fraction of the total number ND of discernable colours of the original distribution. The value of this fraction measures the colour variability of the sample in proportion to the colour variability of the original distribution, and will be denominated polychromaticity index (PI). Therefore, a sample *k* of

a textured object *j*, will have a polychromaticity index of PI_{jk}. Then, the PI of a stimulus will range from 1/ND, for any uniformly coloured stimulus, to 1 for the most naturally textured stimulus (i.e. the texture of the entire object). Table 6.8 lists the PIs for any combination of objects/contour/chromatic conditions.

Stimulus	PI
Apple	1
Banana	1
Carrot	1
Rhomboid – Apple	0.9103
Rhomboid – Banana	0.6428
Rhomboid – Carrot	0.8287
Uniform – Apple	1.205E-05
Uniform – Banana	2.172E-05
Uniform – Carrot	2.312E-05

Table 6.8 – List of the polychromatic indexes for the complete texture presented in the natural shape, the sample of texture presented on the generic (rhomboid) shape, and the uniform conditions for all three objects.

As discussed before, subjects are more accurate (and precise) if the stimulus presented is textured. Therefore it is plausible that an increase of PI implies an increase in the accuracy of the mean. Figure 6.42 and Figure 6.43 show the correlations graphs between these to factors for 3D and 2D, respectively.

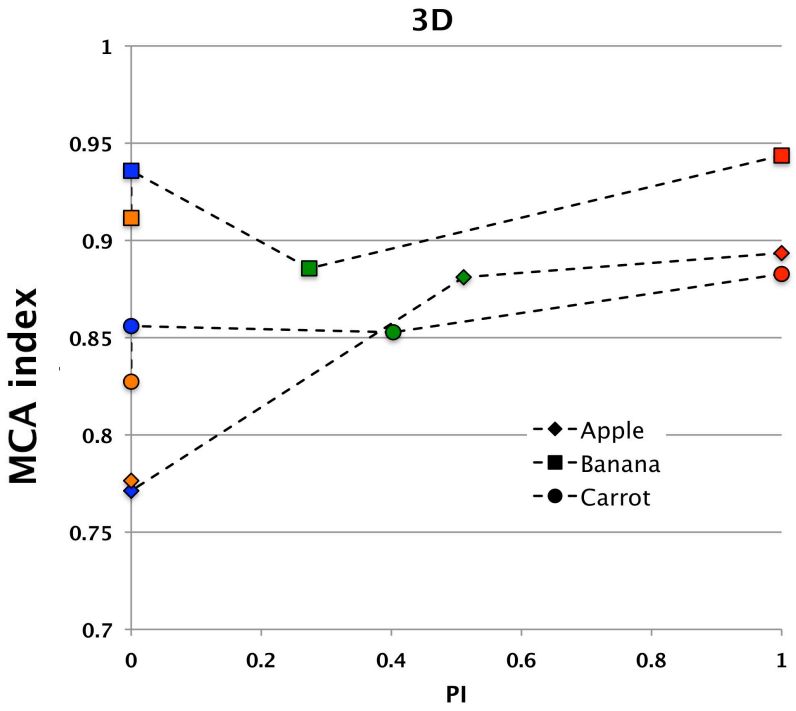

Figure 6.42 – Correlation between the memory colour accuracy (MCA) index and the polychromaticity index for 3D. Diamond: Apple. Square: Banana, Closed circle: Carrot. Red colour=TEX-3DN. Green colour=TEX-3DG. Blue colour=MM-3DN. Orange colour = MM-3DG.

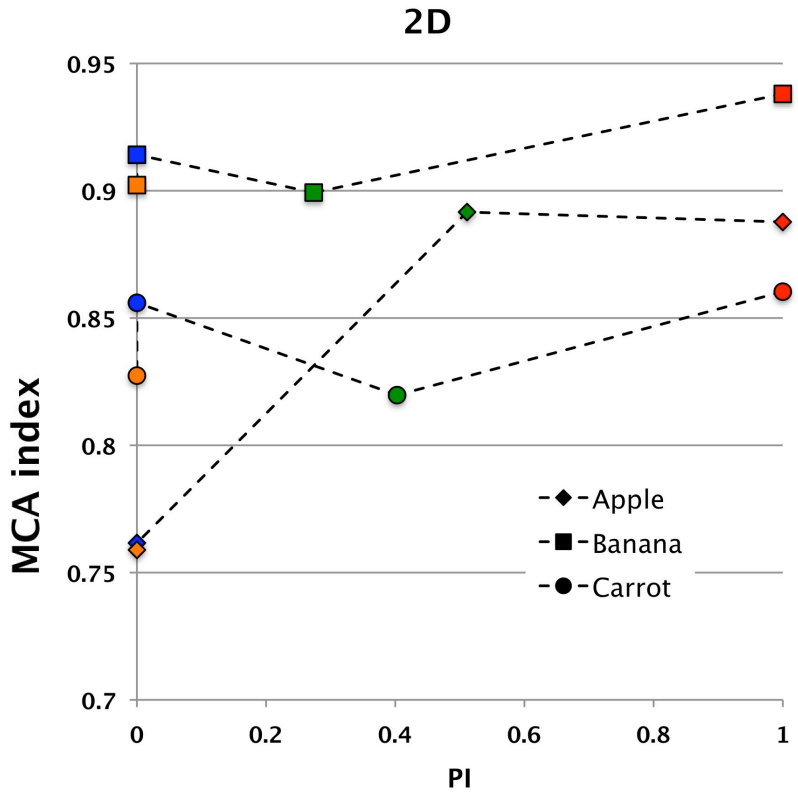

Figure 6.43 - Correlation between the memory colour accuracy (MCA) index and the polychromaticity index for 2D. Diamond: Apple. Square: Banana. Closed circle: Carrot. Red colour=TEX-2DN. Green colour=TEX-2DG. Blue colour=MM-2DN. Orange colour = MM-2DG.

The natural textured objects have the highest chromaticity variation followed by the generic textured objects and then the uniforms (natural or generic). The graphs show that natural textured objects (in red) have in fact better accuracy than all the other cases and textured rhomboids (in green) have better (or similar) accuracy than the uniformly coloured rhomboids (in orange). On the contrary, textured generic objects may or may not have better accuracy than uniformly coloured natural shaped stimuli depending on the object/dimensionality combination. Therefore PI is somewhat positively correlated to memory colour accuracy, in other words, more information about the original spread of the object colour is given better is the subject accuracy in the selection, but other factors interact with its contribution. For example, interactions with contour and dimensionality factor might be to some extent expected. In sections 6.4.2 and 6.4.3 I have proven a positive correlation between recognisability of the identity of the object from shape and memory colour accuracy. Therefore if I hypothesize that the effect of shape on accuracy is stronger than the effect of the chromatic factor, then the incoherence in the effect of PI is justified. More puzzling, instead, is the difference between different objects, prompting the question: "What differences among objects may lead to different memory colour accuracy?" In Chapter 4 and Chapter 5, as well in 6.2.3.1, I have shown that objects, belonging to different categories, possess different ranges of chromatic distributions (or object colour variegation, OV); for example the carrot has a much smaller range of chromaticies than the apple (Table 6.3 in 6.2.3.1). Furthermore the ranges of these distributions, or gamuts, are category dependent. In other words, the gamuts of two objects are

different if the objects belong to two diverse categories, but are similar if they are samples of the same category (if they are presented under the same illuminant). Therefore we can assume that in general a subject has a certain gamut for a specific object and the extent of this gamut varies among object category, but not among different samples of the same category. Thus, a reason for the difference in strength of the effect of PI on memory colour accuracy for different objects could reside in the difference between memory gamuts of the object category.

Based on these observations, I propose that the accuracy of the mean relative to natural distribution should depend on several factors linked to the naturalness of the stimulus and the original object chromatic variegation, i.e. the recognition-by-shape index (RS) and polychromaticity index (PI), and the object colour variegation (OV), as illustrated in the formula below:

$$MCI \propto f(PI_{jk}, RS_j, OV_j, PM_j) \qquad \text{Equation 6-9}$$

where PM is the subject baseline memory for that specific object or better his/her memory accuracy of the object colour without any object's identity cues. Herein PM will be called "pure memory" of the object. It should be notice that in this experiment PM is equal to the subject's memory colour accuracy in the uniform generic condition.

These factors can contribute alone (as a simple sum of variables for example), cooperate with other factors (products or ratios between factors) or be themselves function of other variables. PM might be an example of the latter; in fact, is logical to assume that bigger is the object variation worse is its pure memory. As shown in Figure 6.44 this conjecture is supported by the data. However, if we look at the MS-2D case (orange close circle), the banana shows a bigger accuracy than the carrot, which has a smaller distribution spread; thus the object colour variation is not the only factor influencing the basic accuracy.

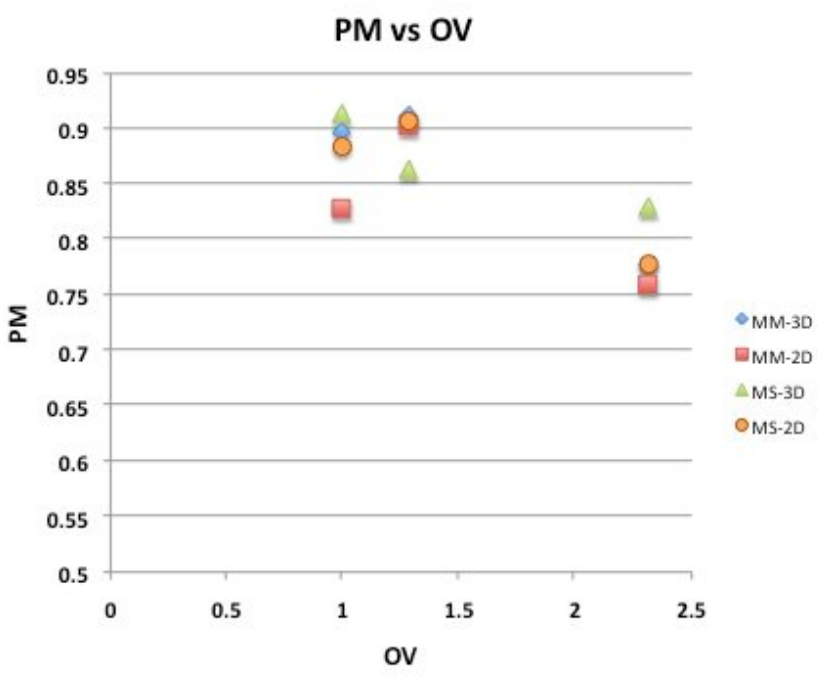

Figure 6.44 – Scatter plot of the object colour variation versus the pure memory accuracy, i.e. the subject accuracy for the uniform case (mean or most saturated condition) in the generic shape (2D or 3D).

Another hypothesis might be that more a shape-texture combination is representative of the original object appearance more the subject is accurate, i.e. implying a multiplicative effect of these two factors. Figure 6.45 plots the product of RS and PI for each object; note that the more this product is close to 1, the more it is representative of the object's natural aspect. The plot shows a direct correlation between the dependent variable "memory colour accuracy" and the independent variable "PI*RS". Each point in the figure represents one object in one condition (TEX-3DN, TEX-2DN, MM-3DN and MM-2DN; note that I impose RS=0 for rhomboid, i.e. RS*PI=0, but other factor do contribute when shape identity is missing, therefore, for clarity, they are not plotted in the figure).

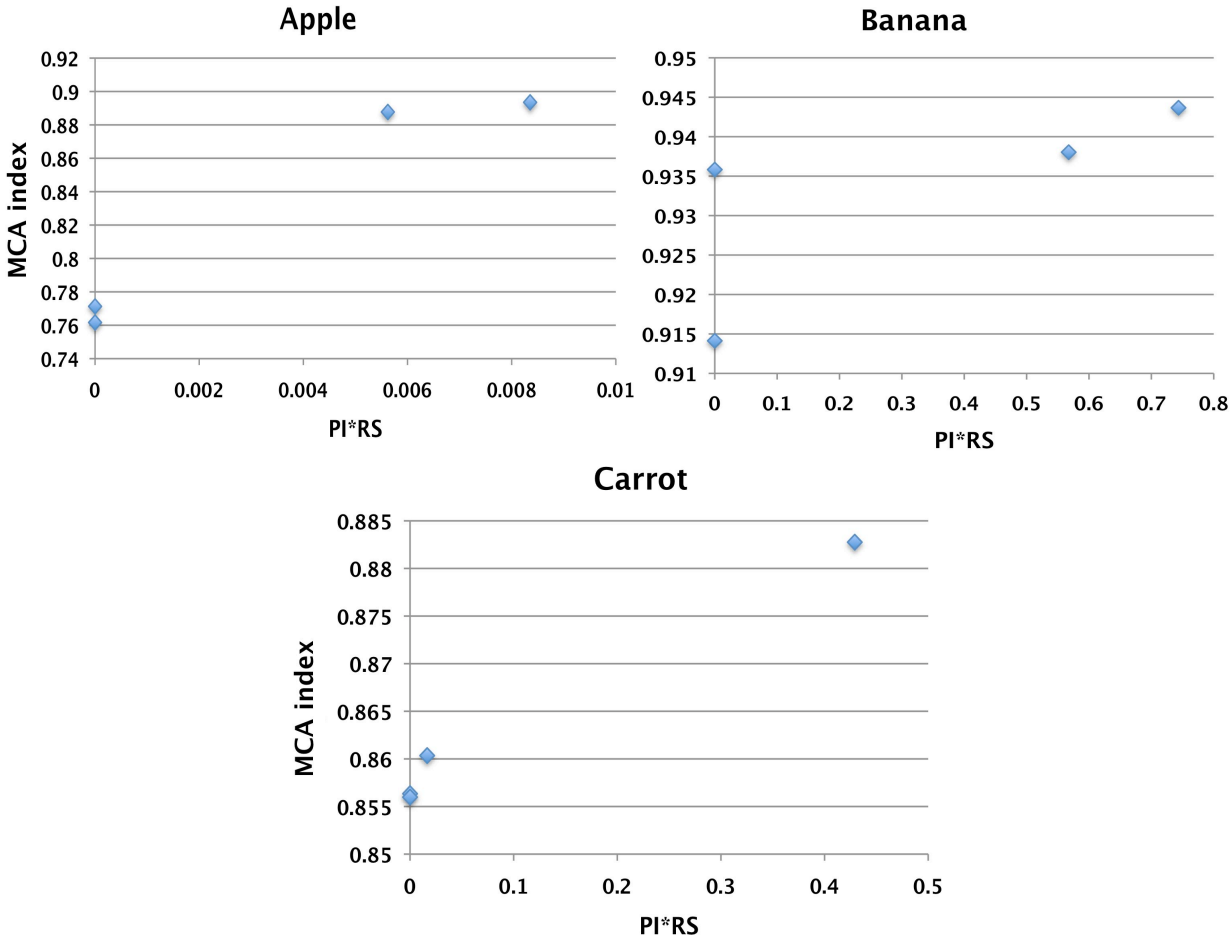

Figure 6.45 – Correlation between memory colour accuracy and the product of recognition-by-shape index and polichromaticity index for: A) Apple; B) Banana; C) Carrot. Each point represents the object in one chromatic (TEX or MM) and shape (3DN or 2DN) combination.

An additional possible combination of factors might be between RS and OV, as seen in Chapter 4. If, for hypothesis, the subject recalls better the memorized object colours' distribution (and therefore its mean) when helped by shape cues, then we can postulate that his/her accuracy in the selection would be enhanced. Consequently the product of RS and OV would be directly proportional to the subject's memory colour accuracy.

Figure 6.46 plots the relationship between the RS*OV for the three object and memory colour accuracy. The graph shows a direct proportionality. Each point represents one object in one condition (TEX-3DN, TEX-2DN, MM-3DN and MM-2DN; note that I imposed RS=0 for rhomboid, ie. RS*OV=0, but other factor do contribute when shape identity is missing, so they are not plotted for clarity).

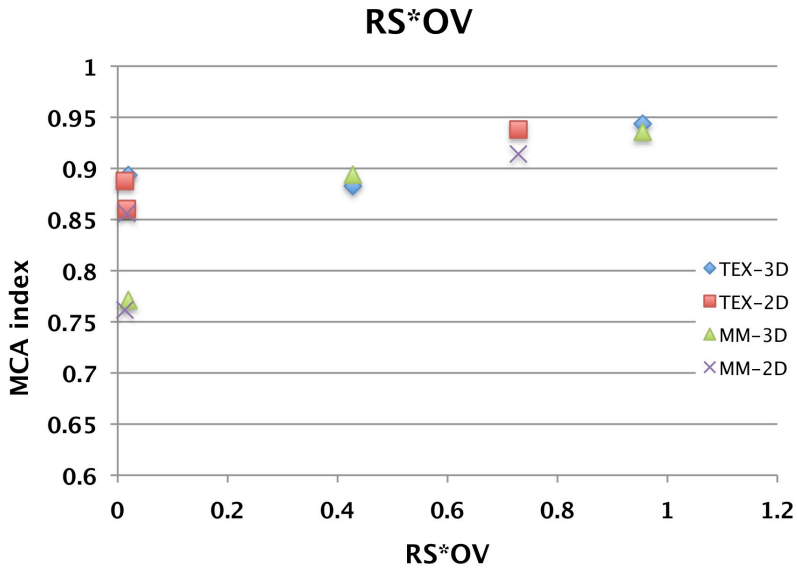

Figure 6.46 - Correlation between memory colour accuracy and the product of recognition-by-shape index and object colour variation index OV. Each point represents one object in one chromatic (TEX or MM) and shape (3DN or 2DN) combination.

In conclusion, memory colour accuracy is a function of the physical appearance of the stimulus (its shape and chromatic variegation), of the original distribution of the object colours alone and of their combinations. Comparing this model with the one presented in Chapter 4, many similarities are evident and we might see this second model as an extension of the previous. Obviously, this "model" is only a simplification and more factors and interactions might be involved. Further investigations may shed some light on other contributions.

6.4.5 EFFECT OF ILLUMINANT CHANGES AND COLOUR CONSTANCY

Colour memory results of the two illuminations are robustly similar, consistently with some recent studies on photographs of textured objects (Olkkonen et al. (2008)). However, the size of the effects varies in relation of the illuminant and factors. The magnitude of effect of chromatic factor on memory colour accuracy differs significantly for all factors. Furthermore, the sizes of effect of dimensionality factor on memory colour precision differ significantly between the two illuminants, while it is similar for chromatic condition and contour factor (for the latter there is a visible difference but is not significant). A plausible reason for these differences can reside in the weight that each factor has in a memory colour "model" relative to its accuracy or precision. In other words, if we consider the illumination as a further variable in the "model" presented in 6.4.4, it is

possible that the factor's weight (i.e. the contribution that each factor gives to memory colour accuracy or precision) is a function of the illumination.

Furthermore, observers showed very high colour constancy, consistent with other recent studies (Hansen et al., 2007; Olkkonen et al., 2008; Murray, Daugirdiene, Vaitkevicius, Kulikowski, & Stanikunas, 2006; Rinner & Gegenfurtner, 2002). Furthermore, chromatically textured surfaces improve the subjects' performance in making colour constant settings, i.e. colour constancy indices are greater for texture condition compared to uniform conditions. This latter result is very strong and significant and it is in line with other studies on chromatically variegated surfaces (e.g. Olkkonen et al., 2008). However, no difference was found on colour constancy between natural and generic contour, or between 3D and 2D configuration.

The question is then: since memory colour is influenced by texture and shape recognisability but constancy only by texture, is memory colour a factor in the colour constancy mechanism? Or does only the number of distinct chromaticities in a scene aid constancy? Based on this study we can assert that: if there is an effect of memory colour on constancy other factors cancel the improving effect of shape. Therefore, the increased performance on memory colour, obtained introducing more identity cues to the object through shape, are lost or deleted by those elements. It is still an open question what additional factors influence such a mechanism.

6.5 CONCLUSIONS

Images of ordinary objects and scenes are normally chromatically very rich. A pixel-by-pixel look at pictures of simple objects generally reveals dramatic chromatic differences across their surfaces. Yet the majority of studies on colour habitually simplify the colour stimuli as basic homogeneous colour patches. Furthermore, we live in a complex shaped three-dimensional environment, in which the shape often induces the recognition of an object category. Again this aspect of the real world is often disregarded by colour studies. This work has demonstrated the need of introducing such cues to obtain a more faithful evaluation of the subjects' memory colour and colour constancy.

As a matter of fact, I found that memory colour is: (1) more accurate, consistent and precise for polychromatic stimuli compared to homogenous stimuli; (2) more accurate, consistent and precise for diagnostic shapes (natural contours) than generic shapes; and (3) more accurate and consistent for 3D than 2D condition. Furthermore the 3D configuration (1) increases the size of the accuracy effect in natural shaped objects relative to generic ones and (2) reduces the precision of memory colour if no reference to the object's actual shape is given (generic contour), otherwise has no effect. *The above findings lead to the conclusion that more cues to the object identity and a more natural stimulus facilitate the observers in accessing their colour information in memory*. I have

also analysed the interaction between chromatic surface representations, shape recognisability, shape dimensionality, and physical proprieties of the object on memory colour accuracy, finding parallels with the experiments in Chapter 4 and the "model" presented there.

Colour constancy was affected by the chromatic complexity of the stimuli (texture) in that colour constancy is higher for texture stimuli relative to uniform stimuli. However, I found no effect of shape diagnosticity or dimensionality. Based on these results, I postulate that the chromatic variegation of the familiar object of the stimuli contributes significantly in colour constancy.

Chapter 7

Surface Discrimination as a Function of Shape, Chromatic Texture, and Familiarity

7.1 INTRODUCTION

In previous chapters I have demonstrated that shape influences the colour appearance of natural objects, as well as the chromatic surface property of the stimulus. For example, Chapter 4 formulated the hypothesis that observers memory colour is affected by shape; the range of typical colours that observers associate with the object increases with the introduction of the natural shape. I have then postulated the theory that humans possess a memory colour distribution (or even a few as suggested in Chapter 5) and not a single memory colour for a familiar object. Furthermore, in Chapter 6, I have shown an effect of shape on memory colour; subjects selected settings matching the typical familiar object colour with more accuracy and precision when the stimulus was displayed in its natural shape. In addition, in the same chapter, when the stimulus represented the full distribution of the object's chromatic surface, the subject was more accurate and precise in selecting the natural object's typical colour. However, a few questions arose from these studies. For example, in both experiments we have explicitly indicated the object category (banana, carrot etc) before performing the task; what effect on colour appearance would this have presenting implicit cues to the object identity as the natural shape or texture of the object? Furthermore, we have seen that the natural shape of the object improved the subject's performance in different tasks compared to a generic non-diagnostic shape; would a shape incongruent with the identity of the object associated with the displayed colour (memory colour of the object) worsen the subject performance? An additional question rised in Chapter 6 was if the improvement induced by the chromatic texture was due to the chromatic variegation of the stimulus or to the specific texture cue to the object identity.

To answer these questions, I performed two experiments in sequence, which examined the effect of shape and chromatic surface cues to object identity. In the first experiment (Experiment 1), I tested the subject ability to simultaneously discriminate between colour stimuli varying in hue. In the second experiment, I replicated the same stimulus conditions as in Experiment 1 in a successive discrimination task. The first experiment can be considered as the base of the second experiment. In fact, Experiment 1 examines the "pure" ability to discriminate between two stimuli while Experiment 2 introduced a memory component testing the ability to recall the object colour and then to match the memory to the displayed stimuli.

Specifically I explore the effects of diagnosticity of shape, diagnosticity of chromatic texture, and colour diagnosticity, where the diagnostic shapes or textures belong to natural fruit and vegetable objects. Generic shapes, randomized textures and uniform colours were used as controls. Results showed a complex pattern and demonstrated that the effects of shape and chromatic conditions are related to high-level or low-level mechanisms depending on task.

7.2 EXPERIMENT 1

The aim of this first experiment is to study chromatic discrimination for solid 3D natural fruit and vegetable objects as a function of the visual cues to the object identity. The subjects were requested to select between two alternatives with respect to a reference present at the same time in the scene. Section 7.2.1 describes the methodology employed; section 7.2.2 presents the results while section 7.2.3 discuss such results. Section 7.2.3 enlists the conclusions drawn from Experiment 1.

7.2.1 METHODS

The experiment was carried out in the set up described in 2.3 in which the surface appearance of a solid object can be digitally manipulated in real time, point-by-point. The projector image (see section) was set at 'sRGB', with contrast 0 and brightness at 0. No side-lights were used in this experiment. All software packages were written in Matlab 7.5 (The Mathworks Inc.) using the Psychophysics Toolbox libraries (Brainard, 1997; Pelli, 1997) and all experiments run on a PC using Microsoft Windows XP.

Six natural, familiar fruit and vegetable objects, bought at an ordinary supermarket, were employed for the first experiment, namely: (1) a Gala apple, (2) a ripe banana, (3) a carrot, (4) a cucumber, (5) a lime, and (6) a potato. Three alternative chromatic configurations (chromatic factor) were evaluated: (1) a full texture of a natural fruit or vegetable object, (2) a randomized texture of a natural fruit or vegetable object, and (3) a uniformly coloured surface whose colour is the mean of a natural fruit or vegetable object. The randomized texture was used as control to disentangle the potential contributions of low-level and high-level effects, as it brings no information on the object identity. In fact, as said in previous chapters the colour appearance of natural objects is influenced by memory colour (see introduction). Hence chromatic discrimination performance might be influenced by the chromatic distributions of the chosen objects (as proven by te Pas&Koenderink (2004) and Zaidi et al. (1998)), but possibly also by the memory colour associated to these familiar objects.

All stimuli in the experiment were presented on a uniform background whose colour defined the adaptation point of the scene. During the experiment the background colour simulated the CIE standard D65 illuminant on a Munsell neutral grey (see section 7.2.1.2).

Object pair (e.g. Apple/Cucumber)	*Chromatic factor*	*Shape factor*
Obj 1 (=Apple)	Obj 1 Natural Texture condition **(Diagnostic texture condition; TEX)**	Obj 1 Natural Shape condition **(Congruent shape; CS)**
		Obj 2 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**
	Obj 1 Randomized Natural texture condition **(Non-diagnostic texture condition; RAND)**	Obj 1 Natural Shape condition **(Congruent shape; CS)**
		Obj 2 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**
	Obj 1 Natural Mean Uniform colour condition **(Uniform condition; MM)**	Obj 1 Natural Shape condition **(Congruent shape; CS)**
		Obj 2 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**
Obj 2 (=Cucumber)	Obj 2 Natural Texture condition **(Diagnostic texture condition; TEX)**	Obj 2 Natural Shape condition **(Congruent shape; CS)**
		Obj 1 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**
	Obj 2 Randomized Natural texture condition **(Non-diagnostic texture condition; RAND)**	Obj 2 Natural Shape condition **(Congruent shape; CS)**
		Obj 1 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**
	Obj 2 Natural Mean Uniform colour condition **(Uniform condition; MM)**	Obj 2 Natural Shape condition **(Congruent shape; CS)**
		Obj 1 Natural Shape condition **(Incongruent shape; IS)**
		Rhomboid Shape condition **(Generic shape; GS)**

Table 7.1 – Example of the combinations of the conditions of the all factors used in both experiments. Object pair: Apple/Cucumber.

Each condition of the chromatic factor, i.e. each chromatic configuration, was displayed in three possible shape configurations: (1) 3D generic shape (rhomboid), (2) 3D natural shape of the natural fruit or vegetable object ***congruent*** with the chromatic factor (i.e. shape and texture or mean

colour belong to the same natural object), and (3) 3D natural shape of the natural fruit or vegetable object ***incongruent*** with the chromatic factor (i.e. shape and texture or mean colour belong to two different natural objects). Hence, each object of the set above was paired with a second of the same set, so as to have 3 pairs of objects. For the first experiment the pairs were: (1) apple/cucumber, (2) banana/lime, and (3) carrot/potato. Table 6.1 lists all the combinations of conditions of the factors (chromatic factor and shape factor) for both experiments and for one pair (for example apple/cucumber). When a texture of an object was displayed on the shape of a different object, i.e. in the incongruent shape condition, the original texture of the first object was synthesized and then shaped as the second object. The synthesising process is described in detail in section 7.2.1.4.2. As illustration, Figure 7.1 shows for one pair (carrot/potato) all the possible combinations. Keep in mind that this is only an example and is represented in 2D and not in 3D as during the experiment. Moreover, due to discrepancy in different medium reproducibility, the colours in the figure are only an approximation of the ones displayed during the experiment.

Figure 7.2 illustrates the comparisons between conditions of different factors: (1) chromatically variegated surfaces (TEX or RAND) versus uniformly coloured (MM) surfaces; (2) diagnostic texture (TEX) versus non-diagnostic texture (RAND); (3) diagnostic shape (CS or IS) versus rhomboid generic shape (GS), and (4) congruent shape versus incongruent shape. The first comparison (Comparison 1) tests the effect of ***chromatic variegation*** of surface on colour appearance. The second comparison (Comparison 2) tests the effect of ***chromatic texture diagnosticity*** on colour appearance. The third comparison (Comparison 3) tests the effect of ***shape diagnosticity*** on colour appearance, while the fourth comparison (Comparison 4) evaluates the effect of ***shape congruency*** on colour appearance.

All objects were photographed under standard daylight at 6,500 K correlated colour temperature (D65) in a Verivide cabinet and processed as in section 3.3.1 (or similarly section 5.2) to obtain the basic object chromatic texture for the experiment. The objects' colours distributions and proprieties are analysed in section 6.2.3.1, while section 7.2.1.4 describes the chromatic condition and explains in detail the texture pre-processing and manipulation technique. Shape configuration characteristics are described in section 6.2.4. Section 6.2.5 describes the experimental procedure.

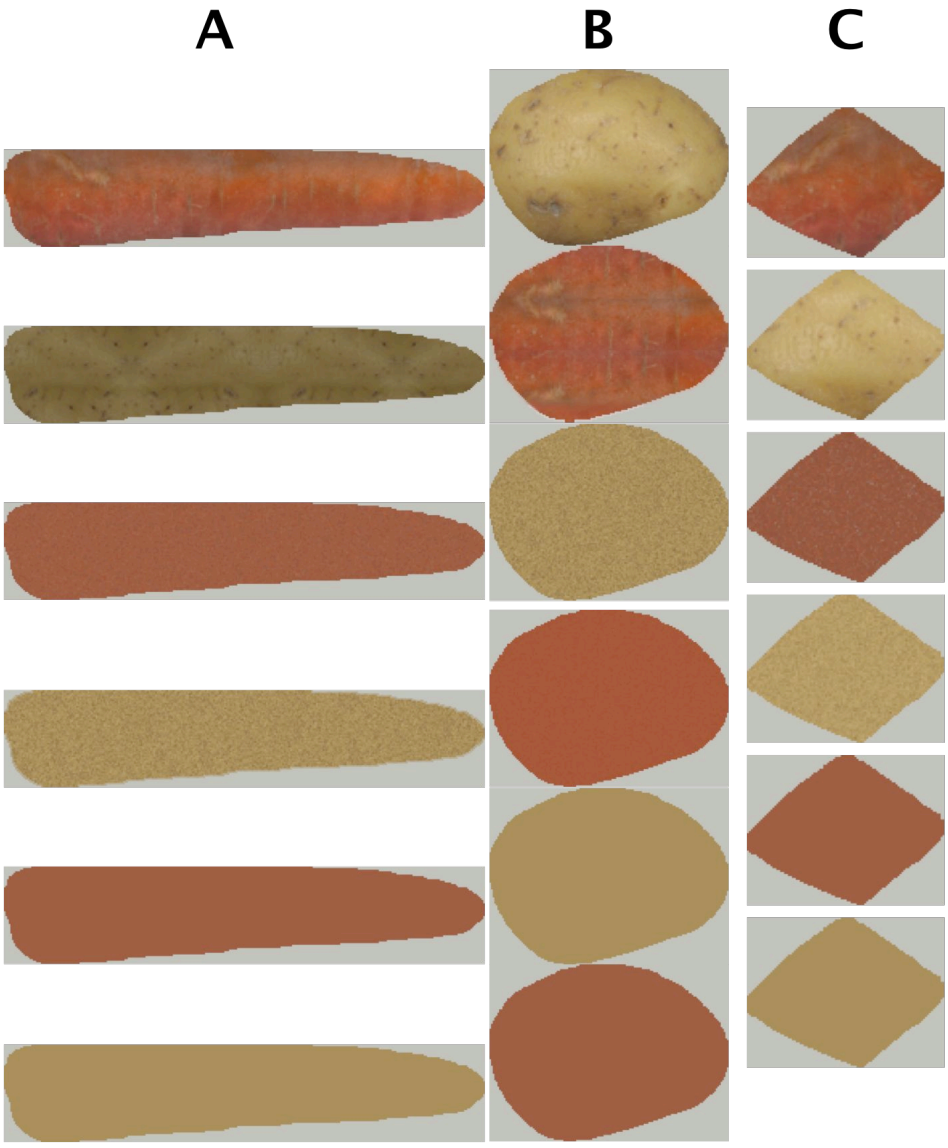

Figure 7.1 – Examples of the stimuli used in this experiment for the object pair carrot/potato. Each subfigure is a combination of the factors' conditions. Column A: carrot shape. Column B: potato shape. Colum C: generic shape. Columns A and B show the following condition combinations, from top to bottom: TEX-CS, TEX-IS, RAND-CS, RAND-IS, MM-CS, and MM-IS. Column C shows for the generic rhomboid shape the following chromatic conditions, from top to bottom: TEX for the carrot, TEX for the potato, RAND for the carrot, RAND for the potato, MM for the carrot, and MM for the potato.

