On Probing Appearance
Testing material-lighting interactions in an image-based canonical approach

Zhang, Fan

DOI
10.4233/uuid:7aab28bf-76cc-46ed-b305-b55d19723a5f

Publication date
2019

Document Version
Proof

Citation (APA)
Zhang, F. (2019). *On Probing Appearance: Testing material-lighting interactions in an image-based canonical approach*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:7aab28bf-76cc-46ed-b305-b55d19723a5f

On Probing Appearance

TESTING MATERIAL-LIGHTING INTERACTIONS IN AN IMAGE-BASED CANONICAL APPROACH

On Probing Appearance

TESTING MATERIAL-LIGHTING INTERACTIONS IN AN IMAGE-BASED CANONICAL APPROACH

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T. H. J. J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
dinsdag 29 oktober 2019 om 12:30 uur

door

Fan Zhang

Master of Science in Robotics,
King's College London, London, United Kingdom,
geboren te Beijing, China.

Dit proefschrift is goedgekeurd door de

promotor: Prof. dr. S. C. Pont
promotor: Prof. dr. H. de Ridder

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. S. C. Pont,	Technische Universiteit Delft
Prof. dr. H. de Ridder,	Technische Universiteit Delft
Onafhankelijke leden:	
Prof. dr. S. F. te Pas	Universiteit Utrecht
Prof. dr. E. Eisemann	Technische Universiteit Delft
Dr. K. Doerschner	Justus-Liebig University Giessen, Germany
Dr. B. Xiao	American University
Prof. dr. ir. K. M. B. Jansen	Technische Universiteit Delft, reservelid
Overig lid:	
Dr. P. Barla	INRIA, University of Bordeaux, France

Pascal Barla, PhD, heeft bijgedragen aan hoofdstukken 5 en 6 van deze dissertatie.

Printed by: Ipskamp Printing

This research was funded by the European Commission through the EU FP7 Marie Curie Initial Training Networks (ITN) project PRISM, Perceptual Representation of Illumination, Shape and Material (PITN-GA-2012-316746).

An electronic version of this dissertation is available at
`http://repository.tudelft.nl/`.

To everyone who has been part of this odyssey.

CONTENTS

1

Introduction

1

1.1. BACKGROUND

We live in a physical world, full of things and stuff (Adelson, 2001). "Things" are often called "objects", for example, a ball, a car, or a cup. "Stuff" usually refers to "materials", i.e. what the objects are made of, such as wood, metal, or glass. Undoubtedly, it is of vital importance for us to correctly recognize things and stuff to interact with the physical world, for instance not to mistake a rubber ball for an apple. And fortunately, the human brain has developed a functional visual system that is powerful and sophisticated to interpret the complex physical world. Within a split of a second, we can easily recognize even things and stuff that are not familiar to us (Sharan, Rosenholtz, & Adelson, 2014) and correctly plan actions to interact with them.

Instead of directly sensing the physical world, our eyes receive images formed by light-rays that are scattered by stuff in the physical world, i.e. two-dimensional (2D) projections of complex optical processes such as specular reflections, mutual reflections, refractions, shad(ow)ing, etcetera. The visual system receives projected images as input, and after intelligent processing, the result is an interpretation of the physical world, i.e. a perception. However, it is impossible to completely capture all optical properties of stuff in a single projected image of a thing, with a certain shape under a certain illumination. As a result, interpretations of stuff interact with interpretations of the shape and illumination of things (Anderson, 2011; Fleming, 2017; te Pas & Pont, 2005), even though these perceptions are usually experienced as unique. We must thus not study material perception without considering these interactions. The main goal of this thesis is to understand the perception of stuff or materials, including its interactions with shape and illumination.

The transformation from a multi-dimensional physical world to a two-dimensional image forms a well-posed optical and mathematical problem, called the forward problem (Poggio, Torre, & Koch, 1985). This problem is addressed in computer graphics and optical modeling. Theoretically, the forward problem is solvable given enough computational resources and knowledge about the optical properties of the world (ecological optics). However, due to the complexity of the optical structure, this forms a hard problem in practice. The inverse problem, inferring the physical properties from images, is addressed in computer and human vision (Poggio & Torre, 1984; Poggio & Koch, 1985; Kaas, 1992). Since any image of a thing can be generated by a range of combinations of materials, shapes and illuminations, the input is in that sense "ambiguous" and there is no unique solution of those properties (the problem is "mathematically underdetermined"). Therefore, in computer vision approaches to material, shape and illumination recognition, simplifications and assumptions are made to restrict the space of solutions, such as assuming that lighting usually comes from the upper left or that the surface is spherical and isotropic (e.g., Romeiro & Zickler, 2010). In this thesis we address the question how the human visual system deals with this underdetermined problem and makes sense of the ambiguous input.

In general, the optical structure of the physical world depends on (at least) three elements: the illumination of the environment (lighting), the three dimensional surface geometry (3D space and shape), and the objects' material scattering characteristics (material). Light can be described as a light field, that is by the luminance as a function of the position in a 3D space (x, y, z), direction (θ, ϕ) and wavelength (λ), i.e. six dimensions (Gershun, 1939). In this thesis we will study the perception of the stuff of a single thing and therefore can safely ignore the dependency on position, limiting the illumination description to a 3D function. The geometry can be simply described by the 3D shape. In order to make the problem tractable we restrict our studies to opaque materials and ignore texture and translucency. Bidirectional Reflectance Distribution Functions (BRDFs) describe how such an opaque material scatters light. A BRDF is defined by four parameters (4D), two for the incident light direction and two for the scattered light direction. Understanding how light scatters at surfaces makes it possible to describe and simulate the optical properties of materials and to understand how things appear. In computer science, BRDF models have been well employed in solving computer graphics and computer vision problems (Newell & Blinn, 1977; Cook & Torrance, 1982; Hapke, Nelson,& Smythe, 1998; Koenderink & Pont, 2003; Koenderink, Van Doorn, Dana, & Nayar, 1999; Nayar & Oren, 1995; Oren & Nayar, 1995; Phong, 1975; Torrance & Sparrow, 1967; Torrance, Sparrow, & Birkebak, 1966; van Ginneken, Stavridi, & Koenderink, 1998; Ward, 1992). To describe how light scatters internally and spatially over a surface for materials that transmit light, such as for translucent or transparent materials, the BRDFs need to be extended with more parameters.

In vision science, when studying the visual perception of materials, traditional feedforward vision theories proposed that the brain explicitly discounts the influence of lighting and shape to estimate material properties (Marr, 1982; Pizlo, 2001; Poggio & Koch, 1985; Poggio, Torre, & Koch, 1985) as a hypothesis. In such a bottom-up hierarchy, the visual system would first measure low-level image structures such as colors, orientations, spatial frequencies, and then group edges into contours and corners before entering mid-level vision. In mid-level vision, optical properties such as surface reflectance and mechanical properties such as softness and roughness would be estimated from the low-level image features and finally materials would be categorized into classes.

However, optical properties do not always simply correspond to perceived properties in visual perception (Fleming, 2014). The projection of such optical properties in images (on the retina or a camera) is confounded with lighting and shape. For example, a smooth surface with specular highlights and a matte-textured surface may result in the same image. The visual system will have to disentangle material, lighting and shape to retrieve information about each of these factors. Reverse modeling of the physical properties of material, lighting and shape, i.e. deconstructing the scene from the input in an "inverse optics" approach will run into "chicken and egg" problems, since inferences of the material (or lighting, or shape) properties require knowledge about the lighting and shape (or shape and material, or material and lighting). Additionally, recent studies revealed that the perception of specific optical and mechanical properties may be triggered in a top-down manner. They found that relationships and associations with familiar categories of materials instead of low- and mid-level image features influenced material perceptions, e.g. silky, cottony, metallic. (Fleming, Wiebel, & Gegenfurtner, 2013; Fleming, 2017). That is to say, the processes between mid-level estimations and high-level perception is not simply one-way.

Figure 1.1 shows a schematic representation of my approach to the perception process. The (objective) plenoptic function (Adelson & Bergen, 1991) describes all there is to see, the final result of the optical interactions between shape, material, and lighting, and thus the confounded optical structure, at one specific scale. The viewing then determines the proximal stimulus for the visual system - sampling just cross-sections of the plenoptic function. The visual system then processes the input, starting already at the retina, via low-level processes, to form mid and high level inferences, which mutually influence each other. It is even likely that the mid and high level inferences will influence viewing behavior and thereby the proximal stimuli and low level input (te Pas et al., 2017).

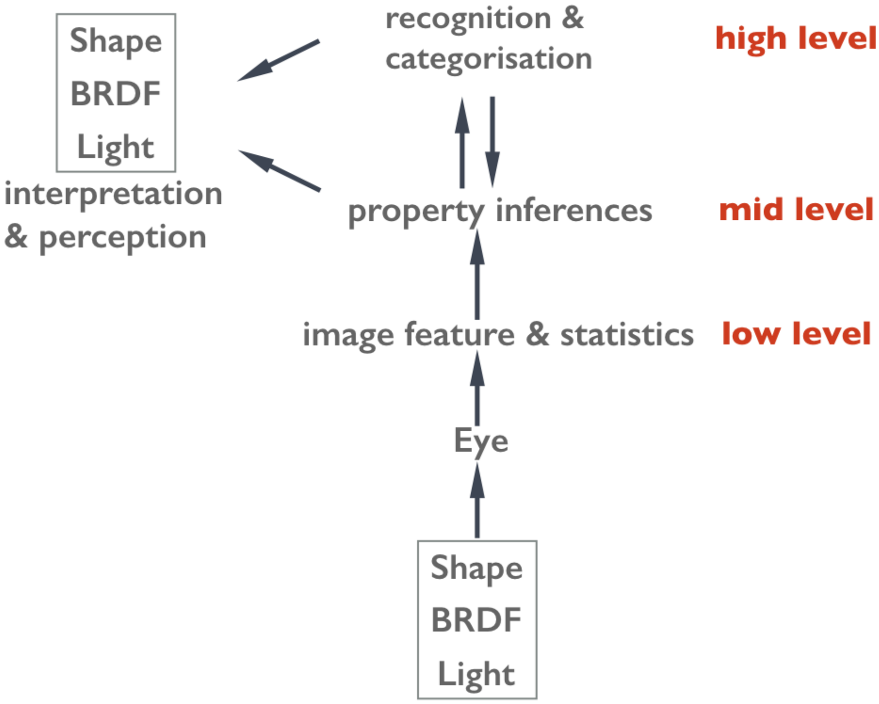

Figure 1.1: A schematic representation of our approach to "things" perception.

Although the study of material perception is still in its infancy compared to other directions in vision science, it has attracted increasing attention in the past decade. The main directions in material perception research are focused on how humans:

- recognize and categorize classes of materials (Sharan, Rosenholtz, & Adelson, 2009, 2014; Fleming, Wiebel, & Gegenfurtner, 2013; Nagai et al., 2015; Balas, 2017)
- estimate physical material properties such as gloss (Fleming et al. 2003; Pont & te Pas 2006; Vangorp et al. 2007; Ho et al. 2008; Xiao & Brainard 2008; Anderson & Kim 2009; Doerschner et al. 2010; Kim & Anderson 2010; Olkkonen and Brainard

2010, 2011; Motoyoshi & Matoba 2012; Marlow et al. 2012; Marlow & Anderson, 2013; Wiebel et al. 2015; Paulun et al., 2016), translucency (Singh & Anderson 2002; Fleming & Bülthoff 2005, Motoyoshi et al. 2005; Robilotto & Zaidi 2004; Faul & Ekroll 2011; Fleming et al. 2011, Schlüter & Faul 2014; Xiao et al., 2014), viscosity (Kawabe et al., 2005; Paulun et al., 2015; van Assen et al., 2016), stiffness (Bouman et al., 2013; Bi & Xiao, 2016; Bi et al., 2018;), etcetera

- attribute high-level concepts or meanings associated with certain materials (e.g. Karana, Hekkert, & Kandachar, 2009)
- understand light-material-shape interactions and perceive lighting effects (Doerschner, Boyaci, & Maloney, 2007; Fleming, Dror, & Adelson, 2003; te Pas & Pont, 2005; Pont & Te Pas, 2006; Vangorp, Laurijssen, & Dutré, 2007; Koenderink, Van Doorn, Wijntjes, & te Pas, 2012; Marlow, Kim, & Anderson, 2012; Zhang, de Ridder, & Pont, 2015).

Progress has been made in studies of both opaque and translucent/transparent materials. In these studies, stimuli were mostly rendered using computer graphics, so that the physical parameters such as surface reflectance and opacity could be controlled systematically and compared to perceptual judgments. An often addressed question in material perception literature concerns how opaque materials with matte-to-glossy variations are perceived. Stimuli with such variations can easily be created to study glossiness perception (e.g. Fleming, Dror, & Adelson, 2003; Marlow, Kim, & Anderson, 2012). One reason is that analytical BRDF models for forward scattering have been well developed. A common approach to model glossy materials is a linear weighted combination of a diffuse scattering and a specular reflection lobe. The weights of the diffuse and specular components can be manipulated to vary the glossiness systematically. When more parameters are included, it is possible to render realistic stimuli that account for a wide range of plastics and metals (Chadwick & Kentridge, 2015; Fleming, 2017). Less well-known are BRDF models for surface scattering and backward scattering. Surface scattering, that is scattering at grazing angles, independent of the illumination direction, happens in velvet and other materials with asperities (hence it is also sometimes called asperity scattering) (Koenderink and Pont, 2003). Backward scattering or retroreflection happens in surfaces with hemispherical pits or other similar surface structures (Pont, & Koenderink, 2002; Ikeuchi, 2014). The micro structure of opaque surfaces modulates the BRDF (Nayar & Oren, 1995) and natural surfaces can show a great variety of BRDF modes (Dana et al., 1999). A major question that arises is thus how to cover this wide range of possible reflectance types in order to increase the ecological validity of material perception studies.

As the starting point of our research, we describe the reflectance of any material as a combination of certain basic reflectance modes, extending the approach to glossy materials as a linear weighted combination of matte and specular modes. In order to cover a wide range of natural materials, the choice of basic modes should be "canonical", that is, consist of typical representations of frequently occurring types of reflectances that are mutually clearly distinct. On the basis of the models described above plus practical realizability, we implemented altogether four canonical material modes, namely, “matte”, “velvety”, “specular” and “glittery”, which altogether span a large part of the BRDF space (Figure 1.2). A backward mode was not implemented due to a lack of backward reflecting green paint. The glittery mode was added to explore what is well-known as a specular micro facet model, plus the occurrence of texture. A "bird" set of both real objects and computer renderings was created to represent the four canonical material modes, which were used as stimuli in our studies. To create the "bird" set of real objects, we covered four of the same bird-shaped objects with different finishes representing the corresponding surface scattering modes. With a simplified 3D model of the bird’s shape and state-of-art BRDF models, we then created the renderings. The canonical modes could next be combined in a linear weighted manner in order to systematically vary the resulting material appearance. This is done via optical mixing (Griffin, 1999): linear weighted superposition of the modes’ images. In this way, we operationalized the basic elements to study material perception in a canonical approach: with a limited number of modes, we can still systematically vary the material appearance such that material perception is explored in a wide ecological perspective.

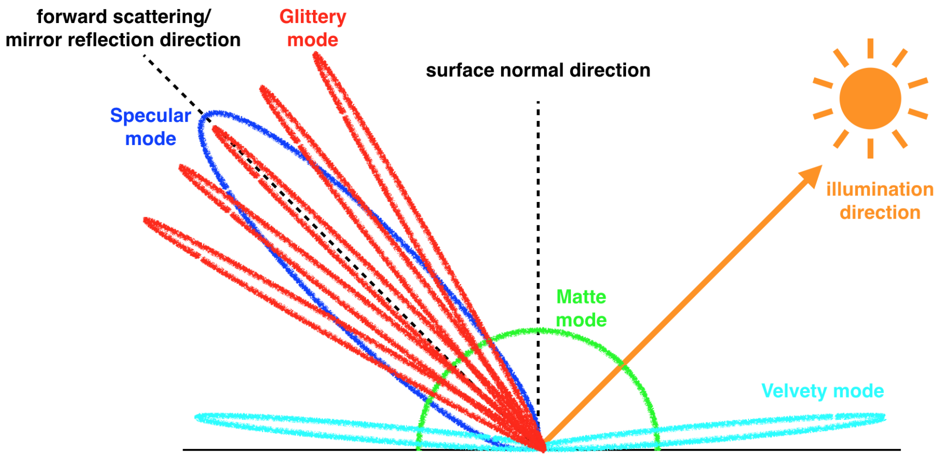

Figure 1.2: A schematic representation of our canonical modes in the BRDF space. Please note that the representation of the glittery mode is strictly not correct; the micro-scale BRDF of the glittery material can be represented by a single, sharply peaked, forward scattering mode; the meso-scale texture results from the fact that the meso-scale facets have different attitudes and thus scatter forward in slightly different directions; and the resulting macro-scale BRDF then depends on that facets’ attitude distribution. The red multi-peaked mode is an attempt to visually communicate these characteristics.

Since lighting, shapes and materials are confounded in the images of things, none of them can be addressed without also addressing the other factors. It would be extremely demanding and time-consuming for observers to test all variables simultaneously in psychophysical experiments. Thus, although progress has been made in research into how we estimate the surface geometry, surface reflectance, or illumination properties such as intensity, direction and diffuseness, little is known about the interactions between them. In order to study such interactions in a systematic way we implemented three canonical lighting modes within a spherical harmonics (SH) and perception based framework (Mury et al., 2009; Xia, Pont, & Heynderickx, 2017; Kartashova et al., 2018; Pont, 2019) that Pont found to be congruent with Kelly's lighting design approach (Kelly, 1952), namely, the "ambient", "focus", and "brilliance" lighting, respectively. Together with the squash tensor these canonical lighting modes form the local light field. Then we combined our canonical materials with the canonical lighting modes for both the real and rendered set to study the perceptual interactions between materials and lighting. Shape variations were tested in the last chapter.