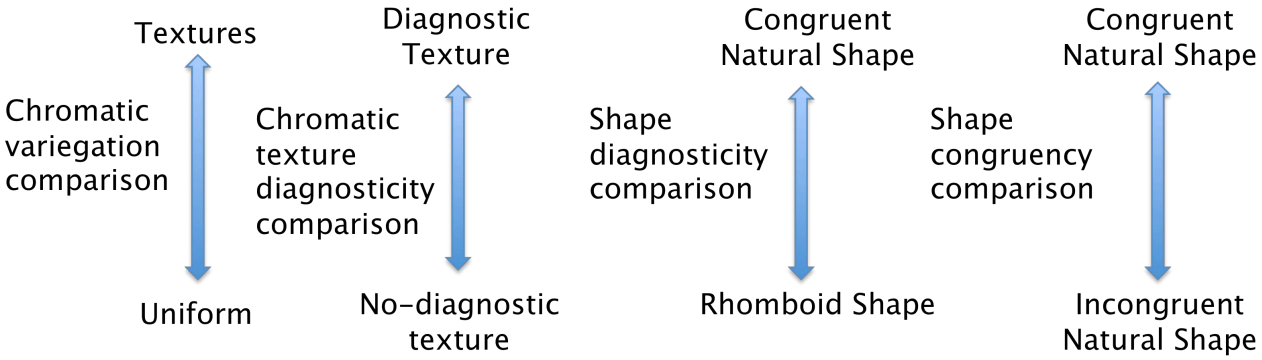

Figure 7.2 - Comparison between conditions of different factors. Comparison 1: chromatically variegated surfaces (TEX or RAND) versus uniformly coloured surfaces (MM). Comparison 2: diagnostic texture (TEX) versus non-diagnostic texture (RAND). Comparison 3: diagnostic shape (CS) versus rhomboid generic shape (GS). Comparison 4: congruent versus incongruent shape.

7.2.1.1 OBSERVERS

Twelve observers, 8 female and 4 male, all in the age range of 19-28 years and all students at Newcastle University, took part in all conditions of Experiment 1. They all tested normal on the Farnsworth-Munsell 100 hue test (mean total error score 35.2 (Kinnear&Sahraie (2002)) and were naïve to the purpose of the experiment. An ethical agreement was obtained before performing each experiment. Before starting each experiment, they acquired familiarity with the task during a 3 minutes practice session.

7.2.1.2 ILLUMINANTS

During each experiment the background colour simulated the CIE standard D65 and CWF illumination as recorded in the Verivide cabinet in which the original photographs of the objects were taken (see Chapter 3). As said in previous chapters, the spectral power distribution of the background is the best approximation to the original illuminant that the projector can generate. In other words, a metamer of the illuminant was employed in these experiments. The CIE Yxy values are shown in the Table 7.2.

Illuminant	Y(cd/m^2)	X	Y
D65 (neutral)	7.8	0.314	0.333

Table 7.2 - CIE Yxy chromaticity values for the two illuminations. On the right: chromaticity coordinates of the background. On the right: mean chromaticity coordinates of the light reflected by the chamber walls.

7.2.1.3 OBJECTS

Photographs of the six familiar objects were recorded under standard daylight illuminant at 6,500 K correlated colour temperature (D65): a Gala apple, a ripe banana, a carrot, a cucumber, a lime, and a potato. The images were subsequently converted to CIE XYZ colour space using the method described in Chapter 3.

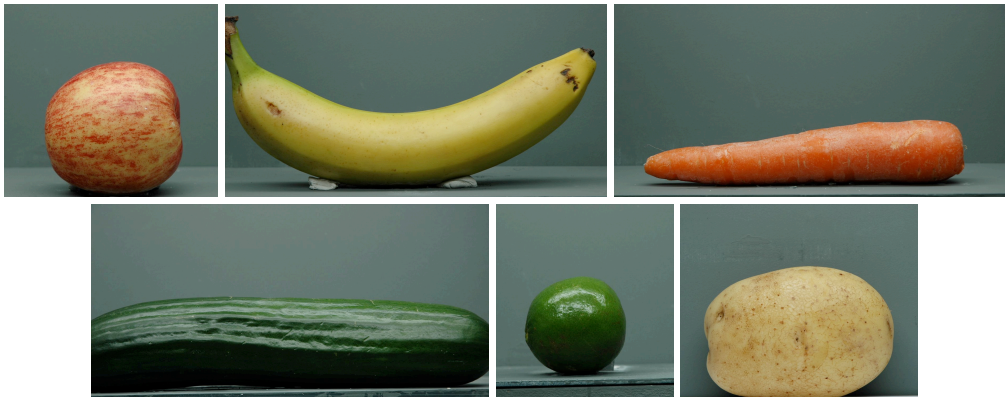

Figure 7.3 – Thumbnails of the six objects' photographs. From left to right, from top to bottom: Gala apple, banana, carrot, cucumber, lime, and potato.

The six photographs of the objects are shown in Figure 6.3. Figure 6.4 illustrates the chromaticity distributions for all objects under D65 in EMG cone contrast space (RG-YB isoluminant plane) obtained after using the camera model transformation.

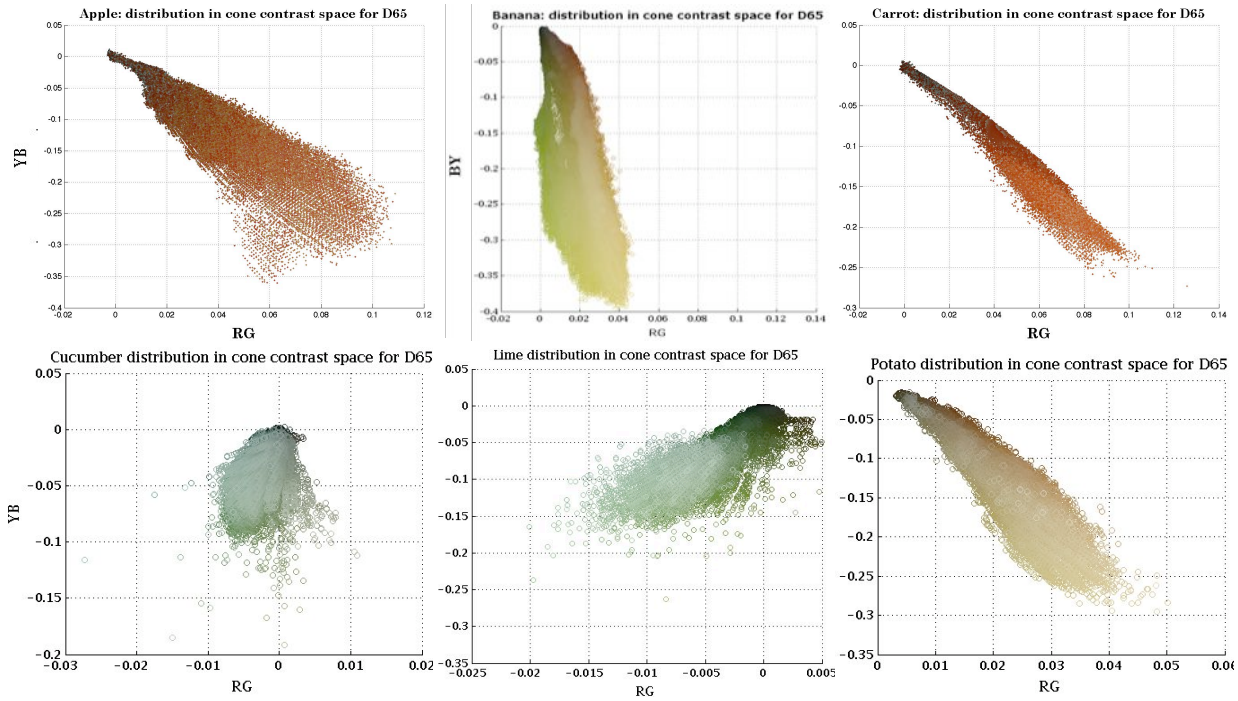

Figure 7.4 – Chromaticity distribution of the 6 objects in EMG cone contrast space under D65. From left to right, from top to bottom: Gala apple, banana, carrot, cucumber, lime, and potato. Note that the colour of each point of the distributions in the figure is only indicative of its real colour. Each point is represented only once even if occurring more frequently in the distribution.

As in previous chapters, it is possible to calculate the area covered by the chromatic distribution of an object in EMG cone contrast space, i.e. the object colour variation indices or OV (see section 4.4), and evaluate the difference between the extent of the objects' spreads (Table 6.3). The distributions can now be manipulated for the purpose of the experiments, as described in the next sections.

	Apple	**Banana**	**Carrot**	**Cucumber**	**Lime**	**Potato**
OV*100	2.309	1.272	0.884	0.459	0.427	0.481

Table 7.3 – List of the object-variation indices (OV).

As shown in Chapter 5, a natural object's chromaticity distribution, when plotted in EMG cone contrast isoluminant plane, forms a distinct signature. In Figure 6.4 we can observe the signatures for the tested objects. As described in Chapter 5 we can compute, for each object under a specific light, the mean hue vector of its surface colours distribution. Table 6.4 lists the object hue angles for the D65 simulated illuminant.

	Apple	Banana	Carrot	Cucumber	Lime	Potato
Hue Angle	70.197	84.133	67.390	175.963	176.457	80.381

Table 7.4 – Hue angles of the six objects under illuminant D65 (hue angles calculated with respect to the RG axis in clockwise direction).

7.2.1.4 CHROMATIC CONDITIONS

During the experiment, I varied the representation of the chromatic proprieties of the surface of the displayed object between: (1) natural chromatic texture of the object (or diagnostic texture; TEX condition), (2) randomized texture of the object's natural texture (or non-diagnostic texture; RAND condition), and (3) mean colour only of the object (or uniform colour; MM condition). The latter condition represents the case in which the colour stimuli were uniformly coloured surfaces. The surface colours were simulated under the D65 illuminant (i.e. the background illumination) as in previous chapters.

Unlike Hansen and colleagues (Hansen et al., 2008), the control non-diagnostic textures used in this experiment were *not* synthetic textures resembling the distribution of the banana generated from white, pink or brown noise, but rather synthetic textures possessing the same chromatic distribution of the objects with which they are compared (section 7.2.1.4.1). Furthermore, since the chromatic texture of one object could be displayed on a second object shape it was necessary to create a method to synthesise the texture of one object in the shape of the object which it is paired. Therefore block 2 of the general framework explained in Chapter 3 (framework in Figure 3.1) is divided into 3 sub-blocks. Figure 7.5 illustrates the flowchart of the block 2 structure and its sub-blocks, namely: (1) sub-block 2A, which treats the synthesis of the randomized texture (section 7.2.1.4.1), (2) sub-block 2B, which considers the synthesis of the incongruent-shape texture (section 7.2.1.4.2), and (3) sub-block 2C, which treats the manipulation to generate the single textured test stimuli (section 7.2.1.4.3). Block 1, block 3, block 4, and block 5 were described in Chapter 3 sections 3.2 to 3.7. Finally section 7.2.1.4.4 describes the generation of the uniformly coloured test stimuli.

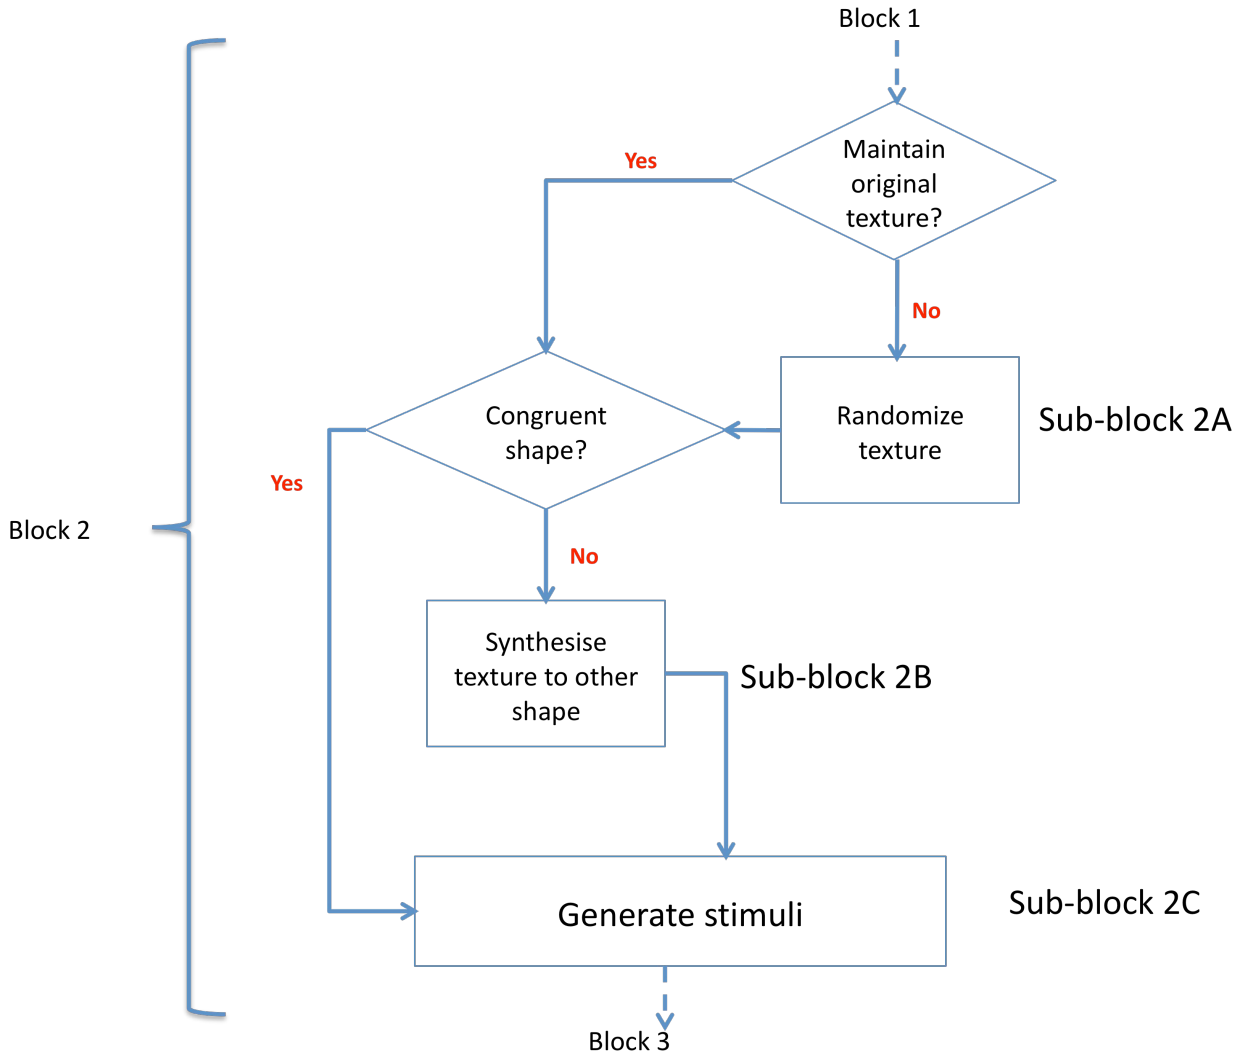

Figure 7.5 - Flowchart of block 2 of the pre-processing technique used in this chapter. Refer to Chapter 3, sections 3.2 to 3.7, for description of the other blocks and to Figure 3.1 for the full framework.

7.2.1.4.1 Randomized texture generation

The main reason for synthesising a non-diagnostic texture was to compare the effect of a generic polychromatic stimulus with a meaningful polychromatic stimulus. However I wanted to give the same chromatic information and shape cue to the observers. The perfect solution would have been to apply a bi-dimensional discrete Fourier transformation on the image and obtain the image in frequency domain. Then we can randomize the phase component of the transformation whilst maintaining the same amplitude at each frequency. This would generate a texture that had the same spatial frequencies and pixel intensity values of the original image, but a broken spatial relationship between pixels, i.e. a different random texture. Note that only the part of the image containing the texture of object needs to be modified (the part within the object contour), i.e. only the content of a non-rectangular region of the image. Alas, a fundamental requirement of the Discrete Fourier transformation dictates that data samples must be contiguous, which, in case of an image, corresponds to having a rectangular region of pixel intensities. It is worth remembering that

a bi-dimensional Fourier transformation consists in applying the one-dimensional Fourier transformation on the rows of the image and then on the column of the image (or vice versa).

A possible solution could be to apply deconvolution algorithm to remove the artefacts introduced by the shape components as proposed by Clark et al. (1999). The basic idea of this technique is to assume that the non-rectangular texture patch is the product of a binary shape mask (the object contour) and a rectangular region of texture, as illustrated by Figure 7.6.

Figure 7.6 – Deconvolution algorithm for non-rectangular texture Fourier analysis. Example: apple.

The caveat of this method is that the deconvolution is an ill-posed problem and has no unique solution. Many techniques have been proposed to solve such problem, amongst which the most common are: Maximum Entropy, Lucy-Richardson, and CLEAN. Although this method was appropriate for our scope, the results of the computation were not visually satisfying and, additionally, it required a laborious and computationally demanding processing. Hence a simpler technique was applied.

I will describe the procedure for one plane of the colour image (here I employed the EMG cone contrast space), remembering that the Fourier transformation is a linear transformation; therefore after transforming all three planes we can apply a 1-D transformation on the third direction for only the point within in the object contour. In the image we can delineate an inside of the object contour (or contour) and an outside. Initially, I selected the points *within* the object contour starting from the pixel at the top-left corner of the image, scanning along one column from the top row to the bottom row across all image from left to right. The selected points were organized in a one-dimensional array in the same order as they were scanned. Then I applied a one-dimension Discrete Fourier transformation (DFT) on such array and replaced the phase of the transformation with the phase of a white noise, while keeping the same amplitude. Let us call the array after transformation "first order transformed array". Then, a second image of the same size as the original was created in which each pixel was or zero or one of element of the first order transformed array. Specifically: (1) the zero values filled the area of the image identical to the area in the original image outside the object contour and (2) each element of the first order transformation array was placed back to the same location as it was in the initial image (ergo in the image contour). Hence we obtained an "image" that has the same structure of the intial image, where the "pixels" inside the contour are the complex values of the 1-D Fourier transformation; let us call this new image the "first order transformed image".

The second stage consisted in applying again a 1-D transformation, but this time on the pixels of the transformed image along its *rows*. Again I selected only the pixels inside the object contour. To do so I selected the points that differ from zero starting from the pixel at the top-left corner of the image, scanning along one row from the left to the right of the image from the top row to the bottom row. The selected points were organized in a one-dimensional array in the same order as they were scanned. Next I reapplied a 1-D DFT and again replaced the phase of the transformation with the phase of a white noise leaving the amplitude unchanged. The third stage was to apply the 1-D transformation this time along the third dimension of the image (RGB) similarly as before. Finally the image was transformed back to the space domain following the inverse process described above.

This technique maintains unchanged the value for each pixel (amplitude) and corrupts the spatial coherence of the image, while partially restoring the original object frequencies. An example of an object (the apple texture in apple shape) before and after randomization is show in Figure 7.7. Additionally, Figure 7.7 shows the frequency components for each RG (red line), BY (blue line) and L (luminance) plane for the natural texture (Figure 7.7 C) and the randomized texture (Figure 7.7D).

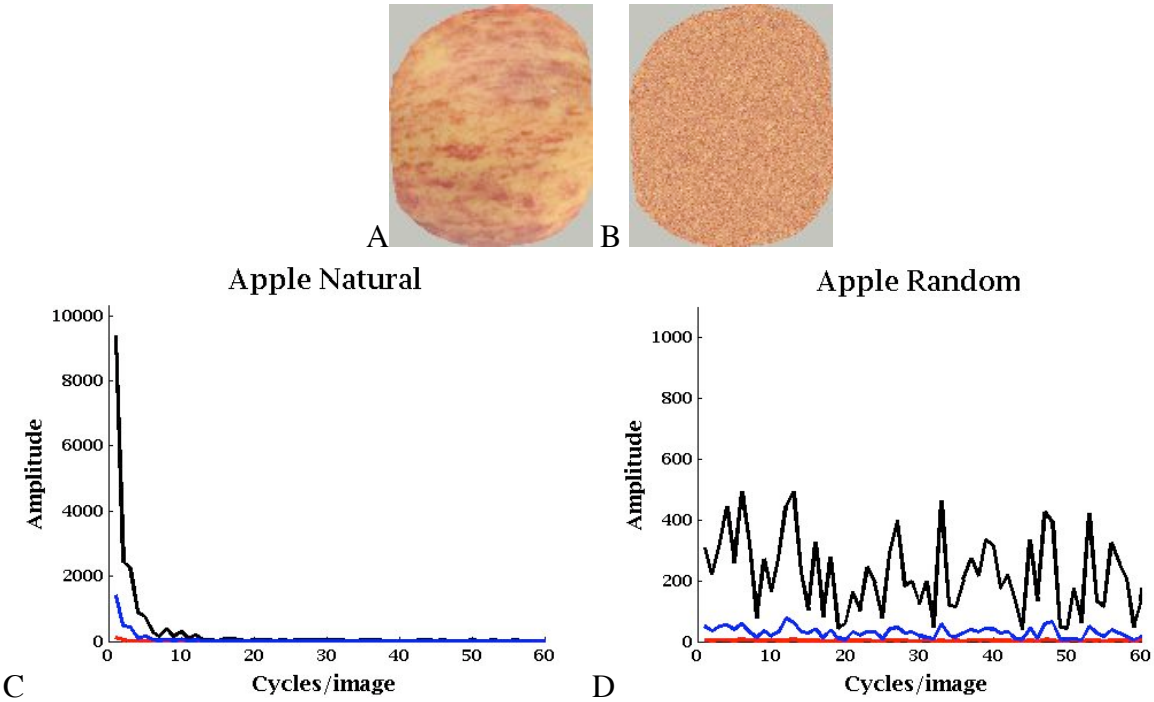

Figure 7.7 – Example of the apple texture in apple shape. Top row: Image of apple before (A) and after (B) randomization. Bottom row: spectral power distribution of the apple before (C) and after (D) randomization. Red line: RG plane. Blue line: BY plane. Black line: luminance (L) plane.

7.2.1.4.2 Texture synthesis

A second type of chromatic stimulus employed in this chapter consisted in an image possessing the same colour (chromaticity and luminance) of the object original texture but a different shape contour. Accordingly it was necessary to develop a method that synthesised the

original texture of the object so as to be applied on a different object, for example the texture of an apple on the shape of a cucumber (Figure 7.9). Clearly, to generate the cucumber/apple image in Figure 7.9, it was necessary to synthesise an apple texture large enough to cover the (much larger) cucumber shape *maintaining* the same perceived texture. In other words, the primary idea of texture synthesis is that, given a sample of any texture, we are able to reproduce a *different* sample reproducing the *same* perceived texture. Evidently this problem is ill-posed in that the sample could have been extracted from a great number of textures. Hence the fundamental assumption is that we are able to select a sample of the original texture large enough that it somehow captures crucial perceptual characteristics of the specific texture. According to the "texton theory" (review in Julesz and Bergen (1983)), textures are discriminated if they differ in the density of certain simple, local statistical textural features, or textons (see Chapter 1).

Assuming that we able to construct a "vocabulary" of textons, i.e. a set of textons that captures the features of the original texture, then we need a method to synthesise a different textured image without incurring in the generation of artefacts and aberrations of the final texture. Hence we need a method that generates a transition between single "textels", where a textel is an element containing one or more textons of the original texture.

Although a great number of studies in computer vision have pursued the solution of such problem and great effort has been put into developing algorithms that successfully apply a solution under certain constraints, alas *no* texture synthesis approach is able to *exactly* reproduce the perception of the original texture. Nonetheless, recently some very successful algorithms have been presented, among which the Efros&Freeman image quilting algorithm stands out (Efros&Freeman (2001), Hays et al. (2006)). This method uses an adaptive search of textels possessing similar features and textels' edge boundary reconstruction. Based on this algorithm I developed a procedure to generate the incongruent texture-shape stimuli.

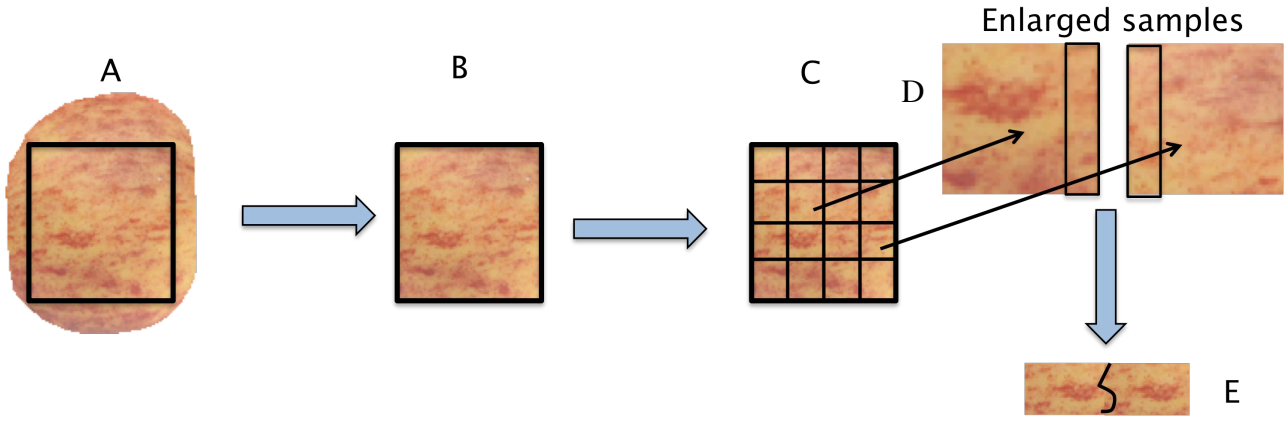

Figure 7.8 – Texture synthesis algorithm: apple texture example. Description in the text.

Figure 7.8 show the framework of the algorithm while Figure 7.9 show the results of the process for the apple texture. Initially I selected a rectangular sample of the original object of the maximum size allowed by the contour of the object in the original picture (Figure 7.8A). Then this

rectangular image is divided into blocks of nxn size, where n is a parameter of the quilting that differs between objects (Figure 7.8B). Next, a first block is randomly chosen among the generated blocks (or alternatives - Figure 7.8C) and tiled on a new image (Figure 7.8D left). Subsequently, a second block is actively searched among the alternative that possesses features similar to a certain area of the first block (Figure 7.8D right). This area is the region on the boundary of the first block and of the second block that will overlap in the tiling process (indicated with black rectangular contours in Figure 7.8D). Therefore we defined an additional parameter of the quilting: the "size of overlap" defined as the number of columns (and/or rows) that will overlap. The block is chosen such that the area of overlap possess the least square error among all alternatives. However the overlap produces a distinctive boundary. Therefore a minimum cost path along the two boundaries is computed in order to obtain a smooth transition between blocks (Figure 7.8E). The same process was repeated until the synthesis texture covered an image of the size of the object. Finally the texture was cut out into the shape of the new object. The image was represented in EMG cone contrast space and all planes were examined simultaneously. The main feature examined in the search for the match was the spatial power spectrum of the area of overlap. Occasionally an object-dependent bandpass filter was applied along the line of transition to allow optimal visual results. Figure 7.9B shows the final results for the apple texture on the cucumber shape.

A  B 

Figure 7.9 - Example of the apple texture in (A) apple shape and (B) cucumber shape (note that the objects' sizes are not scaled by the same amount).

7.2.1.4.3 Textured stimuli generation

As seen in Chapter 6, a convenient method to change the overall colour of the object, and maintain the relative distance between colour points, is to rotate the entire chromatic distribution solidly with its mean vector around the adaptation point. In other words, the different test stimuli for each chromatic condition were obtained by rotating of a certain angular degree the whole distribution around the origin of the EMG cone contrast isoluminant plane. Hence, the luminance value of each point was kept constant and only the RG and BY coordinates changed in the rotation. In order to maximise the displayable gamut for a specific object, the luminance was scaled by a fixed factor. Figure 6.6 shows an example of the gamut of the projector and of the banana colours' distribution with the luminance (L) scaled by 2.5 (note that this is similar to Chapter 6).

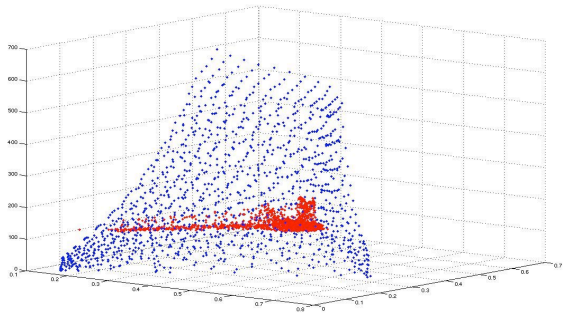

Figure 7.10 – Projector gamut in the setting of this experiment and banana distribution in La*b* colour space. The L value of the banana distribution is scaled by 2.5. Blue dots: sampled projector gamut. Red dots: banana distribution, each colour point shown once.

For each point in the chromatic distribution, the position after rotation was determined as follow. Let $c_i = (r_i, \Theta_i)$ be a point of the initial distribution in the EMG isoluminant plane, given in polar coordinates, Φ_N the angle of the rotation and $c_i' = (r_i', \Theta_i')$ the position of the new adjusted point after rotation. Then, the new position for each point in the distribution was computed, in polar coordinates, as:

$$r'_i = r \qquad \text{Equation 7-1}$$

$$\Theta'_i = \Theta_i + \Phi_N \qquad \text{Equation 7-2}$$

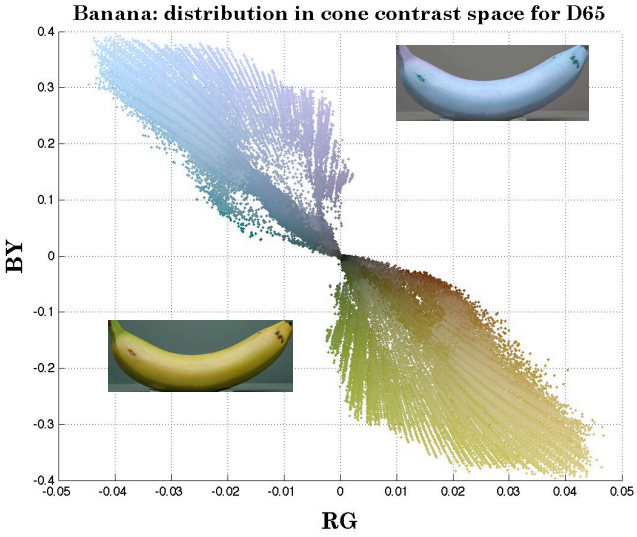

Figure 7.11 – Distribution of a banana in its original configuration (yellow-brown) and in its 180° rotated configuration (blue-violet) in EMG cone contrast space. Note that the colour of each point of the distributions is only indicative of its actual colour in the experiment. Each colour is represented only once even if it occurs more frequently in the distribution. Bottom left insert = original banana image. Top right insert = an altered banana image.