This approach allows us to test whether material constancy can be achieved in different lights and of different shapes. It is known that in some extreme or unusual scenarios the same materials can appear very differently, while different materials may appear to be the same. For example, glossy surfaces appear to be matte in diffuse lighting (Dror, Willsky, & Adelson, 2004; Pont & te Pas, 2006; Zhang, de Ridder, & Pont, 2015; Zhang et al., 2016, 2018). The main reason for this phenomenon lies in the ambiguity of the resulting image: diffuse lighting "diffuses out" the highlights that are the main cue for glossiness perception.

Semantic information might influence estimations of material properties (Fleming, Wiebel, & Gegenfurtner, 2013). Therefore, although without the use of words it becomes especially difficult to test human perception, we prefer to test human visual perception in a purely visual but still quantitative way. To begin with, we created a novel material probe to begin with, using the canonical modes and optical mixing approach. Later, we compared results from our novel probing method and a few other psychophysical experimental methods such as scaling and discrimination, in which the use of words (semantic information) was also involved.

Using our canonical modes method, we implemented psychophysical experiments to measure human visual perception of materials. Psychophysics is one of the main approaches to study the perception of the physical world and how the human visual system represents the optical properties of things and stuff. Although it was developed to solve philosophical questions at first in the 19th Century, it became a powerful tool to study how the brain works in modern science. In psychophysical studies, physical stimuli are presented to human subjects to measure certain psychological responses from the subjects. The analysis and modeling then relates the physical parameters to the human judgments. With regard to material, light and shape perception we then run into the problem that the image is not uniquely defined by the physical parameters. In order to get a grip onto the relationships between those parameters and the perceptual judgments, one thus needs models of the ambiguities present in the image or an alternative approach to the rendering process. To our awareness there is only one such formal model (Belhumeur, Kriegman, & Yuille, 1999) and it is extremely hard to derive such for non-Lambertian (perfectly matte) materials. Our method uses an alternative rendering process instead.

We believe that there are certain key features in the images that trigger perceptions, regardless of confounds between shapes, lighting, and materials. Because the interactions between materials, lightings, and shapes influence such image key features in very complex and nonlinear manners, it is difficult to control the physical parameters of materials such that the resulting appearances change in a systematical way. Using optical mixing of canonical modes we can linearly combine and vary the contributions of image features directly. Therefore we regard this a suitable appearance-based alternative to visually test material perception via direct manipulation of the proximal stimulus for the visual system. This approach bears resemblance to painters' techniques in that it combines "layers" representing different features in a weighted manner and could therefore also be called a painterly approach.

1.2. THIS THESIS

1.2.1. RESEARCH QUESTIONS

The main aim of this thesis is to understand the visual perception of opaque materials of a wide ecological variety and interpret the perceptual results in a manner relating to shape and lighting variations. To start, we developed and tested a novel material probe that allowed optical mixing of our four canonical material modes in a painterly approach, to answer our first research question, which was: **Q1: Can material perception be measured quantitatively in a purely visual way?**

Lighting has an influence on the perception of materials. To systematically study the interplay between the perceptions of materials and lighting, we combined the canonical materials with the canonical lightings framework described above. First, we conducted a matching experiment using our material probe to answer **Q2: Can observers match optically mixed canonical materials while discounting canonical lightings?**

Then, to test and analyze the material-lighting interactions, we further developed our material probe to allow optical mixing of canonical lighting modes, asking **Q3a: Can observers match optically mixed canonical lightings while discounting canonical materials?** In the same study, we implemented another experimenting method on the same set of stimuli, namely a 4-category discriminating task, to study **Q3b: Can observers simultaneously discriminate materials and lightings?** and made a comparison between the results of the two types of experimental methods.

So far, we implemented different experimental methods to study material perception, including a novel probing method. Next, we wanted to further explore the material-dependent lighting effects we found. To be more specific, we found that certain material qualities could be brought out or attenuated by certain canonical lightings, and wondered **Q4: Can we predict material quality effects of light-material interactions?**

Lastly, we addressed the fine details of such effects on perceived material qualities in **Q5: Do light map orientations and the shape of the objects influence the perception of the associated material qualities?** To answer this question, we conducted an experiment in which we systematically varied four independent variables, namely the canonical material mode, the canonical lighting mode, the light map orientation, and the shape.

1.2.2. STRUCTURE OF THIS THESIS

This thesis consists of seven chapters. From Chapters 2 to Chapter 6, each chapter answers a research question listed in section 1.2.1, following a logical sequence such that each chapter answers a question brought up by its previous chapter. Together all chapters in this thesis explore the main mechanisms behind the perception of stuff or materials, including its interactions with shape and illumination. Each Chapter is self-contained and could be read independently.

- Chapter 2 is based on [Zhang-4], in which we developed a novel probing method to measure material perception in a purely visual way without involving semantic information (i.e. not using words) and in a painterly approach. To test the material probe, we integrated it in an interface inspired by a DJ's mixing desk and conducted a matching experiment. The matching task was expected to be difficult for inexperienced observers as they would have to manipulate four materials that all interact with each other. Yet we found that observers were well able to handle the probe. We also found interactions between our canonical material modes and different lightings, which led to the work of Chapter 3.
- Chapter 3 is based on [Zhang-10], and here we implemented the method developed in Chapter 2 to systematically study how canonical lighting modes influence the perception of our canonical material modes. In this study, a matching experiment was conducted using the same interface developed in Chapter 2. We found material-dependent lighting effects.
- Chapter 4 is based on [Zhang-3], in which we further investigated the material-dependent lighting effects, and found asymmetric perceptual confounds between materials and lightings. Three experiments were conducted. The first experiment further developed the material probe to allow optical mixing of canonical lighting modes in a matching task, focused on studying whether observers can match optically mixed lighting modes. In the second experiment, a 4-category discrimination task was conducted, focused on studying whether observers can simultaneously discriminate materials and lightings. The third experiment was conducted to compare and relate the two type of tasks in the previous two experiments.
- Chapter 5 is based on [Zhang-2]. In two experiments we investigated how environment illuminations could evoke certain material qualities using a list of nine terms that are commonly used to describe material qualities. The first experiment focused on the canonical lighting modes we previously implemented, and the second experiment focused on generic environments. Based on the results of the first experiment and novel metrics to quantify lighting properties, we were able to make predictions of the light effects on material appearance for generic light environments. The results of the second experiment validated the predictions we made.
- Chapter 6 is based on [Zhang-1]. In this chapter we investigated how light orientations of the environment illuminations and the shape of the object influence perceived material qualities for corresponding canonical material modes. In a rating

1

experiment, observers were asked to simultaneously rate the associated material quality for fifteen stimuli that differed in light orientations, per lighting and per shape. Results showed different effects of lighting, light orientation, and shape, depending on material modes.

Eventually, Chapter 7 concludes the work presented by interpreting our findings. The limitations of the current work and future directions are also discussed. Lastly, future issues are discussed and some examples are given to show possible future applications of our work.

BIBLIOGRAPHY

Anderson, B. L. (2011). Visual perception of materials and surfaces. Current Biology, 21(24), R978-R983.

Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. Journal of vision, 9(11), 10-10.

van Assen, J. J., Barla, P., & Fleming, R. (2016, May). Cues underlying liquid constancy. In Vision Science Society.

Balas, B. (2017). Children's use of visual summary statistics for material categorization. Journal of vision, 17(12), 22-22.

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. International journal of computer vision, 35(1), 33-44.

Bi, W., Jin, P., Nienborg, H., & Xiao, B. (2018). Estimating mechanical properties of cloth from videos using dense motion trajectories: Human psychophysics and machine learning. Journal of vision, 18(5), 12-12.

Bi, W., & Xiao, B. (2016, July). Perceptual constancy of mechanical properties of cloth under variation of external forces. In Proceedings of the ACM symposium on applied perception (pp. 19-23). ACM.

Bouman, K. L., Xiao, B., Battaglia, P., & Freeman, W. T. (2013). Estimating the material properties of fabric from video. In Proceedings of the IEEE international conference on computer vision (pp. 1984-1991).

Chadwick, A. C., & Kentridge, R. W. (2015). The perception of gloss: A review. Vision research, 109, 221-235.

Cook, R. L., & Torrance, K. E. (1982). A reflectance model for computer graphics. ACM Transactions on Graphics (TOG), 1(1), 7-24.

Dana, K. J., Van Ginneken, B., Nayar, S. K., & Koenderink, J. J. (1999). Reflectance and texture of real-world surfaces. ACM Transactions On Graphics (TOG), 18(1), 1-34.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2007). Testing limits on matte surface color perception in three-dimensional scenes with complex light fields. Vision Research, 47(28), 3409-3423.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. Journal of Vision, 10(4), 8-8.

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

Faul, F., & Ekroll, V. (2011). On the filter approach to perceptual transparency. Journal of Vision, 11(7), 7-7.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. Journal of vision, 3(5), 3-3.

Fleming, R. W., & Bülthoff, H. H. (2005). Low-level image cues in the perception of translucent materials. ACM Transactions on Applied Perception (TAP), 2(3), 346-382.

Fleming, R. W., Jäkel, F., & Maloney, L. T. (2011). Visual perception of thick transparent materials. Psychological science, 22(6), 812-820.

Fleming, R. W., Wiebel, C., & Gegenfurtner, K. (2013). Perceptual qualities and material classes. Journal of vision, 13(8), 9-9.

Fleming, R. W. (2014). Visual perception of materials and their properties. Vision research, 94, 62-75.

Fleming, R. W. (2017). Material perception. Annual review of vision science, 3, 365-388.

Gershun, A. (1939). The light field. Journal of Mathematics and Physics, 18(1-4), 51-151.

van Ginneken, B., Stavridi, M., & Koenderink, J. J. (1998). Diffuse and specular reflectance from rough surfaces. Applied optics, 37(1), 130-139.

Griffin, L. D. (1999). Partitive mixing of images: a tool for investigating pictorial perception. JOSA A, 16(12), 2825-2835.

Hapke, B., Nelson, R., & Smythe, W. (1998). The opposition effect of the moon: coherent backscatterandshadow hiding. Icarus, 133(1), 89-97.

Ho, Y. X., Landy, M. S., & Maloney, L. T. (2008). Conjoint measurement of gloss and surface texture. Psychological Science, 19(2), 196-204.

Ikeuchi, K. (2014). Computer vision: A reference guide. Springer Publishing Company, Incorporated.

Kartashova, T., de Ridder, H., te Pas, S. F., & Pont, S. C. (2018). Visual Light Zones. i-Perception, 9(3), 2041669518781381.

Kass, M. (1992). Inverse problems in computer graphics. In Creating and animating the virtual world (pp. 21-33). Springer, Tokyo.

Kawabe, T., Maruya, K., Fleming, R. W., & Nishida, S. Y. (2015). Seeing liquids from visual motion. Vision research, 109, 125-138.

Kelly, R. (1952). Lighting as an integral part of architecture. College Art Journal, 12(1), 24-30.

1

Kim, J., & Anderson, B. L. (2010). Image statistics and the perception of surface gloss and lightness. Journal of Vision, 10(9), 3-3.

Koenderink, J., & Pont, S. (2003). The secret of velvety skin. Machine vision and applications, 14(4), 260-268.

Koenderink, J. J., Van Doorn, A. J., Dana, K. J., & Nayar, S. (1999). Bidirectional reflection distribution function of thoroughly pitted surfaces. International Journal of Computer Vision, 31(2-3), 129-144.

Marlow, P. J., & Anderson, B. L. (2013). Generative constraints on image cues for perceived gloss. Journal of vision, 13(14), 2-2.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. Current Biology, 22 (20), 1909–1913.

Marr, D. (1982). Vision. San Francisco: W. H. H. Freeman.

Motoyoshi, I., & Matoba, H. (2012). Variability in constancy of the perceived surface reflectance across different illumination statistics. Vision Research, 53(1), 30-39.

Motoyoshi, I., Nishida, S. Y., & Adelson, E. H. (2005, August). Luminance re-mapping for the control of apparent material. In Proceedings of the 2nd symposium on Applied perception in graphics and visualization (pp. 165-165). ACM.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2009). Structure of light fields in natural scenes. Applied Optics, 48(28), 5386-5395.

Nagai, T., Matsushima, T., Koida, K., Tani, Y., Kitazaki, M., & Nakauchi, S. (2015). Temporal properties of material categorization and material rating: visual vs non-visual material features. Vision Research, 115, 259-270.

Nayar, S. K., & Oren, M. (1995). Visual appearance of matte surfaces. Science, 267(5201), 1153-1156.

Newell, M. E., & Blinn, J. F. (1977, January). The progression of realism in computer generated images. In Proceedings of the 1977 annual conference (pp. 444-448). ACM.

Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. Journal of vision, 10(9), 5-5.

Olkkonen, M., & Brainard, D. H. (2011). Joint effects of illumination geometry and object shape in the perception of surface reflectance. i-Perception, 2(9), 1014-1034.

Oren, M., & Nayar, S. K. (1995). Generalization of the Lambertian model and implications for machine vision. International Journal of Computer Vision, 14(3), 227-251.

te Pas, S. F., & Pont, S. C. (2005, August). A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization (pp. 75-81). ACM.

te Pas, S. F., Pont, S. C., Dalmaijer, E. S., & Hooge, I. T. (2017). Perception of object illumination depends on highlights and shadows, not shading. Journal of vision, 17(8), 2-2.

Paulun, V. C., Gegenfurtner, K. R., Goodale, M. A., & Fleming, R. W. (2016). Effects of material properties and object orientation on precision grip kinematics. Experimental brain research, 234(8), 2253-2265.

Paulun, V. C., Kawabe, T., Nishida, S. Y., & Fleming, R. W. (2015). Seeing liquids from static snapshots. Vision research, 115, 163-174.

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

Pizlo, Z. (2001). Perception viewed as an inverse problem. Vision Research, 41 (24), 3145–3161.

Poggio, T., & Torre, V. (1984). Ill-posed problems and regularization analysis in early vision (No. AI-M-773). MASSACHUSETTS INST OF TECH CAMBRIDGE ARTIFICIAL INTELLIGENCE LAB.

Poggio, T., & Koch, C. (1985). Ill-posed problems in early vision: From computational theory to analogue networks. Proceedings of the Royal Society of London B: Biological Sciences, 226 (1244), 303–323.

Poggio, T., Torre, V., & Koch, C. (1985, Setpember 26). Computational vision and regularization theory. Nature, 317 (6035), 314–319.

Pont, S. C. (2019). Light: Toward a Transdisciplinary Science of Appearance and Atmosphere. Annual review of vision science, 5.

Pont, S. C., & Koenderink, J. J. (2002). Bidirectional reflectance distribution function of specular surfaces with hemispherical pits. JOSA A, 19(12), 2456-2466.

Pont, S. C., Koenderink, J. J., van Doorn, A. J., Wijntjes, M. W., & te Pas, S. F. (2012, February). Mixing material modes. In Human Vision and Electronic Imaging XVII (Vol. 8291, p. 82910D). International Society for Optics and Photonics.

Pont, S. C., & te Pas, S. F. (2006). Material—Illumination ambiguities and the perception of solid objects. Perception, 35(10), 1331-1350.

Robilotto, R., & Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. Journal of Vision, 4(3), 5-5.

Romeiro, F., & Zickler, T. (2010, September). Blind reflectometry. In European conference on computer vision (pp. 45-58). Springer, Berlin, Heidelberg. Zhang, F., de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. Human Vision and Electronic Imaging XX, 9394, 93940Q.

Sharan, L., Rosenholtz, R., & Adelson, E. (2009). Material perception: What can you see in a brief glance?. Journal of Vision, 9(8), 784-784.

Sharan, L., Rosenholtz, R., & Adelson, E. H. (2014). Accuracy and speed of material categorization in real-world images. Journal of vision, 14(9), 12-12.

Schlüter, N., & Faul, F. (2014). Are optical distortions used as a cue for material properties of thick transparent objects?. Journal of Vision, 14(14), 2-2.

Singh, M., & Anderson, B. L. (2002). Toward a perceptual theory of transparency. Psychological review, 109(3), 492.

Torrance, K. E., & Sparrow, E. M. (1967). Theory for off-specular reflection from roughened surfaces. Josa, 57(9), 1105-1114.

Torrance, K. E., Sparrow, E. M., & Birkebak, R. C. (1966). Polarization, directional distribution, and off-specular peak phenomena in light reflected from roughened surfaces. JOSA, 56(7), 916-925.

Vangorp, P., Laurijssen, J., & Dutré, P. (2007, August). The influence of shape on the perception of material reflectance. In ACM Transactions on graphics (TOG) (Vol. 26, No. 3, p. 77). ACM.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. Computer Graphics, 26(2), 265-272.

Wiebel, C. B., Toscani, M., & Gegenfurtner, K. R. (2015). Statistical correlates of perceived gloss in natural images. Vision research, 115, 175-187.

Xia, L., Pont, S. C., & Heynderickx, I. (2017). Light diffuseness metric part 1: Theory. Lighting Research & Technology, 49(4), 411-427

Xiao, B., & Brainard, D. H. (2008). Surface gloss and color perception of 3D objects. Visual neuroscience, 25(3), 371-385

Xiao, B., Walter, B., Gkioulekas, I., Zickler, T., Adelson, E., & Bala, K. (2014). Looking against the light: How perception of translucency depends on lighting direction. Journal of vision, 14(3), 17-17.

Zhang, F., de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. Human Vision and Electronic Imaging XX, 9394, 93940Q.