Negative Rotation	ΔEuv difference of the negative rotation	**Positive Rotation**	ΔEuv difference of the positive rotation
-2	2.8460	2	2.8626
-4	2.8274	4	2.8771
-6	2.8070	6	2.8896
-8	2.7849	8	2.9001
-10	2.7611	10	2.9082
-12	2.7358	12	2.9142
-14	2.7092	14	2.9181
-16	2.6815	16	2.9197
-18	2.6527	18	2.9192
-20	2.6232	20	2.9166
-22	2.5931	22	2.9118
-24	2.5627	24	2.905
-26	2.5323	26	2.8961
-28	2.5020	28	2.8853
-30	2.4722	30	2.8725
-32	2.4432	32	2.8580
-34	2.4152	34	2.8416
-36	2.3884	36	2.8236
-38	2.3633	38	2.8039
-40	2.3401	40	2.7827
-42	2.3188	42	2.7601
-44	2.3001	44	2.7361
-46	2.2837	46.25	3.0479
-48	2.2701	48.5	3.0143
-50	2.2594	50.75	2.9792
-52	2.2517	53	2.9427
-54	2.2470	55.25	2.9050
-56	2.2454	57.5	2.8663

Table 7.5 – Angles selected for the potato and the perceptual differences between the means of its distributions for successive angles of rotation.

As an example, Figure 6.7 illustrates the chromaticity distribution of an original banana (0° rotation, named "base distribution") and the same distribution rotated 180° in EMG cone contrast space. Note that due to the uncalibrated reproduction of the colour on paper the colour of each point in the figure is only indicative of the actual colour displayed in the chamber.

For each object, 56 angles of rotation (or "angular steps") were selected. For example, Table 6.5 lists the angular steps for the potato. A complete list for all objects can be found in the Appendix 2). For each object, the perceptual difference between the means of the distributions was in average 2.57 ΔEuv difference. Table 6.5 lists the perceptual difference between successive angles of rotation for the potato; note that there are 23 negative angles, i.e counterclockwise rotation of the distribution, and 23 positive angles, i.e. clockwise rotation of the distribution. In addition, it is worth pointing out that the ΔEuv difference of the first negative or positive angle is calculated with respect to the object's original angle (0° rotation).

To reiterate, as shown in Figure 7.5, depending on the chromatic condition, the distributions of the objects are rotated to generate the test stimuli *after* being processed as in 7.2.1.4.1 and/or 7.2.1.4.2. In any chromatic condition, the distribution with 0° rotation is defined as the "base" distribution. Finally, it is worth pointing out that chromaticity coordinates of all stimuli were simulated with respect to the illuminant D65 generated by the projector.

7.2.1.4.4 Uniformly coloured stimuli

In addition to the chromatic texture (TEX and RAND), for each angle of rotation of the distribution described above, a uniformly coloured stimulus was generated whose chromaticity was equal to the mean chromaticity of the rotated distribution. The luminance of the uniform stimuli was equal to the mean luminance of the object's distribution; thus I obtained 57 uniform stimuli per object.

7.2.1.5 SHAPE CONDITIONS

Three copies of seven solid objects were used in this experiment, of which six were 3D artificial replicas of the fruit and vegetable objects listed before and one was an unfamiliar handmade object (generic object). Specifically, the three apples and the three bananas were copies in polystyrene of the real object shape, while the other objects were hand-made as described in 6.2.4.1. The three generic objects were identical irregular rhomboid-based pyramids made using air-drying modelling clay (DAS white by GIOTTO/FILA). For the purpose of this thesis it will be called "3D rhomboid" or 3D generic shape, alternatively. All objects were finished off with matt white acrylic-based paint. As a result, these objects preserve all the shape cues of the test objects,

but can be seen as a "blank canvas" on which different test colours can be projected. Figure 4.3 shows the 3D white-painted objects and all their copies for Experiment 1 (note that in the figure is also present three copies of a pear that will substitute the lime in Experiment 2).

Figure 7.12 – Photograph of the objects used in Experiment 1 (see text) and Experiment 2.

Thus, for each object, the chromatic stimulus can be presented on: (1) the 3D shape of the object from which the chromatic stimulus was generated (congruent object shape, CS), (2) the 3D shape of the object *paired* to the object from which the chromatic stimulus was generated (incongruent object shape, IS), or (3) the 3D rhombic shape (generic shape, GS). Table 6.6 lists the visual area covered by each object when placed in the centre of the screen.

Objects	Degrees
Apple	3.814x3.148
Banana	7.360x3.101
Carrot	8.718x1.670
Cucumber	12.27x2.243
Lime	2.958x2.624
Potato	4.384x2.862
Generic	3.005x2.148

Table 7.6 – Visual degrees of angle covered by the objects used in the experiment.

As for the experiments in previous chapters, the object cut out in the projected image exactly aligns with the visible surfaces of individual objects placed in the box, segmented from the background. It is important to say that the texture condition of the generic object (or rhomboid) consists of a cut out of the original object chromatic texture that *subsequently* is manipulated as

described in 7.2.1.4.3. In the following, each combination of the condition of the factors (or "factors combination") will be indicated as "chromatic label-shape label", e.g. the natural uncorrupted texture condition combined with the 3D generic shape will be referred as TEX-GS.

7.2.1.5.1 Object casting

The sequence of object casting is shown in Figure 6.11. First a soft mould of the object was made using algae gelatine. Note that the objects photographed for these experiments and the objects used to make the die are the same. Before the gelatine settled the object was pressed inside the mixture. After solidification, the object was slowly removed. Lightweight white plaster of Paris was poured into the form. Before this, the mould was sprayed with synthetic oil to allow easier extraction of the cast after the compost had hardened. The same die was used to obtain 3 identical copies of the object.

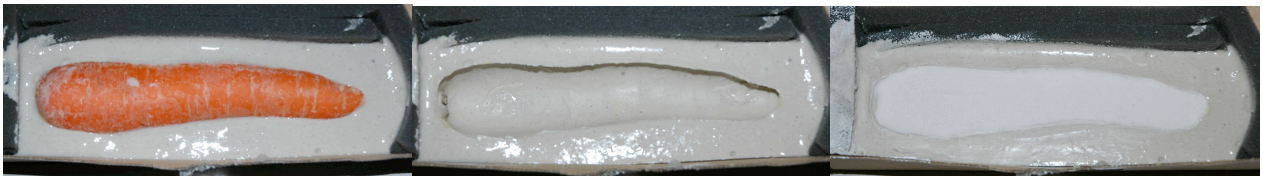

Figure 7.13 – Casting process sequence. From left to right: from mould making to final cast.

7.2.1.6 EXPERIMENTAL PROCEDURE

In the first experiment the subject was asked to perform a discrimination task between two alternatives with respect to a central reference. The general scene displayed two sets of three objects placed on a uniform background. Each set of objects had the same shape, S1 or S2. One object of each set was placed in central position, while the other two copies were placed one to its left and one to its right side. An example of the placement of two sets of the objects in the scene is shown in Figure 7.14. The central object was designated as the reference object and the two side objects as tests. In each experimental trial, one chromatic configuration was displayed on shape S1 or S2. In other words, for each trial, the subject observed a scene consisting of three solid identically shaped objects whose chromatic surfaces belong to one of the chromatic conditions (TEX, RAND or MM). The central stimulus was always the base distribution (the distribution not rotated) in the specific object/shape/chromatic configuration combination (see Table 7.1). Only one of the two test stimuli was identical in colour to the reference while the chromaticity distribution of the other was rotated as described in 7.2.1.4.3; the latter test stimulus was called "target stimulus". The position of the target stimulus was randomly alternated between the two test object positions.

The observer was required to sit on a chair, stabilize his/her head on the chin rest and look inside the experimental chamber; then s/he was handed a 3-button mouse. The experimental room was in complete darkness. The observers had to complete 4 blocks of trials, one for each combination of objects (3 pairs) and one for the generic shape. The generic shape was always

presented first, while the sequence of other three object combinations was pseudorandomised, with ~16.5% of the participants shown the same order. This order ensures a gradual introduction of more identity cues and it controls for bias.

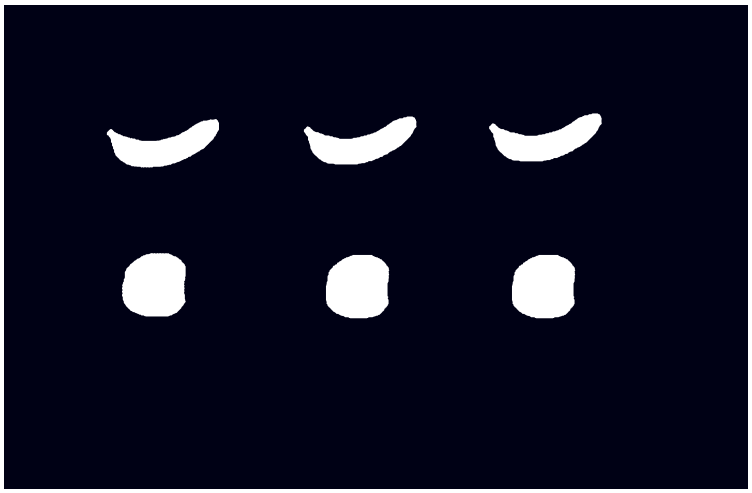

Figure 7.14 - Typical experimental scene. Objects: banana (top row) and apple (bottom row).

Each block was divided into 2 sub-blocks: in one the chromaticity distribution of target stimulus was rotated negatively with respect to reference stimulus, while in the other it was rotated positively. In each sub-block, the shape/chromatic condition combinations were randomly displayed on the first or the second set of objects, also in random order. The paired objects (apple/cucumber, banana/lime, and carrot/potato) were not presented in the same scene. Specifically, the objects' shapes simultaneously presented in the scene are: apples with bananas, carrots with limes, and cucumbers with potatoes. Therefore, in each sub-block (or block) 12 different chromatic/shape/object combinations were randomly displayed, with the exception of the generic shape in which all chromatic conditions for all objects were displayed in the same sub-block (18 interleaved staircases, see below). Thanks to this expedient the observer was unable to learn the reference colour. The sequence of each sub-block within the block was random and its duration was 45-55 minutes. The subject was allowed to take a break between sub-blocks. Note that a double staircase (i.e. negative and positive rotation in the same staircase) was not employed in this experiment because the subject's attention dropped after 1 hour. Therefore, if we wanted to maintain the same length of the sub-block we would have had to reduce the number of chromatic/shape/object combinations per sub-block, with the caveats of (1) presenting the same reference stimulus twice more than in the chosen procedure and (2) interleaving with a lower number of combinations. Since I wanted to avoid the possibility that the subjects learned the reference stimulus, the use of a double staircase was discarded.

At the start of each experimental sub-block, the subject fixated for 70 seconds a uniform background whose chromaticity was equal to the selected illuminant value (adaptation phase). Then the first stimulus was presented. The observer task was to select, as quickly and accurately as possible, which of the two test stimuli was different from the reference stimulus by pressing the

right or left button of the mouse. The stimuli were displayed for as long as the subject needed to make his/her judgement. Once the decision was made, the selection and the time of response were recorded together with the other variables of the trial (see below). Next, a fixation cross was presented for 0.5 sec and the next trial initiated. No feedback was given after each response. The general protocol is presented in Figure 6.12.

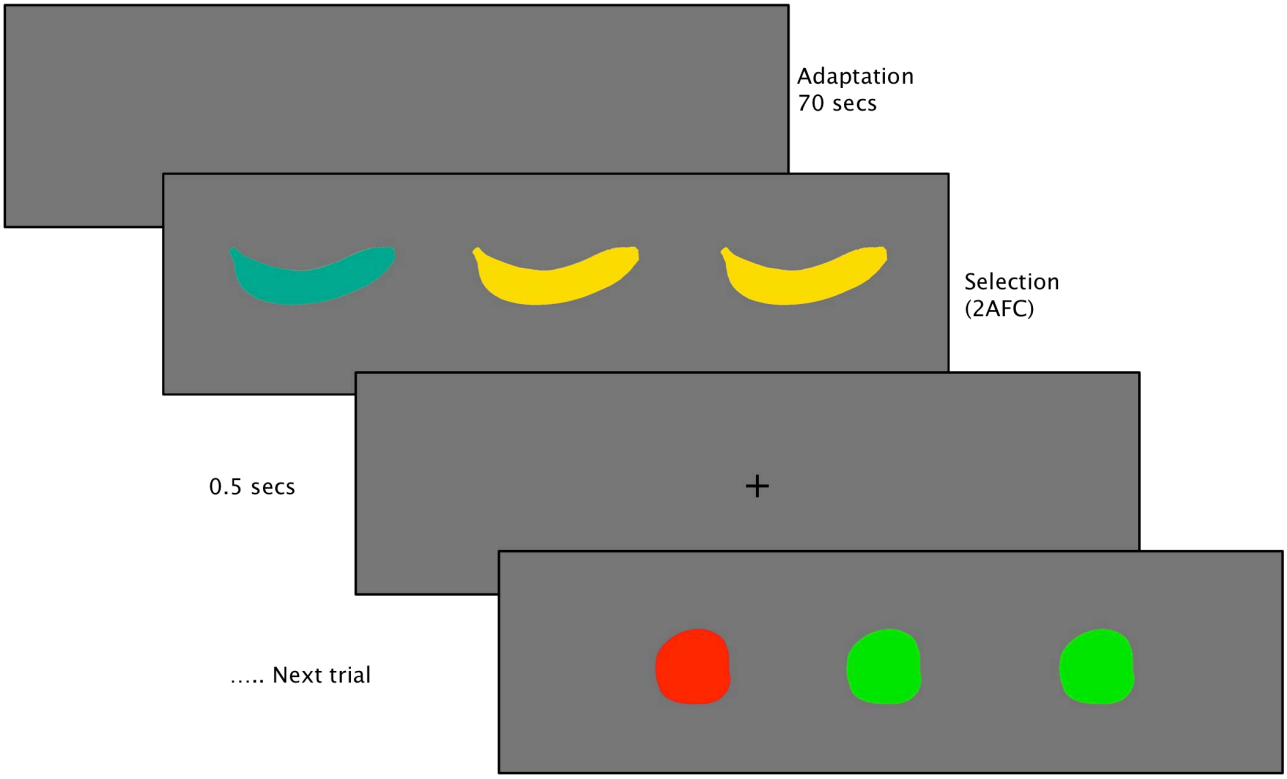

Figure 7.15 – General procedure: example of the experiment images sequence. The displayed colours are only illustrative.

The increment or decrement of the angular difference from reference stimulus was determined by an adaptive staircase procedure (QUEST - Watson& Pelli, 1983; Pelli & Farell, 1995; in Psychophysics toolbox), which aimed for a discrimination yielding 75% correct responses. After 3 consecutive correct responses, the angular difference between the target and the reference stimulus was decreased by one angular step; after one incorrect response, it was increased by 3 angular steps. Each interleaved staircase started from a 20 angular steps away from the reference stimulus (positively or negatively) and was given 100 trials to converge. Each staircase terminated after 5 reversals. At the end of each staircase, 4 *other* variables were recorded: (1) the objects' combination, (2) the factors' combination, (3) the number of trials performed, and (4) the number of correct and incorrect responses for the specific staircase.

7.2.2 RESULTS

The primary aim of Experiment 1 was to explore the effect of diagnostic shape and chromatic texture cues on chromatic surface discrimination. The results were analysed on two levels, organized here into two subsections: (1) the discrimination thresholds of the objects

(subsection 7.2.2.1), and (2) the facility in performing the task in a specific condition (subsection 7.2.2.2). The latter analysis evaluated: (1) the mean time to respond at one stimulus in one trial belonging to a specific condition (subsection 7.2.2.2.1), and (2) the mean number of incorrect responses (subsection 7.2.2.2.2). The significances of the results were measured using a four-way repeated measures ANOVA. Finally I compared the results with the physical proprieties of the natural object.

7.2.2.1 DISCRIMINATION THRESHOLD

The subject could select between 57 angle of rotation: the "zero" angle which is the original hue angle of object, 23 angles obtained by rotating clockwise the distribution (positive angles) and 23 angles obtained by rotating counter clockwise the distribution (negative angles). The difference between the zero angle and the other angles is defined as the "angular step", e.g there are three angular steps between the zero and the third positive or negative angle chosen for a certain observer. Thus the ability to discriminate between two stimuli is described in angular steps and the discrimination threshold is expressed in angular steps (i.e. one angular step is one unit of discrimination). Table 6.5 gives an example of the angular steps for the potato.

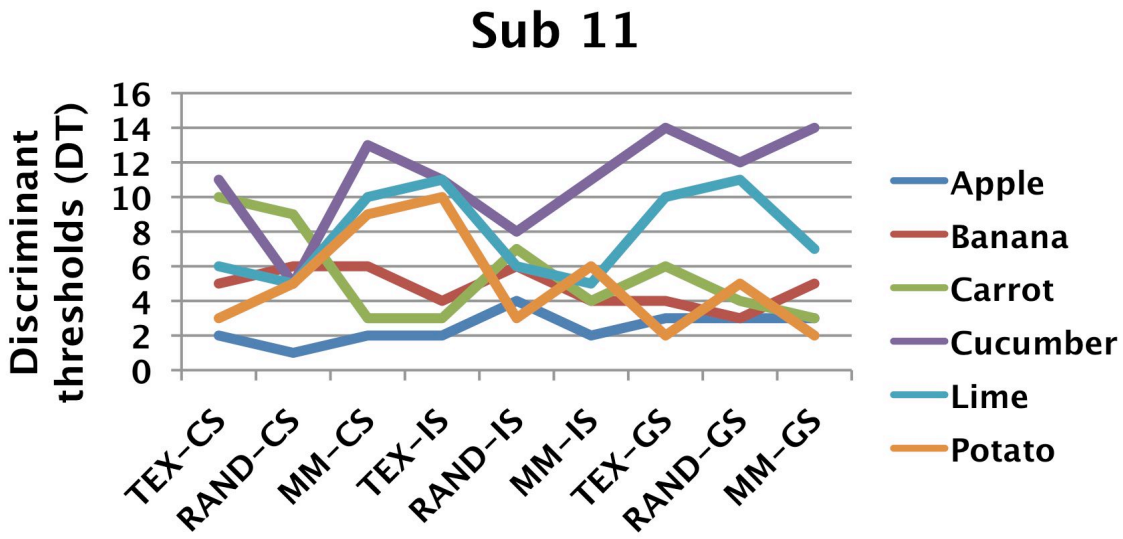

Figure 7.16 – Discrimination thresholds for one subject (sub 11) for all objects and condition combinations. The discrimination unit is one angular step. Each line represents one object for all conditions. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition.

Figure 7.16 shows an example of discrimination thresholds (DT) for one subject (subject 11) for all objects and condition combinations (negative rotation), while Figure 7.17 shows the discrimination thresholds for one condition combination (TEX-CS) for all subjects and objects (negative rotation). Furthermore Figure 7.18 shows the discrimination thresholds for one object (Cucumber) for 4 subjects (subjects 1 to 4) and condition combinations (negative rotation). Through visual inspection we can observe from both figures that there is: (1) great variability between

objects of discrimination threshold (Figure 7.16), (2) great variability between subjects of discrimination threshold (Figure 7.17), and (3) a strong correlation between the 4 subjects' discrimination thresholds in different conditions for the cucumber object (Figure 7.18).

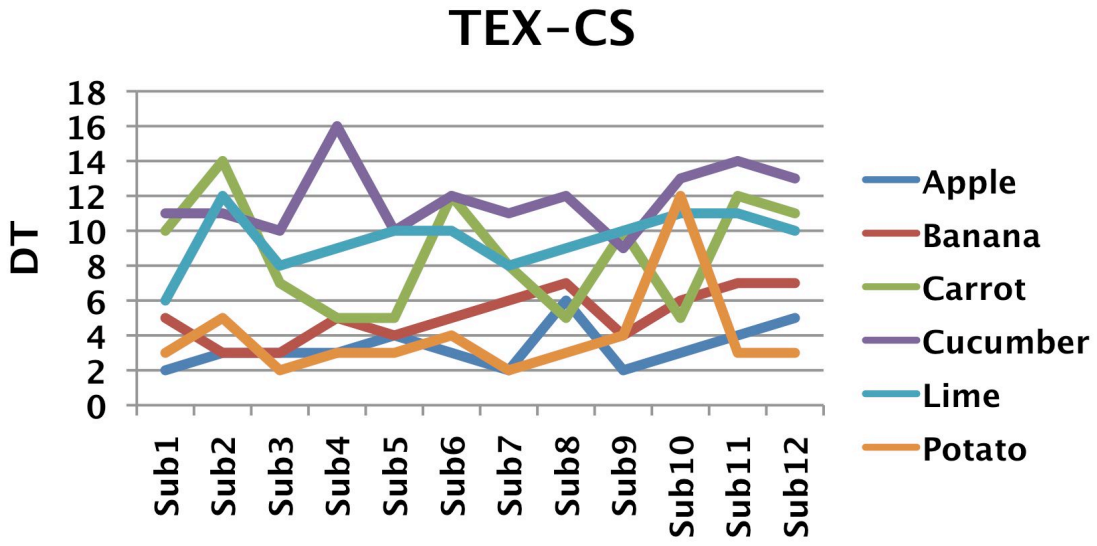

Figure 7.17 - Discrimination thresholds for one condition combination (TEX-CS) for all objects and subjects (N=12). The discrimination unit is one angular step. Each line represents one object for all conditions. TEX-CS: natural or diagnostic texture (TEX) condition combined with congruent shape (CS) condition.

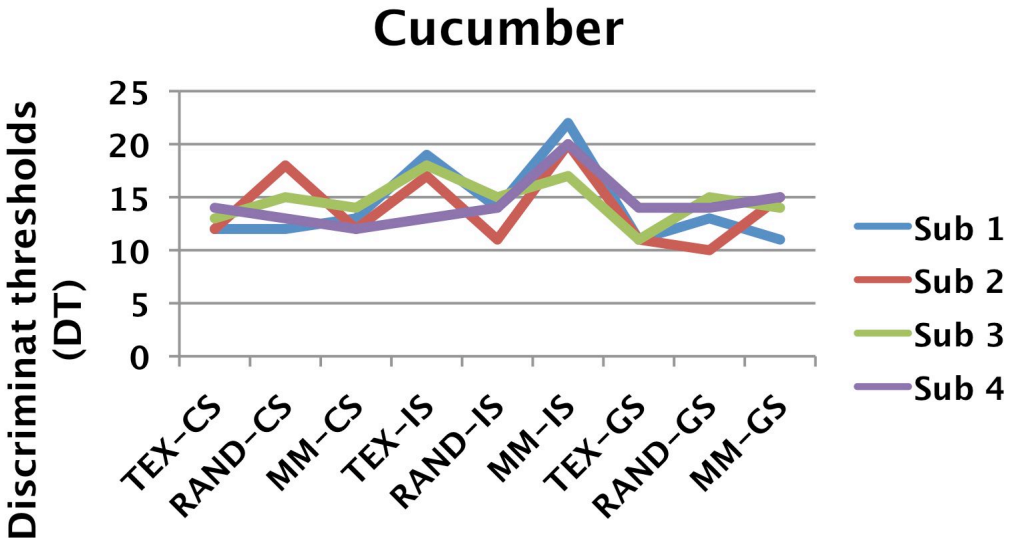

Figure 7.18 - Discrimination thresholds for the cucumber object for 4 subjects and all condition combinations. The discrimination unit is one angular step. Each line represents one subject for all conditions. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition.

To visually compare different factor conditions, I define "mean discrimination threshold" (MDT) for a combination X as the average across all trials belonging to the combination X of shape/chromatic condition over all objects and subjects. MDT can be express by the following formula:

$$MDT_X = \frac{\sum_{i=1}^{n}\sum_{i=1}^{n} DT_{ij}\Big|_X}{N * M}$$

Equation 7-3

where DT_{ij} is the discrimination threshold of subject *i* for the object *j* in combination X, N is the number of subjects (N=12), M is the number of objects (M=6). Figure 7.19 depicts the mean discrimination threshold for all 9 combinations.

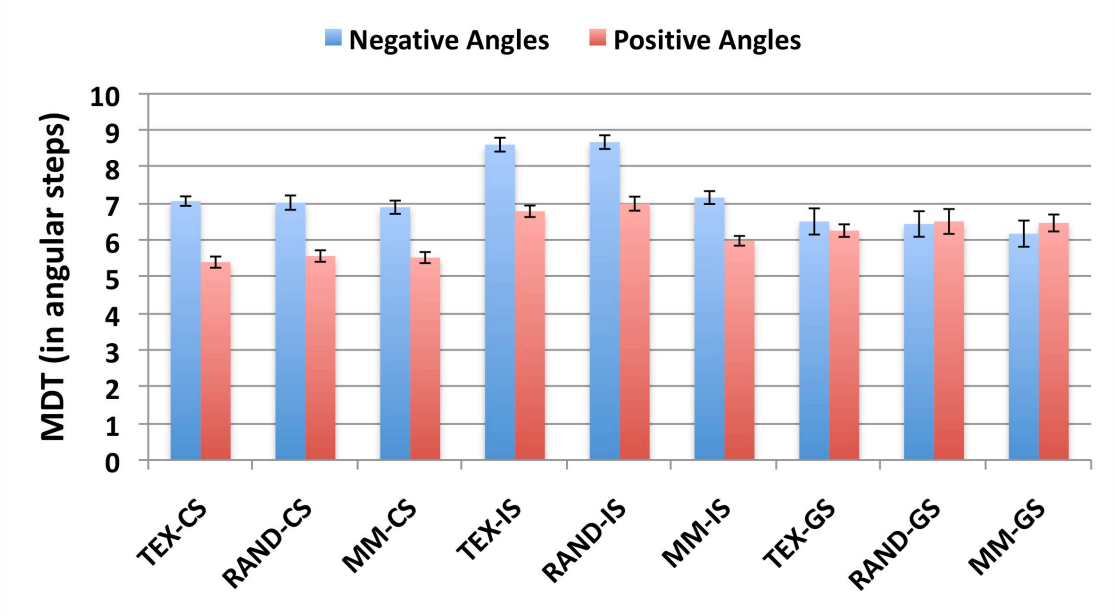

Figure 7.19 – Mean discrimination thresholds over all subjects and objects. Unit: angular steps. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

A 4-way repeated measures ANOVA was performed on the mean response time of all subject for all objects and both directions of rotation. The factors considered were shape, chromatic surface property, object, and directions of the rotation. Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using the using Greenhouse-Geisser estimates. The ANOVA reported as significant the main effect of shape ($F(16.61, 1.51)=16.29$, $p<0.0001$) and chromatic factor ($F=(15,97,1.452)=24.671$, $p<0.0001$) on discrimination thresholds. Note that, if not explicitly stated, the result analyses are for positive and negative rotation combined.

Figure 7.20A shows the average across all discrimination thresholds (DT) belonging to one chromatic condition (TEX, RAND or MM) over all subjects and objects (DTC). Results show that discrimination thresholds for natural texture (or diagnostic texture, TEX) condition are on average 1.41% lower than the random texture (RAND) condition and 6.32% higher than the uniform (MM) condition (μ_{TEX}= 6.77±0.286, μ_{RAND}= 6.87±0.319, and μ_{MM}= 6.37±0.396). Contrasts reported as significant the difference on mean discrimination threshold between TEX and MM condition ($F(1,11)=18.29$, $p<0.002$) and between RAND and MM condition ($F(1,11)=88.36$, $p<0.0001$) implying a significant difference between chromatically variegated surfaces and uniformly coloured surfaces (Figure 7.21A). In fact, mean discrimination thresholds for all chromatically variegated stimuli (combined TEX&RAND) are on average 6.46% significantly higher than for uniform stimuli; Figure 7.21A compares the mean discrimination thresholds of chromatically variegated

stimuli against uniform stimuli. Visually inspecting Figure 7.21A we can notice that all points lie above or on the unity line indicating that in general mean discrimination thresholds are higher for chromatically variegated surface. Conversely, no difference was found between diagnostic and non-diagnostic texture ($F(1,11)=1.752$, $p=0.212$) as illustrated by Figure 7.21B. This graph compares the mean discrimination thresholds of diagnostic textured stimuli (TEX) against non-diagnostic texture stimuli (RAND). By eye we can observe that all points in this figure lie on or about the unity line. *In brief, the discrimination thresholds for the chromatic conditions follow these relationships: MM<TEX<RAND, where MM<TEX and MM<RAND are significant inequalities. This indicates a significant effect of chromatic variegation, but no significant effect of texture diagnosticity.* Note that each subfigure in Figure 7.21 depicts one comparison as described in 7.2.1 and illustrated in Figure 7.2. Specifically Figure 7.21A represents Comparison 1 whereas Figure 7.21B represents Comparison 2 for the mean discrimination threshold.

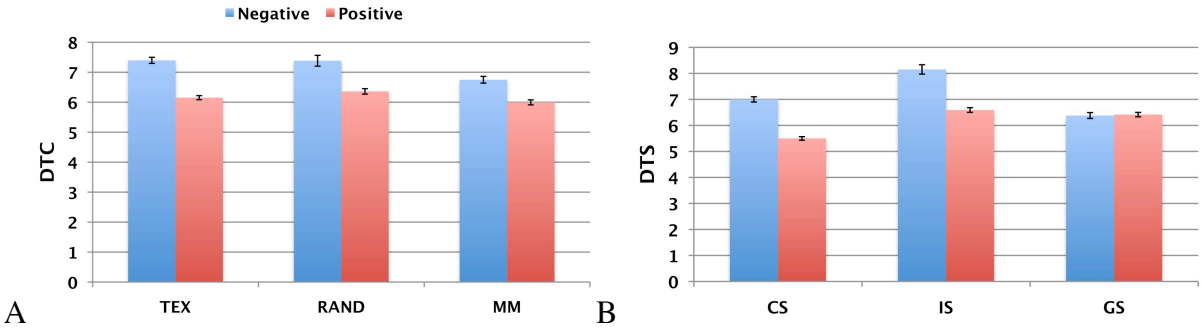

Figure 7.20 – A) Average of the discrimination threshold across all trials belonging to one chromatic condition (TEX, RAND or MM) over all subjects and objects (DTC). B) Average of the discrimination threshold across all trials belonging to one shape condition (CS, IS or GS) over all subjects and objects (DTS). Unit: angular step. Blue bar: negative rotation. Red bar: positive rotation. Error bars represent one standard error of the mean (s.e.m).