Zhang, F., de Ridder, H., Fleming, R. W., & Pont, S. (2016). MatMix 1.0: Using optical mixing to probe visual material perception. Journal of vision, 16(6), 11-11.

Zhang, F., de Ridder, H., & Pont, S. C. (2018). Asymmetric perceptual confounds between canonical lightings and materials. Journal of vision, 18(11), 11-11.

List of Publications

International peer-reviewed journal papers

Zhang-1 **Zhang, F.**, de Ridder, H., Barla, P., & Pont, S. (2019). Effects of light directions on the visual perception of canonical materials. *under revision*

Zhang-2 **Zhang, F.**, de Ridder, H., Barla, P., & Pont, S. (2019) A systematic approach to testing and predicting light-material interactions. *Journal of vision,19*(4), 1-22. doi:10.1167/19.4.11

Zhang-3 **Zhang, F.**, de Ridder, H., & Pont, S. (2018) Asymmetric perceptual confounds between canonical lightings and materials. *Journal of vision,18*(11), 1-19. doi:10.1167/18.11.11

Zhang-4 **Zhang, F.**, de Ridder, H., Fleming, R. W., & Pont, S. (2016). MatMix 1.0: Using optical mixing to probe visual material perception. *Journal of vision,16*(6), 1-18. doi:10.1167/16.6.11

Other publications

Zhang-5 **Zhang, F.**, de Ridder, H., Barla, P., & Pont, S. (2017). Material dependent appearance effects brought out by natural light environments. Poster presentation at *European Conference on Visual Perception 2017*, Berlin, Germany.

Zhang-6 **Zhang, F.**, de Ridder, H., van Egmond, R., & Pont, S. (2017). The interplay between material qualities and lighting. *Journal of vision,17*(10), 228-228. doi:10.1167/17.10.228

Zhang-7 **Zhang, F.**, de Ridder, H., & Pont, S. (2016). Canonical Material and Illumination Confounds. In *PERCEPTION*(Vol. 45, pp. 130-130).

Zhang-8 **Zhang, F.**, de Ridder, H., & Pont, S. (2016). Can people match optically mixed canonical lighting modes? *Journal of Vision,16*(12), 642-642. doi:10.1167/16.12.642

Zhang-9 **Zhang, F.**, de Ridder, H., & Pont, S. (2015). Matmix 1.0, a noval material probe for quantitatively measuring visual perception of materials. *Journal of Vision,15*(12), 824-824. doi:10.1167/15.12.824

Zhang-10 **Zhang, F.**, de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. In *Proc. SPIE 9394, Human Vision and Electronic Imaging XX*, 93940Q

(March 17, 2015). doi:10.1117/12.2085021

Zhang-11 **Zhang, F.**, de Ridder, H., & Pont, S. (2014). Towards an interactive probe for material studies. In *Proceedings Experiencing light 2014: international conference on the effects of light on wellbeing*, Eindhoven, The Netherlands, 10-11 November 2014(p. 119).

2

MatMix 1.0: Using optical mixing to probe visual material perception

Abstract

MatMix 1.0 is a novel material probe we developed for quantitatively measuring visual perception of materials. We implemented optical mixing of four canonical scattering modes, represented by photographs, as the basis of the probe. In order to account for a wide range of materials, velvety and glittery (asperity and mesofacet scattering) were included besides the common matte and glossy modes (diffuse and forward scattering). To test the probe, we conducted matching experiments in which inexperienced observers were instructed to adjust the modes of the probe to match its material to that of a test stimulus. Observers were well able to handle the probe and match the perceived materials. Results were robust across individuals, across combinations of materials, and across lighting conditions. We conclude that the approach via canonical scattering modes and optical mixing works well, although the image basis of our probe still needs to be optimized. We argue that the approach is intuitive, since it combines key image characteristics in a "painterly" approach. We discuss these characteristics and how we will optimize their representations.

Published as: Zhang, F., de Ridder, H., Fleming, R. W., & Pont, S. C. (2016) MatMix 1.0: Using optical mixing to probe visual material perception. Journal of Vision, 16(6).

2.1. INTRODUCTION

Natural materials scatter light in various manners. Even if we limit ourselves to the main scattering characteristics of opaque materials, we probably still need about a dozen scattering types or canonical modes to represent most materials. Bidirectional reflectance distribution functions (BRDFs) provide a physical description of how opaque material surfaces scatter light. Knowing how light scatters from surfaces makes it possible to simulate materials using parametric BRDF models in computer renderings (Newell & Blinn, 1977; Cook & Torrance, 1982; Hapke, Nelson, & Smythe, 1998; Koenderink & Pont, 2003; Koenderink, Van Doorn, Dana, & Nayar, 1999; Nayar & Oren, 1995; Oren & Nayar, 1995; Phong, 1975; Torrance & Sparrow, 1967; Torrance, Sparrow, & Birkebak, 1966; van Ginneken, Stavridi, & Koenderink, 1998; Ward, 1992). Generally speaking, if the scattering properties or optical characteristics of materials can be accurately described, the so-called forward rendering problem can be solved. However, it is very unlikely that these optical characteristics correspond to the representation of the visual attributes in the brain (Fleming, 2014). In other words, we do not see BRDFs. On the one hand, a BRDF combined with various object shapes and lighting conditions can result in different images of the same material (we consider an image as the resulting optical structure projected on a picture or the retina). On the other hand, different combinations of BRDF, object shape, and lighting can result in similar images. In other words, the so-called inverse problem does not have a unique solution. Thus, images contain ambiguities of material, shape, and light (Belhumeur, Kriegman, & Yuille, 1999), and consequently material, light, and shape perception are confounded (Anderson, 2011). In the present study, we will exploit this metamerism and investigate key ingredients in the images that trigger our material perceptions.

Most investigations into material perception are confined to matte–glossy variations. That is, the perception of materials varying from matte to shiny has been intensively studied on perceived glossiness (Anderson, 2011; Anderson & Kim, 2009; Fleming, 2012; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2006, 2008; Kim, Marlow, & Anderson, 2011; Marlow, Kim, & Anderson, 2012; Motoyoshi, Nishida, Sharan, & Adelson, 2007; Nishida & Shinya, 1998; Pellacini, Ferwerda, & Greenberg, 2000; Vangorp, Laurijssen, & Dutré, 2007; Wiebel, Toscani, & Gegenfurtner, 2015; Wijntjes & Pont, 2010). In our research, we want to address material perception not only within the matte–glossy continuum but also for as wide a range of natural materials as possible. To date, little is known about the visual perception of materials outside the matte–glossy dimension, such as velvetiness (Koenderink & Pont, 2003; Nishida, Sawayama, & Shimokawa, 2015; te Pas & Pont, 2005) or other material dimensions (Fleming, Wiebel, & Gegenfurtner, 2013; Sharan, Rosenholtz, & Adelson, 2014). The main problem seems to be the lack of a tool to test purely visually (without referring to physical parameters or attributes) and quantitatively what material is perceived, for a wide range of materials.

We hereby present MatMix 1.0, a novel material probe using optical mixing, which will be explained in the next section. MatMix 1.0 is meant to account for a wide range of opaque materials. We use optical mixing of four canonical scattering modes as a tool for quantitatively measuring visual perception of materials. In our main study, we integrated the probe into a MATLAB graphical user interface and conducted two matching experiments without (Experiment 1) and with (Experiment 2) variation of the illumination and viewpoint conditions. Images of real objects were used as a basis set. Before the experiments started, we expected the task to be difficult for inexperienced observers, as they would have to simultaneously manipulate four different scattering modes. Surprisingly, we found that all participants could handle MatMix 1.0 well, as indicated by the finding that they performed far above chance level within reasonable amounts of time. In an additional study, we replaced the images with renderings and conducted a similar matching experiment (Experiment 3). Again, participants performed far above chance level, demonstrating that the approach works well with both real and simulated materials. In the General discussion and the Conclusion we address the relationships between a few key image characteristics and the results.

2.2. MatMix 1.0: A novel material probe

2.2.1. Optical mixing: A painterly approach

Many arbitrary materials can be represented by linearly combining surface scattering distributions (Matusik, Pfister, Brand, & McMillan, 2003; Pellacini et al., 2000; Ward, 1992). Instead of directly combining reflectance functions, we propose to linearly superpose images of objects with the same shape but finished with different materials. This image-combination process, called optical mixing, was introduced by Griffin (1999), who also described the mathematics behind the opticalmixing method and showed that it could be used as a tool for visual-perception studies. In Brainard's lab, Griffin's partitive mixing method was applied to reduce the number of stimuli to be rendered for their experiments (Olkkonen & Brainard, 2010; Radonjić, Cottaris, & Brainard, 2015; Xiao & Brainard, 2008). Although applying image mixing was not the main purpose of those studies, it can still be concluded from them that implementing optical mixing in psychophysical studies is indeed feasible and efficient. However, it has not been implemented yet for variations other than matte–glossy.

The optical-mixing procedure shows an interesting analogy with how a painter renders materials in a scene. Most painters do not think about image statistics or BRDFs when they paint. Instead, their approach is more similar to optical mixing of key visual ingredients layer by layer. A frequently observed recipe for oil paintings (Wallert, 1999) is to first draw the contour of an object, then apply the matte layer (the diffuse body scattering), and finally add highlights or a bright contour to render specular or velvety elements (forward or asperity scattering). We reasoned that optical mixing of nonspherical objects of arbitrary scattering modes should work because it similarly combines key image ingredients that trigger our perceptions—even though it may be physically incorrect.

Pont, Koenderink, Van Doorn, Wijntjes, and te Pas (2012) generated optical mixtures of three canonical scattering modes (matte, velvety, and specular) by optically mixing real objects in a viewing box. The task for the observers was to rate perceived material qualities such as glossiness, warmth, hardness, and softness. In a follow-up study, observers performed the same task, but now the stimuli were optically mixed images of matte, velvety, and specular materials displayed on a screen (Pont, te Pas, & Wijntjes, 2014). They obtained robust and systematic ratings for material qualities as a function of the weights of the three modes in both experiments. On the basis of these studies, we hypothesized that observers should be able to match the perceived material of a certain object if they have the opportunity to create a mixture with desired material attributes via control over the weights of the underlying canonical material modes in that mixture. This forms the basis of our proposed new materials probe, MatMix 1.0.

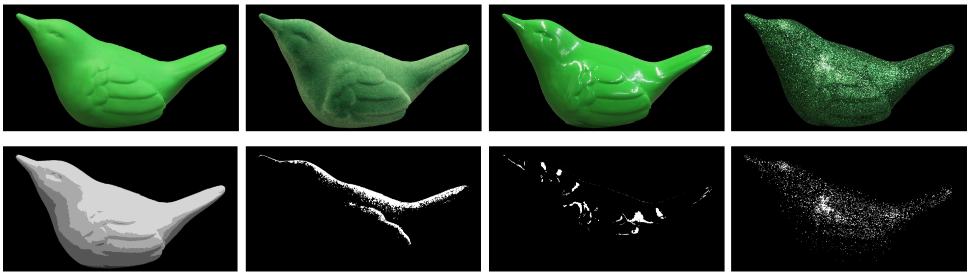

Figure 2.1: The top row shows the images of the birdlike object with the four materials representing the chosen canonical scattering modes: diffuse, asperity, forward, and meso-facet scattering (from left to right). These modes are represented by matte, velvety, specular, and glittery materials. The bottom row shows the prototypical image characteristics of each material. Note that the reflectance components are not only in different directions in BRDF space but result in characteristics of the proximal image in different regions too. For the image of the matte bird, the green channel was posterized from 255 to six levels. For the velvety, specular, and glittery bird images, we performed red-channel thresholding at the 50% level. These extremely simple processes resulted in smooth shading from the top to the bottom of the object for the matte object, bright contours for the velvet object, highlights at specular points for the specular object, and bright speckles all over the surface for the glittery object..

For MatMix 1.0, we employed four basis materials by finishing four bird-shaped objects with matte, velvety, specular, and glittery materials (Figure 2.1). These materials represent four canonical scattering modes, namely diffuse, asperity, forward and meso-facet scattering modes. The scattering distributions of these four canonical scattering modes together span a large part of the BRDF space. Because the main scattering directions of these scattering modes are different, the key characteristics in the images of the corresponding objects will end up in different locations on the object too. This means that the reflectance components not only are complementary in BRDF space but also allow the user to adjust different characteristics of the proximal image. With simple image analyses, the prototypical image characteristics of the four materials can be easily distinguished from each other, as shown in Figure 2.1. Note that there are many alternative image manipulations that would give similar results; the examples just serve to demonstrate the main idea.

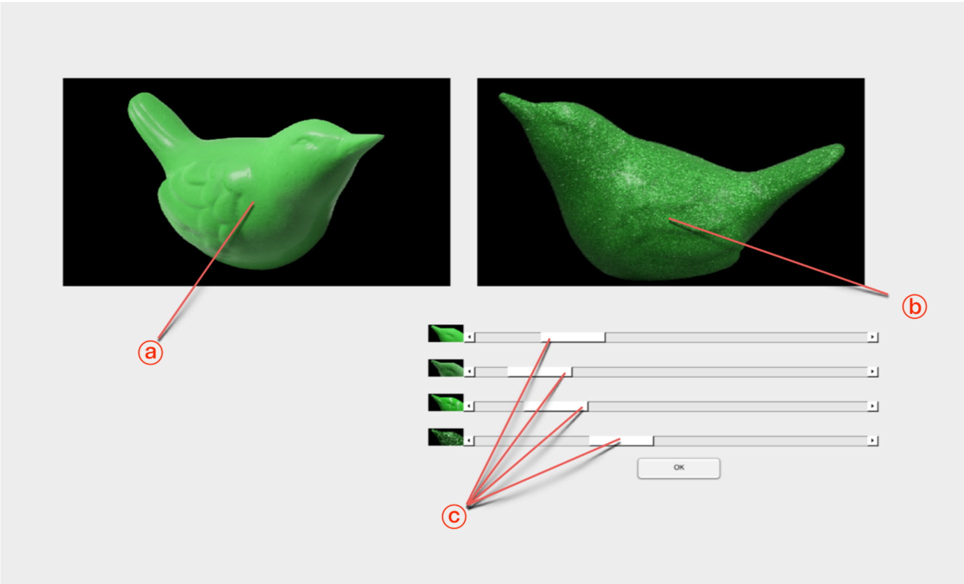

Figure 2.2: The interface of MatMix 1.0. (a) Stimulus image. (b) Material probe, generated by linear weighted superposition of the four images representing the canonical scattering modes. (c) Four sliders, with the position of each slider bar representing the selected weight value per material mode, ranging from 0 to 1.2 (left to right). The icon on the left of each slider visualizes the corresponding material component. The task of the observers was to change the material of the probe to match the stimulus. They could take as much time as they needed. Observers could click the OK button below the sliders to finish the matching procedure. Here, it is obvious that the two materials do not match.

2.2.2. The interface of MatMix 1.0

Inspired by audio-mixing desks, we built a user interface consisting of four sliders, a stimulus window, and a probe window (Figure 2.2). During each matching trial, the stimulus image and the probe image were simultaneously presented to the observers in the corresponding windows, with the stimulus on the lefthand side and the probe on the right-hand side. The four sliders were positioned directly underneath the probe window. In order to give purely visual information, we avoided the use of terms like "matte," "velvet," and so on in the interface. Instead, we put cropped images (the head parts of the bird images) in front of each slider, representing the material modes. The position of each slider bar represents the selected weight value per material mode, varying between 0 and 1.2.

Inspired by audio-mixing desks, we built a user interface consisting of four sliders, a stimulus window, and presented to observers on a linearly calibrated Apple 15-in. Retina display.

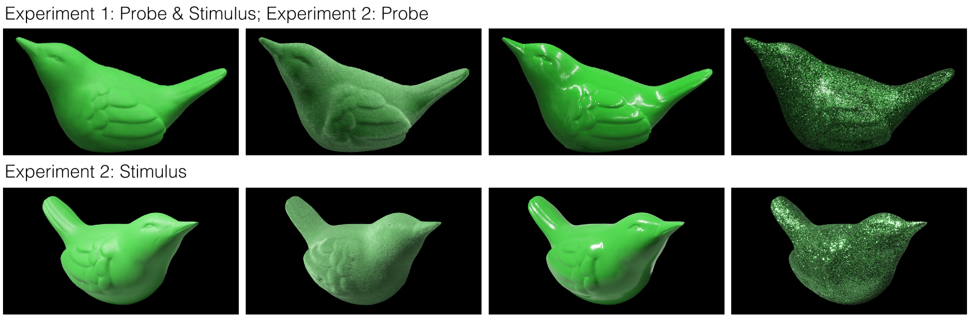

Figure 2.3: Basis images. From left to right, the columns represent the matte, velvety, specular, and glittery modes, respectively. The images in the top row were taken under office lighting and were used as basis images for the probe in the main study. These images were also used as the basis for the stimulus images in Experiment 1. The images in the bottom row were taken in studio lighting and from a different viewing angle than the first set, and were used as the basis for the stimulus in Experiment 2.