Furthermore, Figure 7.20B shows the average across all discrimination thresholds (DT) belonging to one shape condition (CS, IS or GS) over all subjects and objects (DTS). Results show that discrimination thresholds for congruent shape (CS) condition are on average 15.2% lower than the incongruent shape (IS) condition and 2.31% lower than the generic shape (GS) condition[16]. Contrasts revealed that no differences on discrimination thresholds between CS and GS is significant ($F(1,11)=0.35$, $p=0.56$). However it is worth pointing out that Figure 7.21C shows that the positive and negative direction of rotation show opposite trends on discrimination thresholds indicating an interaction between these two factors (see below) and explaining that the lack of a significance that instead is clearly visible. Finally, contrast reported significant the difference on discrimination thresholds between CS and IS condition ($F(1,11)= 63.439$, $p<0.0001$) as shown in Figure 7.21D. *In brief, the discrimination thresholds for the chromatic conditions follow these relationships: CS<GS<IS where CS<IS is significant. These findings indicate that there is a significant effect of shape congruency (*Figure 7.21*D, Comparison 4) on discrimination thresholds.*

[16] Discrimination threshold mean for CS, IS and GS respectively: $\mu_{CS} = 6.250\pm0.249$, $\mu_{IS}= 7.372\pm0.343$, and $\mu_{GS}=6.398\pm0.396$

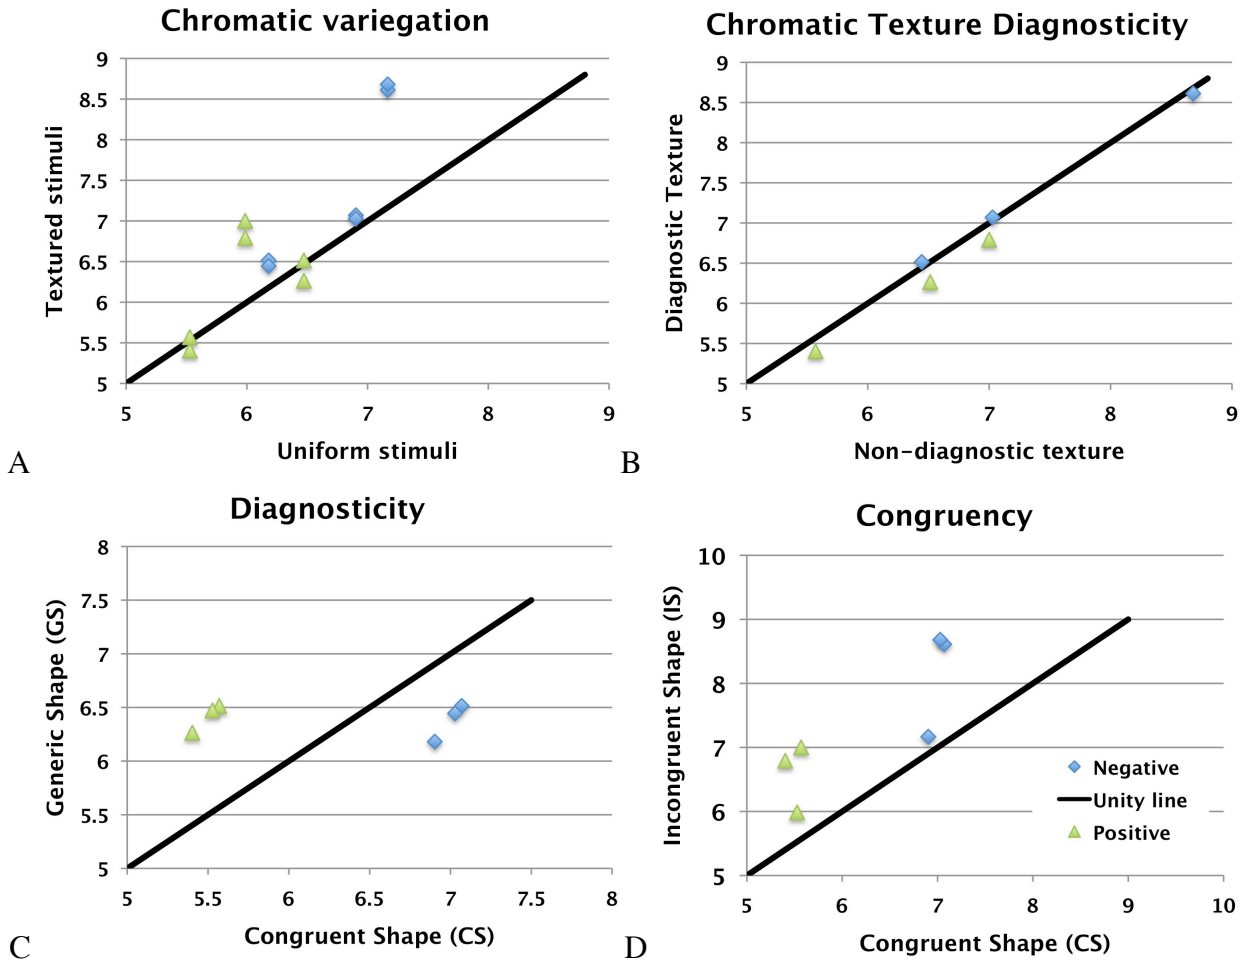

Figure 7.21 – Contrast between mean discrimination thresholds as a function of different chromatic or shape factor conditions. Unit: angular step. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation. Note that the origin of the graphs is [5,5]. Error bars represent the standard error of the mean.

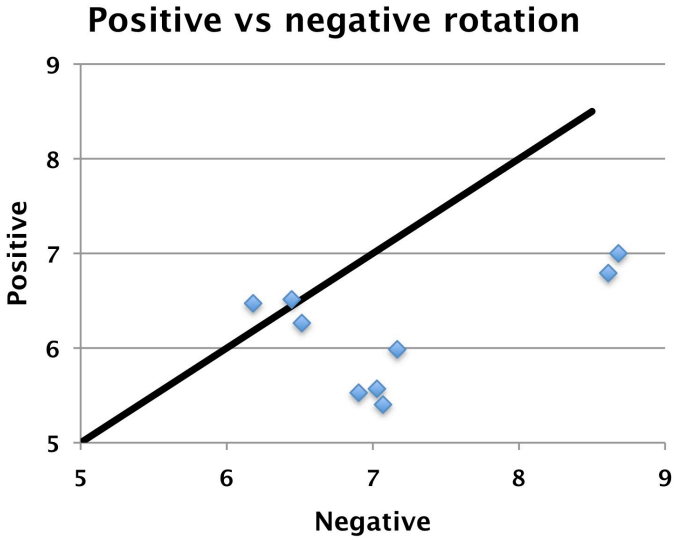

Figure 7.22 – Relationship between discrimination thresholds in opposite directions of rotation. Each point represents one shape/chromatic condition combination. Note that the origin of the graphs is [5,5].

In addition, the ANOVA reported significant the main effect of the direction of the angular rotation (or rotation direction) on discrimination thresholds ($F(1,11)= 30.13$, $p<0.0001$), as depicted in Figure 7.22 (see also Figure 7.20). Moreover there was a significant interaction between rotation direction and chromatic factor ($F(1.89, 20.849=4.711$, $p<0.05$). Three 3-way ANOVAs were performed for each chromatic condition for positive and negative rotation on discrimination thresholds. Results showed significant effects of direction on DT for all three conditions[17]. This indicates that the mean discrimination thresholds of different chromatic conditions are different between positive rotation and negative rotation. To explore further this interaction, contrasts were performed between rotation directions for the comparisons between: (1) diagnostic texture versus non-diagnostic texture (or texture diagnosticity), (2) diagnostic texture and uniform, and (3) non-diagnostic texture and uniform. Contrasts revealed as significant the difference between rotation direction only for the comparison between diagnostic texture and uniform ($F(1,11)=10.688$, $p<0.01$)[18]. Hence mean discrimination thresholds significantly differ for positive and negative rotation only when comparing diagnostic texture with uniform (Figure 7.23A).

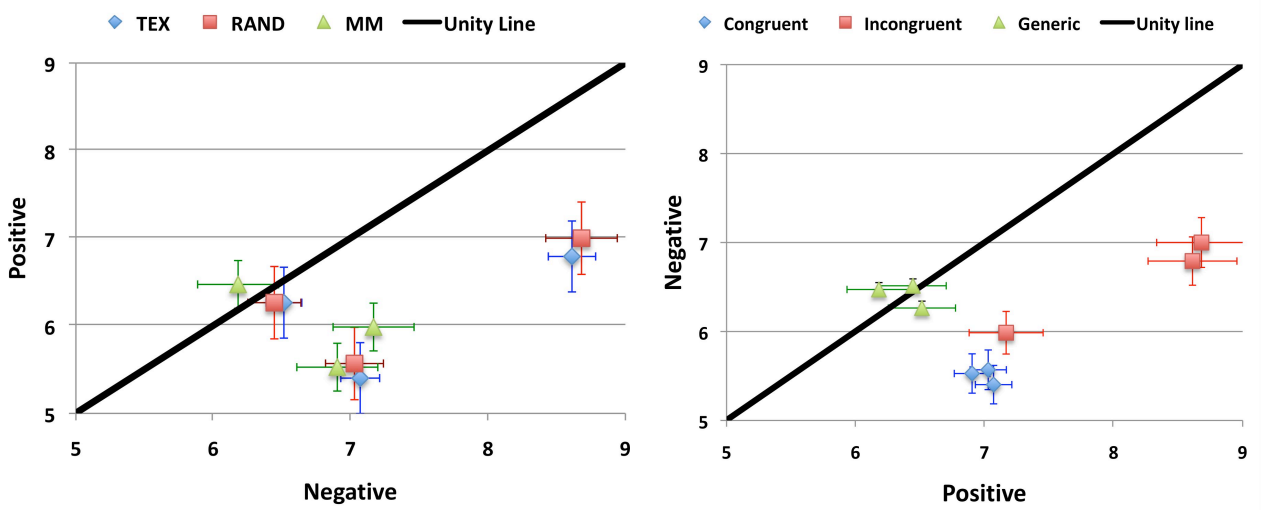

Figure 7.23 – Comparison between mean discrimination thresholds of the positive and negative direction of angular rotation as a function of chromatic condition (A) and shape (B). Each point is one shape/chromatic condition combination. Note that the origin of the graph is [5,5].

Furthermore there is a significant interaction between direction and shape ($F(1.19, 13.08)=16.702$, $p<0.005$) as illustrated in Figure 7.23B. Figure 7.21D shows that mean discrimination thresholds for the congruent shape and incongruent shape condition are significantly higher for negative than rotation. Nonetheless the difference between positive and negative rotation effect on DTM is similar for non-diagnostic shape and diagnostic shape ($F(1,11)= 0.098$, $p=0.76$). Interestingly there is no significant difference on discrimination thresholds between positive and negative rotation for generic shape ($t=0.150$, $p=0.883$; 3-way ANOVA, $F(1,11)= 0.014$, $p=0.907$),

[17] ($F(1,11)=31.39$ for TEX, $F(1,11)=19.028$ for RAND and $F(1,11)=25.51$ for MM, and for all $p<0.005$.

[18] For thoroughness, $F(1,11)= 1.56$, $p=0.238$ for contrast with the comparison diagnostic texture versus non-diagnostic texture, and $F(1,11)=3.174$, $p=0.102$ the comparison non-diagnostic texture and uniform.

whereas such difference is significant for congruent condition (t= 7.344, p<0.0001; 3-way ANOVA, F(1,11)= 32.488, p<0.001) and incongruent shape (F(1,11)= 27.622, p<0.001). *Thus, the discrimination thresholds for generic shape are similar for both direction of rotation, whereas significantly vary for congruent or incongruent shape (*Figure 7.23*B)*. Besides, contrasts revealed no significant interaction between the CS and GS (F(1,11)= 0.098, p=0.76), *indicating that the size of the difference between positive and negative angles is the same between CS and GS*.

Additionally, *there is a significant interaction between shape and chromatic factor* (F(2.809, 30.901)=13.368, p<0.). To investigate this interaction, contrasts were performed on pairs of shape conditions against pairs of chromatic conditions. These analyses reported a *significant interaction between shape congruency and chromatic variegation comparisons on discrimination thresholds*, i.e. Comparison 1 versus Comparison 3[19]. The p-values were respectively for TEX vs MM and RAND vs MM chromatic combinations: p<0.0001 (F(1,11)=27.779) and p<0.05 (F(1,11)= 19.398). Conversely no significant interaction was found between shape congruency and texture diagnosticity comparisons (i.e. Comparison 2 versus Comparison 3; F(1,11)= 0.24, p=0.634). As expected from Figure 7.21, no significant interaction was found for shape diagnosticity (CS versus IS) and chromatic factor comparisons. The p-values were respectively for TEX vs MM, RAND vs MM, and TEX vs RAND condition combinations: p=0.81 (F(1,11)=0.06), p=0.691 (F(1,11)=0.167), and p=0.885 (F(1,11)=0.022).

Finally the ANOVA demonstrated that the discrimination thresholds significantly differ between objects (F(1.825, 20.077)=227.173, p<0.0001; Figure 7.16).

7.2.2.2 DIFFICULTY OF IN PERFORMING THE TASK

An alternative level of analysing the observer's response is to evaluate the facility in performing the task in a specific condition. This can be explored examining for each condition combination: (1) the mean of time that the subject requires to respond per trial (section 7.2.2.2.1), and (2) the number of mistakes made by the subject (7.2.2.2.2) in the specific condition combination. The combination of these two variables is directly related to the difficulty in performing the task. Next sections will describe the factors' effects on these two variables.

7.2.2.2.1 Mean response time

Let's define the "response time" as the time that a subject required to select a match from the instant the specific stimuli was presented in a certain trial. The average of the response times

[19] In other words: [chromatic variegated texture versus uniform] against [diagnostic versus non-diagnostic shape]

(RS) over all trials belonging to the same factor condition performed by *one* observer is defined as the mean response time (MRS) of the observer, and can be expressed by the following formula:

$$MRS_X = \frac{\sum_{j=1}^{K} RS_{jX}}{K} \qquad \text{Equation 7-4}$$

where RS_{iX} is the response time of one subject to trial *i* in condition X and K is the number of trial for that condition. In addition we can average the mean response time across all subjects and objects for one condition X obtaining the "global response time for condition X". Figure 7.24 illustrates the global response time (RT) for all conditions. Note that in the following figures RT is measured in seconds.

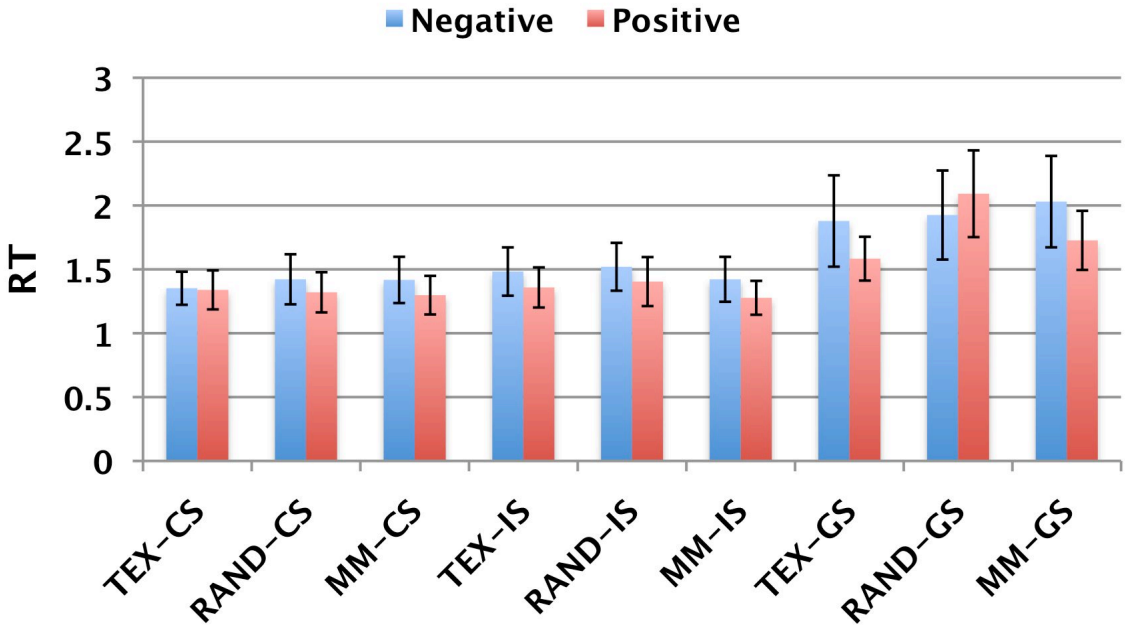

Figure 7.24 – Average of the mean response times across all subjects and objects. Unit: seconds. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

A 4-way repeated measures ANOVA was performed on the mean response time of all subject for all objects and both directions of rotation. The factors considered were shape, chromatic surface property, object, and direction of the rotation. Figure 7.25 shows the 4 comparisons as described in 7.2.1 and illustrated in Figure 7.2. Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using the using Greenhouse-Geisser estimates. The ANOVA reported as significant the main effect of shape on mean response times ($F(1.059, 11.652)=6.79$, $p<0.05$). Results showed that mean response times for generic shapes are 30.15% higher[20] than for congruent shapes ($F(1,11)=7.59$, $p<0.05$) and 27.45% higher than for incongruent shape ($F(1,11)=6.17$, $p<0.05$). Figure 7.25C compares the effect on global mean

[20] Global mean response time CS=1.358±0.16, IS=1.41±0.17, GS=1.94±0.35.

response time of congruent shape against generic shape condition; note that all points lie above the unity line. On the other hand, Figure 7.25C compares the effect on global mean response time of congruent shape against incongruent shape condition showing that the two condition effects do not differ ($F(1,11)=0.358$, $p=0.562$). *In brief, the mean response times as a function of the shape conditions follow this relationship: CS<IS<GS, where CS<GS and IS<GS are significant inequalities.*

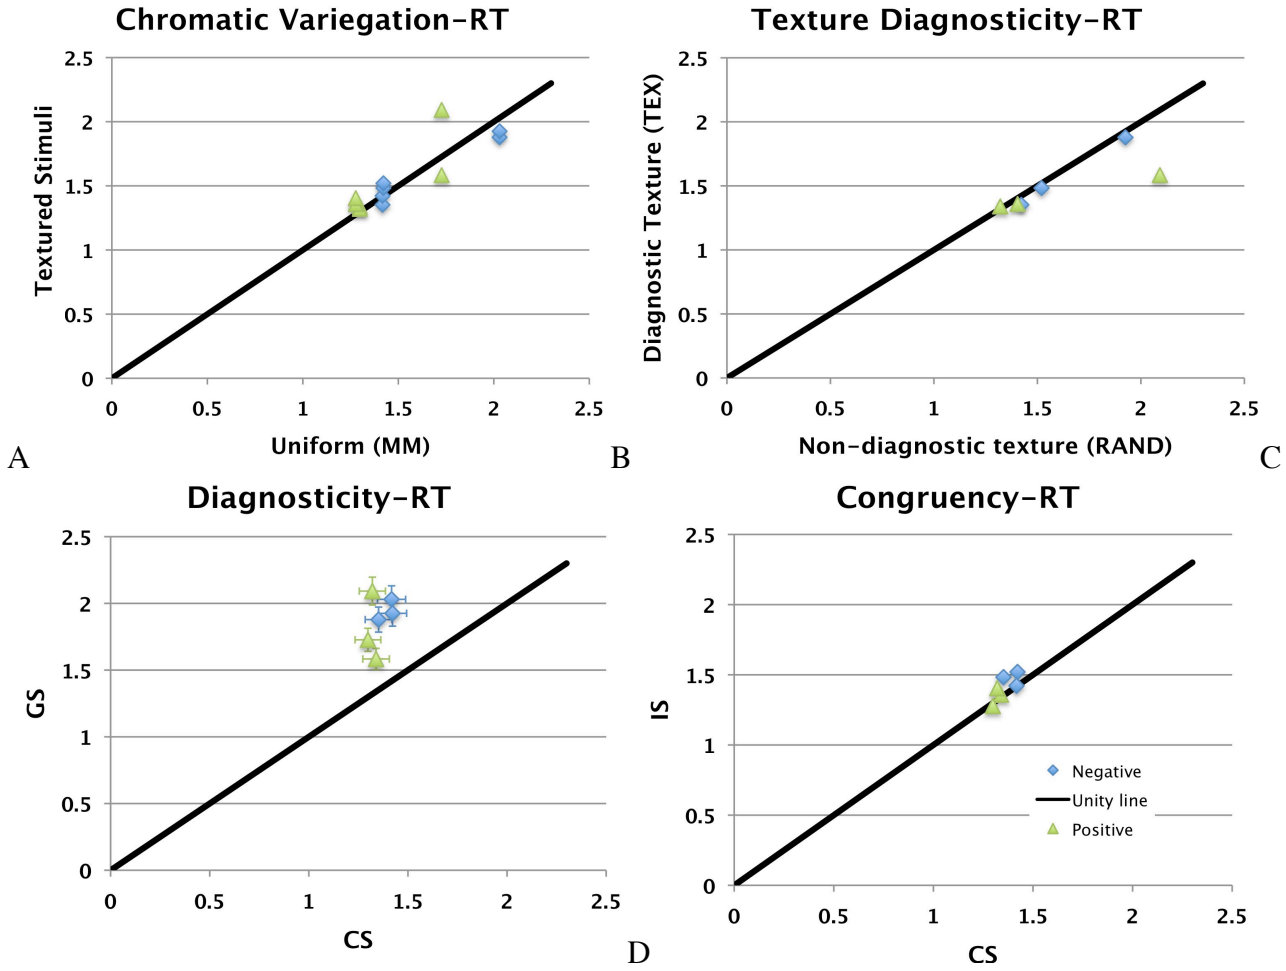

Figure 7.25 – Contrast between global response times as a function of different chromatic or shape factor conditions. Unit: seconds. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation. Error bars represent the standard error of the mean.

Conversely, there is *no significant main effect of chromatic factor* on mean response times ($F(1.079, 11.874)=1.113$, $p=0.318$), as shown in Figure 7.25A and B. Contrasts revealed no interactions.

Finally, second level contrasts reported a *significant interaction between shape diagnosticity and texture diagnosticity on mean response time* ($F(1,11)=5.36$, $p<0.05$) indicating that these two effects influence each other (see Discussion 7.2.3.3). Lastly, mean response times are significantly

affected by the main effect of object (F(1.589,17.48)=8.943, p<0.005), whereas there is no effect of rotation direction (Figure 7.26, F(1,11)= 1.655, p=0.225).

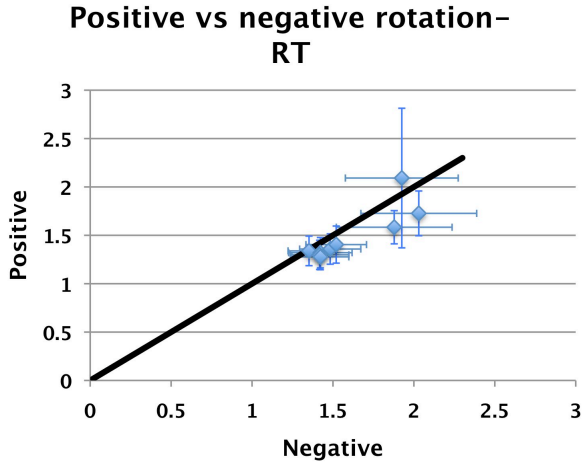

Figure 7.26 – Relationship between response time means in opposite directions of rotation. Each point represents one shape/chromatic condition combination.

7.2.2.2.2 Number of incorrect responses

An additional way to analyse the response is to calculate the number of trials in which a subject made an incorrect response per condition combination ("number of incorrect responses", IR). The mean number of incorrect responses over all subjects and objects for one shape/chromatic condition combination is defined as the mean number of incorrect responses (NM). Figure 7.27 shows the mean number of incorrect responses (NM) for all 9 combinations.

Figure 7.28 compares the effects on mean number of incorrect responses for all the comparisons described in 7.2.1 and illustrated in Figure 7.2. A 4-way ANOVA reported as significant the main effect of chromatic factor (F(1.864, 20.507)=19.207, p<0.05). Specifically there is a significant effect of chromatic variegation (Figure 7.28A, Comparison 1) on the number of incorrect responses. In fact, contrast analysis reported p-values below 0.05 for both TEX versus MM and RAND versus MM (F(1,11)= 8.91, p<0.05 and F(1,11)=14.586 6.773, p<0.01). On average the number of incorrect responses for uniform condition (μ=5.02±0.47) is 5.25% significantly lower than diagnostic texture (μ=5.28±0.53) and 7.97% significantly lower than the non-diagnostic texture condition (μ=5.42±0.45). Beside, the number of incorrect responses for diagnostic texture is 2.52% lower than non-diagnostic texture condition. Nevertheless this difference is not significant, F(1,11)= 1.477 and p=0.25 (Figure 7.28B, Comparison 2). *In brief, the number of incorrect responses as a function of the chromatic conditions follows these relationships: MM<TEX<RAND where MM<TEX and MM<RAND are significant inequalities*. This indicates a significant effect of chromatic variegation, whereas no significant effect of texture diagnosticity. However, additional analyses showed that this effect is mainly due to the contribution of the incongruent shape (3-way ANOVA p>0.05 for CS and GS and p<0.02 for IS).

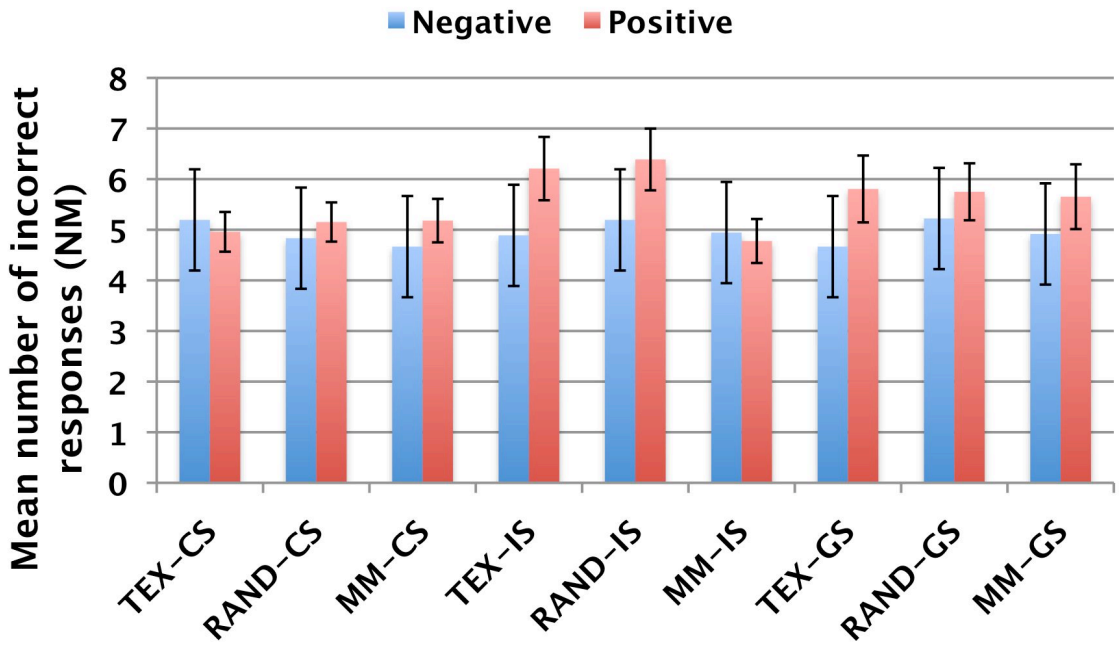

Figure 7.27 – Mean number of incorrect responses over all subjects and objects. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

Furthermore, there is no significant main effect of shape ($F(1.63, 17.94)=0.263$, $p=0.107$) on the number of incorrect responses, although contrasts showed a significant effect of congruency ($F(1,11)= 7.556$, $p<0.02$)[21]. Thus the number of incorrect responses for congruent shapes is on average 20.68% significantly smaller than incongruent shapes and 19.17% smaller than generic shape ($F(1,11)=7.022$, $p<0.05$). *In brief, for the number of incorrect responses, the inequality CS<IS is significant.*

No interaction was found between shape and chromatic factor on number of incorrect responses ($F(2.899, 31.888) =2.475$, $p=0.081$). However there is a significant interaction between shape congruency and the comparison RAND versus MM ($F(1,11)= 5.737$, $p=0.036$).

There is a significant main effect of direction of rotation ($F(1,11)= 18.309$, $p<0.05$). However there is no effect of rotation on chromatic factor ($F(1.564, 17.204)=1.078$, $p=0.346$) or on shape factor ($F(1.976, 21.737)=1.87$, $p=0.178$). *This indicates that there is no difference between direction rotation for different shapes or different chromatic conditions,* but there is an interaction between direction of rotation and certain combinations of shape and chromatic factors ($F(2.596, 28.554)=3.795$ $p<0.05$)[22]. Finally there is a significant main effect of object ($F(4.232, 46.556)=4.037$, $p<0.01$)

[21] NM - CS=4.283±;0.45 IS=5.400±0.5; GS=5.335±0.3

[22] For example, congruent shape and diagnostic texture induce a difference in the effect of rotation on number of incorrect responses with respect to incongruent shape and uniform ($F(1,11)= 15.959$, $p<0.05$).

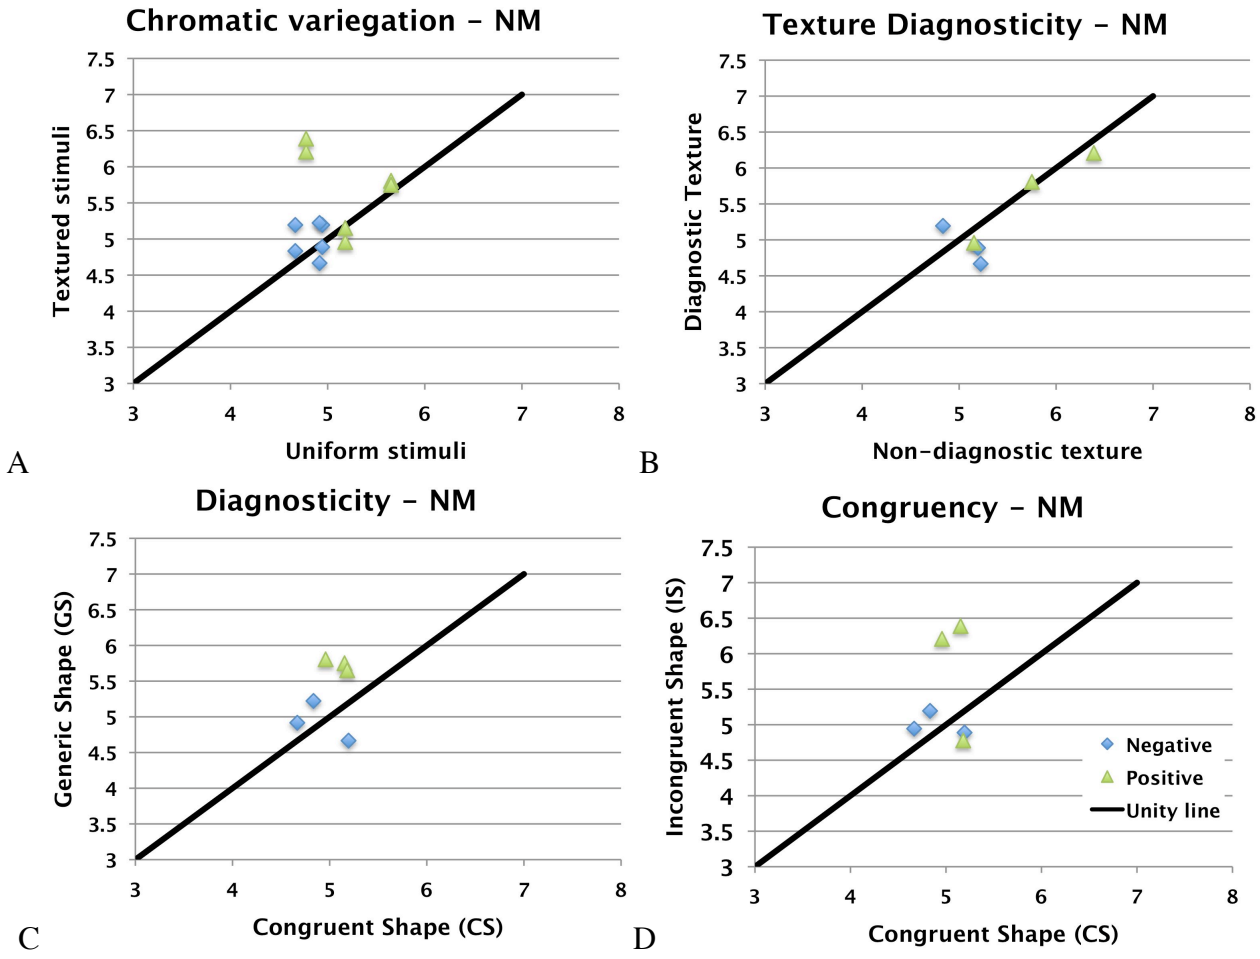

Figure 7.28 - Contrast between numbers of incorrect responses as a function of different chromatic or shape factor condition. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation. Note that the origin of the graphs is [3,3]. Error bars represent the standard error of the mean.