2.2.3. BASIS IMAGES

The surfaces of four physical objects with identical shapes were finished with matte, velvety, specular, and glittery materials. The bird-shaped objects were purchased in a shop. They were originally made of ceramic and had exactly the same shape. The matte and specular birds were created by spray-painting them with matte and glossy paint, respectively (both color RAL 6018). The glittery bird was created by repeatedly sprinkling green glitter over a layer of spray glue. The velvety bird was finished by a factory using a technique called flocking (color RAL 6018). These materials represent diffuse, asperity, forward, and meso-facet scattering modes, respectively. We took photos of the objects under office lighting and under studio lighting from different viewing angles. The office lighting consisted of multiple fluorescent tubes in the ceiling of a room without daylight. The studio lighting consisted of a halogen spotlight from the left side of the object. The camera settings were kept constant per lighting condition and we used raw imaging in order to photometrically gauge the basis images. Furthermore, to allow superposition of the basis images, we placed each object in exactly the same position. To do so, we drew their cast shadows and base outlines on their groundings as references. Next we adjusted the white balance of the raw images using Adobe Photoshop so that the highlights were all white.We did this in the same manner for all images per lighting condition. Then we segmented the images using the shared contours of the birds and made the background black for all images. Last, in order to avoid color interactions, we set the hue value to 0.33 (green) for all images usingMATLAB. Because the birds were pure green, this transformation had a negligible influence on the images (Figure 2.3). The saturation of the colors was not adjusted, because the saturation as a function of lighting and viewing angles can vary strongly. For instance, specular reflections lower the saturation of highlights. This effect depends on the type of scattering (Klinker, Shafer, & Kanade, 1987; Koenderink et al., 1999; Koenderink & Pont, 2008; Shafer, 1985; Wolff, 1994). This is why the different modes have substantial differences in saturation.

2.2.4. The probe: MatMix 1.0

The probe is a linearly superposed optical mixture of the basis images. The mixing process can be illustrated by Equation 2.1:

$$I_{probe} = w_m \cdot I_m + w_v \cdot I_v + w_s \cdot I_s + w_g \cdot I_g, \tag{2.1}$$

where subscripts {m, v, s, g} denote the four scattering modes matte, velvety, specular, and glittery, representing the four canonical scattering modes (diffuse, asperity, forward, and meso-facet scattering); $\{w_m,\ w_v,\ w_s,\ w_g\}$ are the weight values corresponding to the positions of the slider bars, ranging from 0 to 1.2 (see Figure 2.2); and $\{I_m,\ I_v,\ I_s,\ I_g\}$ are the basis images under office lighting (top row in Figure 2.3) for Experiments 1 and 2 in the main study. The linearly mixed image I_{probe} plus the interface forms the probe MatMix 1.0, which allows real-time dynamic and interactive variation of a visual presentation of material through adjustments of the slider bars.

2.3. Main study (optical mixing with images of real objects)

2.3.1. Introduction

In the main study we tested the material probe we developed. The study consisted of two experiments, which mainly differed in the illumination and viewpoint conditions under which the photos of real objects were taken. In Experiment 1, we created the stimuli by mixing the basis images shown in the top row in Figure 2.3. In Experiment 2, we created the stimuli by mixing the basis images shown in the bottom row in Figure 2.3. Thus, in Experiment 1 the stimuli and probe were mixed from the same basis, while in Experiment 2 the stimuli were mixed from a different basis than the probe.

2.3.2. Method

Stimuli

We tested 15 weight combinations of the four scattering modes, as shown in Table 2.1. The basis images were linearly superposed, implementing Equation 2.1in the form:

$$I_{stimulus} = \left(w_m{}' + x_m\right) \cdot I_m + \left(w_m{}' + x_m\right) \cdot I_v + \left(w_m{}' + x_m\right) \cdot I_s + \left(w_m{}' + x_m\right) \cdot I_g, \tag{2.2}$$

where $\{\ w_m{}',\ w_v{}',\ w_s{}',\ w_g{}'\ \}$ are the weights of the scattering modes; $\{x_m,\ x_v,\ x_s,\ x_g\}$ are randomly generated offsets in a range from -0.1 to 0.1 that were added to the nonzero weights only; and $\{I_m,\ I_v,\ I_s,\ I_g\}$ are the stimulus basis images shown in Figure 2.3 (top row for Experiment 1, bottom row for Experiment 2). The resulting linearly mixed image is the stimulus image $I_{stimulus}$. The complete set of stimulus images for Experiments 1 and 2 is shown in Figure 2.4.

Stimulus	$w_m{}'$	$w_v{}'$	$w_s{}'$	$w_g{}'$
1	1	0	0	0
2	0	1	0	0
3	0	0	1	0
4	0	0	0	1
5	0.5	0.5	0	0
6	0.5	0	0.5	0
7	0.5	0	0	0.5
8	0	0.5	0.5	0
9	0	0.5	0	0.5
10	0	0	0.5	0.5
11	0.33	0.33	0.33	0
12	0.33	0.33	0	0.33
13	0.33	0	0.33	0.33
14	0	0.33	0.33	0.33
15	0.25	0.25	0.25	0.25

Table 2.1: Overview of the weight combinations of the four material modes that were used to generate the stimulus images. There were 15 stimuli in total.

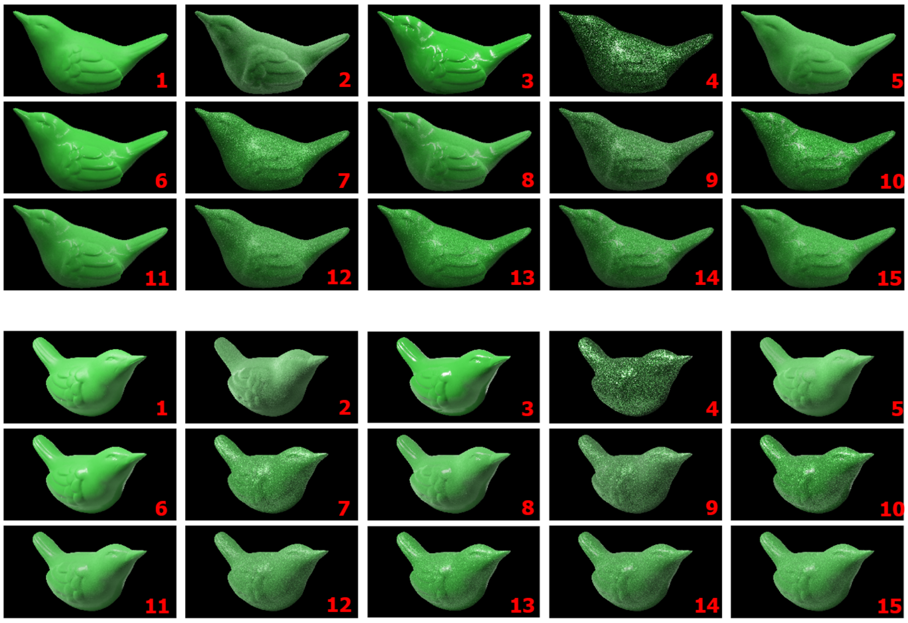

Figure 2.4: The stimuli. The top set represents the test stimuli in Experiment 1. The bottom set represents the test stimuli in Experiment 2. The randomly generated offsets $\{x_m, x_v, x_s, x_g\}$ were set to 0 to generate these images. The numbers in the images correspond to the stimulus numbers in Table 1.

Observers

There were eight paid inexperienced participants in total (four men and four women, aged 23 to 30), with normal or corrected-to-normal vision. All of them participated first in Experiment 1, and a few days later in Experiment 2. Participants read and signed a consent form before the experiments. The experiments were conducted in agreement with the Declaration of Helsinki and local ethical guidelines and approved by the Human Research Ethics Committee of the Delft University of Technology.

Procedure

The positions of the slider bars (i.e., the initial weights of the probe) were randomly initialized in each trial. In Experiments 1 and 2, each stimulus weight combination in Table 1 was repeated three times. Three repeats combined with 15 different weight combinations, making each experiment 45 trials in total. The trials were presented in pseudo-random order. At the start of the experiment, the interface (Figure 2) was shown to the observers. The observers were instructed that their main task was to move the sliders to adjust the material of the bird in the top right window (probe) until it appeared to be made of the same material as the bird in the top left window (stimulus), and that they could take as much time as they needed. Once observers finished a matching trial, they pressed the “OK” button, after which only the stimulus and probe images were presented on the screen. The observers were asked to indicate to what extent they were satisfied with thematching result.After they pressed the "Next" button, the next matching trial started. Three trials were performed as practice trials before the experiment formally started. In the practice trials, participants were told that they could move the slider bars by dragging the mouse or pressing the left and right arrow keys on the keyboard. Moving the slider bars by dragging the mouse resulted in bigger steps, while pressing the arrow keys resulted in smaller steps and more gradual changes in the probe.

2.3.3. RESULTS

OVERVIEW

In order to test the usability of the method and evaluate the general matching results per experiment, we will first fit one single linear equation to the complete set of weights of the stimuli and probe adjustments. Then we will analyze the satisfaction ratings and the durations of the matching processes. After that we will look into the details of the interactions of the four canonical material modes to analyze the perceptual effects in detail.

MATCHING RESULTS

The general results of the matching experiments were evaluated by solving the linear factor matrix A in Equation 2.3

$$[\mathbf{Y}]_{4\times 360} = [\mathbf{A}]_{4\times 4} \cdot [\mathbf{X}]_{4\times 360} + [\mathbf{E}]_{4\times 360}\,, \tag{2.3}$$

where, $[\mathbf{X}] = \begin{bmatrix} w_m{}' \\ w_v{}' \\ w_s{}' \\ w_g{}' \end{bmatrix}$, $[\mathbf{Y}] = \begin{bmatrix} w_m \\ w_v \\ w_s \\ w_g \end{bmatrix}$, and the residual $[\mathbf{E}] = \begin{bmatrix} e_m \\ e_v \\ e_s \\ e_g \end{bmatrix}$.

For each trial, one column in matrix **X** represents the weights of the four scattering modes in the stimulus image, and the corresponding column in matrix **Y** represents the weights of the four scattering modes in the probe image, i.e., the values represented by the positions of the four sliders set by the participant. We consider all eight participants together. Thus, there are 45 trials for 8 participants = 360 columns in matrix **X**, matrix **Y**, and matrix **E** (the residuals). The 4×4 linear factor matrix A was solved using a least-squares fit in MATLAB, and then the matrix E was simply calculated as the difference between **Y** and **A** · **X**. If observers were to move all sliders so that the weights in matrix **Y** would be exactly equal to the corresponding weights in matrix **X** (i.e., the matching would be veridical), then **A** would be a 4×4 identity matrix and **E** would be a zero matrix.

The resulting matrix **A** of Experiment 1 is surprisingly close to an identity matrix (see Table 2.2). To be more specific, the nondiagonal values are 0.18 or lower and close to 0, and the diagonal values are 0.78, 0.89, 0.91, and 1.08 for the matte, velvety, specular, and glittery modes, respectively. In the resulting matrix for Experiment 2 the first three diagonal elements decreased to 0.65, 0.69, and 0.63 for the matte, velvety, and specular

	${w_m}'$	${w_v}'$	${w_s}'$	${w_g}'$
Experiment 1				
w_m	**0.78**	0.14	0.16	0.00
w_v	0.18	**0.89**	0.03	0.00
w_s	0.18	0.04	**0.91**	0.04
w_g	-0.02	0.09	0.02	**1.08**
Experiment 2				
w_m	**0.65**	0.25	0.32	0.02
w_v	0.14	**0.69**	0.10	0.00
w_s	0.30	0.24	**0.63**	0.19
w_g	0.01	0.12	-0.00	**1.09**

Table 2.2: Linear factor matrices **A** for Experiments 1 and 2.

modes, respectively. The diagonal value for the glittery mode is 1.09, which is similar to that of Experiment 1. The nondiagonal values that represent the interactions between the scattering modes are larger for Experiment 2 than for Experiment 1. To be more specific, $\{w_m, {w_v}'\}$—the value between w_m and ${w_v}'$ in matrix **A**—was 0.14 in Experiment 1, which means that occasionally velvety contributions in the stimuli w_v were perceived to match with a matte contribution in the probe w_m. The value increased from 0.14 to 0.25 in Experiment 2, showing that the chance increased of perceiving velvety contributions in the stimuli to match a matte contribution in the probe. Similarly, for the combination $\{w_m,\ {w_s}'\}$ the value increased from 0.16 to 0.32; for $\{w_s,\ {w_m}'\}$ it increased from 0.18 to 0.30; for $\{w_s,\ {w_v}'\}$ it increased from 0.04 to 0.24; and for $\{w_s,\ {w_g}'\}$ it increased from 0.04 to 0.19. Thus, overall, a comparison of the off-diagonal elements between the two experiments shows that the interactions between perceptions of matte, velvety, and specular modes became stronger when stimulus and probe were under different lighting and viewing conditions.

Another measure of general performance is the ratio between the sum of the four diagonal values in matrix **A** and the sum of all values in matrix **A**. This ratio can vary from 0 to 1, with veridical behavior at 1 (identity matrix) and chance level at 0.25 (all values in matrix **A** being equal). For each individual, we solved the linear factor matrix **A** with 45 trials per observer per experiment and calculated the ratios. As shown in Figure 2.5, in Experiment 1 the ratios for the observers were 0.80, 0.85, 0.72, 0.83, 0.80, 0.77, 0.70, and 0.80 ($M = 0.78, SD = 0.05$). In Experiment 2 these ratios were 0.47, 0.70, 0.58, 0.76, 0.51, 0.55, 0.64, and 0.69 ($M = 0.61, SD = 0.10$). Overall, all observers performed far above chance level.

We also analyzed how close the residuals (matrix **E**) were to 0. We first took the absolute values of the 4×360 matrices, and then calculated the mean of all elements in each 4×45 matrix, for each observer. The results were quite similar between observers per experiment. As shown in Figure 2.6, in Experiment 1 the means of the residuals' abso-

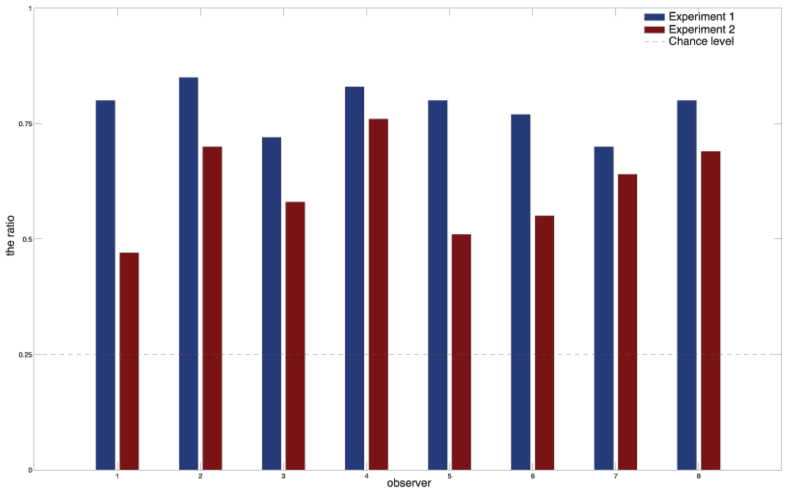

Figure 2.5: The ratio between the sum of the four diagonal values in matrix **A** and the sum of all values in matrix **A**. All eight observers performed far above chance level in Experiment 1 (blue) and Experiment 2 (red).

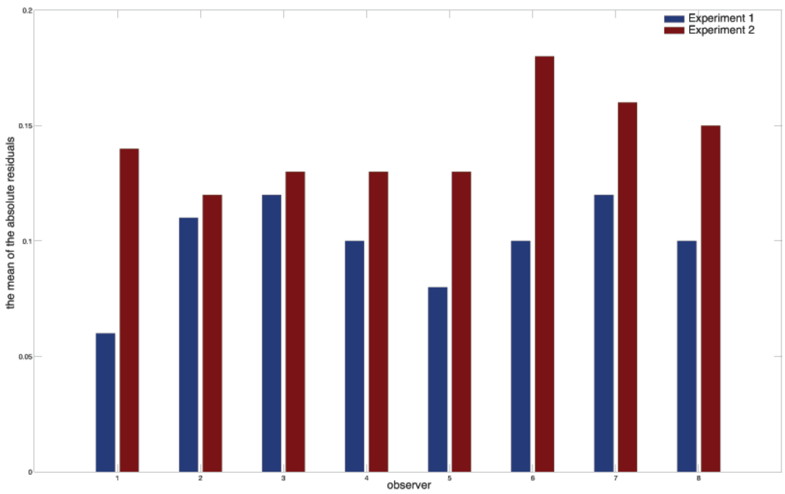

Figure 2.6: The mean of the absolute residuals of each observer in Experiment 1 (blue) and Experiment 2 (red)

lute values for the eight observers were 0.06, 0.11, 0.12, 0.10, 0.08, 0.10, 0.12, and 0.10 ($M = 0.10, SD = 0.02$). In Experiment 2 these values became 0.14, 0.12, 0.13, 0.13, 0.13, 0.18, 0.16, and 0.15 ($M = 0.14, SD = 0.02$). We can conclude that the least-squares fit method properly solved the linear Equation 2.3.

2.3.4. DURATIONS AND SATISFACTION RATINGS

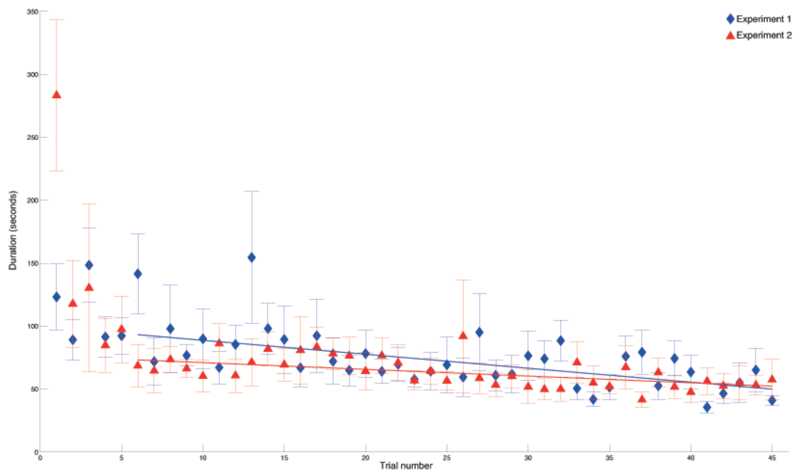

Figure 2.7: Mean trial duration as a function of trial number, averaged across all observers, for Experiment 1 (blue) and Experiment 2 (red). Error bars of each data point represent one standard error of the sample mean.

In Figure 2.7 we plotted the mean duration of the matching trials over all observers as a function of trial number. We fitted the data for both experiments simultaneously by means of multiple linear regression with one dummy variable to directly compare the slopes and establish a possible shift between the two regression lines. The first five trials of each experiment were excluded in the linear regression because we found the duration data in those trials to vary wildly, probably because observers were still exploring the possibilities of the interface. After the first five trials, the pattern of trial durations stabilized (see Figure 2.7). For Experiment 1, the slope of the regression line was found to deviate significantly from 0 $\left(-1.11 \pm 0.21, p < 0.001\right)$. The difference between the two slopes was also significant (0.58 ± 0.3, $p = 0.05$) resulting in a slope of -0.53 for Experiment 2. The offset for Experiment 1 (99.85 ± 6.0, $p < 0.001$) was higher than that for Experiment 2 (76.35; difference equals -23.49 ± 8.46, $p < 0.001$). These results imply that the duration for Experiment 1 started at a higher level than for Experiment 2, and afterward the durations of both experiments systematically decreased with trial number, converging to the same level at the final trials. In conclusion, the main effect is a gradual but small decrease of trial duration as a function of trial number. On average, the duration was slightly above 1 min per matching trial.

The satisfaction ratings were defined to range from 0 (not satisfied with the matching) to 1 (satisfied with the matching). Subsequently, we took the average over all observers per trial. Excluding the first five trials, data were again fitted by multiple linear

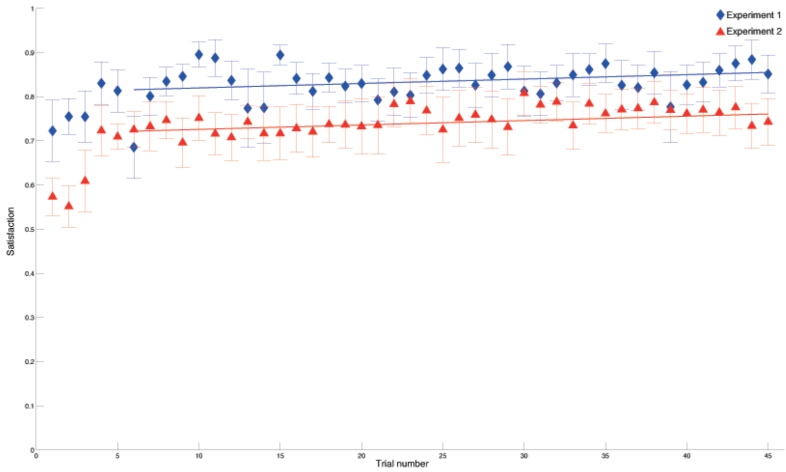

Figure 2.8: Average satisfaction ratings over all observers as a function of trial number, for Experiment 1 (blue) and Experiment 2 (red). Error bars of each data point represent one standard error of the sample mean.

regression with one dummy variable (Figure 2.8). The only significant effects were the offset for Experiment 1 (0.81 ± 0.02, $p < 0.001$) and the difference between the two offsets (-0.09 ± 0.03, $p < 0.001$). Both slopes (0.001 ± 0.01, $p = 0.17$) did not significantly deviate from 0. We can conclude that the participants generally found the matching task feasible, as the average satisfaction is quite high, but that changing the illumination and viewpoint conditions significantly decreased the satisfaction ratings.

2.3.5. SUM OF WEIGHTS

Here we analyze the sum of the four weights in the probe—i.e., the sum of the four slider values—per trial. Our interface and mixing algorithm is based on additive mixing. The sums theoretically can vary from 0 to 4.8 (each slider ranges from 0 to 1.2). However, if the observers were to adjust the image as a partitive mixture constraining the overall brightness of the probe, the sum would be 1 (Griffin, 1999). Similar to what we did when analyzing the durations and the satisfaction ratings, we considered only the last 40 trials per experiment for all eight observers, so in total there were 320 values per experiment. We found that the averages of these sums were 1.13 ± 0.06 and 1.18 ± 0.09 in Experiments 1 and 2, respectively. Because of the randomly generated offsets xm, xv, xs, xg in the stimuli, the sums of the weights in the stimuli were very close to 1 but not exactly equal to 1. The averages of the sums in the stimuli were actually 1.00 ± 0.01 in Experiment 1 and 0.99 ± 0.00 in Experiment 2. We calculated the differences between the sums in the probe and the sums in the stimuli and found that these differences significantly deviated from 0 (one-sample t test, $p < 0.001$ for both experiments), with the sums of the weights in the probe being larger than those of the stimuli in both experiments. We also found a significant difference between the two experiments (paired twosample t test, $p < 0.001$), with the average sum of the weights in the probe of Experiment 2 being larger than that of Experiment 1.

2.3.6. Interactions between scattering modes

In Figure 2.9 we visualized the interactions between each combination of two scattering modes by means of ellipses representing 1 *SD* of bivariate normal distributions fitted to the 24 data points (8 observers × 3 repetitions) for each stimulus. Every data point represents the settings of two of the four sliders in the probe in one trial. For clarity of presentation the data points themselves were rendered invisible in the plots. Each subplot contains 6 ellipses, which are the results of three different weight combinations in the stimuli in the two experiments. The crosses depict the corresponding stimulus weight combinations. This provides a means to visualize the extent to which participants would trade off - or confuse - the weights of different reflectance modes.

To give an example, in the top left subplot the blue ellipses depict the variations of the weights of the matte and velvety modes in the probe for matches to stimulus number 5 (half matte and half velvety in the stimulus, as in Table 2.1 and Figure 2.4). The solid plot represents the result in Experiment 2, which is centered close to the veridical value (the blue cross). The dashed plot represents the result in Experiment 1, which is slightly shifted upward—i.e., in these trials the matte slider was set around its veridical value, while the velvety slider setting was set larger than the weight of the velvety mode in the stimulus. This indicates that in our office lighting, the half-matte, half-velvety mixture was perceived as a match to mixtures of half-matte and more-than-half-velvety components.

Another way of interpreting Figure 2.9 is to look at how the ellipses are oriented and shifted from their veridical centers. To be more specific, in both Experiment 1 (dashed lines) and Experiment 2 (solid lines), the matte and specular contributions strongly interacted with each other, as seen by the ellipses oriented and shifted diagonally in the middle left subplot. For the velvety and matte (top left) and velvety and specular contributions (middle), we also find diagonal shifts for Experiment 2, while for Experiment 1 there are primarily horizontal or vertical shifts. The glittery contributions were all set around their veridical values in both experiments, and the ellipses in the three subplots at the bottom primarily shifted horizontally from their veridical centers. To conclude, in Experiment 1 we found interactions primarily between the matte mode and the specular mode. In Experiment 2 the matte, velvety, and specular modes interacted strongly with each other. The glittery mode remained quite independent in both experiments.

2.4. Validation study (optical mixing with rendered images)

To cross-validate the method, we conducted an additional experiment with MatMix 1.0 in which we used computer-rendered images as the basis images for the mixtures of the

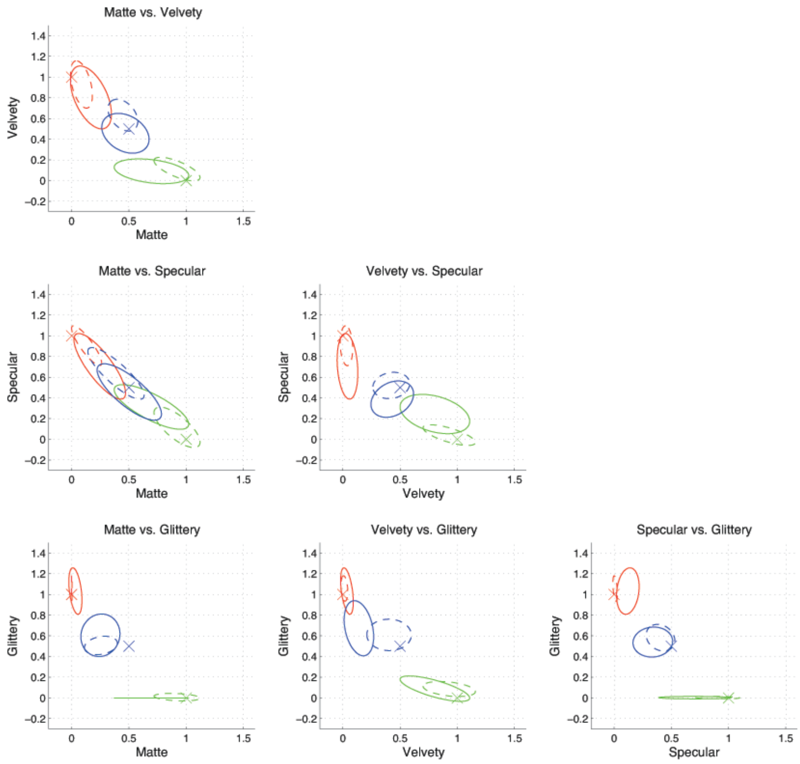

Figure 2.9: A visualization of the interactions between each combination of two scattering modes. Different colors correspond to different weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent 1 *SD* of bivariate normal distributions fitted to the data. Dashed plots represent the data of Experiment 1 and solid plots represent the data of Experiment 2.

stimuli and the probe. In Figure 2.10 we show the basis images of Experiment 3. To generate these basis images, we built a 3-D model of a birdshaped object in Blender (Figure 2.11). We then applied four different "materials" to the object in MaxwellRender (Figure 2.12). The parameters of the materials in MaxwellRender can be obtained from its online material library. We carefully adjusted the parameters of the four materials to represent matte, velvety, specular, and glittery finishes. High-dynamic-range image-based lighting was used as the illumination environment in rendering. For the basis images of the probe, we used Debevec's "Grace Cathedral" environment map (Debevec, 1998). For the basis images of the stimuli, we used Debevec's "Eucalyptus Grove" environment map.

Experiment 3 was conducted at the University of Giessen, Germany. Five paid, inexperienced participants with normal or corrected-to-normal vision participated in the experiment. Participants read and signed a consent form before the experiment. The

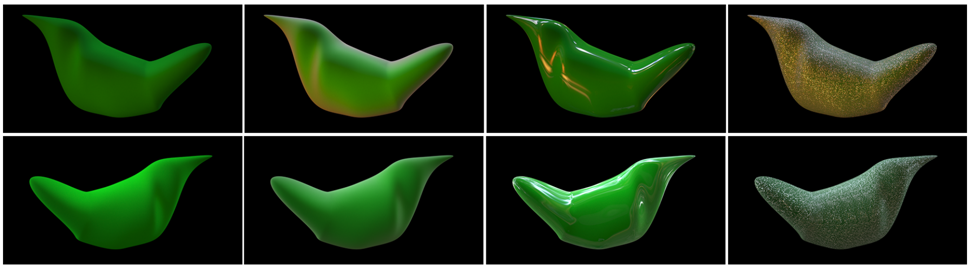

Figure 2.10: The basis images for Experiment 3. The images in the top row were used as basis images for the probe. The images in the bottom row were used as basis images for the stimuli. From left to right, columns represent matte, velvety, specular, and glittery modes, respectively.

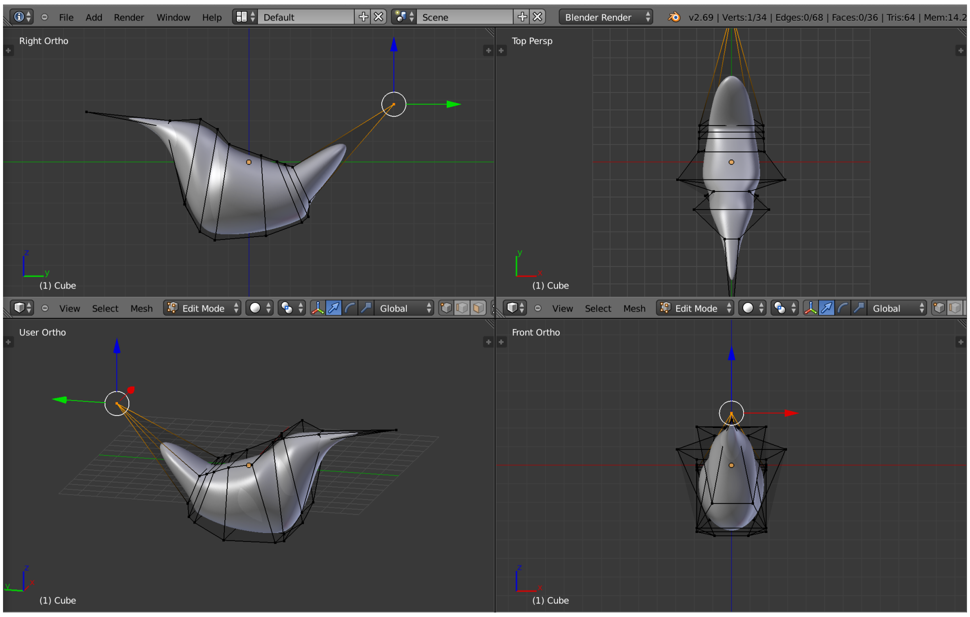

Figure 2.11: A screenshot of Blender during the 3-D modeling process. The model is mirror symmetric. Note that the model is a simplified version of the shape we used in Experiments 1 and 2.

experiment was done in agreement with the Declaration of Helsinki and local ethical guidelines.

We asked the observers to perform the matching task with MatMix 1.0 using the rendered images in Figure 2.10 instead of the photographs in Figure 2.3. Observers spent 50–100 s per trial, which is similar to the durations in Experiments 1 and 2.The satisfaction ratingswere 0.76 on average, which is similar to what we found in Experiment 2. Thus, using renderings as the basis images in MatMix 1.0 does not influence the time costs or the satisfaction ratings of observers in the matching experiment.

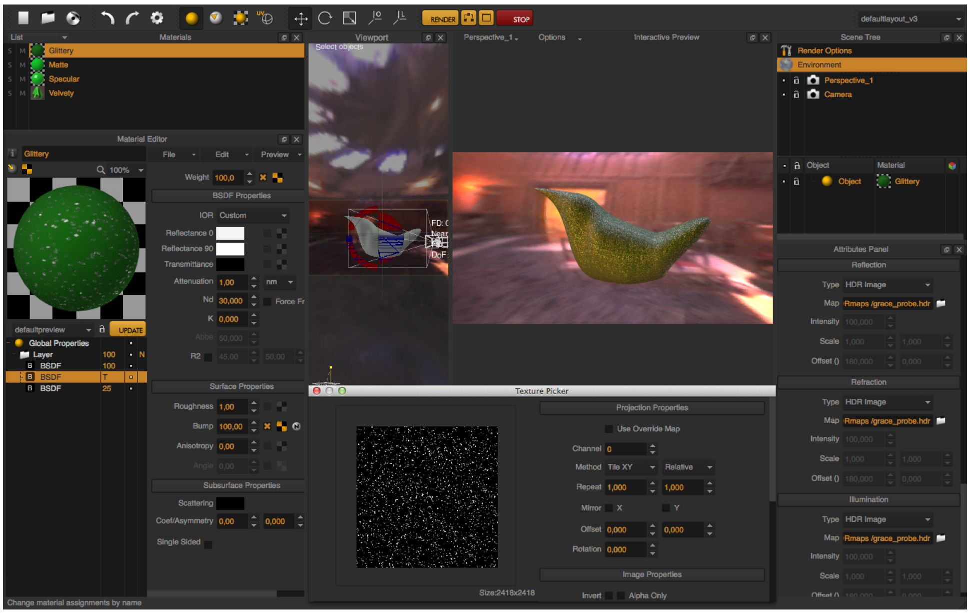

Figure 2.12: A screenshot of MaxwellRender during the rendering process. In this figure, glittery material was assigned to the object, and the Grace Cathedral image was selected as the environment map.

	${w_m}'$	${w_v}'$	${w_s}'$	${w_g}'$
w_m	**0.94**	0.46	0.24	-0.07
w_v	0.35	**0.62**	0.20	0.04
w_s	0.05	0.08	**1.01**	-0.07
w_g	0.00	0.03	0.02	**1.00**

Table 2.3: Linear factor matrices **A** for Experiments 3.

The linear factor matrix **A** for Experiment 3 is shown in Table 2.3. It is very close to an identity matrix, except for the values that represent the perception of the velvety mode. The ratios between the sum of the diagonal values and all values for the five observers were 0.79, 0.71, 0.62, 0.70, and 0.60 ($M = 0.68$, $SD = 0.08$), and thus far above chance level (0.25). The nondiagonal values, specifically 0.46 for $\{w_m, {w_v}'\}$ and 0.35 for $\{w_v, {w_m}'\}$, indicate that the perception of the velvety mode strongly interacted with the perception of the matte mode in Experiment 3. The residuals were all close to 0. The averages of the absolute value of the residuals were 0.16, 0.12, 0.14, 0.16, and 0.16 for the five observers, and thus similar to those of Experiment 2.

To sum up, we find that (a) MatMix 1.0 could be implemented by replacing the basis images of the stimuli and the probe with rendered images and (b) with renderings as the basis images, observers can still perform the task well. However, we found increased interactions between the matte mode and the velvety mode. This probably reflects limitations in the current simulations of such reflectance properties.