Note that the number of correct responses increase directly with the number of incorrect responses as for each incorrect match the subject had to perform 9 correct trials to reach the same point in which s/he made an error. Thus, the measure of incorrect responses is effectively a measure of the number of excursions around the threshold that the subject had to perform before converging on the threshold, or the number of total steps taken on the staircase

7.2.3 DISCUSSION

The main goal of Experiment 1 was to assess the effect of shape and surface chromatic variegation on chromatic discrimination. Performance was evaluated on two levels: (1) discrimination thresholds and (2) difficulty in performing the task. Here I discuss the results, and compare them with previous studies and with other chapters in this thesis.

7.2.3.1 *EFFECT OF SURFACE CHROMATIC VARIEGATION ON CHROMATIC DISCRIMINATION*

The effect of surface chromatic variegation of an object was investigated in Experiment 1. One group of stimuli consisted of chromatic textures, while a second group consisted of uniformly coloured surfaces. The chromatic variegated stimuli were textures of natural familiar objects or spatially-randomized versions of such textures. Results showed that discrimination thresholds were significantly higher for chromatically variegated stimuli than for uniformly coloured stimuli. The magnitude of this difference depended on the type of the object (for discussion see section 7.2.3.4).

This finding is in line with the studies on chromatic discrimination using chromatic texture stimuli (te Pas & Koenderink (2004), Hansen et al. (2008), and Giesel et al. (2009)). Unfortunately only one study evaluates the effects of natural textures or of textures resembling a natural chromatic distribution (Hansen et al. (2008)) and can in some extent be compared to Experiment 1. However a confirmation of the validity of the effect of chromatic variegation found in Experiment 1 can be indirectly extrapolated also from the other few studies. te Pas and Koenderink compared the discrimination thresholds of different distributions along the red-green, blue-yellow, cyan-magenta axis in RGB space. They reported higher discrimination thresholds for textured stimuli compared to uniform stimuli. However, due to the difference in colour space (the RGB space entangles luminance and pure chromatic variation) and the choice of the direction of variation, a direct comparison with Experiment 1 in this chapter is difficult. Similarly Hansen et al (2008) and Giesel et al. (2009) demonstrated a significant effect of chromatic texture compared to uniform colours. In the second study, they tested two groups of stimuli consisting of: (1) uniformly coloured disks or (2) chromatically variegated disks. The chromaticitiy distribution of the variegated stimuli varied along one of two perpendicular directions (45°–225° or 135°-315°) in DKL space. Their results showed that the effect of chromatic texture on discrimination thresholds was stronger in the direction of the variation of the chromatic distribution of the stimuli when the test location was at the adaptation point. An effect was also found away from the adaptation point where the chromatic distribution "thresholds were dominated by an elevation of thresholds in the direction of the shift away from the adaptation point" (Giesel et al. (2009)). Similar results were obtained in Hansen et al. (2008) for images of natural textured object and synthetic chromatic textures, although in this study the effect of texture at test points away from the adaptation point did not show a significant effect. A major difference between these studies and ours is that the entire distribution was rigidly shifted to the test points and not rotated as in our case. Thus although both studies compared discrimination thresholds measured as the distance between the means of two distributions, the axis of the distribution varied in different ways around the means in the two studies. Furthermore, we

are testing thresholds orthogonal to the axis of variation of the object distribution, i.e. orthogonal to the angle of rotation that has components along both the RG and BY axis. Despite the above differences between studies, all experiments, *including the one presented in this chapter*, demonstrate that the discrimination threshold for chromatically textured stimuli is significantly higher relative to uniformly coloured stimuli.

Furthermore the task resulted more difficult for altered chromatic texture than uniform texture in that the observers made more mistakes in the latter case. Facing a decision the observer is more uncertain for a richer stimulus (chromatic variegated condition) than for a simpler, thus he/she had to be exposed to the same stimuli more times before converging. In other words, if X is the subject discrimination threshold for a specific condition, if the subject makes 5 successive incorrect responses exactly at position X, we will obtain exactly 5 reversals and the specific staircase would end. Clearly this subject for this condition is very precise as s/he fails to see a difference always at the same point. On the contrary the same observer in another condition would be very imprecise if s/he makes mistakes in different points, although mostly at his/her threshold. As the subject was instructed to respond as quickly as possible, the increase in the variability of the response reflects that the stimulus required more decisions in the same short time. Keep in mind that the observers were also keen in completing the task quickly suggesting the reason for similar mean time response across conditions.

7.2.3.1.1 Diagnostic chromatic texture or just chromatic variation?

Why are discrimination performances for chromatic textures poorer? This is one of the crucial questions that arises at the end of Experiment 1. Is it because more chromaticities are involved and therefore more comparisons? Or is that we are performing a surface-matching task instead of a direct colour matching task, and we accept that the same surface might have intrinsic variability in its surface texture which allows for inexact matches between patches of the same surface? Or is there a cognitive factor involved? Experiment 1 proved that, if we corrupt the spatial relationships between chromatic elements of the texture but maintain its overall chromatic distribution, discrimination thresholds between the natural and the altered texture conditions do not differ. According to these results then there is no effect of diagnostic chromatic texture on discrimination, but only an effect of surface polychromaticity. This result is also in line with the work of Hansen et al. (2008), in which no difference was found for discrimination between photographs of naturally textured objects and synthesised chromatic textures. In view of these considerations the difference between chromatic discrimination performances cannot be attributed to higher-level mechanisms such as the object's memory colour. Most likely instead is due to low-

level features as chromatic distribution spread and the spectral power distribution of the stimulus alone.

7.2.3.2 THE EFFECT OF SHAPE

The most intriguing result of Experiment 1 is the effect of shape on discrimination performance. Two groups of shapes were examined in this experiment: (1) a generic non-diagnostic shape and (2) a set of 6 natural shapes. Each chromatic stimulus was projected onto the generic shape and onto two of the natural shapes. A natural shape X was defined as congruent when the mean chromaticity of its colour stimulus was generated from an object having the same shape X; otherwise the shape was defined as incongruent. The data demonstrated a strong effect of shape congruency on discrimination threshold. Discrimination thresholds increased by 15% relative to the congruent shape when the colour stimuli were displayed on the incongruent shape. On the contrary there is no difference between congruent and generic shape. Since the generic shape was presented before any other stimuli we may consider the discrimination threshold for this shape as the baseline for our group of observers. Hence, we concluded that the ***incongruent*** *"object colour-object shape" stimulus deteriorates the subject's ability to discriminate between chromatic surfaces relative to the baseline*.

Equally we might conclude that, since there is no overall difference in discrimination thresholds between congruent and generic shape, then the presence of a congruent shape does not facilitate the task of surface discrimination. This conclusion is invalidated, though, by the strong interaction between shape diagnosticity and direction of the angular rotation. While the positive and negative rotations for the generic shape lead to the similar values of discrimination thresholds (Figure 7.21C), for congruent shapes (and similarly incongruent shapes. Figure 7.23) the discrimination thresholds vary depending on direction. The reason for these puzzling results might be that the presence of a natural shape introduces a memory colour effect.

First, it should be pointing out that the chromatic distributions of the objects in this chapter are located in the third or fourth quadrant of the EMG colour space (see Figure 6.4 and Figure 7.29). As Figure 7.29 shows, in EMG isoluminant plane the colours along the x-axis vary between red (or positive x-axis) and green (negative) and along the y-axis vary between purple (or positive y-axis) and yellow-greenish (negative). Therefore a small counter-clockwise rotation (i.e. a negative rotation) of the objects in the third quadrant will result in a change of the hue of each point of the distribution toward the yellow-greenish, while a small clockwise rotation gives a change towards the greenish-blue. Note that the objects in the third quadrant are the cucumber and lime whose memory colour is usually greenish. Similarly a counter-clockwise (negative) rotation of object distributions in the fourth quadrant such as the carrot's will result in a change in hue towards the

red, while a clockwise (positive) rotation gives a change towards the yellow-green. Again note that the memory colour of a carrot is orange-red. Hence on average, for these objects, a positive rotation moves away from the memory colour of the object and therefore leads to easier discrimination between the two stimuli. Figure 7.30 illustrates this point by plotting the distribution of discrimination thresholds for positive and negative directions for the TEX-CS case as a function of object. Note that all objects, except the potato, show a significant difference between directions of rotation (T-test, $p<0.05$). Specifically, for the apple, carrot, cucumber and lime, the negative rotations have higher thresholds and, for the banana, positive rotation has higher threshold. Finally the potato shows no significant difference as it has a rather neutral memory colour.

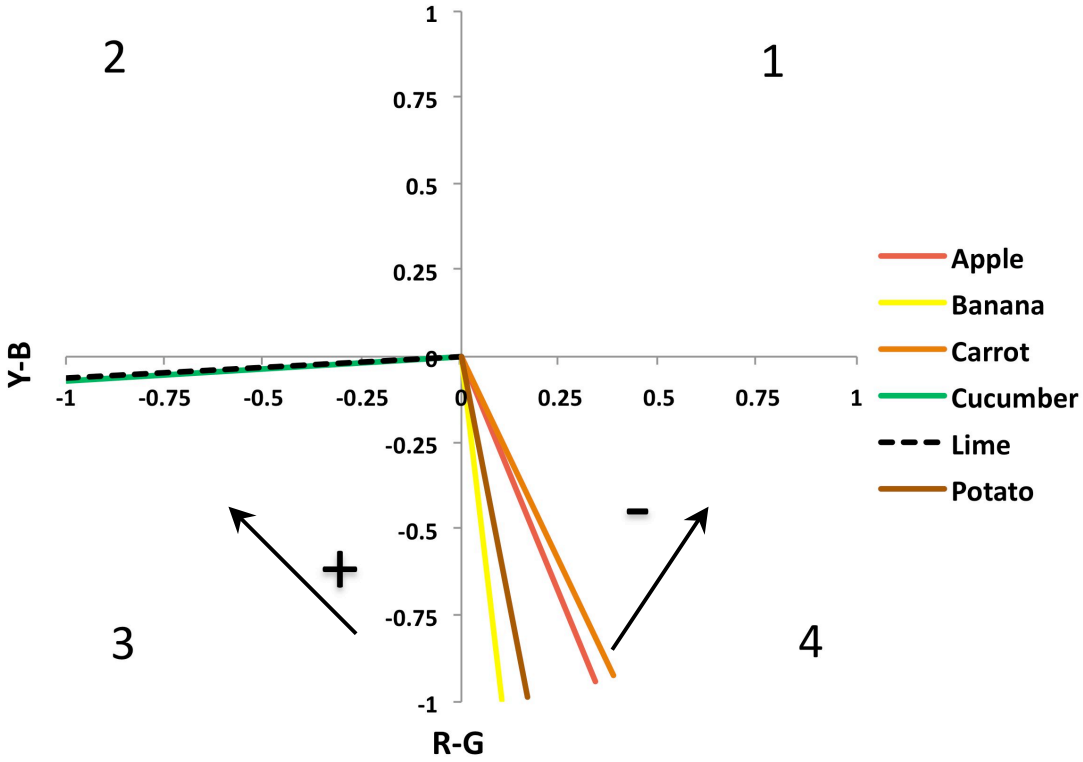

Figure 7.29 – Schematic of the 4 quadrants, hue angles of the objects and directions of the rotation.

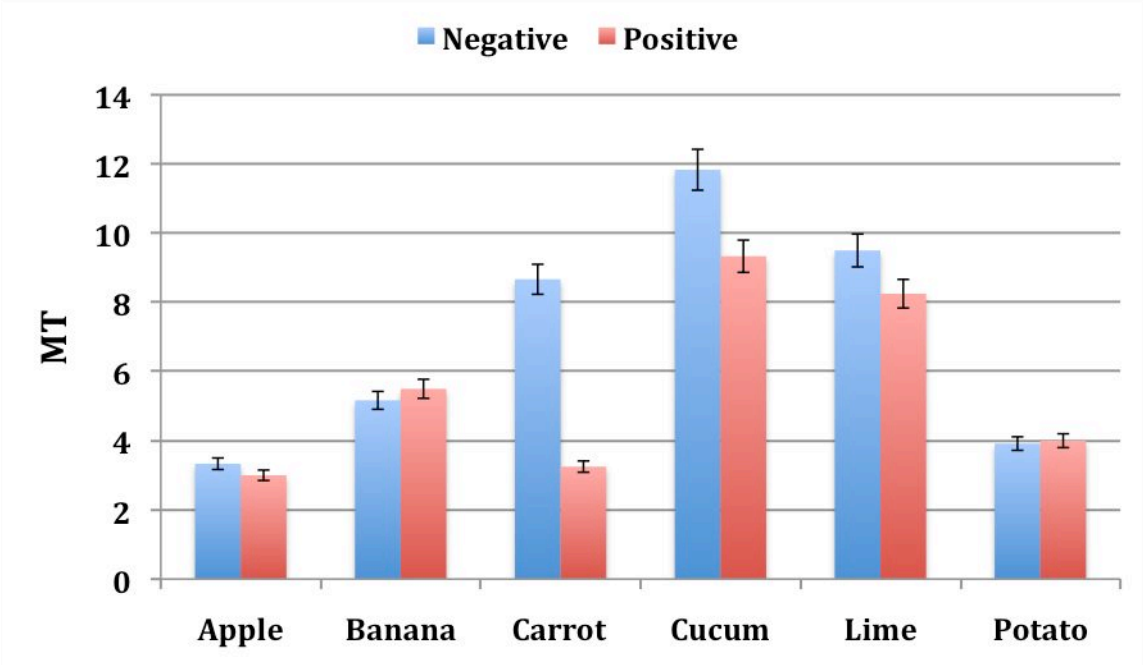

Figure 7.30 – Positive (red bars) and negative (blue bars) discrimination thresholds in the TEX-CS combination for all six objects. Error bars: s.e.m.

The task is usually harder in the generic shape condition and easier for the natural shapes; the mean response time for congruent shapes is the smallest. Analogously, the number of incorrect responses is about 20% smaller for congruent shapes compared to incongruent shapes. Therefore the subject's performance is significantly facilitated by the presence of the object's original shape. This is independent of the direction of the angular rotation. Therefore the recognition of the object shape facilitates the discrimination of its colour.

Finally note that all objects in the experiment were presented in 3D. In other words, all the cues to the real object shape were given to the observer. Can we assume that similar results will be found in a 2D presentation? Provided with the knowledge I have acquired from the previous experiments in this thesis, I hypothesise that if the 2D stimulus provides an identifiable cues to the actual shape of the object in the same way and amount as the 3D shape then this assumption is correct.

7.2.3.3 I SAY TOMATO, YOU SAY POTATO. OR ABOUT THE INCONGRUENCY BETWEEN CHROMATIC TEXTURE AND SHAPE CUES TO OBJECT IDENTITY.

An additional important finding of Experiment 1 is the interaction between natural chromatic texture and shape. In fact, *there is a significant interaction between shape and chromatic factor on thresholds, specifically* between shape congruency and chromatic variegation on discrimination threshold. When two opposite cues of the object identity (its chromatic distribution

and its shape) compete then the subject's chromatic discrimination threshold increases (Figure 7.19). Moreover, the results on mean response time showed an interaction between shape diagnosticity and texture diagnosticity. As a matter of fact, mean response times are lower for the congruent object shape-texture combination that any other chromatic-generic combination Therefore I postulate that congruency between identity cues enhance the observer ability to discriminate between stimuli, whereas the incongruency between such cues reduce the subject discrimination performance. Experiment 2 will confirm these results, showing an even stronger effect of texture diagnosticity (see General discussion, 7.4.2). Thus I can conclude that high-level mechanisms are involved in the observer behaviour.

In previous chapters I have cited various studies (e.g. Duncker (1939)) that show the effect of diagnostic shape alone on colour appearance as the memory colour of the object influenced the observer perception. Additionally, I have shown in Chapter 6 that the texture of the stimuli affects the subject colour perception. However, to the author's knowledge, this is the first study that contrast shapes and textures of natural object, presenting the same textures on different natural shapes and evaluating the effect on the subject visual colour perception.

7.2.3.4 EFFECT OF OBJECT COLOUR

It is well known that colour and hue discrimination thresholds vary across colour space (e.g. MacAdam (1942), Wright&Pitt (1934), Newtow&Eskew (2003), Eskew et al. (2009)). For example, "green-blue" colours are poorly discriminated while yellow-red colours are more easily discriminated. Although in choosing the object stimuli I aimed to select the largest range of objects in colour and shape, nonetheless the majority of the object colour distributions are located in the third or fourth quadrant of the EMG cone contrast space (note that there are fruit and vegetable objects which have a blue-violet typical colour, e.g. blackberries). Therefore our objects' colours range from green(ish) for the cucumber and lime to reddish-yellow for the apple. It is then appropriate to see how and in what degree this variation in colour has affected the object thresholds. Figure 7.31 shows the object mean discrimination threshold averaged across positive and negative directions of rotation as a function of the hue angle (Table 6.4). As expected, the discrimination thresholds of the object vary with its hue angle, with the "green" objects having on average the highest thresholds.

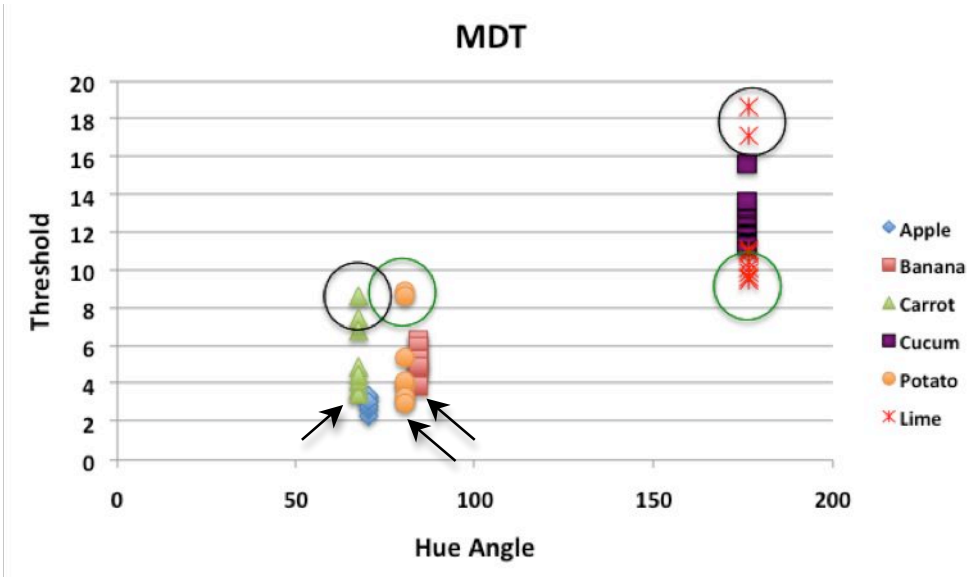

Figure 7.31 – Mean discrimination threshold for all objects in all conditions averaged across direction as a function of the object hue angle. Black arrows indicate the MM-GS condition for the banana, the carrot and the potato. Black and green open circles are explained in the text.

Nonetheless Figure 7.31 is interesting for several reasons. First it shows that objects of slightly different hues might have rather different thresholds in the same condition or, vice versa, objects of rather different hues might have similar thresholds in the same condition (black arrows indicates for the MM-GS condition the banana, the carrot and the potato). Besides different objects with very different hue angles can have the same threshold in different conditions (e.g. green open circles in Figure 7.31: TEX-IS potato and MM-GS lime). In addition, and most importantly *the discrimination threshold difference between two identical conditions for one object can be almost twice as larger as than for other object*. In fact, the black open circles in Figure 7.31 indicate the TEX-IS and RAND-IS condition for both carrot and lime; note that their differences with respect to the MM-GS thresholds are very different, the lime having a larger range than the carrot. Lastly, objects such as the apple have a much smaller variation in threshold. Therefore it is clear that other properties of the object colour play a role in discrimination. For example, we could plot the object thresholds as a function of their original chromatic variation index OV of the object (Table 6.3). Figure 7.32 shows that there is a direct relationship between OV and MDT in that object having larger original chromatic range have smaller variation between conditions. Note that before we have said that polychromaticity of the *stimuli* increase the observer discrimination threshold. Here instead I am postulating that the polychromaticity of the original object narrows the difference between condition thresholds (the object between condition variability). In other word, it is possible that the memory of the original object distribution linked to the object colour may play a role in this discrimination task.

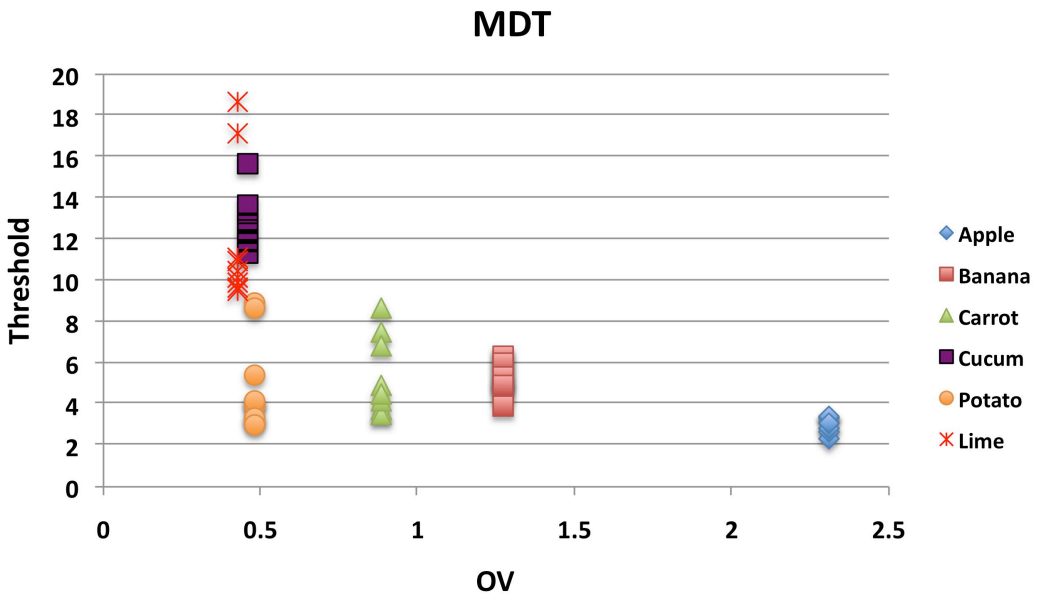

Figure 7.32 - Mean discrimination threshold of all objects in all conditions average across direction as a function of the as a function of object variation index OV.

In a previous experiment we have tested the discriminability between two chromatic surfaces using a rapid successive surface discrimination procedure (Hurlbert et al. (2009)); here I compare the results obtained in that experiment with the finding of Experiment 1. Figure 7.33 compares the discrimination thresholds for natural texture conditions combined with the generic shape in Experiment 1 (red bars, TEX-GS) with the discrimination thresholds of the rapid successive surface discrimination experiment presented in Hurlbert et al. (2009) (blue bars) for 5 object category used in both experiments: apple, banana, carrot, lime, and potato. Note that the objects used in these two experiments are not identical exemplars but they belong to the same categories and, as discussed in Chapter 5, their chromatic surface proprieties possess similar characteristics relative to the purpose of this analysis. For comparison I have multiplied by 5 the threshold from Hurlbert et al. (2009) to fit the scale from Experiment 1. Both experiments return the similar trends.

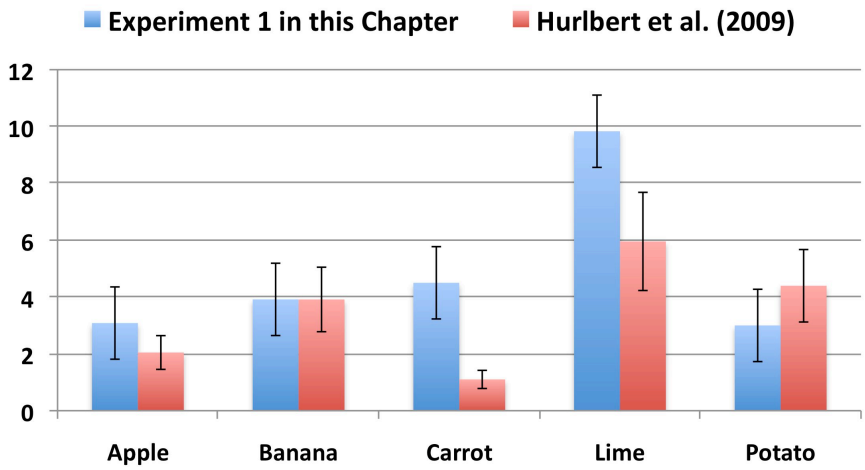

Figure 7.33 – Comparison between Experiment 1 and Hurlbert et al. (2009) study. Error bars: s.e.m.

7.2.4 CONCLUSION

Experiment 1 demonstrated that surface discrimination performance declined (i.e. discrimination thresholds increased) when the stimulus was chromatically variegated compared to uniform and also declined if it was presented on a shape incongruent with the object's colour (a red lime for example). In addition, *generally*, the presence of a natural shape facilitated the task, especially if the shape was congruent. However there was no significant difference between discrimination performances for natural diagnostic textures and randomized non-diagnostic textures, which leads to the conclusion that this effect was most likely due to low level features of the stimuli.

We can consider Experiment 1 as the ground floor of a building. In Experiment 1 we tested at the lowest level the effect of chromatic variegation and shape on surface discrimination. Experiment 2, in the next section, was designed to test the observer's ability to discriminate between two alternatives when a delay between the reference stimuli and the alternatives is introduced, i.e. when memory is involved.

7.3 EXPERIMENT 2

In Experiment 1, I have demonstrated the effect of polychromaticity and shape on simultaneous discrimination performance. In Experiment 2, I investigate the same factors and replicate the same set of stimuli that in Experiment 1 using a delayed matching-to-sample paradigm. Hence the main difference between the two experiments is that in Experiment 2 the observer had to memorize and *then* match the object chromatic surface. This section thus follows the same structure as section 7.2 (Experiment 1).

Unlike Experiment 1, in Experiment 2 I evaluate the effect of shape and texture on a successive discrimination task, which involves working memory.

7.3.1 METHODS

In this second experiment I have used objects belonging to the same categories that Experiment 1, with the exception of the lime that was substituted with a pear (object 5). Specifically the cucumber and the potato were the same object used in Experiment 1, while a different apple, banana and carrot were used in this experiment. Figure shows the photographs of the objects used in Experiment 2. Chromatic distributions, hue angles and object variation indices of the apple, banana and carrot were similar to the one used in Experiment 1 (as expected, see Chapter 5). The hue angles and indices for the 4 new objects (pear, apple, banana, and carrot) are listed in Table 7.7.

The chromatic distribution of the pear is shown in Figure 7.35. The artificial replica of the pear covered a visual angle of 6.14x2.86 degree and is shown in Figure 4.3.

Figure 7.34 - Thumbnails of the six objects' photographs under illuminant D65. From left to right, from top to bottom: apple, banana, carrot, cucumber, pear, and potato.

	Apple	Banana	Carrot	Pear
Hue Angle	72.60	84.20	67.27	89.11
OV*100	1.801	1.042	0.556	0.423

Table 7.7 - Hue angles and object-variation index for the four new objects in Experiment 2 computed as described in Chap 5.

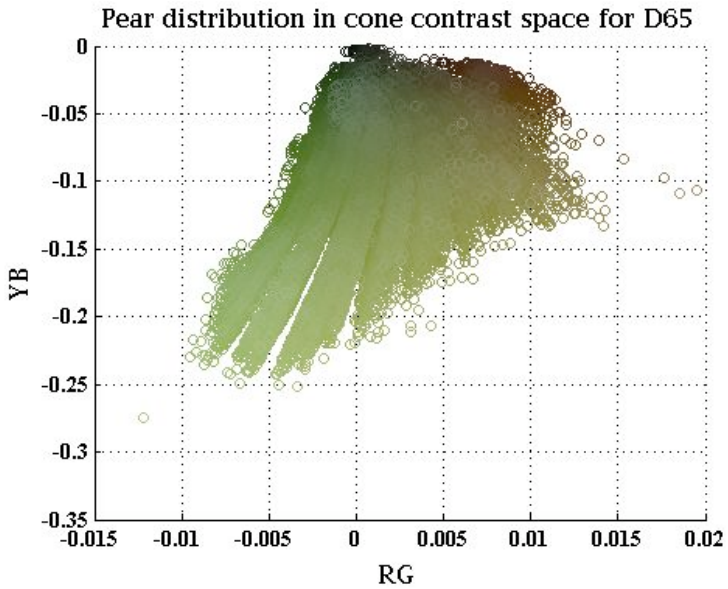

Figure 7.35 – Chromaticity distribution of the pear used in Experiment 2 in EMG cone contrast space under D65. Note that the colour of each point of the distributions in the figure is only indicative of its real colour. Each point is represented only once even if occurring more frequently in the distribution.

The set up employed is the same as in Experiment 1. The identical procedure used to generate the stimuli in Experiment 1 was employed for Experiment 2. Conversely the objects were paired as follows: (1) apple/pear, (2) banana/potato, and (3) carrot/cucumber. The chromatic and shape conditions are the same as in Experiment 1 and listed in Table 7.1. Figure 7.36 gives an

example of the factors' combinations for the pear (shape or texture); the same comparisons as in Experiment 1 can be performed in Experiment 2 (Figure 7.2). Note this figure shows 12 combinations and not 9 to additionally illustrate the case of apple shape/pear surface. Since the apparatus used in this experiment was the same as in Experiment 1 we can make the same consideration as in 7.2.1.2 and 6.2.3.1.

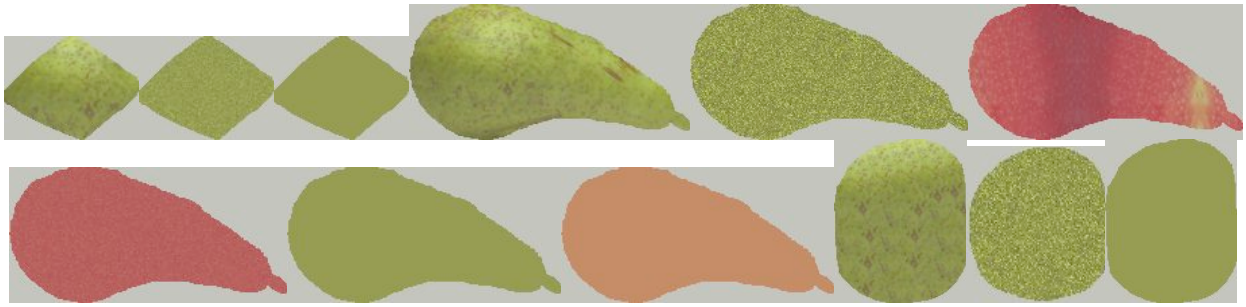

Figure 7.36 – Example of the combinations of all conditions for the pear. From left to right, from top to bottom: pearTEX-GS, pearRAND-GS, pearMM-GS, pearTEX-CS, appleRAND-IS, appleRAND-IS, pearMM-CS, appleMM-IS, pearTEX-IS, pearRAND-IS, and pearMM-IS.

7.3.1.1 OBSERVERS

Nine observers, 8 female and 1 male, three of which participated in Experiment 1, took part in all conditions of Experiment 2. They were all in the age range of 19-24 years and students at Newcastle University. All tested normal on the Farnsworth-Munsell 100 hue test (mean total error score 23.6) and were naïve to the purpose of this experiment. Before starting the experiment, they acquired familiarity with the task during a 4 minutes practice session.

7.3.1.2 PROCEDURE

The general scene shown to the observers was identical to Experiment 1. Two sets of three objects of the same shape were placed on a uniform background. The observer was positioned on a chair, and asked to stabilize his/her head on the chin rest and look inside the experimental chamber; then s/he was handed a 3-button mouse. The experimental room was in complete darkness. The observers had to complete 4 blocks of trials, one for each combination of objects (3 pairs) and one for the generic shape. The generic shape was always presented first, while the sequence of other three object combinations was pseudorandomised, with ~16.5% of the participants shown the same order.