2.5. Discussion

A major finding in this study is our demonstration that the interface (Figure 2.2) and the probe enable accurate and robust measurements of material perception. Although observers were asked to manipulate four canonical reflectance modes simultaneously, they could do the task within a reasonable amount of time and felt satisfied about their matching results. Moreover, the general matching results were found to be far above chance level in all experiments. In Experiment 1, the illumination and viewpoint conditions were the same for both stimuli and probe images. Observers might have simply compared the two images on a pixel-topixel basis and searched for the perfect match. In order to avoid this possibility, we implemented Experiments 2 and 3. Results showed that observers were able to match the materials even if the stimulus and probe images did not correspond. The similarities of the results of Experiments 2 and 3 further convinced us that observers were not just doing a best possible image match. Specifically, in Experiment 3 the light fields of stimulus and probe were quite different, but results were similar to those of Experiment 2. This suggests that the observers were indeed matching perceived materials.

Unlike in Griffin's study (1999), the weights of the material modes do not necessarily add up to 1 when the mixing is performed. Compared to Griffin's partitive mixing method, MatMix 1.0 implements additive mixing. As a result, observers had the freedom to manipulate each of the material modes independently, so that changing the weight of one material does not affect the weight of the others. Theoretically, the sum of the four slider settings could range from 0 to 4.8. We calibrated the luminances of all basis images per set in the same manner so that their relative luminances corresponded with the physical relations. So in order to generate a probe image that was neither too bright nor too dark, the sum of the four weights should be around 1. We found that the sums were somewhat higher than 1, which might be an overall effect of the velvety and glittery basis images having a somewhat lower luminance than the matte and specular basis images. An alternative approach could be equalizing the average luminance of all basis images. However, since the lightness of the resulting images is dependent on material, shape, and illumination, it is more logical to calibrate the physical inputs of the different materials by applying the same camera settings. Nevertheless, the finding that the sum was close to 1 suggests that participants can approximately match the overall magnitude of reflectance (or albedo) while simultaneously reporting precise differences in the quality of the reflectance.

We found systematic shifts in material perception between Experiments 1 and 2 (Figure 9), showing how the perception of material was influenced as the object orientation changed and the lighting changed from office light to studio light. These shifts can be related to the changes of the values in matrix A. For example, in Figure 9 the subplot for matte versus specular (middle left) shows that matte–specular interactions increased in Experiment 2 compared to those in Experiment 1, which corresponds to an increase of the nondiagonal values $\{w_m, w_s'\}$ and $\{w_s, w_m'\}$ in Table 2.2. Material–lighting interactions have been addressed by many researchers (Dror, Willsky, & Adelson, 2004; Fleming et al., 2003; Hunter, 1975; Marlow et al., 2012; Motoyoshi & Matoba, 2012; Olkkonen & Brainard, 2010, 2011; Pont & te Pas, 2006; te Pas & Pont, 2005). In a recent study we combined our canonical material modes with three canonical lighting modes, and in this manner we were able to systematically investigate material–lighting interactions for a broader range of materials and lightings (Zhang, de Ridder, & Pont, 2015). We found systematic effects that depended on lighting and material.

However, whether the type of interface we used is the most suitable one remains to be seen. MatMix 1.0 was designed and tested with a limited basis set consisting of four materials. In the future we want to include more material modes to span a wider gamut of the BRDF space, such as backscattering, split-specular scattering, and so on. In order to do this well, we need knowledge about which canonical materials have to be included to cover the perceptual space of natural opaque materials, and about how redundancies between modes could elicit formal ambiguities. But in order to generate such knowledge we would need an extended probe. Moreover, the interface needs to be optimized using knowledge about the perceptual space (e.g., using nonlinear rescaling of the sliders to make the adjustment steps perceptually uniform). These issues are currently still chicken-and-egg problems. We will approach these issues in future research via several iterations in typical design loops (van Boeijen, Daalhuizen, Zijlstra, & van der Schoor, 2013): redesign (extend interface with extra modes), test and evaluate (via formal psychophysical experiments), and adjust the design (on the basis of the experimental outcomes). Other techniques, such as psychophysical scaling methods (Knoblauch & Maloney, 2008; Maloney & Yang, 2003) may also aid with the scaling and selection of the reflectance components.

In the Introduction we made an analogy between optical mixing and painting. In order to analyze our results qualitatively in terms of image characteristics, we did some simple image analysis of the basis images of Experiments 1 (Figure 1), 2 (for results, see Figure 2.13A), and 3 (for results, see Figure 2.13B, C). In general, similar to what was shown in Figure 1, we find smooth shading to be typical for the diffuse scattering component (matte material), bright contours for the asperity scattering component (velvety material), highlights at specular points for the forward scattering component (specular material), and bright speckles all over the surface for meso-facet scattering (glittery material). Such key image characteristics may well form the main triggers for general material perception, in a weighted-mixture manner. So across illuminations and object orientations, the diffuse mode typically yields smooth variations, whereas the asperity mode tends to yield bright contours, the specular mode localized highlights, and the glitter mode high-spatial-frequency variations in the image. We find that this indeed allows the user to adjust different aspects of the proximal image. For highlights, many authors have already shown how their specific characteristics influence perception of glossiness (see e.g., Anderson, 2011; Giesel & Zaidi, 2013; Motoyoshi et al., 2007). Perception of velvetiness, glitter, and sparkle concern undeveloped topics. We argue that such understanding of separate modes, together with our findings about how these characteristics combine and interact, will eventually lead to in-depth understanding of any opaque material.

In Experiment 3, strong interactions between the velvety and matte modes were found. In Figure 2.13C, the rendered basis images of the velvety mode are just very subtly different from those of the matte mode. This might be due to both the rendering functions and the illumination environment (Giesel & Zaidi, 2013).

A

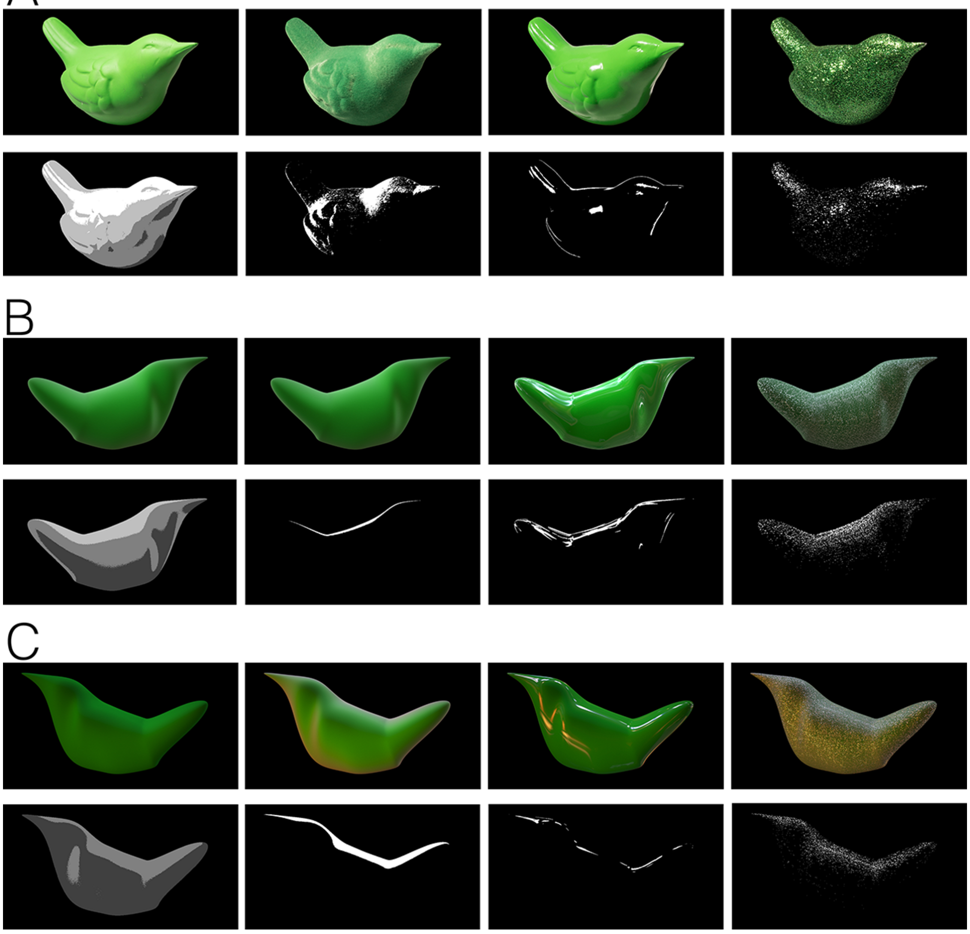

Figure 2.13: The basis images and some related image characteristics. In each subfigure, the first row shows the basis images; the bottom row shows the prototypical image characteristics of each material: the green channel of the basis images after posterization from 255 to six levels for the matte mode, the red-channel thresholding at a somewhat arbitrary 50% level for the velvety, specular, and glittery modes, respectively. Columns from left to right: Representations of matte, velvety, specular, and glittery modes. (A) Photographed set in studio lighting. (B) Rendered image set in Debevec's "Eucalyptus Grove" image-based lighting. (C) Rendered image set in Debevec's "Grace Cathedral" image-based lighting.

A connected question is how to represent each scattering mode properly. For instance, it might be better to mix only the highlights of the glossy bird to represent the specular mode, instead of using the green glossy bird which actually also includes a diffuse mode (see specular mode in Figure 2.13). Similarly, it might be better if only the bright contours were added to represent the velvety mode. This might avoid some interactions between the matte mode and other modes, and thus make them more independent. An analogy is that in a painterly approach, after drawing the contour, the body color is usually painted first in a diffuse manner, after which highlights are added to make the material look glossy or bright contours added to make it look velvety (Wallert, 1999). Additionally, in future studies we want to investigate whether color variation will affect material perception. Currently, green is used disproportionately in material-perception research, for no clear reason (Fleming et al., 2003; Marlow et al., 2012; Marlow & Anderson, 2013). Thus, in a novel version of our MatMix probe we will include color variations accordingly. This will also allow optical mixtures of differently colored modes. For example, specular plastic materials have white highlights, while metals have highlights in the color of their diffuse reflectance, and thus we need color variations in the specular modes to cover both plastics and metals.

Many computer-graphics systems also include sliders to allow the user to alter the material parameters. However, MatMix 1.0 is different in several respects from the standard approach found in computer-graphics interfaces. First, we are able to combine (photographs of) real materials, which can exhibit subtle effects that cannot yet be modeled by computer graphics. Second, even experts find computer-graphics interfaces sometimes overwhelming, especially when there is no realtime feedback on appearance. In typical computergraphics interfaces there are a large number of parameters to adjust, and it is often not intuitive how they are related to the proximal-image result. In contrast, in our approach the basis is limited to a smaller number of canonical visuals, and we have shown that the task is natural and intuitive for inexperienced observers. Third, the purpose of the method is to probe perceptual judgments rather than to design materials from scratch. Thus, the observer will typically have a reference object whose appearance they are trying to match. Finally, the bases are selected to provide perceptually intuitive means for altering proximal-image properties rather than parameters of a physical model, which may not have any distinctive perceptual correlate.

2.6. CONCLUSION

We tested a novel approach to probe material perception in a quantitative and purely visual manner. In the main study (Experiments 1 and 2), we took photos of real materials under different illuminations and implemented them as basis images in MatMix 1.0. In an additional study (Experiment 3), we rendered four materials similar to the ones we used in the first two experiments and used them as basis images to perform another matching experiment with our probe. We found that (a) no matter how difficult the task was or how satisfied participants were, it cost them on average around the same amount of time per matching trial; (b) the participants were matching the probe to the stimuli on the basis of perceived materials instead of simply matching the two images pixel to pixel;

and (c) participants performed well above chance level. In conclusion, it was found that the participants were able to handle the MatMix 1.0 interface well, and our MatMix 1.0 probe was shown to form a robust and intuitive method to test visual material perception. We believe that our painterly optical-mixing approach is promising, because it reflects how weighted mixtures of key ingredients for material representations can trigger our perceptions.

BIBLIOGRAPHY

Anderson, B. L. (2011). Visual perception of materials and surfaces. *Current Biology, 21*(24), R978—R983.

Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. *Journal of Vision, 9*(11):10, 1—17, doi:10.1167/9.11.10.

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision, 35*(1), 33—44.

van Boeijen, A., Daalhuizen, J., Zijlstra, J., & van der Schoor, R. (Eds.). (2013). *Delft design guide: Design methods.* Amsterdam, The Netherlands: BIS Publishers.

Cook, R. L., & Torrance, K. E. (1982). A reflectance model for computer graphics. *ACM Transactions on Graphics, 1*(1), 7—24.

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

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. *Journal of Vision, 4*(9):11, 821—837, doi:10. 1167/4.9.11.

Fleming, R. W. (2012). Human perception: Visual heuristics in the perception of glossiness. *Current Biology, 22*(20), R865—R866.

Fleming, R. W. (2014). Visual perception of materials and their properties. *Vision Research, 94*, 62—75.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision, 3*(5):3, 347—368, doi:10.1167/3.5.3.

Fleming, R. W., Wiebel, C., & Gegenfurtner, K. (2013). Perceptual qualities and material classes. *Journal of Vision, 13*(8):9, 1—20, doi:10.1167/13.8.9.

Giesel, M., & Zaidi, Q. (2013). Frequency-based heuristics for material perception. *Journal of Vision, 13*(14):7, 1—19, doi:10.1167/13.14.7.

van Ginneken, B., Stavridi, M., & Koenderink, J. J. (1998). Diffuse and specular reflectance from rough surfaces. *Applied Optics, 37*(1), 130—139.

Griffin, L. D. (1999). Partitive mixing of images: A tool for investigating pictorial perception. *Journal of the Optical Society of America A, 16*, 2825—2835.

Hapke, B., Nelson, R., & Smythe, W. (1998). The opposition effect of the moon: Coherent backscatter and shadow hiding. *Icarus, 133*(1), 89—97.

Ho, Y. X., Landy, M. S., & Maloney, L. T. (2006). How direction of illumination affects visuallyperceived surface roughness. *Journal of Vision, 6*(5): 8, 634—648, doi:10.1167/6.5.8.

Ho, Y. X., Landy, M. S., & Maloney, L. T. (2008). Conjoint measurement of gloss and surface texture. *Psychological Science, 19*(2), 196—204.

Hunter, R. S. (1975). *The measurement of appearance*. New York: John Wiley.

Kim, J., Marlow, P., & Anderson, B. L. (2011). The perception of gloss depends on highlight congruence with surface shading. *Journal of Vision, 11*(9): 4, 1—19, doi:10.1167/11.9.4.

Klinker, G. J., Shafer, S. A., & Kanade, T. (1987). Using a color reflection model to separate highlights from object color. In *Proceedings First international conference on computer vision*. New York: IEEE.

Knoblauch, K., & Maloney, L. T. (2008). MLDS: Maximum likelihood difference scaling in R. *Journal of Statistical Software, 25*(2), 1—26.

Koenderink, J., & Pont, S. (2003). The secret of velvety skin. *Machine Vision and Applications, 14*(4), 260–268.

Koenderink, J. J., & Pont, S. C. (2008). Material properties for surface rendering. *International Journal for Computational Vision and Biomechanics, 1*(1), 43—53.

Koenderink, J. J., Van Doorn, A. J., Dana, K. J., & Nayar, S. (1999). Bidirectional reflection distribution function of thoroughly pitted surfaces. *International Journal of Computer Vision, 31*(2—3), 129–144.

Maloney, L. T., & Yang, J. N. (2003). Maximum likelihood difference scaling. *Journal of Vision, 3*(8): 5, 573—585, doi:10.1167/3.8.5.

Marlow, P. J., & Anderson, B. L. (2013). Generative constraints on image cues for perceived gloss. *Journal of Vision, 13*(14):2, 1—23, doi:10.1167/13.14.2.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology, 22*(20), 1909—1913.

Matusik, W., Pfister, H., Brand, M., & McMillan, L. (2003). Efficient isotropic BRDF measurement. In *Proceedings of Eurographics/SIGGRAPH workshop rendering* (pp. 241—248). Aire-la-Ville, Switzerland: Eurographics Association.

Motoyoshi, I., & Matoba, H. (2012). Variability in constancy of the perceived surface reflectance across different illumination statistics. *Vision Research, 53*(1), 30—39.

Motoyoshi, I., Nishida, S. Y., Sharan, L., & Adelson, E. H. (2007, May 10). Image statistics and the perception of surface qualities. *Nature, 447*(7141), 206—209.

Nayar, S. K., & Oren, M. (1995, Feb 24). Visual appearance of matte surfaces. *Science, 267*(5201), 1153—1156.

Newell, M. E., & Blinn, J. F. (1977). The progression of realism in computer generated images. In *Proceedings of the 1977 Annual Conference* (pp. 444—448). New York: ACM.

Nishida, S. Y., Sawayama, M., & Shimokawa, T. (2015). Material-dependent shape distortion by local intensity order reversal. *Journal of Vision, 15*(12): 940, doi:10.1167/15.12.940.

Nishida, S. Y., & Shinya, M. (1998). Use of imagebased information in judgments of surface-reflectance properties. *Journal of the Optical Society of America A, 15*(12), 2951—2965.

Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. *Journal of Vision, 10*(9):5, 1—19, doi:10.1167/10.9.5.

Olkkonen, M., & Brainard, D. H. (2011). Joint effects of illumination geometry and object shape in the perception of surface reflectance. *i-Perception, 2*(9), 1014—1034.

Oren, M., & Nayar, S. K. (1995). Generalization of the Lambertian model and implications for machine vision. *International Journal of Computer Vision, 14*(3), 227—251.

te Pas, S. F., & Pont, S. C. (2005). A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In *Proceedings of the 2nd symposium on applied perception in graphics and visualization* (pp. 75—81). New York: ACM.

Pellacini, F., Ferwerda, J. A., & Greenberg, D. P. (2000). Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th annual conference on computer graphics and interactive techniques* (pp. 55—64). New York: ACM Press/Addison-Wesley Publishing Co.

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

Pont, S. C., Koenderink, J. J., Van Doorn, A. J., Wijntjes, M. W. A., & Te Pas, S. F. (2012). Mixing material modes. *Proceedings of SPIE, 8291*, 82910D, doi:10.1117/12.916450.

Pont, S. C., & te Pas, S. F. (2006). Material-illumination ambiguities and the perception of solid objects. *Perception, 35*(10), 1331—1350.

Pont, S., te Pas, S., & Wijntjes, M. (2014). Qualities of optically mixed real materials and photographs— Towards a material probe. *Journal of Vision, 14*(10): 458, doi:10.1167/14.10.458.

Radonjić, A., Cottaris, N. P., & Brainard, D. H. (2015). Color constancy supports cross-illumination color selection. *Journal of Vision, 15*(6):13, 1—19, doi:10. 1167/15.6.13.

Shafer, S. A. (1985). Using color to separate reflection components. *Color Research & Application, 10*(4), 210—218.

Sharan, L., Rosenholtz, R., & Adelson, E. H. (2014). Accuracy and speed of material categorization in real-world images. *Journal of Vision, 14*(9):12, 1—24, doi:10.1167/14.9.12.

Torrance, K. E., & Sparrow, E. M. (1967). Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America A, 57*(9), 1105—1112.

Torrance, K. E., Sparrow, E. M., & Birkebak, R. C. (1966). Polarization, directional distribution, and off-specular peak phenomena in light reflected from roughened surfaces. *Journal of the Optical Society of America A, 56*(7), 916—924.

Vangorp, P., Laurijssen, J., & Dutré, P. (2007). The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics, 26*(3), 1—9.

Wallert, A. (1999). *Still lifes: Technique and style—The examination of paintings from the Rijksmuseum.* Amsterdam: Rijksmuseum.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. *Computer Graphics, 26*(2), 265—272.

Wiebel, C. B., Toscani, M., & Gegenfurtner, K. R. (2015). Statistical correlates of perceived gloss in natural images. *Vision Research, 115*(B), 175—187.

Wijntjes, M. W., & Pont, S. C. (2010). Illusory gloss on Lambertian surfaces. *Journal of Vision, 10*(9):13, 1–12, doi:10.1167/10.9.13.

Wolff, L. B. (1994). Relative brightness of specular and diffuse reflection. *Optical Engineering, 33*(1), 285—293.

Xiao, B., & Brainard, D. H. (2008). Surface gloss and color perception of 3D objects. *Visual Neuroscience, 25*(3), 371—385.

Zhang, F., de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. *Proceedings of SPIE, 9394*, 93940Q, doi:10.1117/12.2085021.

3

THE INFLUENCE OF LIGHTING ON VISUAL PERCEPTION OF MATERIAL QUALITIES

Abstract

We studied whether lighting influences the visual perception of material scattering qualities. To this aim we made an interface or "material probe", called MatMix 1.0, in which we used optical mixing of four canonical material modes. The appearance of a 3D object could be adjusted by interactively adjusting the weights of the four material components in the probe. This probe was used in a matching experiment in which we compared material perception under generic office lighting with that under three canonical lighting conditions. For the canonical materials, we selected matte, velvety, specular and glittery, representing diffuse, asperity, forward, and specular micro facet scattering modes. For the canonical lightings, we selected ambient, focus and brilliance lighting modes. In our matching experiment, observers were asked to change the appearance of the probe so that the material qualities of the probe matched that of the stimuli. From the matching results, we found that our brilliance lighting brought out the glossiness of our stimuli and our focus lighting brought out the velvetiness of our stimuli most similarly to office lighting. We conclude that the influence of lighting on material perception is material-dependent.

Published as: Zhang, F., de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. *Proceedings of SPIE, 9394*, 93940Q, doi:10.1117/12.2085021.

3.1. INTRODUCTION

As part of the EU Marie Curie PRISM project, one of our goals is to understand the relations between material and lighting perception of real objects in natural scenes. Novel aspects of our approach are that i) we developed a probing method to test perceptual material qualities in a purely visual way, and ii) we considered generic materials including not only matte-glossy variations but also other opaque materials, such as velvety and glittery materials. In the current paper we studied how lighting influences material perception. In order to systematically address this question in a generic sense, we approach “lighting” and “material” via a representation in canonical modes. Canonical modes refer to standard or archetypal components that are essential ingredients to mix into generic representations.

In 2012, Pont et al. showed that optical mixtures (Griffin, 1999) of canonical BRDF modes (such as matte, velvety and specular) can be used to systematically vary visual material qualities of real objects in a real setup and on a computer screen (Pont, Koenderink, et al., 2012; Pont, van Doorn et al., 2012). On the basis of this study, we developed a novel material probe by optically mixing four materials that have clear features representing diffuse, asperity, forward and micro facet scattering modes. We selected these four canonical reflectance modes because they together span a large part of the BRDF space. The BRDF is a function that represents the amount of light scattered in every direction as a function of the amount of light arriving from every direction. Generally, if the main directions of the BRDF lobes are different, the main image features will appear in different locations on the object too. For instance, forward scattering will cause highlights at specular attitudes while asperity scattering causes bright contours. Thus, we propose to take the BRDFs as the criteria, and believe that when the BRDFs span the whole BRDF space, the image features will also span the whole space of possible visual appearances (Thompson et al., 2011).

Photographs of the probe were made under generic office lighting. In order to test the probe, we integrated it into an interface and performed matching experiments (Zhang, de Ridder, & Pont, 2014). The stimuli in the interface were also created by optically mixing images of the four materials. However, in order to systematically study the influence of lighting on material perception, the photographs for the stimuli were taken using three canonical modes of lighting. We chose so-called ambient lighting, focus lighting and brilliance lighting since these are commonly applied in lighting design and architecture for building up lighting plans (Ganslandt & Hofmann, 1992). Moreover, these modes were proven to be basic components of physical decompositions of the light field (Pont, 2013; Mury, Pont, & Koenderink, 2009), and it has been shown that human observers are sensitive to these modes (Pont, 2013; Xia, Pont, & Heynderickx, 2013). Ambient lighting is represented by a uniform light distribution illuminating an object in it equally intense from all directions. A combination of snow and fog is an example of a condition in which ambient lighting can be encountered in nature. Focus lighting can be created by a single collimated light source that produces hard shadows, like the sun. Brilliance lighting is represented by the higher order statistics of the light field or, in other words, high angular frequencies of the light. In real scenes, the local light field is normally a combination of the canonical lighting modes in various ratios. However, in this study we will test the influence of the canonical lightings in isolation.

3.2. METHODS

3.2.1. BASIS IMAGES

Figure 3.1: The setup for the photography of the object under different lighting conditions. Left: The object and the camera were fixed onto a frame attached to a tripod and placed on a small trolley. This setup could be rolled into the three canonical lighting setups of which one is shown on the right. The ambient lighting scene was created using a white photo tent – which was almost closed during the photography.

In order to systematically vary the visual perception of material qualities in the stimuli, a series of basis images was created from photos of four bird-shaped objects. First, the objects were finished with matte, velvety, glossy and glittery materials to represent the four canonical reflectance modes, which were diffuse, asperity, forward and micro facet scattering, respectively. Second, we took photos of the objects under different lighting conditions by using the equipment shown in Figure 3.1. The camera and one of four objects were fixed on a frame, which was attached to a tripod standing on a cart. We then rolled the whole setup between three canonical lighting setups, namely ambient lighting (by using a white photo tent as shown in Figure 3.1 – which was almost closed at the front side during the photography but left open for visualization purposes for this figure), focus lighting (by using a spot light from the left side of the object), and brilliance lighting (by hanging a LED-strip containing 150 small sources around the object). In order to align the bird images, we had to make sure each object was placed in exactly the same position and orientation on the frame and in the lighting setups. Their shadows and base outline on their grounding were drawn as references in order to align the imagery, and the cart positions were marked on the floor.

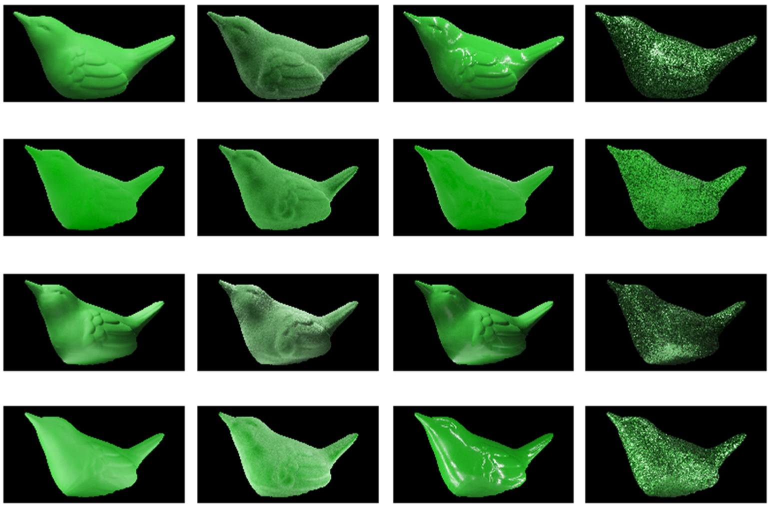

Figure 3.2: Basis images. From left to right each column shows a canonical reflectance mode, namely diffuse, asperity, forward and micro facet scattering, represented by matte, velvety, specular and glittery materials, respectively; From top to bottom the first row shows the probe basis images which were shot under generic office lighting, below which each row shows stimulus basis images for a canonical lighting mode, namely ambient lighting, focus lighting and brilliance lighting (from second to fourth row, respectively).

Next, we edited the photos to find the shared contour, and made the backgrounds outside the contour black for all images. To avoid color interactions the images were processed in MATLAB to set the hue values to 0.33 (green) – which only had a minor influence on the images since the birds were pure green. In addition we photometrically calibrated the luminances, to arrive at a linearly calibrated optical mixing system. The resulting basis images are shown in Figure 3.2.

3.2.2. THE STIMULI AND THE PROBE

The stimuli were linearly superimposed optical mixing results of the basis images taken under any of the three canonical lighting conditions. We made 15 weight combinations of four canonical reflectance modes as shown in Table 3.1. The mixing was done per lighting mode and can be illustrated by equation 1.1:

$$I_{stimulus} = w_m \cdot I_m + w_v \cdot I_v + w_s \cdot I_s + w_g \cdot I_g, \tag{3.1}$$

where subscripts {m, v, s, g} denote matte, velvety, specular and glittery, representing the four canonical reflectance modes (diffuse, asperity, forward and micro facet scattering); $\{w_m, w_v, w_s, w_g\}$ are the weights of the reflectance modes (Table 3.1); $\{I_m, I_v, I_s, I_g\}$ are the processed basis images of the different materials (Row 2 to row 4 in Figure 3.2). As a result, 15 mixed images for each lighting mode were created. In combination with the three lighting conditions this makes 45 stimuli in total, as shown in Figure 3.3.

The probe was a linearly optically mixed result of basis images taken under generic office lighting. Similarly to equation 1.1), the result can be described by equation 3.2

$$I_{probe} = w_m{}' \cdot I_m + w_v{}' \cdot I_v + w_s{}' \cdot I_s + w_g{}' \cdot I_g, \tag{3.2}$$

where$\{ w_m{}', w_v{}', w_s{}', w_g{}' \}$ are the weight values corresponding to the positions of the slider bars in the corresponding sliders (see Figure 3.4);$\{I_m{}', I_v{}', I_s{}', I_g{}'\}$are the processed basis images under office lighting (first row in Figure 3.2).

3.2.3. PROCEDURE

An interface (Figure 3.4) containing two images and four sliders was shown to the observers. The task for the observers was to move the sliders to adjust the appearance of the bird in the top-right window (the probe) until it appeared to be made of the same material as the bird in the top-left window (the stimulus). In order to avoid the use of

Stimulus	w_m	w_v	w_s	w_g
1	1	0	0	0
2	0	1	0	0
3	0	0	1	0
4	0	0	0	1
5	0.5	0.5	0	0
6	0.5	0	0.5	0
7	0.5	0	0	0.5
8	0	0.5	0.5	0
9	0	0.5	0	0.5
10	0	0	0.5	0.5
11	0.33	0.33	0.33	0
12	0.33	0.33	0	0.33
13	0.33	0	0.33	0.33
14	0	0.33	0.33	0.33
15	0.25	0.25	0.25	0.25

Table 3.1: Overview of the weight combinations of the four material modes that were used to generate the stimulus images. There were 15 stimuli in total.

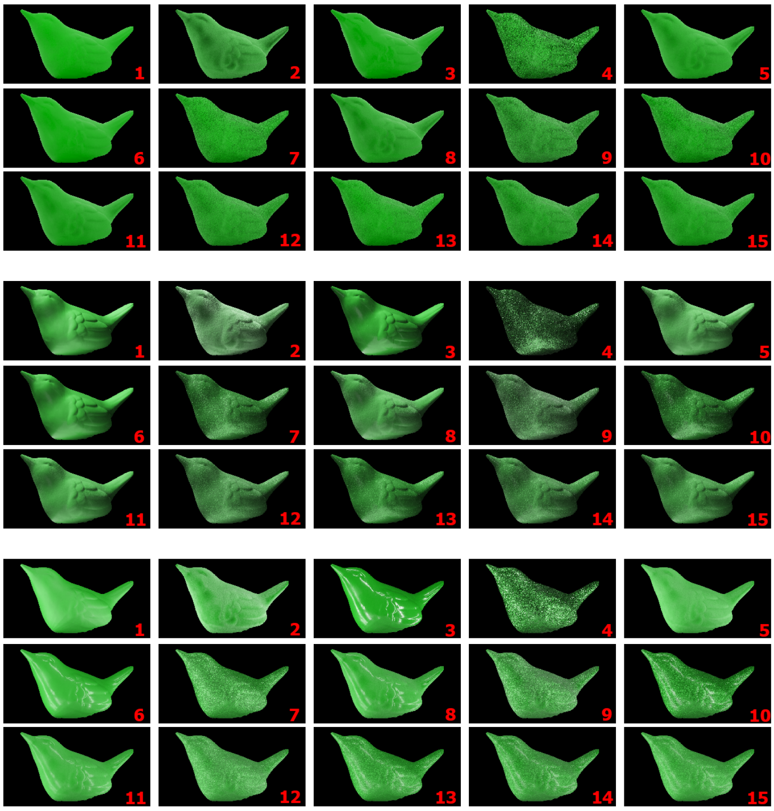

Figure 3.3: The stimuli. Rows 1, 2, 3 show the stimuli for ambient lighting; rows 4, 5, 6 show the stimuli for focus lighting; rows 7, 8, 9 show the stimuli for brilliance lighting. The number in each image corresponds to the weight combination number in Table 1.

terms and provide purely visual information, we put a cropped image in front of each slider, representing one of the material components. The 45 stimuli were presented in a randomized order. In each trial, the positions of the sliders were randomly initialized as well. Once observers finished a trial, they pressed a button, after which only the stimulus image and the probe were presented on the screen, and they were asked to indicate on a scale how satisfied they were about the matching results. Five practice trials were performed before the experiment formally started. The interface was developed using Graphic User Interfaces (GUIs) in MATLAB R2014a, and presented to observers on a linearly calibrated Apple Inc. 15-inch Retina display.

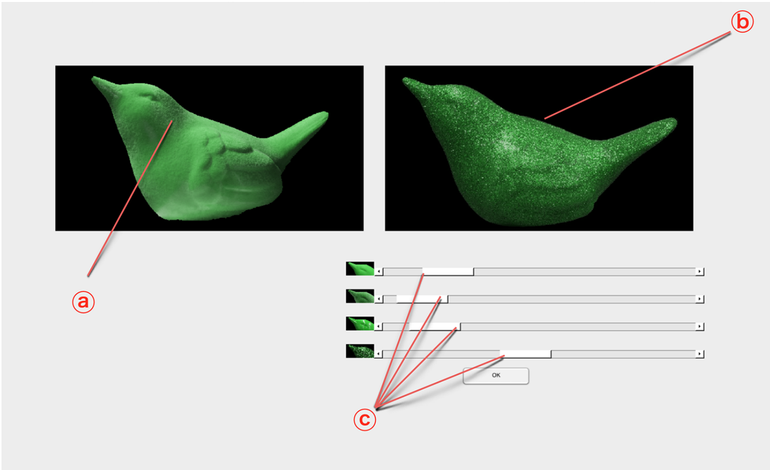

Figure 3.4: The interface. a): A stimulus image. b): The material probe. c): Four sliders. The position of each slider bar represents a weight corresponding to a reflectance mode, ranging from 0 to 1.2. The task of the observers was to match the materials of the probe and stimulus. Here, the two materials obviously do not match.

3.2.4. OBSERVERS

In total, 8 unpaid observers participated in the experiment (5 women and 3 men). They all have normal or corrected-to- normal vision. The study was approved by the TUDelft Human Research Ethics Committee. Before the experiment, we explained to observers that their participation was voluntary, and asked them to read and sign the consent form and introduction of the study.

3.3. RESULTS

The overall performance across all 8 observers per lighting mode was evaluated by solving the linear factor matrix A of equation 3.3

$$[\mathbf{Y}]_{4\times120} = [\mathbf{A}]_{4\times4} \cdot [\mathbf{X}]_{4\times120} + [\mathbf{E}]_{4\times120}\,, \tag{3.3}$$

where, $[\mathbf{X}] = \begin{bmatrix} w_m \\ w_v \\ w_s \\ w_g \end{bmatrix}$, $[\mathbf{Y}] = \begin{bmatrix} w_m{}' \\ w_v{}' \\ w_s{}' \\ w_g{}' \end{bmatrix}$, and the residual $[\mathbf{E}] = \begin{bmatrix} e_m \\ e_v \\ e_s \\ e_g \end{bmatrix}$.