The observers were required to look at a reference stimulus placed in central position on a uniformly coloured background and memorize its chromatic surface. After 5 seconds the reference disappeared (that is, the reference surface appearance reverted to neutral grey, although the solid 3D object remained in place) and a countdown of 5 seconds was displayed on a uniform background (delay time). During the entire experiment the chromaticity of the background remained metameric to illuminant D65 (as in Experiment 1). Then the observer was presented with two alternatives (test

stimuli) respectively placed to the left and right of center. As in Experiment 1, one of the alternatives was identical to the reference stimuli while the other's chromaticity distribution was rotated as in 7.2.1.4.3. The observer's task was to select, as quickly and accurately as possible, the alternative that was *identical* to the reference (such alternative is in this experiment the target stimulus). No time constraints to the selection were given nor feedbacks on the correctness of the response. The stimuli were presented and selected as in Experiment 1 and interleaved adaptive staircases, with the same specifications as in Experiment 1, were employed to determine the subject's performance threshold for each condition of the task. For each trial, the selection and the response time were recorded together with information about the stimulus condition as in Experiment 1. Each half block of the experiment lasted on average between 55 and 75 minutes. In general the two experiments differ only in the task performed by the observer and on the method of presentation of the reference stimulus with respect to the target stimulus (i.e. successive presentation instead of simultaneous presentation of reference stimulus and target stimulus). As in Experiment 1 I will compare different factor conditions.

7.3.2 RESULTS

The primary aim of Experiment 2 was to explore the effect of shape and chromatic texture cues to the identity of an object on successive discrimination (or delayed matching-to-sample) of chromatic surfaces.

As in Experiment 1 the results are analysed on two levels, organized here into two subsections: (1) the matching performance (sub-section 7.3.2.1), and (2) the facility in performing the task in a specific condition (subsection 7.3.2.2). The latter analysis evaluates: (1) the mean time to respond to one stimulus in one trial belonging to a specific condition (subsection 7.3.2.2.1), and (2) the mean number of incorrect responses (subsection 7.3.2.2.2). Here I tested the significance of the effects using a 4-ways repeated measures ANOVA.

7.3.2.1 MEMORY PERFORMANCE

As in experiment 1 the subjects were given 57 distinct angles of rotation as alternatives (the original angle and 56 alternatives) and I can define the difference between stimuli in steps between adjacent angular alternatives. Note that, as each staircase starts from 20 steps away from the original angle, the subject saw at least 41 of these angles for each object. Here I define the performance of a subject as the angular step at which the adaptive staircase converges for 75% of correct responses, and indicated as "memory threshold".

Figure 7.37 shows an example of memory thresholds (MT) for one subject (subject 4) for all objects and condition combination for the negative rotation direction, while Figure 7.38 shows the

memory thresholds for one condition combination (TEX-CS) for all subjects and objects, again for the negative rotation direction.. Furthermore Figure 7.39 shows the memory thresholds for one object (potato) for 4 subjects (subjects 1 to 4) and condition combinations (negative rotation). Through visual inspection we can observe from both figures that there is great variability between objects (Figure 7.37), subjects (Figure 7.38), and different conditions (Figure 7.39) on memory thresholds.

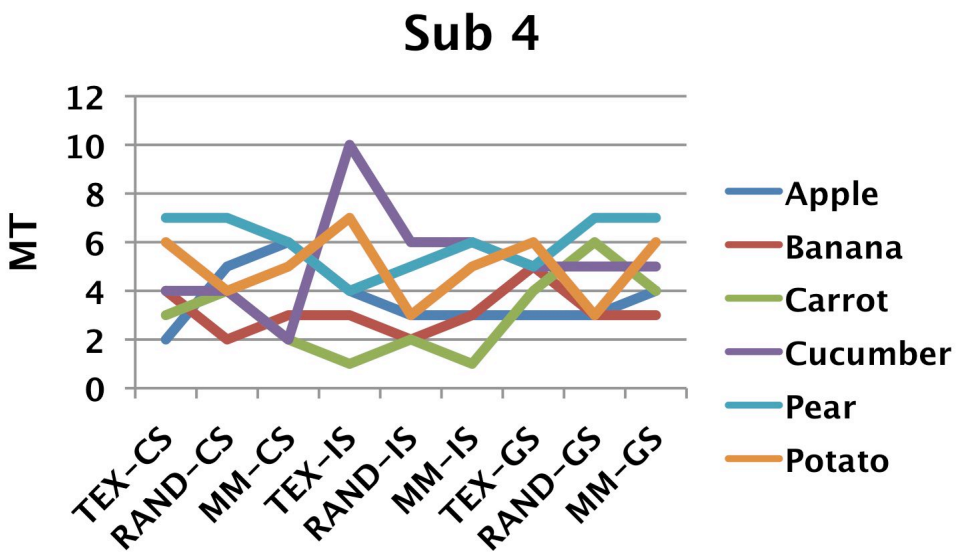

Figure 7.37 – Memory thresholds for one subject (sub 4) for all objects and condition combinations. Each solid line represents one object for all conditions. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition.

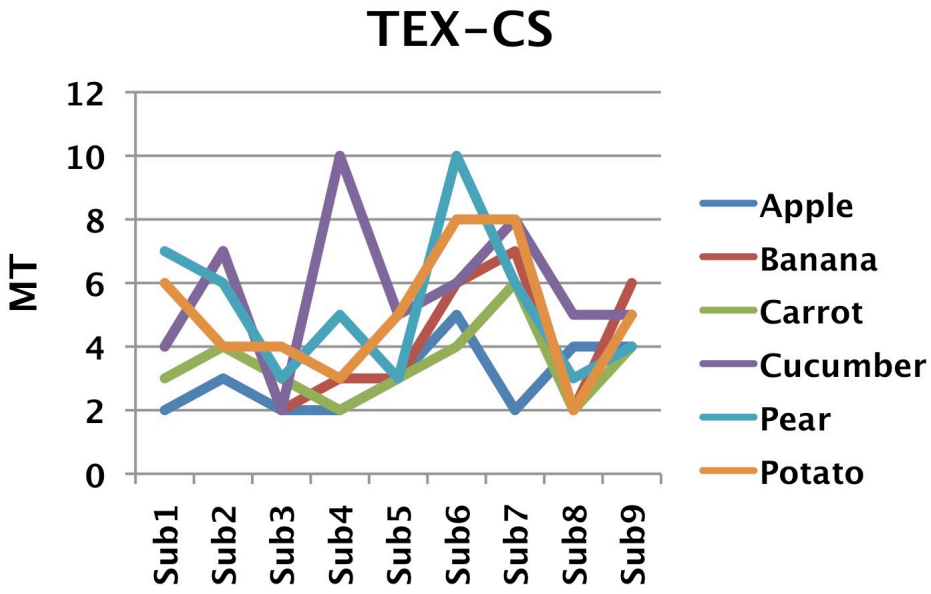

Figure 7.38 - Memory thresholds for one condition combination (TEX-CS) for all objects and subjects (N=12). Each line represents one object for all conditions. TEX-CS: natural or diagnostic texture (TEX) condition combined with congruent shape (CS) condition.

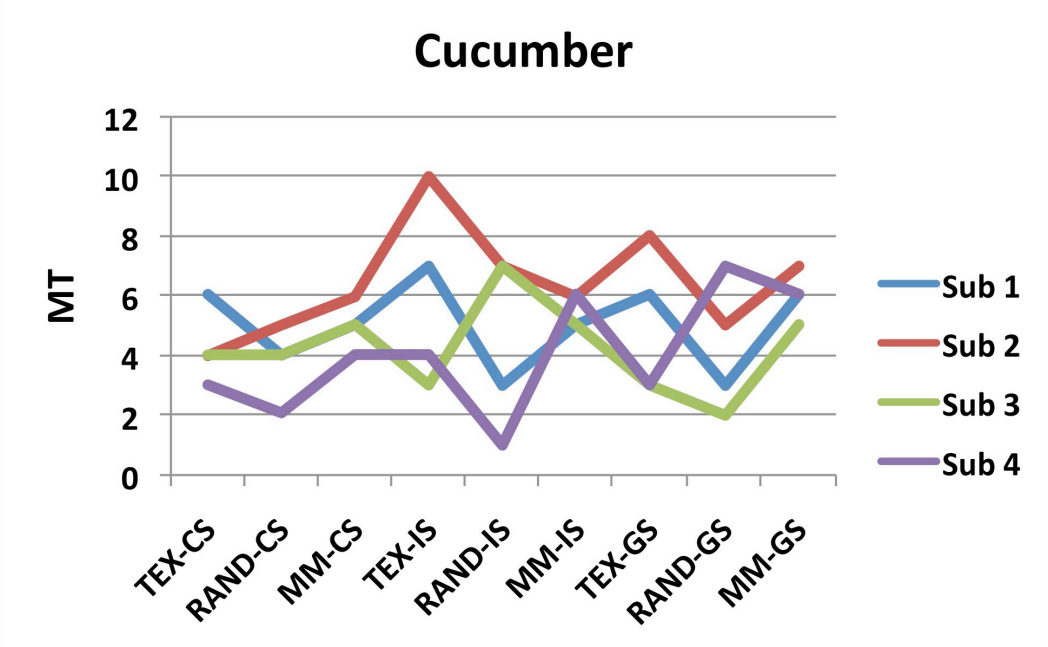

Figure 7.39 - Memory thresholds for the potato object for 4 subjects and all condition combinations. Each line represents one subject for all conditions. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition.

To visually compare different factor conditions, let us define the mean memory threshold (MMT_X) for a combination X as the average of the memory thresholds across all trials belonging to the combination X of shape/chromatic condition over all objects and subjects. This can be expressed by the following formula:

$$MMT_X = \frac{\sum_{i=1}^{N}\sum_{q=1}^{Q} MT_{iq}\Big|_X}{N * Q}$$

Equation 7-5

where MT_{iq} is the memory threshold of subject *i* for the object *q* in combination X, N is the number of subjects (N=12), Q is the number of objects (Q=6). Figure 7.40 depicts the mean discrimination threshold for all 9 combinations. Through visual inspection it is possible to notice that TEX-CS has the smallest threshold. Let us examine these results quantitatively.

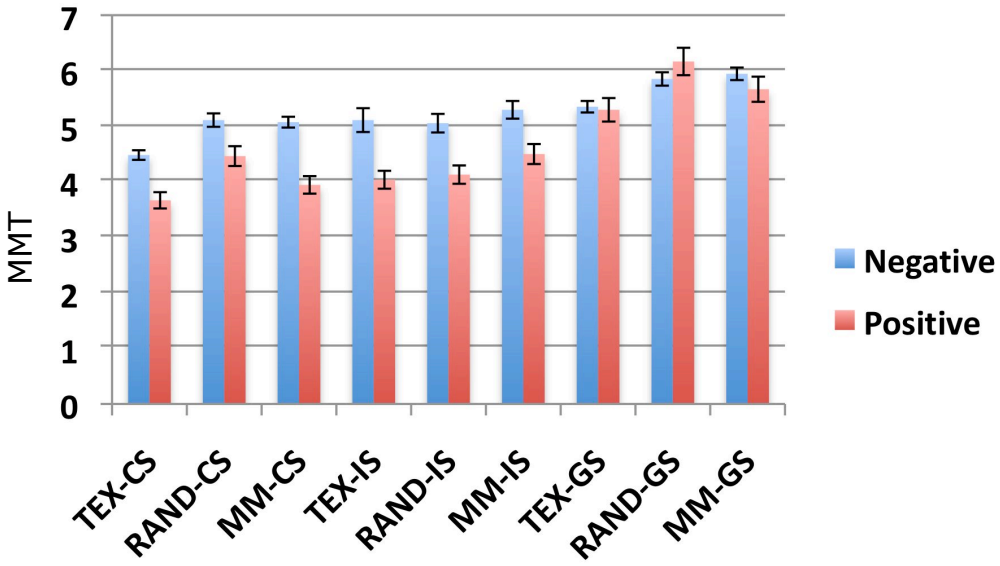

Figure 7.40 – Mean memory threshold over all subjects and objects. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

A 4-way repeated measures ANOVA was performed on the memory thresholds of all subject for all objects and both directions of rotation. The factors considered were shape factor, chromatic factor, object factor, and direction factor. Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using the using Greenhouse-Geisser estimates. The ANOVA reported as significant the main effect of the chromatic factor ($F(1.511, 12.091)=5.823$, $p<0.0001$) on memory threshold. Analogously to Experiment 1, Figure 7.41 shows the comparison between mean memory threshold of different factor's conditions as describe in 7.2.1 and illustrated in Figure 7.2. Specifically each subfigure represents one comparison (Figure 7.41A = Comparison 1; Figure 7.41B = Comparison 2; Figure 7.41C = Comparison 3; Figure 7.41D = Comparison 4).

Contrasts show that memory thresholds for diagnostic texture (TEX) condition are on average 8.16% significantly smaller than the random texture ($F(1,8)=7.11$, $p<0.01$) condition and 9.21% significantly smaller than the uniform condition ($F(1,8)=6.37$, $p<0.01$; $\mu_{TEX}=4.65\pm0.34$, $\mu_{RAND}= 5.13\pm0.39$, and $\mu_{MM}= 5.07\pm0.36$). On the contrary, no significant difference was found between non-diagnostic texture and uniform condition on memory thresholds ($F(1,8)=0.34$, $p=0.575$). These results can be visualized in Figure 7.41A and B. Note that all points in subfigure A and B lie below or on the unity line. *In brief, the memory thresholds for the chromatic conditions follow this* ***significant*** *relationship: TEX<MM<RAND. This implies that there is an effect of texture diagnosticity and chromatic variegation.*

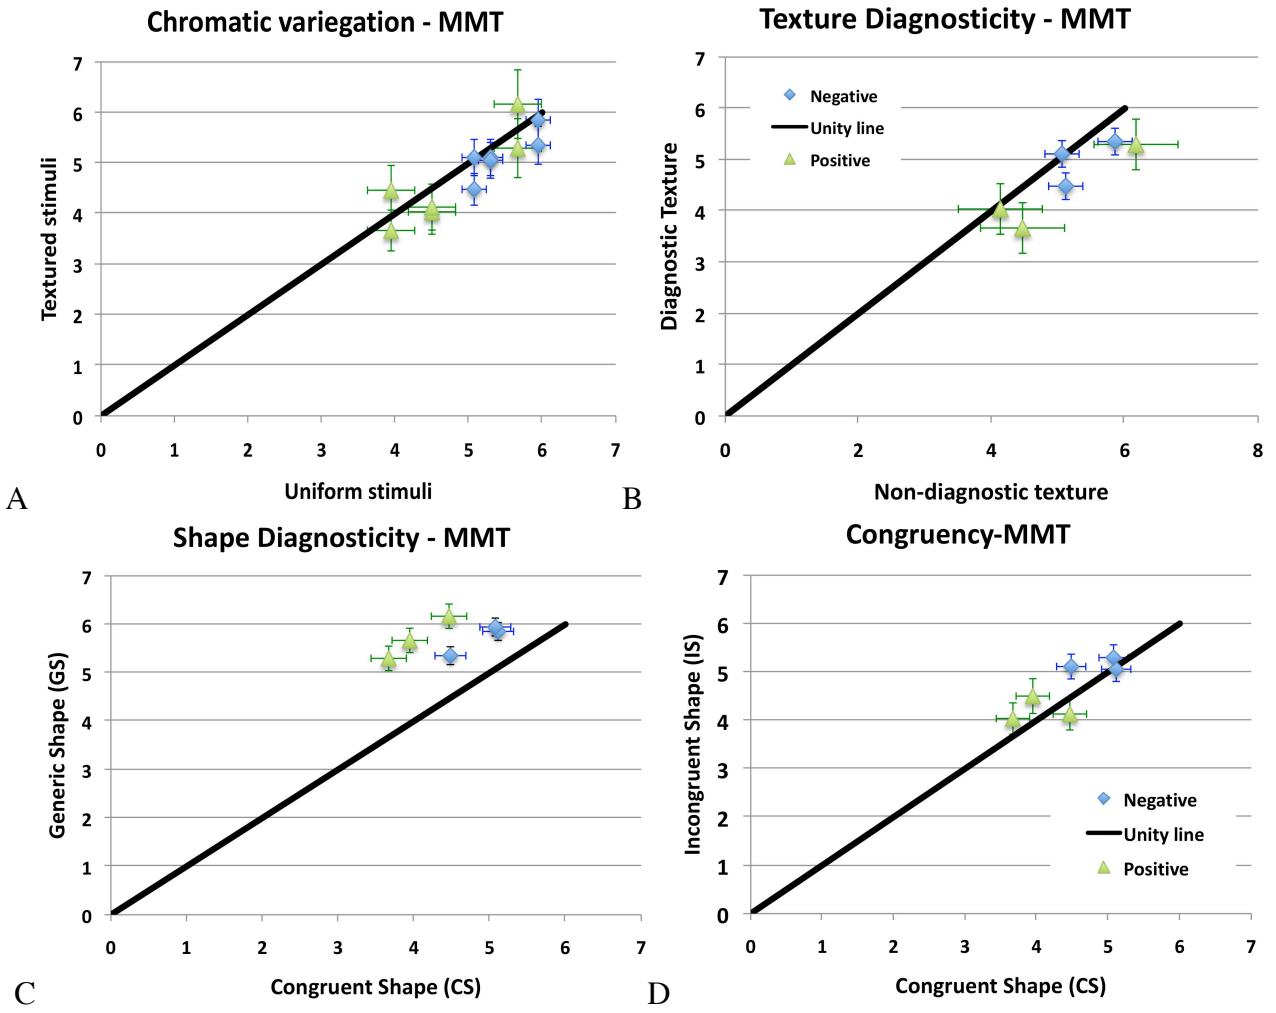

Figure 7.41 – Contrast between mean memory thresholds as a function of different chromatic or shape factor condition. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation. Error bars represent the standard error of the mean.

Furthermore, there is a significant main effect of shape factor on memory thresholds ($F(1.774, 14.19)=47.261$, $p<0.0001$). Yet analysis of contrasts reveal that memory thresholds are significantly different between congruent and generic shapes ($F(1,8)=82.188$, $p<0.0001$), but not between congruent and incongruent shapes ($F(1,8)=4.14$, $p=0.076$). Hence there is an effect of shape diagnosticity (Figure 7.41C), but not of shape congruency (Figure 7.41D). Specifically, memory thresholds for the congruent shape condition are 22.0% significantly lower than generic shape and 4.94% lower than for the incongruent shape (overall mean: μ_{CS} = 4.45±0.315, μ_{IS}= 4.69±0.0.35, and μ_{GS}=5.71±0.41). Besides, memory thresholds are significantly different between incongruent and generic shape conditions ($F(1,8)=6.95$, $p<0.05$). *In brief, the memory thresholds for the shape conditions follow these relationships: CS<IS<GS where CS<GS and IS<GS are significant inequalities. This implies that there is a significant effect of shape naturalness.*

In addition the ANOVA shows a significant main effect of object ($F(2.703, 21.621)=22.678$, $p<0.0001$) and direction of angular rotation ($F(1,8)=15.6$, $p<0.05$). However there is no significant interaction between direction of rotation and chromatic condition ($F(1.5,12.04)=0.94$, $p=0.391$) or shape ($F(1.5,12.2)=4.4$, $p=0.068$) on memory thresholds. Second level contrast analyses revealed significant only the interaction between the comparison incongruent and generic shape and direction of rotation ($F(1,8)= 6.925$).

Finally, the ANOVA reported no significant main interaction between chromatic factor and shape ($F(2.797, 22.379)=2.233$, $p=0.116$). *Nonetheless there is a significant interaction between shape congruency and texture diagnosticity on memory threshold* ($F(1,8)= 7.353$). This indicates the TEX-CS is significantly different from, for example, RAND-shape and TEX-shape, with TEX-CS being the smallest. A significant effect was found between incongruency (IS vs GS) and direction of rotation ($F(1,8)= 6.925$). All other interactions were not significant.

7.3.2.2 DIFFICULTY IN PERFORMING THE TASK

As in Experiment 1, I evaluate the difficulty in performing the task on two levels: (1) the mean duration that the subject requires to respond per trial (section 7.3.2.1), and (2) the number of mistakes made by the subject (section 7.3.2.2) in a specific condition combination. (7.3.2.2.1) (7.3.2.2.2). As said the combination of these two variables is directly related to the difficulty in performing the task.

7.3.2.2.1 Mean response time

The mean response time of the observer is defined as in Equation 7-4. Similarly to Experiment 1, the global mean response time (RT) is the average of all the mean response times across all subjects and objects. Figure 7.42 illustrates the global mean response time (RT) for all conditions.

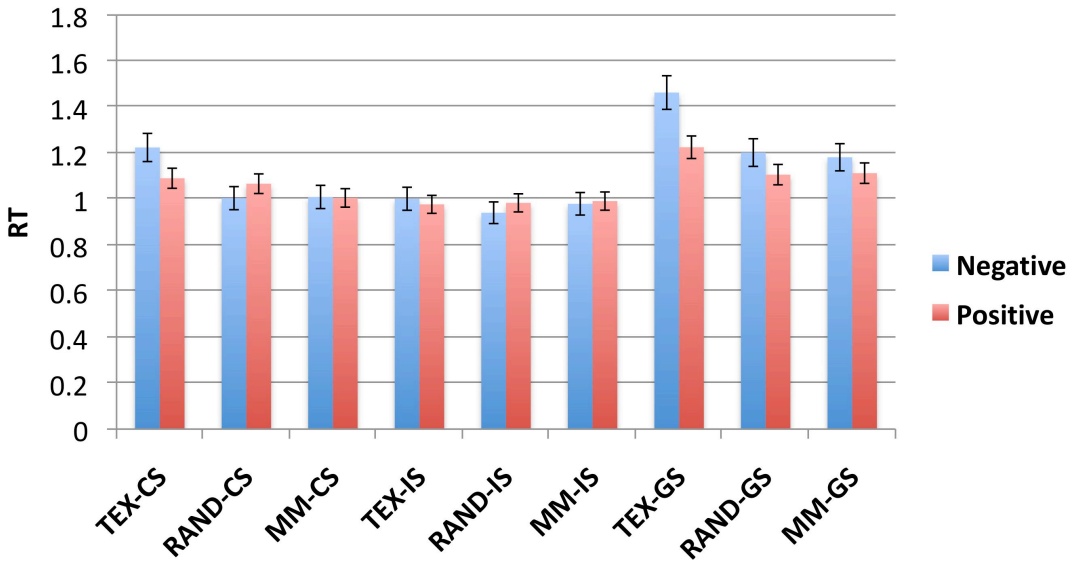

Figure 7.42 – Average of the mean response times across all subjects and objects. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

A 4-way repeated measures ANOVA was performed on the mean response times of all subjects for all objects and both directions of rotation (factors: shape, chromatic variegation, object, and direction). Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using the using Greenhouse-Geisser estimates. The ANOVA reported a significant main effect of chromatic factor on mean response times (F 1.146,10.31)=6.05). Contrasts revealed as significant the effect of texture diagnosticity ($F(1,8)= 6.26$, $p<0.05$) and chromatic variegation ($F_{TEX}(1,8)=5.42$, $p<0.05$). Specifically, mean response times for diagnostic texture condition are 12.71% higher than non-diagnostic texture condition and 9.39% higher than the uniform colour condition[23]. As in Experiment 1, I can compare the effect of different factors' conditions as described in 7.2.1 and illustrated in Figure 7.2 (Figure 7.43). Firstly, note that all the points in Figure 7.43A (Comparison 1 – chromatic variegation) and Figure 7.43B (Comparison 2 – texture diagnosticity) lie above or near the unity line. However no difference was found between non-diagnostic texture condition and uniform condition ($F(1,8)=0.008$, $p=0.929$). *In brief, the mean response times for the chromatic conditions follow these relationships: TEX>MM>RAND where TEX>MM and TEX>RAND are significant inequalities. This implies that there is an effect of texture diagnosticity and of chromatic variegation of the natural texture.*

[23] Global mean response time: TEX=1.17; RAND=1.04; MM=1.072 (RAND<MM<TEX)

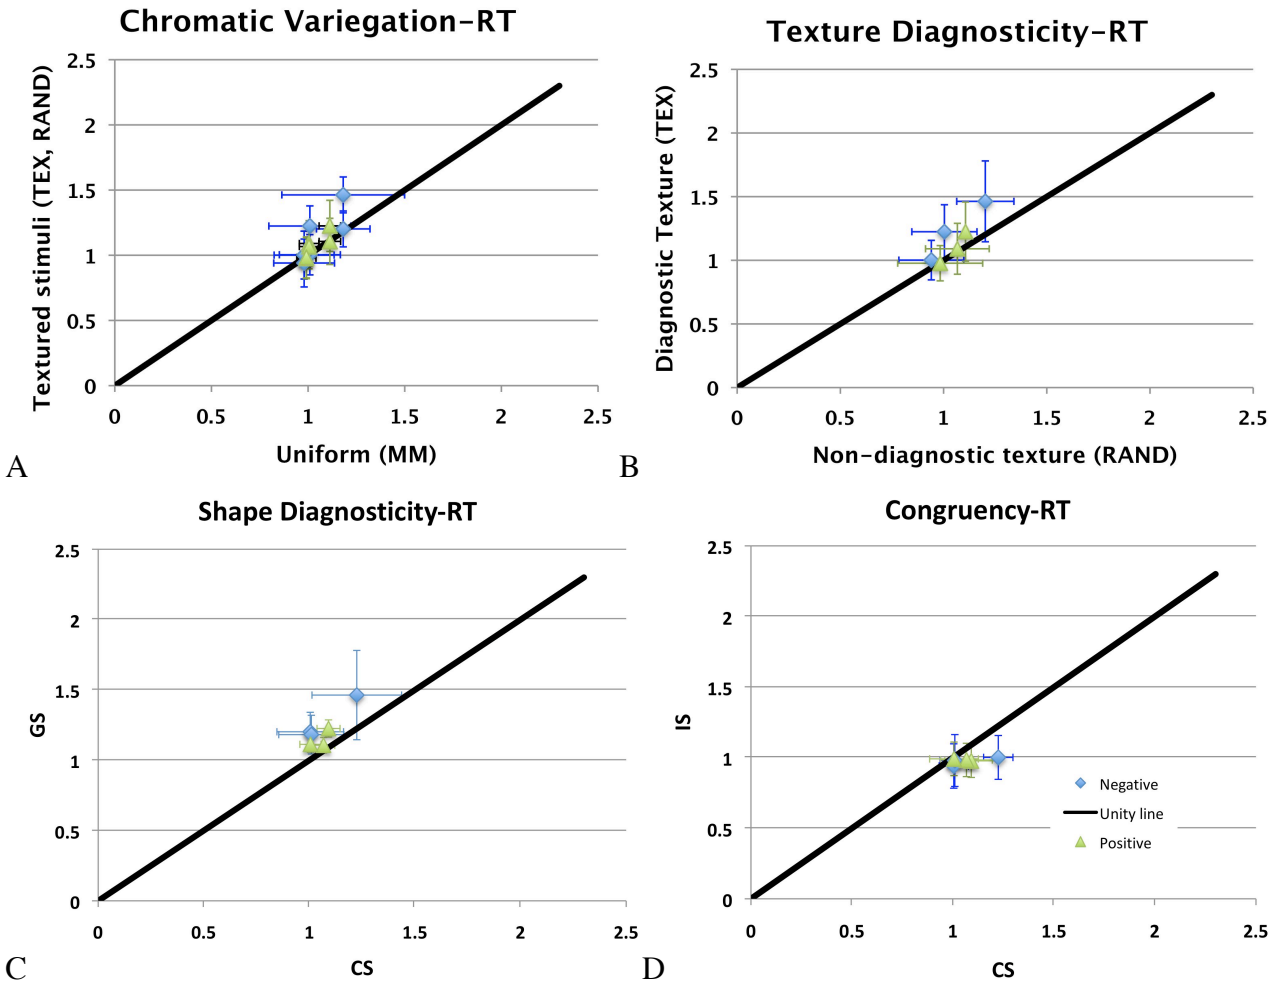

Figure 7.43 – Contrast between global mean response times as a function of different chromatic or shape factor condition. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation. Error bars represent the standard error of the mean.

Furthermore, the ANOVA reported a significant main effect of shape factor on mean response times ($F(1.311,10.49)=5.108$, $p<0.05$). Figure 7.43C compares the effect on global mean response time of congruent shape against generic shape condition (Comparison 3 – shape diagnosticity); note that all points lie above the unity line (each point represent one chromatic-GS condition combination against one chromatic-GS condition combination for positive and negative rotations). Lastly, Figure 7.43D shows that on average mean response times are lower for congruent shape than incongruent shapes. Specifically, contrast showed a significant effect of object congruency with mean response times 9.18% higher for congruent shapes than for incongruent shapes ($F(1,11)=6.36$, $p<0.05$). Conversely, the mean response times are 12.3%[24] lower for the congruent shape than for the generic shape ($F(1,11)=7.5$, $p<0.05$). *In brief, the mean response times for the shape conditions follow these significant relationships: GS>CS>IS.*

[24] Global mean response time: CS=1.07±0.1; IS=0.98±0.3; GS=1.22±0.2 (IS<CS<GS)

Finally, mean response times are significantly affected by the main effect of object ($F(2.58, 20.64)=9.484$, $p<0.005$), whereas there is ***no*** effect of rotation direction ($F(1,11)=0.999$, $p=0.349$).

7.3.2.2.2 Number of incorrect responses

An additional way to analyse the response is to calculate the number of trials in which a subject made an incorrect response per condition, or the number of incorrect responses for condition combination (IR). As in Experiment 1, the mean number of incorrect responses over all subject and objects for one shape/chromatic condition combination is defined as mean number of incorrect response (NM). Figure 7.44 show NM for all conditions and direction.

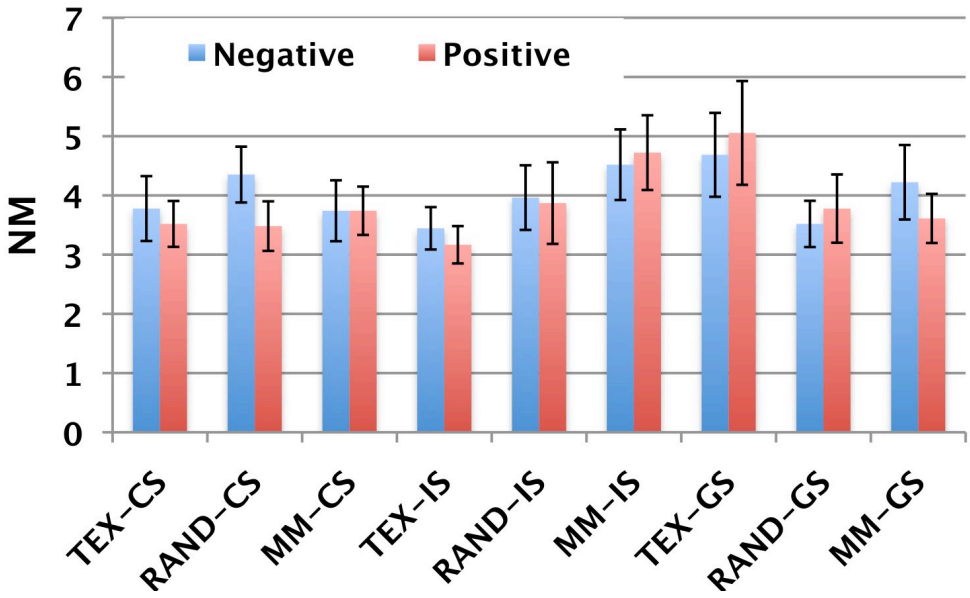

Figure 7.44 – Mean of number of incorrect responses over all subjects and objects. Blue bar: Negative angular rotation. Red bar: positive angular rotation. TEX: natural texture; RAND: randomized texture; MM: uniform colour. CS: congruent shape; IS: incongruent shape; GS: generic shape. Note that for example RAND-GS is the combination RAND chromatic condition and GS shape condition. Error bars: standard error of the mean (s.e.m.).

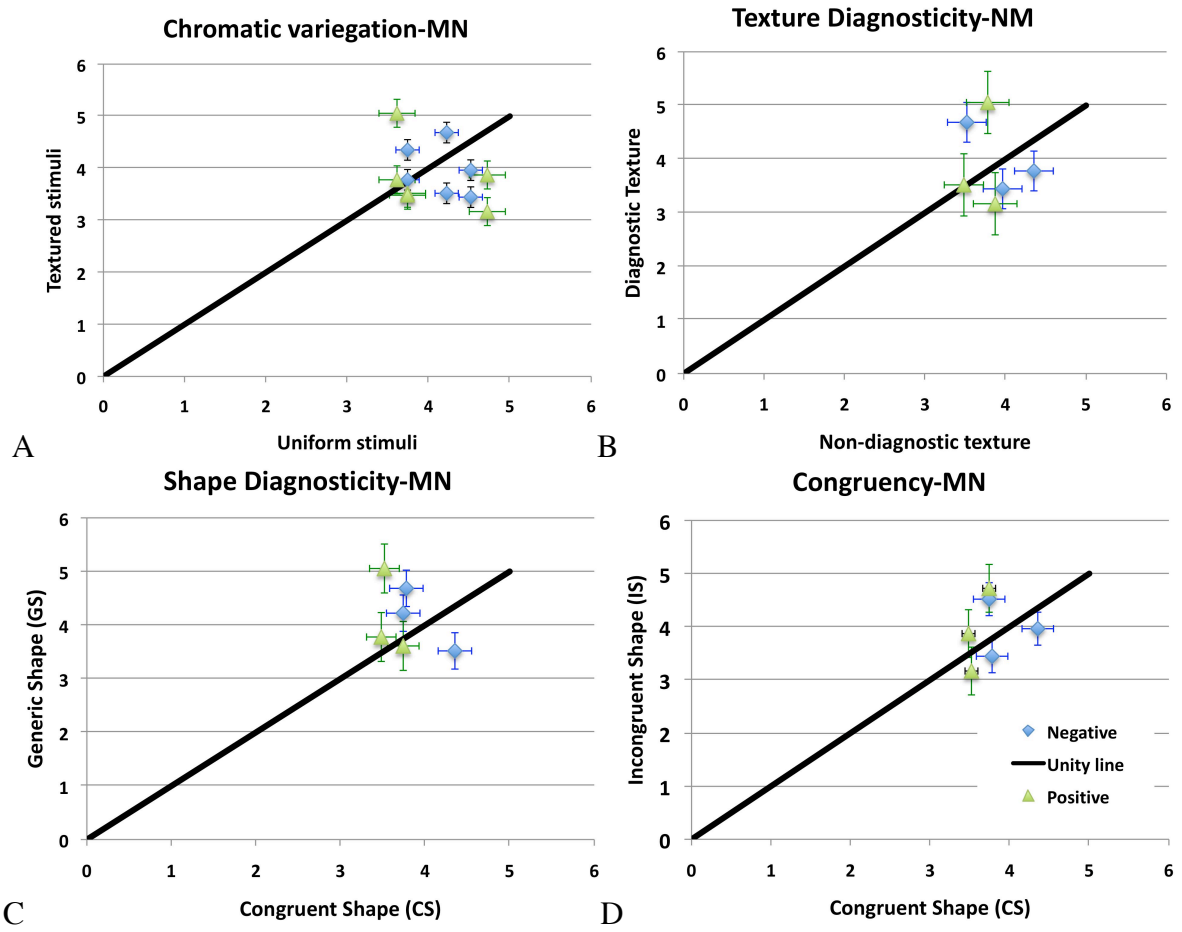

Figure 7.45 – Contrast between mean numbers of incorrect responses as a function of different chromatic or shape factor condition. (A) Chromatically variegated stimuli against uniformly coloured stimuli (Comparison 1). Each point in panel A contrasts one TEX-shape combination against one MM-shape combination (same shape paired) or one RAND-shape combination against one MM-shape combination. (B) Diagnostic (TEX) against non-diagnostic (RAND) stimuli (Comparison 2). Each point in panel B represents one TEX-shape combination against one RAND-shape combination (same shape paired). (C) Congruent shape against generic shape (Comparison 3). Each point in panel C represents one chromatic-GS combination against one chromatic-CS combination (same chromatic condition paired). (D) Congruent against incongruent shape (Comparison 4). Each point in the panel D represents one chromatic-IS combination against one chromatic-CS combination. Blue diamond: negative angular rotation. Green triangle: positive angular rotation.

A 4-way repeated measures ANOVA was performed on the mean response times of all subjects for all objects and both direction of rotation as in Experiment 1. Mauchly's test indicates that the assumption of sphericity was violated and the ANOVA was corrected using the using Greenhouse-Geisser estimates. For the number of incorrect responses the ANOVA reported no significant main effect of (1) chromatic factor, (2) shape, or (3) direction of the rotation. No main interaction was found significant. *Contrast showed a significant interaction between shape diagnosticity and the comparison TEX vs MM on number of incorrect responses*, indicating that TEX-CS is significantly different from TEX-GS and MM-CS from MM-IS ($F(1,8)=20.93$, $p<0.005$). The ANOVA reported significant the main effect of object on number of incorrect responses ($F(2.559, 20.473)=8.64$, $p<0.005$). Figure 7.45 shows in each subfigure one of the 4 usual comparisons (as in Figure 7.2). Note that in each subfigure all points gravitate around the unity line (i.e. they are both above and below the unity line, indicating no specific trend).

7.4 GENERAL DISCUSSION

The main goal of Experiment 2 was to assess the effect of shape and surface chromatic variegation on a delayed match-to-sample task. The results of Experiment 2 lead to the following four conclusions. Firstly, natural polychromatic textures are better discriminated, but more slowly, in a delayed match-to-sample task than either synthesized textures or uniformly coloured surfaces when all appear as surfaces of 3D shapes. Secondly, synthesized polychromatic textures are more difficult but quicker to discriminate than uniformly coloured surfaces in the same delayed match-to-sample task. Thirdly, the surfaces of natural shapes are more easily and quickly discriminated than generic shapes. Lastly, shape-congruent surfaces are more rapidly discriminated than shape-incongruent surfaces. Consequently, the results of this study raise four interesting questions to which present knowledge provides only partial answers. The following sections present these questions and their possible explanations, comparing Experiment 2 with previous studies, previous chapters, and Experiment 1.

7.4.1 THE EFFECT OF CHROMATIC FACTOR

In this successive discrimination task, natural textures are better discriminated than the synthesised textures, although the two have identical chromaticity distributions and appear on the same shapes. What factors, then, cause this difference in discriminability? Why are the two types of textures discriminated differently with respect to a uniformly coloured surface of the same mean colour and shape – i.e. natural textures better, synthesised textures worse? To answer these two questions it is first worth pointing out that there are only two differences between the two types of texture: (1) one represents a diagnostic texture while the second a non-diagnostic texture, and (2) they possess different spatial frequency power spectra (SFPD). In fact, as said in section 7.2.1.4.1 the method of texture randomization generates a non-diagnostic texture that maintains in general only the medium-high spatial frequencies of the diagnostic texture (TEX). Initially I will examine this second difference and its potential effect on discrimination.

Several studies have measured the chromatic contrast sensitivity of a stimulus as a function of spatial frequency (van der Horst et al. (1967); Hilz&Cavonius (1970)), time (de Lange (1958); Kelly&van Norren (1977); Wisowaty (1981); Swanson et al. (1987)), and space-time combinations (van der Horst and Bouman (1969); Noorlander et al. (1981)). These studies report lowpass sensitivity generally for equiluminous chromatic variations, with steep fall-off at high spatial frequencies, and bandpass sensitivity for luminance variation (e.g. Tiippana et al. (2000), Delahunt et al. ((2005)). Based on such results, one might well expect a significant difference in the simultaneous discrimination thresholds between diagnostic and non-diagnostic textures in

Experiment 1 based on the differences in their spatial frequency components. But in Experiment 1 I have demonstrated that there is no such significant difference. *I therefore postulate that the difference between SFPD of the two distributions is sufficiently similar* ***not*** *to affect the discrimination threshold also in a delayed paradigm.*

Therefore I argue that the only relevant difference between the two textures is in the information about the identity of the object. Few recent studies have investigated the cognitive effects of natural texture on colour appearance. For example, in a recent work (Hurlbert et al., 2009) we have shown a relationship between texture and colour in object representation. Specifically, using a rapid successive discrimination task, we have proven that discrimination thresholds for "abnormal" chromatic textures of natural object are higher than for "normal" chromatic textures of natural objects. The "abnormal" texture was created by rotating the typical chromatic distribution of a familiar fruit or vegetable object (the "normal" texture) around the scene adaptation point (as in Experiment 1) to a hue angle opposite or significantly away from the original angle. We concluded that colour diagnosticity might play a role in chromatic discrimination. Furthermore, Hansen et al. (2006) have shown that the memory colour of an object affects its colour appearance. In fact, when an observer was instructed to adjust to neutral (i.e. until it looked achromatic) the colour of a photograph of a natural object, his/her match always deviated from the neutral point (i.e. the adaptation point) in the direction opposite to the object memory colour. Moreover, in Chapter 6 I have shown that cues to object identity (such as its natural texture) affect colour appearance. Specifically, I have proven that natural chromatic textures improve colour constancy and the accuracy of the observer's selection of the object typical colour. I then concluded that the memory colour of the object was a function of the cues to its identity influencing our perception. Considered together with the above studies, the findings in Experiment 2 may be explained by a possible cognitive effect of the "memory colour texture" of the object.

In other words, a natural familiar object possesses a distinctive luminance texture and chromatic texture in addition to a distinctive, typical mean colour. In Experiment 1 and 2, I have corrupted both of these components of the object texture and, in this way, completely removed this meaningful information. Since in Experiment 2 the observer had to hold the chromatic information in memory during the delay before the matching, I propose that the typical representation of the object texture holds more easily as it is linked to a more stable memory of the object -- its memory colour distribution. As already discussed in previous chapters, the main theory introduced in this thesis is that humans possess a range of memory colours for one object, differing to the common opinion of single memory colour. However Experiment 2 proves that *we not only possess a memory of the object colour distribution, but also its spatio-chromatic structure*. Besides, these two elements are linked. Thus this suggests that simply presenting the same distribution does not

activate the same memory. On the contrary, the random texture promps a new memory colour texture for this object; thus we access a more recent and inaccurate memory needing less time with lower performance (short-term memory). On the other hand, the typical texture is already known, it is solid information, and we have just to recall it; we access a long-term memory employing more time, but with better performance.

Further, I argue that uniform surfaces lie in between these two categories. They may be seen as (1) references, eliciting unbiased, baseline discrimination thresholds, ***or*** (2) affected by the memory colour too in the classical sense, i.e. influenced by the *mean* typical colour associated with a familiar object. From the results of Experiment 1 and 2 it is difficult to determine which of the two statements above are true and which memory mechanism is then involved. It is most likely that both mechanisms are involved. We may infer this conclusion probably from the interaction between shape diagnositicity and chromatic surface diagnosticity (texture vs uniform) on the number of incorrect responses. However these results are insufficient to determine whether there is an effect of memory colour on uniformly coloured surfaces.

In an ongoing study we have separated luminance texture and chromatic texture and altered only the spatial coherence of the first component. With this experiment we aim to understand to what degree the two components influence texture diagnosticity.

For thoroughness, note that, in Experiment 1, the subject performed the discrimination task in 2.5 seconds on average. A number of studies (Vimal (2002) Tiippana et al. (2000)) have shown 100% performance on a detection task after 500ms-1500ms at any spatial frequency at maximum contrast. Therefore we can safely assume that the observer has fully detected and processed the reference stimulus in the 5 seconds in which it was presented regardless of stimulus type (textured or uniform). The same consideration may be made for the two alternatives.

7.4.2 THE EFFECT OF SHAPE

In Experiment 1 I have shown that the presence of incongruent shape deteriorates the subject's ability to discriminate between two chromatic surfaces. Furthermore in the same experiment, I have showed that the presence of the congruent shape may induce a memory colour effect on discrimination performance. In Experiment 2, instead, memory thresholds are similar for congruent and incongruent shape, although congruent shapes have slightly lower thresholds. Then why are memory thresholds for natural shapes lower than for generic shape? Why do we not find the same significant difference as in Experiment 1 between congruent and incongruent shapes? A first argument could be that, as long as the subject can anchor to a natural object shape to remember the stimulus colour, his/her performance are similar. A second argument could be that all natural shape shared a similar feature opposite to the generic shape. For example, while objects as the lime

and the apple have similar retinal angular size to the rhomboid, the banana or the cucumber have a much bigger size. Hence is possible to argue that contrary to Experiment 1, only low level mechanism due to the difference in size between the natural objects and the generic object influenced the subject's response. However, the consistent difference between congruent and incongruent shape, suggests that we cannot discard the possibility that the object actual shape facilitate the memory task. In support of this last statement are also the results on the mean response time. In fact, the subject responds significantly more rapidly when the chromatic surface is displayed on a congruent shape than an incongruent or a generic shape. Therefore I suggest that shape cues to the object identity enhance the observer performance in a successive discrimination task, although the effect due to high-level mechanism (i.e. memory) is small with respect to low-level effect (i.e. object size).

In addition, the interaction between shape congruency and texture diagnosticity on memory thresholds suggest an influence between these two previous "knowledge" of the object (shape and texture) on colour appearance. As a matter of fact, when the stimulus is presented in its natural shape and natural texture, the subject demonstrated the best successive discrimination performance compared to all other condition combinations. Based on these considerations, I postulate that congruency between cues to the object identity facilitate the subject's ability to discriminate as the object surface appearance is congruent with the memory that the subject has of object and, therefore, of its colour.

7.5 GENERAL CONCLUSIONS

In this chapter I have presented 2 experiments examining the effect of shape and chromatic proprieties of a surface on simultaneously and successive discrimination of chromatic surfaces of natural objects.

While the first experiment (Experiment 1) demonstrates that the discrimination threshold performances were only due to low-level mechanism of the stimulus, Experiment 2 proved that, when memory is involved, high-level mechanism are established. In fact, the subject discriminates best between two stimuli for the diagnostic texture condition than for the non-diagnostic. Thus, I proposed that the typical representation of the object texture holds more easily as it is linked to the memory colour distribution of the object.

In addition, the two experiments showed different effects of shape cues to object identity. Experiment 1 showed a modulatory effect of memory on the discrimination performance. Compared to the generic shape (the baseline), incongruent shapes inhibited the discrimination performance of the subjects, whereas for congruent shapes the subject performance is enhanced when the match is made in the opposite direction of the memory colour of the object and inhibited

when it is made in the same direction. Experiment 2 showed no such pronounced high-level effect, which mainly manifested itself in the quicker response for congruent shape with respect to the incongruent and generic shape.

Finally, the strong interaction between chromatic cues and shape cues to the object identity on both simultaneous and successive discrimination lead to the conclusion that high level mechanisms linked to the recognition of the object (congruency between object shape and texture) facilitated both tasks.

Chapter 8

SUMMARY AND CONCLUSIONS

The primary goal of this work was to investigate colour perception in a natural environment and to contribute to the understanding of how cues to familiar object identity influence colour appearance. A large number of studies on colour appearance employ 2D uniformly coloured patches, and only few consider the context in which actually the observer is immersed. However, our world is much more complex, and perceptual cues such as binocular disparity, 3D luminance shading, mutual reflection, surface texture, and glossy highlights are integral parts of a natural scene. Furthermore natural objects possess specific cues that help us recognize them, such as shape, texture or colour distribution. Starting from the most obvious cue to object identity, i.e. shape, I examined how also other cues such as binocularity, 3D luminance shading, and surface texture cues to object identification affected colour appearance.

In the first main experiment presented in this thesis (Chapter 4), I aimed to understand the effect of shape on colour perception, specifically on (1) memory colour under constant and varying illumination and on (2) colour constancy. I employed uniformly coloured stimuli in three shape configurations: (a) disk, (b) 2D silhouette of a familiar object, and (c) solid (3D) replica of such object. The set up used allowed the observer to view 3D solid objects changing colour. To this end, I designed a software package that performed geometric and light calculation and that allowed real-time manipulation of the stimuli. Subjects were asked to select all the colours that they judged appropriate for the named object. The results demonstrated the existence of a range of memory colours associated with a familiar object, the size of which was strongly object-shape-dependent. This suggested the theory of the "memory colour distribution", subsequently further explored in this thesis. To understand the object-shape-dependent effect on memory colour range I designed two controls: (1) a shape recognisability experiment and (2) a chromatic analysis of the original object texture. Provided with these control studies I proposed a relationship between memory colour range and (1) shape cues to object identity and (2) natural chromatic distribution of the object. Furthermore, I suggested that shape cues to the object identity affect directly the range of memory colour proportionally to the original object chromatic distribution. Results also showed that 3D shapes, if more indicative of the object identity, increased the subject performance. These results were consistent across all illuminants. Moreover, for all objects and illuminations, memory retrieval was significantly faster when the shape of the object was presented. However, no significant difference was found between shape conditions on colour constancy, probably due to the limitation of the set up performance.

For this reason and to display stimuli that represented the actual chromatic surface texture of a natural object I design a virtual environment to allow the interactive adjustment and presentation of the stimuli while preserving the natural cues to 3D shape. The new apparatus was accessorised

with a software package designed in three stages. In the first stage, I developed a number of techniques to generate images that accurately represented the chromatic information of the natural object, i.e. a method that accurately recorded an image and converted each pixel in CIE XYZ coordinates. In the second stage (Chapter 5), I analysed and modelled characteristics of the object chromatic distribution in a physiologically based colour space. Results from this analysis revealed the characteristics of two physical indices of the chromatic distribution of natural objects: (1) the illuminant-invariance of a "hue vector" defined in this space and (2) the object-independent correlation between chromatic distribution of objects belonging to the same category under same or different lilluminants, and (3) the illuminant-dependent correlation between chromatic distribution of objects belonging to the same category. In a third stage, I developed a technique to extract only the chromatic texture of the object and to manipulate and synthesise this texture for other purposes. Provided with these tools, I successively performed three other experiments with the purpose of investigating the effect of chromatic variegation and texture diagnosticity together with shape diagnosticity and dimensionality.

In Chapter 6, I examined the subject accuracy and precision in adjusting a stimulus colour to its typical appearance. I employed four shape configurations (2D generic shape, 2D object silhouette, 3D generic shape, and solid 3D replica of the familiar object) and three chromatic conditions (full chromatic texture, mean colour and most saturated colour of the distribution) under two illuminants. Here I tested the accuracy and precision of the subjects' memory colour and colour constancy as a function of shape and chromatic variegation of the stimuli. Results showed that memory colour accuracy and precision were enhanced for: (1) chromatic texture with respect to uniform stimuli, (2) diagnostic shapes with respect to generic shape, and (3) 3D configuration with respect to 2D. A strong interaction was found between diagnosticity and dimensionality of the shape. These findings lead to the conclusion that more cues to the object identity and a more natural stimulus facilitate the observers in accessing their colour information from memory. In addition, I demonstrated a direct interaction between chromatic surface representation, physical properties of the object, shape identificability and shape dimensionality on memory colour accuracy. Colour constancy was affected by the chromatic complexity of the stimuli (texture) in that colour constancy is higher for texture stimuli relative to uniform stimuli. However no effect of shape diagnosticity or dimensionality was observed. Based on these results, I postulated that the chromatic variegation of the familiar object of the stimuli contributes significantly in colour constancy. Lastly, I suggested a relationship between colour memory and colour constancy due to the robustness of the effects under changes of illumination, however other elements might play a role on such a mechanism for familiar colour-diagnostic objects. This experiment demonstrates the effect of chromatic texture

and/or shape on memory colour and colour constancy accuracy and precision, as well as their combined contributions.

In the third and fourth experiments, I tested the subject's ability to discriminate between two chromatic stimuli in a simultaneous (in Chapter 7 indicated with Experiment 1) or successive (Experiment 2) 2AFC task, respectively. In both experiments I employed three shapes (congruent, incongruent and generic) and three chromatic conditions (full chromatic texture, full randomized chromatic texture, uniform colours) and evaluated combination of them for 3 objects pairs. Discrimination thresholds, response time and number of incorrect response were recorded. Results showed a strong effect of object, shape diagnosticity, shape naturalness, texture diagnosticity and chromatic variegation. Specifically, Experiment 1 demonstrates that the simultaneous discrimination threshold performances for polychromatic surfaces were only due to low-level mechanism of the stimulus, whereas Experiment 2 proved that, when memory is involved, high-level mechanism are established. In fact, the subject discriminates best between two stimuli for the diagnostic texture condition than for the non-diagnostic. Thus, I proposed that the typical representation of the object texture holds more easily, as it is linked to the memory colour distribution of the object.

In addition, the two experiments showed different effects of shape cues to object identity. Experiment 1 showed a modulatory effect of memory on the discrimination performance. Compared to the generic shape (the baseline), incongruent shapes inhibited the discrimination performance of the subjects, whereas for congruent shapes the subject performance is enhanced when the match is made in the opposite direction of the memory colour of the object and inhibited when it is made in the same direction. Experiment 2 showed no such pronounced high-level effect, which mainly manifested itself in the quicker response for congruent shape with respect to the incongruent and generic shape.

Finally, the strong interaction between chromatic cues and shape cues to the object identity on both simultaneous and successive discrimination lead to the conclusion that high level mechanisms linked to the recognition of the object (congruency between object shape and texture) facilitated both tasks.

In conclusion this work demonstrated that colour appearance of natural object is influenced by the visual cues to the object identity. Shape diagnosticity and dimensionality modulate the size of the memory colour distribution as well as memory colour accuracy and precision. Furthermore, shape affects the discrimination of chromatic surfaces. Textures enhance the observer's colour constancy performance but decrease simultaneous discrimination. Moreover, diagnostic textures increase the successive discrimination performance of the observers.

Hence, the current thesis presents new findings on memory colour and colour constancy presented in a natural context and demonstrates the effect of high-level mechanisms in chromatic discrimination as a function of cues to the object identity such as shape and texture. This work contributes to a deeper understanding of the colour perception and object recognition in the natural world.

REFERENCES

Arend, L. and A. Reeves (1986). "Simultaneous color constancy." J Opt Soc Am A **3**(10): 1743-1751.

Arend, L. J., A. Reeves, et al. (1991). "Simultaneous color constancy: papers with diverse Munsell values." Journal of the Optical Society of America A.

Bannai, N., R. B. Fisher, et al. (2007). "Multiple color texture map fusion for 3D models." Pattern Recognition Letters **28**(6): 748-758.

Bartels, A. and S. Zeki (1998). "The theory of multistage integration in the visual brain." Proceedings of the Royal Society - Biological Sciences (Series B) **265**(1412): 2327-2332.

Bartleson, C. (1960). "Memory colors of familiar objects." Journal of the Optical Society of America **50**(1): 73-77.

Bartleson, C. (1961). "Color in memory in relation to photographic reproduction." Phot Sci Eng **5**(6): 327-331.

Beck, J., A. Sutter, et al. (1987). "Spatial frequency channels and perceptual grouping in texture segregation." Computer Vision, Graphics, and Image Processing **37**: 299-325.

Beeckmans, J. (2009). "How chromatic phenomenality largely overflow its cognitive accessibility." Consciousness and Cognition **18**(4): 917-928.

Bloj, M., D. Kersten, et al. (1999). "Perception of three-dimensional shape influences colour perception through mutual illumination." Nature **402**(6764): 877-879.

Bonnier, N., C. Leynadier, et al. (2008). "Improvements in Spatial and Color Adaptive Gamut Mapping Algorithms." CGIV 2008.

Bonnier, N., F. Schmitt, et al. (2007). Spatial and color adaptive gamut mapping: A mathematical framework and two new algorithms. Proceedings of the 15th IS&T/SID Color Imaging Conference, Albuquerque, NM
.

Brainard, D. (1998). "Color constancy in the nearly natural image. 2. Achromatic loci." Journal of the Optical Society of America A: Optics and Image Science, and Vision **15**(2): 307-325.

Brainard, D. (2004). "Colour constancy." The Visual Neurosciences: 948-961.

Brainard, D., W. Brunt, et al. (1997). "Color constancy in the nearly natural image. I. Asymmetric matches." Journal of the Optical Society of America A: Optics and Image Science, and Vision **14**(9): 2091-2110.

Brainard, D. and W. Freeman (1997). "Bayesian color constancy." Journal of the Optical Society of America A **14**(7): 1393-1411.

Brainard, D. and B. Wandell (1992). "Asymmetric color matching: how color appearance depends on the illuminant." J. Opt. Soc. Am. A **9**(9): 1433–1448.

Braun, G. J. and M. D. Fairchild (2000). "General-purpose gamut-mapping algorithms: Evaluation of contrast-preserving rescaling functions for color gamut mapping." Journal of Imaging Science and Technology **44**: 343-350.

Bruner, J., L. Postman, et al. (1951). "Expectation and the perception of color." The American Journal of Psychology **64**: 216-227.

Caelli, T. and D. Reye (1993). "On the classification of image regions by colour, texture and shape." Pattern Recognition **26**(4): 461-470.

CIE (2004). "Guidelines for the evaluation of Gamut Mapping Algoithms."

Clark, A., B. Thomas, et al. (1999). "Texture deconvolution for the fourier-based analysis of non-rectangular regions." British Machine Vision Conference: 193–202.

Cornelissen, F. and E. Brenner (1995). "Simultaneous colour constancy revisited: An analysis of viewing strategies." Vision Research **35**(17): 2431-2448.

Craven, B. and D. Foster (1992). "An operational approach to color constancy." Vision Research **32**: 1359-1366.

Cremens, D., M. Rousson, et al. (2007). "A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape." Int J Comput Vision **72**(2): 195-215.

Dana, K., S. Nayar, et al. (1997). Reflectance and texture of real-world surfaces. CVPR97**:** 151-157.

de Fez, M., P. Capilla, et al. (2001). "Asymmetric colour matching: Memory matching versus simultaneous matching." Color Res. Appl. **26**(6): 458-468.

Delk, J. and S. Fillenbaum (1965). "Differences in perceived color as a function of characteristic color." American Journal of Psychology **78**: 290-293.

Derrington, A., J. Krauskopf, et al. (1984). "Chromatic mechanisms in lateral geniculate nucleus of macaque." Journal of Physiology **VOL. 357**: 241-265.

Duncker, K. (1939). "The influence of past experience upon perceptual properties." American Journal of Psychology **52**: 255-265.

Efros, A. and W. Freeman (2001). "Image quilting for texture synthesis and transfer." Proceedings of SIGGRAPH 2001.

Efros, A. and T. Leung (1999). "Texture synthesis by non-parametric sampling." Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on **2**.

Eskew, R., J. McLellan, et al. (1999). "Chromatic detection and discrimination." Color Vision: From Genes to Perception: 345–368.

Eskew, R. T. (2009). "Higher order color mechanisms: A critical review." Vision Research **49**(22): 2686-2704.

Fairchild, M. and P. Lennie (1992). "Chromatic Adaptation to Natural and Incandescent

Illuminants." Vision Research **32**(11): 2077-2085.

Fairchild, M. D. (2005). Colour Appearance models, John Wiley & Sons Ltd.

Finlayson, G., M. Drew, et al. (1994). "Color constancy: generalized diagonal transforms suffice." J Opt Soc Am A **11**(11): 3011-3019.

Gagalowicz, A., S. De Ma, et al. (1986). Efficient models for color textures. 8th ICPR.

Gegenfurtner, K. (2003). "Cortical mechanisms of colour vision." Nature Reviews Neuroscience.

Gegenfurtner, K. and J. Rieger (2000). "Sensory and cognitive contributions of color to the recognition of natural scenes." Current Biology.

Giesel, M., T. Hansen, et al. (2009). "The discrimination of chromatic textures." Journal of Vision.

Golz, J. and D. MacLeod (2002). "Influence of scene statistics on colour constancy." Nature **415**(6872): 637-640.

Gonzales, R. C. and R. E. Woods (2007). Digital Image Processing, Harlow: Prentice Hall.

Goshtasby, A. (1986). "Piecewise linear mapping functions for image registration." Patter Recognition **19**: 459-466.

Graeme R. Cole, T. H., and William McIlhagga (1993). "Detection mechanisms in L-, M-, and S-cone contrast space." JOSA A **10**(1): 38-51.

Green, P. and L. MacDonald (2008). Colour gamut mapping, Ed. Wiley.

Hansen, T., M. Giesel, et al. (2008). "Chromatic discrimination of natural objects." Journal of Vision.

Hansen, T., M. Olkkonen, et al. (2006). "Suplememntary: Memory modulates color appearance." Nat Neurosci(Suplementary).

Hansen, T., M. Olkkonen, et al. (2006). "Memory modulates color appearance." Nat Neurosci **9**(11): 1367-1368.

Hansen, T., S. Walter, et al. (2007). "Effects of spatial and temporal context on color categories and color constancy." Journal of Vision **7**(4): 2.

Haralick, R. M., K. Shanmugam, et al. (1973). "Textural features for image classification." IEEE Transactions on Systems, Man, and Cybernetics **SMC-3**(6): 610-621.

Helmholtz, H. v. (1892). Die Grundempfindungen in normalen und anomalen Farbsystemen", Zeitschrift für Psychologie 4, pp. 241-347 (1892); E. G. Boring, Sensation and Perception in the History of Experimental Psychology», New York 1892.

Hering, E., L. Hurvich, et al. (1878). "Outlines of a theory of the light sense (Republished in English translation) 1964."

Horn, B. (1974). "Determining lightness from an image." Computer Graphics and Image Processing **3**: 277-299.

Hurlbert, A. (1986). "Formal connections between lightness algorithms." Journal of the Optical Society of America A **3**: 1694-1693.

Hurlbert, A., D. Bramwell, et al. (1998). "Discrimination of cone contrast changes as evidence for colour constancy in cerebral achromatopsia." Experimental Brain Research **123**(1-2): 136-144.

Hurlbert, A. and Y. Ling (2004). "Colour constancy of real 3-D objects." Perception **33**(ECVP Abstract Supplement).

Hurlbert, A. and Y. Ling (2006). "The colour appearance of chromatically textured natural surfaces." CIE Expert Symposium on Visual Apparence.

Hurlbert, A., I. Pietta, et al. (2009). "Surface discrimination of natural objects: When is a blue kiwi off-colour?" JOV **9**(8): 330a.

Jin, E. and S. Shevell (1996). "Color memory and color constancy." Journal of the Optical Society of America A: Optics and Image Science, and Vision **13**(10): 1981-1991.

Judd, D., D. MacAdam, et al. (1964). "Spectral distribution of typical daylight as a function of correlated color temperature." J. Opt. Soc. Am. **54**(8): 1031-1040.

Julesz, B. and J. Bergen (1983). "Textons, the fundamental elements in preattentive vision and perception of textures." Bell Systems Technological Journal **62**: 1619-1645.

Kandel, E., J. Schwartz, et al. (2000). Principles of neural science.

Kinnear, P. R. and A. Sahraie (2002). "New Farnsworth-Munsell 100 hue test norms of normal observers for each year of age 5-22 and for decades 30-70." Br J Ophthalmol. **86**: 1408-1411.

Kraft, J. and D. Brainard (1999). "Mechanisms of color constancy under nearly natural viewing." Proceedings of the National Academy of Sciences **96**(1): 307.

Krantz, D. (1968). "A theory of context effects based on cross-context matching." Journal of Mathematical Psychology **5**: 1-48.

Kusunoki, M., K. Moutoussis, et al. (2006). "Effect of background colors on the tuning of color-selective cells in monkey area V4." Journal of Neurophysiology **95**(5): 3047-3059.

Kutas, G., K. Gocza, et al. (2004). "Colour size effect." CGIV 2004 - Second European Conference on Color in Graphics, Imaging, and Vision and Sixth International Symposium on Multispectral Color Science: 70-73.

Land, E. (1964). "The retinex." American Scientist **52**(2): 247-264.

Land, E. (1974). "The retinex theory color vision." Proceedings of the Royal Institute of Great Britain **47**: 23-58.

Lanterman, A., U. Grenander, et al. (2000). "Bayesian Segmentation via Asymptotic Partition Functions." IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) **22**(4): 337-347.

Lee, H. (1986). "Method for computing the scene-illuminant chromaticity from specular highlights." J Opt Soc Am A **3**(10): 1694-1699.

Lennie, P. and M. D'zmura (1988). "Mechanisms of color-vision." Critical Reviews in Neurobiology **3**(4): 333-400.

Ling, Y. (2005). YazhuThesis. School of Biology. Newcastle upon Tyne, Newcastle University. **PhD:** 1-182. Thesis

Ling, Y. and A. Hurlbert (2004). "Color and size interactions in a real 3D object similarity task." Journal of Vision **4**(9): 721-734.

Ling, Y. and A. Hurlbert (2005). "Colour constancy is as good as colour memory allows--A new colour constancy index." Perception **34**(ECVP Abstract Supplemen).

Ling, Y. and A. Hurlbert (2008). "Role of color memory in successive color constancy." J Opt Soc Am A **25**(6): 1215-1226.

Linhares, J., P. Pinto, et al. (2008). "The number of discernible colors in natural scenes." Journal of the Optical Society of America A **25**(12): 2918-2924.

Liu, F. and W. Picard (1996). "Periodicity, directionality, and randomness: World features for image modelling and retrieval." EEE Trans. on PAMI **18**(7): 722-733.

Malik, J. and P. Perona. (1990). "Preattentive texture discrimination with early vision mechanisms." Journal of the Optical Society of America **7**: 923-932.

Maloney, L. and B. Wandell (1986). "Color constancy: A method for recovering surface spectral reflectance." J Opt Soc Am A **3**(1): 29-33.

Mamassian, P., M. Landy, et al. (2002). "Bayesian modeling of visual perception." Probabilistic Models of the Brain: 13-36.

Manjunath, B. S. and W. Y. Ma (1996). "Texture features for browsing and retrieval of image data." IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) **18**(8): 837-842.

McCann, J., S. McKee, et al. (1976). "Quantitative studies in retinex theory. A comparison between theoretical predictions and observer responses to the 'color mondrian' experiments." Vision Research **16**(5): 445-458.

Millán, M. S. and E. Valencia (2006). "Color image sharpening inspired by human vision models." Applied Optics **45**(29): 7684-7697

Moore, A., G. Fox, et al. (1990). "A VLSI neural network for color constancy." Advances in Neural Information Processing System **3**.

Morovic, J. and M. R. Luo (2001). "The Fundamentals of Gamut Mapping: a Survey." Journal of Imaging Science and Technology **45**(3): 283-290.

Morovic, J. and Y. Wang (2003). A multi-resolution, full-colour spatial gamut mapping algorithm. Proceedings of IS&T and SID's 11th Color Imaging Conference: Color Science and Engineering: Systems, Technologies, Applications.

Murray, L. I., A. Daugirdiene, et al. (2006). "Cone contrasts do not predict color constancy." Visual Neuroscience **23**: 543-547.

Nascimento, S., F. Ferreira, et al. (2002). "Statistics of spatial cone-excitation ratios in natural scenes." Journal of the Optical Society of America A **19**(8): 1484-1490.

Nascimento, S. and D. Foster (1997). "Detecting natural changes of cone-excitation ratios in simple and complex coloured images." Proceedings of the Royal Society - Biological Sciences (Series B) **264**(1386): 1395-1402.

Newton, J. R. and R. T. Eskew (2003). "Chromatic detection and discrimination in the periphery: A postreceptoral loss of color sensitivity." Visual Neuroscience **20**: 511-521.

Noorlander, C. and J. J. Koenderink (1983). "Spatial and temporal discrimination ellipsoids in color space." J. Opt. Soc. Am. **73**: 1533-1543.

Oliva, A. and P. Schyns (2000). "Diagnostic Colors Mediate Scene Recognition* 1." Cognitive Psychology **41**(2): 176-210.

Olkkonen, M., T. Hansen, et al. (2008). "Color appearance of familiar objects: Effects of object shape, texture, and illumination changes." Journal of Vision.

Panjwani, D. K. and G. Healey (1995). "Markov random field models for unsupervised segmentation of textured color images." IEEE Trans. Patt. Anal. Mach. Intell **17**(10): 939-954.

Perez-Carpinell, J., R. Baldovi, et al. (1998). "Color memory matching: Time effect and other factors." Color Res. Appl. **23**(4): 234-247.

Perez-Carpinell, J., M. de Fez, et al. (1998). "Familiar objects and memory color." Color Res. Appl. **23**(6): 416-427.

Provenzi, E., L. Carli, et al. (2005). "Mathematical definition and analysis of the Retinex algorithm." J. Opt. Soc. Am. A **22**(12).

Reed, T. and J. M. Hans Du Buf (1993). "A review of recent texture segmentation and feature extraction techniques." CVGIP: Image Understanding **57**(3): 359-372.

Reinhard, E., E. A. Khan, et al. (2008). Color imaging: fundamentals and applications, A. K. Peters Ltd.

Rinner, O. and K. Gegenfurtner (2002). "Cone contributions to colour constancy." Perception **31**(6): 733-46.

Robilotto, R. (2004). "Limits of lightness identification for real objects under natural viewing conditions." JOV **4**(9): 779-797.

Robilotto, R. and Q. Zaidi (2006). "Lightness identification of patterned three-dimensional, real objects." J. Vision **6**: 18-36.

Ruettiger, L., D. Braun, et al. (1999). "Selective color constancy deficits after circumscribed unilateral brain lesions." Journal of Neuroscience **19**(8): 3094-3106.

Russ, J. C. (1995). The Image Processing Handbook, 2nd Ed. Boca Raton, Florida, IEEE Press.

Sankeralli, M. J. and K. T. Mullen (1996). "Estimation of the L-, M-, and S-cone weights of the postreceptoral detection mechanisms." J. Opt. Soc. Am. A **13**: 906-915.

Siple, P. and R. Springer (1983). "Memory and preference for the colors of objects." Perception & Psychophysics **34**: 363-370.

Tanaka, J., D. Weiskopf, et al. (2001). "The role of color in high-level vision." Trends in Cognitive Sciences **5**(5): 211-215.

te Pas, S. and J. Koenderink (2004). "Visual discrimination of spectral distributions." Perception **33**: 1483-1497.

Tiippana, K., J. Rovarno, et al. (2000). "Contrast matching across spatial frequencies for isoilluminant chromatic gratings." Vision Research **40**(16): 2159-65.

Tomazevic, D., B. Likar, et al. (2002). "Compartive evaluation of retrospective shading correction methods." Journal of Microscopy **208**(pt 3 december): 212-223.

Tominaga, S. and B. Wandell (2002). "Natural scene-illuminant estimation using the sensor correlation." P Ieee **90**(1): 42-56.

Troost, J. and C. de Weert (1991). "Naming versus matching in color constancy." Perception and Psychophysics **50**(6): 591-602.

Van der Horst, G. and M. Bouman (1969). "Spatiotemporal chromaticity discrimination." Journal of the optical Society of America **59**(11): 1482.

van Ginneken, B. and J. Koenderink (1999). "Texture histograms as a function of irradiation and viewing direction." Int J Comput Vision **31**(2/3): 169-184.

Vimal, R. (2002). "Spatial frequency discrimination: a comparison of achromatic and chromatic conditions." Vision Research **42**(5): 599-611.

Von Kries, J. (1902). "Chromatic adaptation." Sources of Color Science: 109-119.

Voorhees, H. and T. Poggio (1988). "Computing texture boundaries from images." Nature **333**: 364-367.

Webster, M. and K. D. Valois (1990). "Orientation and spatial-frequency discrimination for luminance and chromatic gratings." J. Opt. Soc. Am. A **7**(6): 1034-1049.

Weickert, J. (1999). "Coherence-enhancing diffusion of colour images." Image Vision Comput **17**: 201-212.

Wichmann, F., L. Sharpe, et al. (2002). "The contributions of color to recognition memory for natural scenes." Learning, Memory **28**(3): 509-520.

Wolf, K. and A. Hurlbert (2003). "The effect of global contrast distribution on colour appearance." Normal and Defective Color Vision: 239-247.

Wright, W. and F. Pitt (1934). "Hue-discrimination in normal colour-vision." Proceedings of the Physical Society **46**: 459-473.

Wu, Y., M. Li, et al. (2007). "Multiple features data fusion method in color texture analysis." Applied Mathematics and Computation **185**(2): 784-797.

Xiao, B. and D. Brainard (2006). "Color perception of 3D objects: Constancy with respect to variation in surface gloss." ACM Symposium on Applied Perception in Graphics and Visualization.

Yang, C. and W. Runsheng (2006). Texture segmentation using independent component analysis of gabor features. International Conference on Pattern Recognition. **147-150**.

Young, M. and M. Ellefson (2003). "The joint contributions of shape and color to variability discrimination* 1." Learning and Motivation **34**(1): 52-67.

Zaidi, Q. (1998). "Identification of illuminant and object colors: Heuristic-based algorithms." Journal of the Optical Society of America A: Optics and Image Science, and Vision **15**(7): 1767-1776.

Zaidi, Q. (2001). "Color constancy in a rough world." Color Res Appl **26**(SUPPL.).

Zeki, S., J. Watson, et al. (1991). "A direct demonstration of functional specialization in human visual cortex." Journal of Neuroscience **11**(3): 641-649.

APPENDIX 1. SUPPLEMENTAL GRAPHS FOR CHAPTER 4

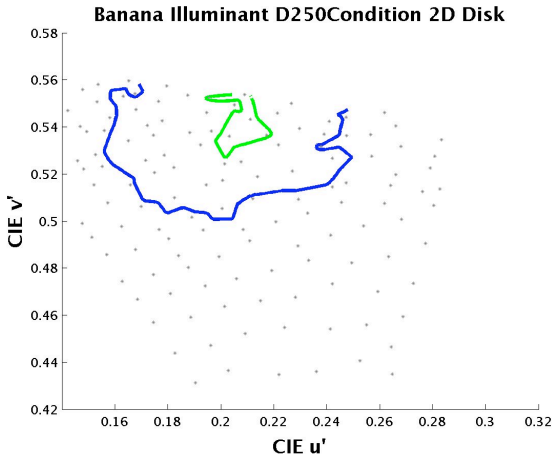

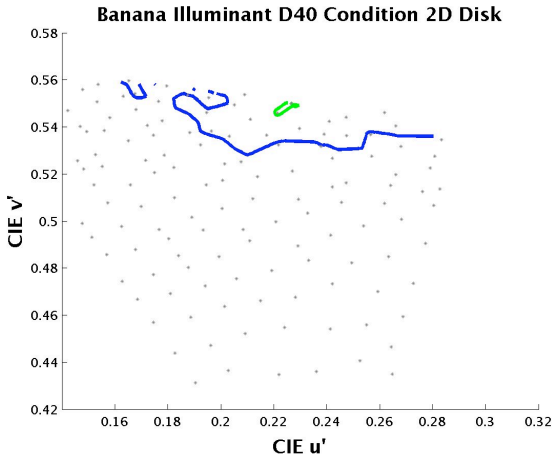

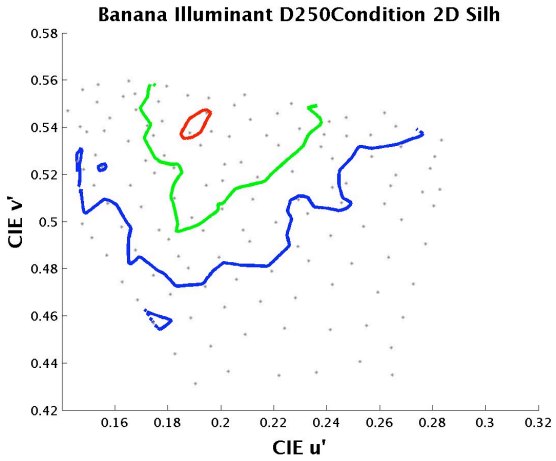

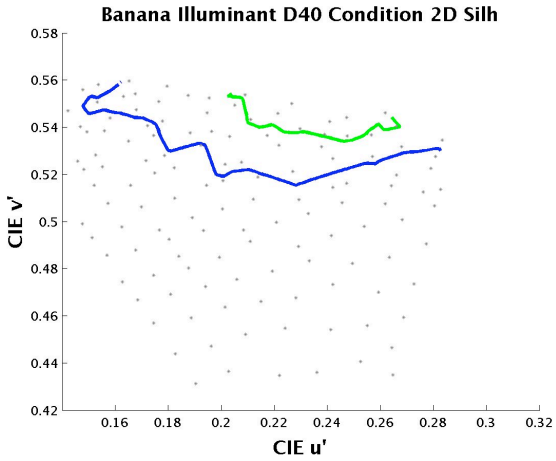

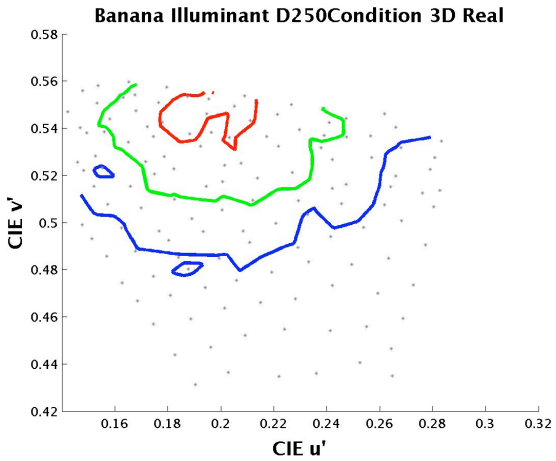

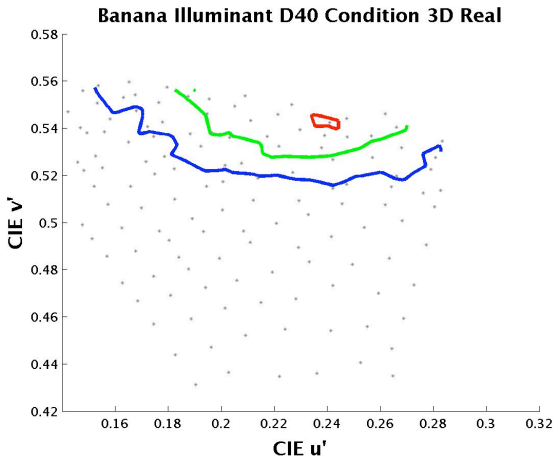

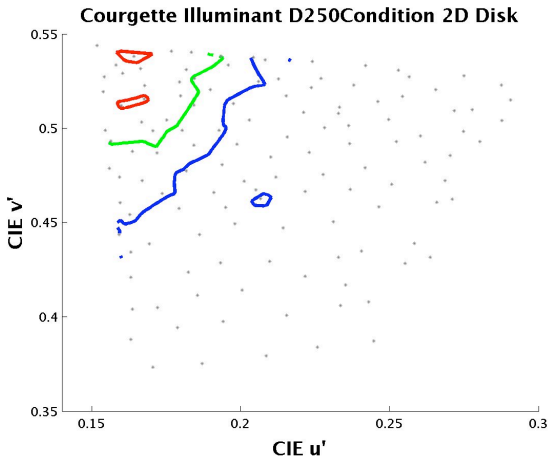
Courgette Illuminant D250Condition 2D Disk
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

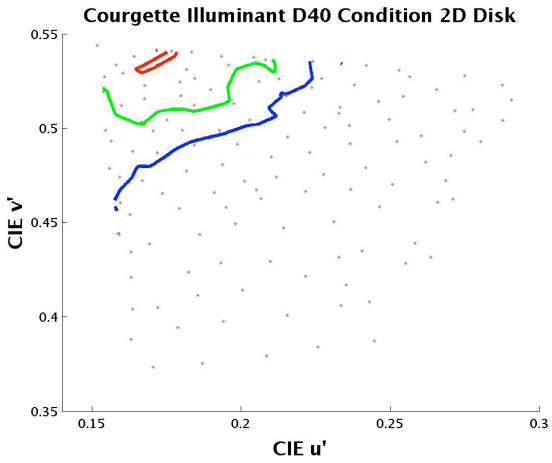
Courgette Illuminant D40 Condition 2D Disk
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

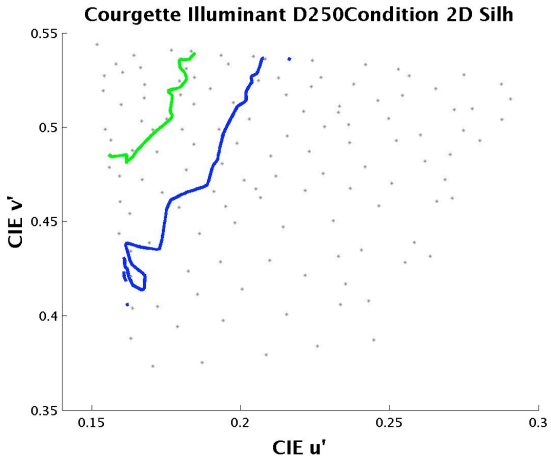
Courgette Illuminant D250Condition 2D Silh
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

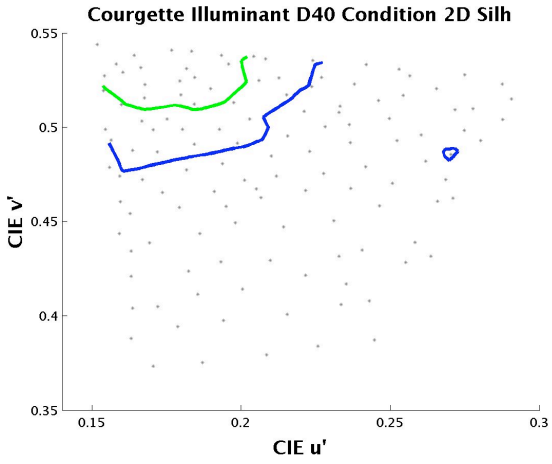
Courgette Illuminant D40 Condition 2D Silh
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

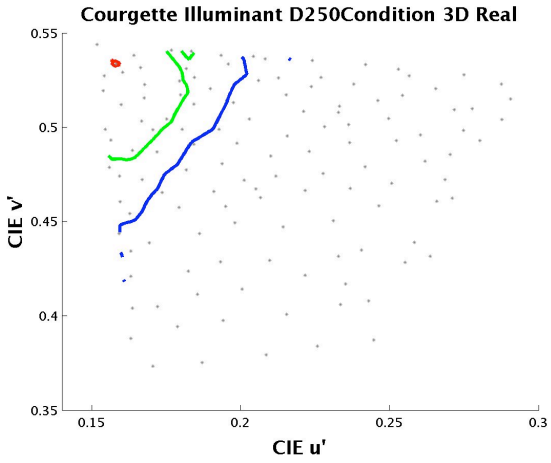
Courgette Illuminant D250Condition 3D Real
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

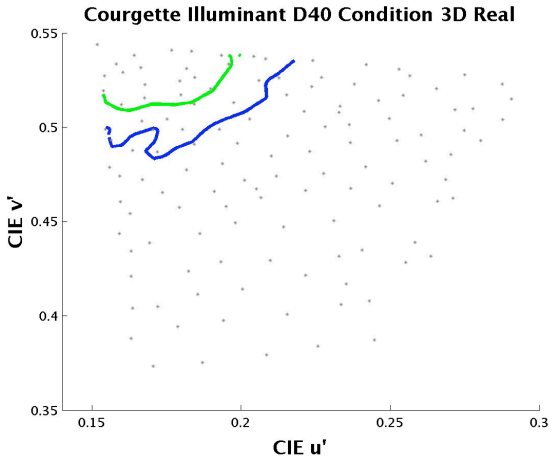
Courgette Illuminant D40 Condition 3D Real
CIE v'
CIE u'
0.55
0.5
0.45
0.4
0.35
0.15
0.2
0.25
0.3

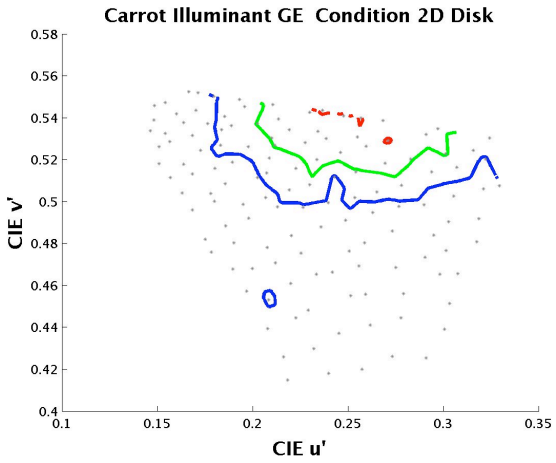
Carrot Illuminant GE Condition 2D Disk
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

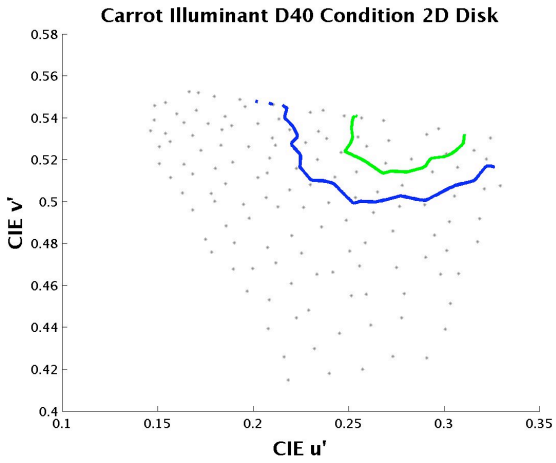
Carrot Illuminant D40 Condition 2D Disk
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

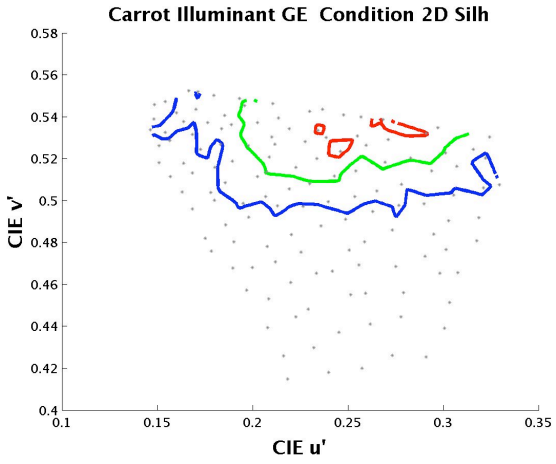
Carrot Illuminant GE Condition 2D Silh
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

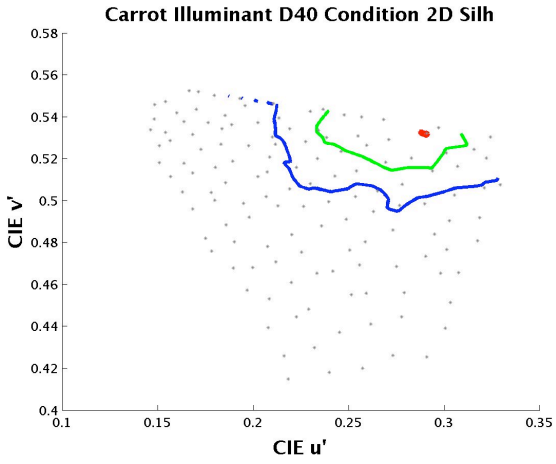
Carrot Illuminant D40 Condition 2D Silh
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

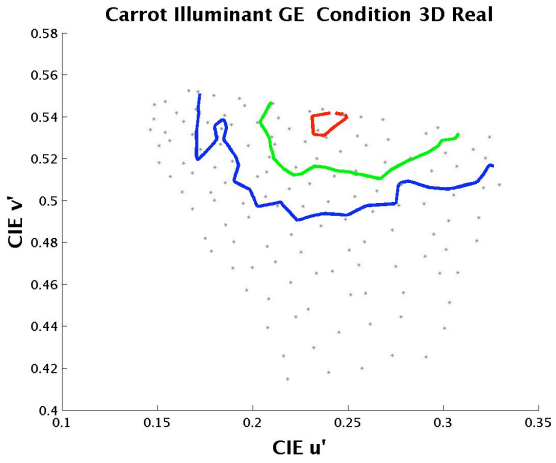
Carrot Illuminant GE Condition 3D Real
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

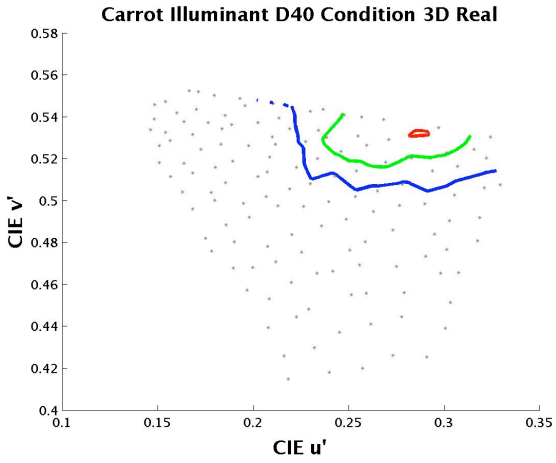
Carrot Illuminant D40 Condition 3D Real
CIE v'
CIE u'
0.58
0.56
0.54
0.52
0.5
0.48
0.46
0.44
0.42
0.4
0.1
0.15
0.2
0.25
0.3
0.35

APPENDIX 2. SUPPLEMENTAL TABLES FOR CHAPTER 7

Apple

Angle	ΔEuv
-35	
-33.75	2.8173
-32.5	2.8325
-31.25	2.847
-30	2.8607
-28.75	2.8735
-27.5	2.8854
-26.25	2.8963
-25	2.9063
-23.75	2.9153
-22.5	2.9233
-21.25	2.9303
-20	2.9363
-18.75	2.9412
-17.5	2.9451
-16.25	2.948
-15	2.9499
-13.75	2.9507
-12.5	2.9505
-11.25	2.9493
-10	2.9471
-8.75	2.9439
-7.5	2.9397
-6.25	2.9346
-5	2.9285
-3.75	2.9215
-2.5	2.9136
-1.25	2.9048
0	2.8952
1.4	3.2301
2.8	3.2159
4.2	3.2005
5.6	3.1842
7	3.1668
8.4	3.1484
9.8	3.1291
11.2	3.109
12.6	3.088
14	3.0662
15.4	3.0438
16.8	3.0206
18.2	2.9967
19.6	2.9723
21.2	3.3663
22.8	3.3329
24.4	3.2988
26	3.2642
27.6	3.229
29.2	3.1933
30.8	3.1572
32.4	3.1208
34	3.0842
35.6	3.0473
37.35	3.2907
39.1	3.2464
40.85	3.2021
42.6	3.158

Banana

Angle	ΔEuv
-49	
-47.25	2.0142
-45.5	2.0113
-43.75	2.0105
-42	2.0118
-40.25	2.0152
-38.5	2.0208
-36.75	2.0284
-35	2.038
-33.25	2.0495
-31.5	2.0629
-29.75	2.078
-28	2.0947
-26.25	2.1129
-24.5	2.1323
-22.75	2.1529
-21	2.1745
-19.25	2.1968
-17.5	2.2198
-15.75	2.2433
-14	2.267
-12.25	2.2909
-10.5	2.3147
-8.75	2.3384
-7	2.3617
-5.25	2.3846
-3.5	2.4069
-1.75	2.4285
0	2.4492
1.75	2.469
3.5	2.4878
5.25	2.5055
7	2.522
8.75	2.5372
10.5	2.5511
12.25	2.5637
14	2.5749
15.75	2.5846
17.5	2.5929
19.25	2.5996
21	2.6049
22.75	2.6087
24.5	2.6109
26.25	2.6117
28	2.611
29.75	2.6088
31.5	2.6051
33.25	2.6
35	2.5935
36.75	2.5857
38.5	2.5765
40.25	2.566
42	2.5542
43.75	2.5413
45.5	2.5272
47.25	2.5119
49	2.4956

Carrot

Angle	ΔEuv
-24	
-23	2.6185
-22	2.6139
-21	2.6083
-20	2.6017
-19	2.5943
-18	2.586
-17	2.5768
-16	2.5667
-15.25	2.6679
-14.5	2.6615
-13.85	2.3011
-13.2	2.2958
-12.55	2.2903
-11.9	2.2845
-11.2	2.4536
-10.5	2.4465
-9.75	2.6131
-9	2.6043
-8.25	2.5953
-7.5	2.586
-6.65	2.9193
-5.8	2.9066
-4.9	3.0634
-4	3.0485
-3	3.3692
-2	3.3497
-1	3.3297
0	3.3093
1	3.2885
2	3.2672
3	3.2456
4	3.2236
5	3.2013
6	3.1787
7	3.1557
8	3.1326
9	3.1091
10	3.0855
11.05	3.2141
12.1	3.1876
13.3	3.4102
14.5	3.5751
15.7	3.5399
16.9	3.5044
18.1	3.4689
19.3	3.4333
20.5	3.3976
21.7	3.3619
22.85	3.1884
24	3.1557
25.25	3.3931
26.5	3.3547
27.75	3.3165
29	3.2785
30.25	3.2407
31.5	3.2032

Cucumber

Angle	ΔEuv
-180	
-174.5	1.4109
-169	1.3945
-162.5	1.6424
-156	1.6552
-149.5	1.6864
-143	1.7337
-136.5	1.7939
-130	1.8632
-123.5	1.9378
-117	2.0142
-110.5	2.0891
-104	2.1593
-97.5	2.2223
-91	2.2758
-84.5	2.3181
-78	2.3475
-71.5	2.3632
-65	2.3643
-58.5	2.3508
-52	2.3228
-45.5	2.2814
-39	2.2279
-32.5	2.1645
-26	2.0939
-19.5	2.0198
-13	1.9464
-6.5	1.8785
0	1.8213
6.5	1.7795
13	1.7571
19.5	1.7565
26	1.7779
32.5	1.8194
39	1.8775
45.5	1.9471
52	2.0233
58.5	2.101
65	2.1755
71.5	2.2431
78	2.3007
84.5	2.3457
91	2.3766
97.5	2.3924
104	2.3925
110.5	2.3773
117	2.3472
123.5	2.3036
130	2.2479
136.5	2.1822
143	2.1089
149.5	2.0308
156	1.9511
162.5	1.8733
169	1.8012
174.5	1.4752
180	1.4383

Lime

Angle	ΔEuv
-85.2	
-82.15	2.8909
-79.1	2.9063
-76.05	2.9181
-73	2.9259
-69.95	2.9297
-66.9	2.9294
-63.85	2.925
-60.8	2.9164
-57.75	2.9037
-54.7	2.887
-51.65	2.8663
-48.6	2.8418
-45.55	2.8137
-42.5	2.7823
-39.45	2.748
-36.4	2.7111
-33.35	2.6721
-30.3	2.6316
-27.25	2.5901
-24.2	2.5483
-21.15	2.507
-18.1	2.4668
-15.05	2.4288
-12	2.3936
-9	2.3236
-6	2.2974
-3	2.2761
0	2.2604
2.75	2.0634
5.5	2.0601
8.25	2.0619
11	2.0686
13.75	2.0803
16.5	2.0967
19.25	2.1175
22	2.1424
24.75	2.1708
27.5	2.2022
30.25	2.2362
33	2.272
35.75	2.3093
38.5	2.3472
41.25	2.3855
44	2.4234
46.75	2.4606
49.5	2.4965
52.25	2.5308
55	2.5631
57.75	2.5931
60.5	2.6204
63.25	2.6449
66	2.6664
68.75	2.6846
71.5	2.6995
74.25	2.7109
77	2.7189

Potato

Angle	ΔEuv
-56	
-54	2.2454
-52	2.247
-50	2.2517
-48	2.2594
-46	2.2701
-44	2.2837
-42	2.3
-40	2.3188
-38	2.34
-36	2.3633
-34	2.3884
-32	2.4152
-30	2.4432
-28	2.4722
-26	2.502
-24	2.5323
-22	2.5627
-20	2.5931
-18	2.6232
-16	2.6527
-14	2.6815
-12	2.7092
-10	2.7358
-8	2.7611
-6	2.7849
-4	2.807
-2	2.8274
0	2.846
2	2.8626
4	2.8771
6	2.8896
8	2.9
10	2.9082
12	2.9142
14	2.9181
16	2.9197
18	2.9192
20	2.9166
22	2.9118
24	2.905
26	2.8961
28	2.8853
30	2.8725
32	2.858
34	2.8416
36	2.8236
38	2.8039
40	2.7827
42	2.7601
44	2.7361
46.25	3.0479
48.5	3.0143
50.75	2.9792
53	2.9427
55.25	2.905
57.5	2.8663

APPENDIX 3. ANSWER SHEET FOR THE CONTROL EXPERIMENT IN CHAPTER 4

Name: _____ Age:_____ Gender:_____

Object 1

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).

Object 2

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).

Object 3

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).

Object 4

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).

Object 5

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).

Object 6

What's the object? _____

How certain are you? _____

(Scale from 1-7, 1 not certain at all, 7 100% sure).