Each column in Matrix X contains weights of 4 reflectance modes that define a stim-

	w_m	w_v	w_s	w_g
Ambient lighting				
${w_m}'$	**0.86**	0.49	0.76	0.07
${w_v}'$	0.21	**0.50**	0.13	0.33
${w_s}'$	0.13	0.01	**0.26**	0.18
${w_g}'$	0.00	-0.01	0.02	**0.93**
Focus lighting				
${w_m}'$	**0.81**	0.08	0.43	-0.08
${w_v}'$	0.20	**1.19**	0.08	0.21
${w_s}'$	0.09	-0.03	**0.51**	0.03
${w_g}'$	-0.01	0.07	0.01	**0.96**
Brilliance lighting				
${w_m}'$	**0.88**	0.42	0.10	-0.01
${w_v}'$	0.15	**0.84**	-0.07	0.20
${w_s}'$	0.24	0.13	**1.17**	0.00
${w_g}'$	0.02	0.09	0.01	**1.17**

Table 3.2: The linear factor matrices of the model (the matrix A in Eq. 3.3).

ulus image in one trial. The 4 values in Matrix Y in the corresponding column record the probing results. For each lighting mode, there are 15 weight combinations, which makes 15 trials * 8 observers = 120 datapoints in total per lighting mode. We solved the matrix A (a 4 by 4 matrix) using a least squares method (in MATLAB). If all observers would have matched the material mixtures perfectly, i.e. moved the sliders to exactly those positions that would have made matrix Y to be equal to matrix X, we would have gotten a 4 by 4 identity matrix for matrix A and a zero matrix for the residual matrix E. In the limiting case in which observers would have randomly adjusted the sliders according to some uniform distribution, all 16 elements in matrix A would be of equal value, with the value being dependent on the boundary conditions (e.g. for the sum of the weights, which is related to the overall brightness). In our analysis we will address the results for matrix A and not for matrix E, since we found that the residuals were negligible.

The linear factor matrices of the model (the matrix A in Equation 3.3) are shown in Table 3.2. These matrices represent the perceptual relations between the weight combinations in the stimuli and the probing results. In this model, the relations between the weights of the same material component in both the stimuli and the probe, i.e.

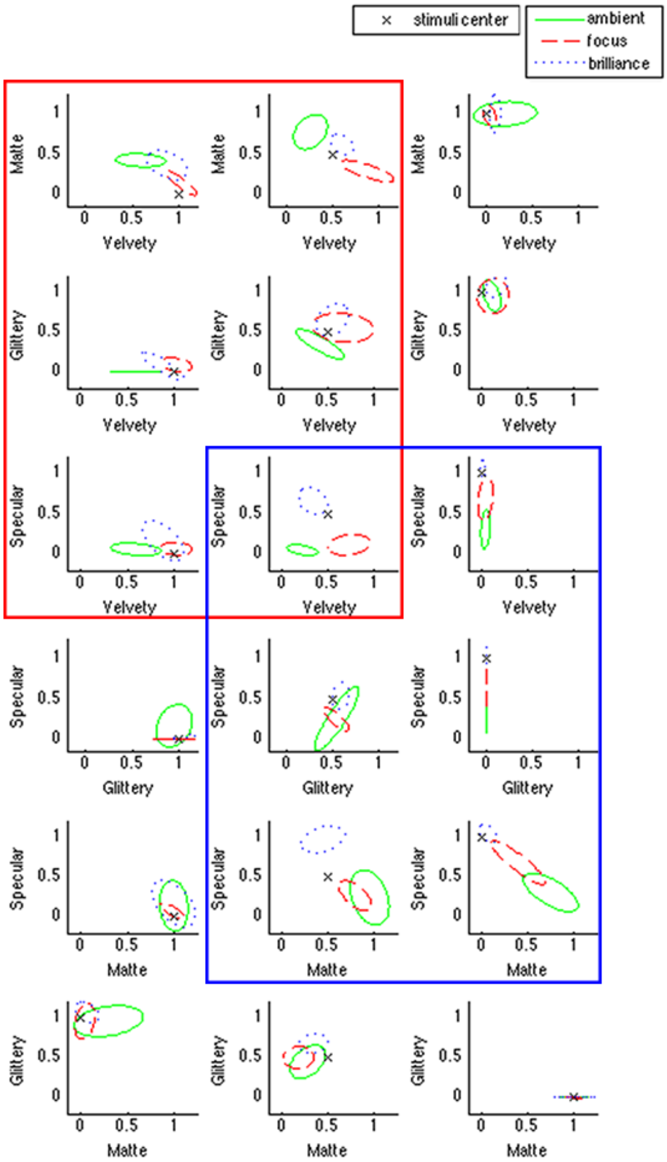

Figure 3.5: Interactions between each two of the reflectance modes. In each subplot, the matching results of 8 observers are fitted into one standard deviation bivariate ellipses. The black cross represents the weight combination of the stimulus in each subplot. In the top-left (red) block of six subplots, perceived velvetiness of the stimuli in comparison with the probe under office lighting decreases as the lighting changes from focus lighting to brilliance lighting to ambient lighting. In the middle-right (blue) block of another six subplots, perceived specularity of the stimuli in comparison with our probe decreases as the lighting changes from brilliance lighting to focus lighting to ambient lighting.

the perceptual relations between each material in one of the canonical lightings and the office lighting, are shown in the diagonal elements. For example, 86% of the matte component in the ambient lighting was matched with the matte component in the office lighting, while only 26% of the specular component in the ambient lighting was matched with the specular component in the office lighting. The perceptual relations between the weights of different material components in the stimuli and the probe are shown in the off-diagonal elements. As a result, the weight of the matte component perceived in the probe ($w_m{}'$) in the office lighting can be predicted by the sum $w_m{}' = 0.86 \cdot w_m + 0.49 \cdot w_v + 0.76 \cdot w_s + 0.07 \cdot w_g + e_m$, when being matched to the stimuli in the ambient lighting.

For the ambient lighting mode we find dramatic influences, as matrix A deviates from an identity matrix to a great extent. Specifically, the diagonal elements of the matrix are 0.86 for matte, 0.50 for velvety, only 0.26 for specular, and 0.93 for glittery. For the focus lighting mode, the element in the diagonal that represents the specular component is also small (0.51), while the elements for the matte, velvety and glittery components remain around 1. For the brilliance lighting mode, the matrix is close to an identity matrix.

Besides the diagonal elements in matrix A, which show the overall matching results of all observers for each canonical lighting mode, we are also interested in the perceptual interactions between reflectance and lighting modes. If we compare the diagonal elements between the three matrices we can see that the diagonal elements for velvetiness in the matrices can be sorted in the following ascending order: ambient lighting (0.50) – brilliance lighting (0.84) – focus lighting (1.19), while for specularity, the order changed to ambient lighting (0.26) – focus lighting (0.51) – brilliance lighting (1.17).

In Figure 3.5 we plotted the bi-variations of the matching results per material combination (in the rows) and lighting mode (represented by the drawn, dashed and dotted lines in each graph). The crosses depict the stimulus weights combinations. The bivariate normal distribution of the sets of individual datapoint was shown as ellipse with one standard deviation. We can easily see that in the stimuli containing a glittery component, very few interactions are found as the ellipses' centers deviate very little from the stimulus positions (the crosses) when glittery is involved. In contradistinction to the glittery component, the matte, velvety and specular components systematically interact with each other.

3.4. CONCLUSION

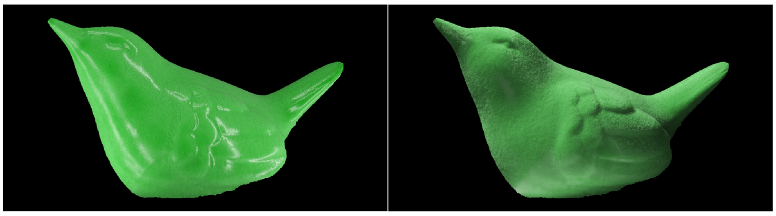

Figure 3.6: Stimulus 8, an optical mixture of velvety and specular materials, looks specular under the brilliance lighting (left image), while it looks velvety under the focus lighting (right image).

Our results showed that the visual perception of materials was systematically (un-)affected by the canonical lighting modes, depending on the material mode. Firstly, the diagonal elements for matte and glittery components in all matrices remained at around approximately 0.9 and 1, respectively, which indicates that the matte and glittery components in our stimuli were perceived similarly to the matte and glittery components in the probe and hardly influenced by the lighting modes. Secondly, the magnitude of the deviations from veridical for the velvety and specular modes differed per lighting mode. The ambient lighting mode caused the largest influence on the perception of velvetiness and specularity. The focus lighting mode, surprisingly, caused almost half of the specular component in the stimuli being confused with the matte component in the probe. The brilliance lighting mode had the least influence on the perception of the four material modes. Finally, the diagonal elements in the matrices representing velvetiness increased as the lighting modes changed in the order ambient lighting - brilliance lighting - focus lighting, while for specularity, it increased according to the order ambient lighting – focus lighting – brilliance lighting. Thus, our brilliance lighting brought out the specularity most similarly to our office lighting, while our focus lighting brought out the velvetiness most similarly to the office lighting. Figure 3.6 illustrates this effect; it concerns stimulus 8 (weights of velvety and specular contributions both 0.5) that looks rather velvety under focus lighting (left image), while it looks rather specular under brilliance lighting (right image). This result indicates that specific, material-dependent lighting modes are required to bring out the characteristics of specific materials. Since our experiment only addressed relative judgments or comparisons we cannot draw conclusions about, for instance, which of our lightings brought out the glossiness or velvetiness best. This question we will address in future experiments.

The phenomenon that glossy materials look like matte materials under ambient lighting has already been reported (Hunter, 1987; Dror, Willsky, &Adelson, 2004; Pont & te Pas, 2006). Here we found the same effect and presented it in a quantitative way, i.e. 0.76 of the specular component in the stimuli in our ambient lighting scene contributed to the perception of the matte component in the probe in the office lighting. In addition, we found that 0.49 of the velvety component in the ambient lighting contributed to the perception of the matte one in our office lighting and 0.33 of the glittery component in the ambient lighting contributed to the perception of the velvety one in our office lighting.

In our focus lighting scene, we also found a small decrease of the glossiness contribution. As shown in Table 2, 0.43 of the specular mode in the stimuli was confused with the matte component in the probe. When considering all elements in the first row for the focus lighting mode, we can find that the total amount of matte component in the matching results consisted of one third of the specular component and two thirds of the matte component in the stimuli. We believe this was due to the relative position of the light source and the object in the focus lighting scene, resulting in a less glossy appearance, compared to the probe which was taken in office lighting. In Figure 3.3 it can be seen that the number of high contrast highlights was much smaller in focus lighting and more comparable to the office lighting and brilliance lighting cases. Since highlight coverage, sharpness and contrast are the main cues for perceived glossiness in non-disparity conditions (Marlow Kim, & Anderson, 2012), this decrease of sharp, high contrast highlights probably explains our finding.

Comparing to the office lighting environment in the probe, our brilliance lighting mode is the most similar one among the three canonical lighting modes we implemented. Due to the limitations of the laboratory conditions, there were also a little focus lighting mode and ambient lighting mode present in the brilliance lighting mode. These altogether probably have led to the fact that the brilliance lighting caused the smallest effects.

In conclusion, a matching experiment was conducted in a quantitative and purely visual manner by using the material probe we previously developed. Three canonical lighting modes (ambient, focus and brilliance) and four canonical reflectance modes (matte, velvety, specular and glittery) were included. The ambient lighting had the largest impact in comparison to office lighting, especially on the perception of velvetiness and specularity. The focus lighting had a strong influence on the perception of specularity. The brilliance lighting had the least strong influence on the perceived reflectance modes. These findings suggest that due to complex material-lighting interactions, perceived material qualities will depend on both lighting and material. Thus, this means that in lighting design (for architecture, computer rendering, retail design, webshop photography, etc.) one needs to be aware of such interactions and explore per material which lighting will bring out the desired material qualities maximally or eliminate undesired material qualities.

BIBLIOGRAPHY

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. *Journal of Vision, 4*(9):11, 821—837, doi:10. 1167/4.9.11.

Ganslandt, R., & Hoffmann, H. (1992). *Handbook of lighting design.* ERCO GmbH.

Griffin, L. D. (1999). Partitive mixing of images: A tool for investigating pictorial perception. *Journal of the Optical Society of America A, 16*, 2825—2835.

Hunter, R. S. (1975). *The measurement of appearance.* New York: John Wiley.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology, 22*(20), 1909—1913.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2009). Structure of light fields in natural scenes. *Applied Optics, 48*(28), 5386-5395.

Pont, S. C. (2013). Spatial and Form-Giving Qualities of Light. *Handbook of experimental phenomenology: Visual perception of shape, space and appearance*, 205-222.

Pont, S. C., van Doorn, A. J., te Pas, S. F., Wijntjes, M. W. A., Koenderink, J. J., "Perceptual qualities of optically mixed materials", In *Predicting Perceptions: Proceedings of the 3rd International Conference on Appearance*, 9-12 (2012).

Pont, S. C., Koenderink, J. J., van Doorn, A. J., Wijntjes, M. W., & te Pas, S. F. (2012,). Mixing material modes. In *Human Vision and Electronic Imaging XVII* (Vol. 8291, p. 82910D). International Society for Optics and Photonics.

Pont, S. C., & te Pas, S. F. (2006). Material–illumination ambiguities and the perception of solid objects. *Perception, 35*(10), 1331—1350.

Thompson, W., Fleming, R., Creem-Regehr, S., & Stefanucci, J. K. (2016). *Visual perception from a computer graphics perspective.* AK Peters/CRC Press.

Xia, L., Pont, S. C., & Heynderickx, I. (2013). Probing light in real scenes using optical mixtures. In *Proceedings of the ACM Symposium on Applied Perception* (pp. 137-137).

ACM.

Zhang, F., de Ridder, H., Pont, S. C. (2014). Towards an interactive probe for material studies. In *Proceedings Experiencing light 2014: international conference on the effects of light on wellbeing*, Eindhoven, The Netherlands, 10-11 November 2014(p. 119).

4

Asymmetric perceptual confounds between canonical lightings and materials

Abstract

To better understand the interactions between material perception and light perception, we further developed our material probe MatMix 1.0 into MixIM 1.0, which allows optical mixing of canonical lighting modes. We selected three canonical lighting modes (ambient, focus, and brilliance) and created scenes to represent the three illuminations. Together with four canonical material modes (matte, velvety, specular, glittery), this resulted in 12 basis images (the "bird set"). These images were optically mixed in our probing method. Three experiments were conducted with different groups of observers. In Experiment 1, observers were instructed to manipulate MixIM 1.0 and match optically mixed lighting modes while discounting the materials. In Experiment 2, observers were shown a pair of stimuli and instructed to simultaneously judge whether the materials and lightings were the same or different in a four-category discrimination task. In Experiment 3, observers performed both the matching and discrimination tasks in which only the ambient and focus light were implemented. Overall, the matching and discrimination results were comparable as (a) robust asymmetric perceptual confounds were found and confirmed in both types of tasks, (b) performances were consistent and all above chance levels, and (c) observers had higher sensitivities to our canonical materials than to our canonical lightings. The latter result may be explained in terms of a generic insensitivity for naturally occurring variations in light conditions. Our findings suggest that midlevel image features are more robust across different materials than across different lightings and, thus, more diagnostic for materials than for lightings, causing the asymmetric perceptual confounds.

Published as: Zhang, F., de Ridder, H., & Pont, S. (2018) Asymmetric perceptual confounds between canonical lightings and materials. *Journal of vision,18*(11), 1-19. doi:10.1167/18.11.11

4.1. INTRODUCTION

The appearance of an illuminated object is determined by its surface geometry (shape), its surface reflectance characteristics (material), and the illumination (lighting). With arbitrary combinations of material, shape, and lighting, the outcomes are difficult to predict. In computer graphics, given models for the shape, illumination, and material and enough computational power, an object can be precisely rendered by calculating the amount of illumination received by the hypothetical camera ("forward optics"). One classic approach that explains how the human visual system estimates physical properties is called "running physics in reverse" or "inverse optics" (Marr, 1982; Pizlo, 2001; Poggio & Koch, 1985; Poggio, Torre, & Koch, 1985). For material perception, using such an approach, the visual system would need to discount the lighting and shape while estimating the material. To do so, the visual system also would need to discount the material before it could estimate the lighting or the shape. Thus, this is a "chicken and egg" problem. Instead, we take as a given that shape, material, and lighting perception are perceptually confounded. Separate studies have been done on how humans visually perceive shapes, materials, or lightings, yet little is known about the interactions between shape, material, and lighting perception. Varying one of the three elements could result in systematic changes of appearance and, thus, could trigger systematic changes of light, material, and shape perceptions, and varying two or three of the elements simultaneously could result in similar appearances and, thus, trigger ambiguities (Dror, Adelson, & Willsky 2001; Morgenstern, Murray, & Harris, 2011; Pont & te Pas, 2006; te Pas & Pont, 2005; Zhang, de Ridder, & Pont, 2015). In this study, we focus on the interactions between lighting perception and material perception. In order to simplify the problem, we kept the shape of our stimuli constant, limited the study to opaque materials, and systematically varied materials and lightings.

4.1.1. CANONICAL LIGHTING MODES

Unlike in physics, light in space and the visual perception of its properties have not been intensively studied in psychophysics (Schirillo, 2013). Koenderink, Pont, van Doorn, Kappers, and Todd (2007) introduced a light probe to measure light perception. They placed a gauge object into a scene and asked observers to adjust the appearance of the probe such that it visually fit into the scene. Ever since then, progress has been made in measuring how humans estimate illumination properties, such as the relative intensity, direction, diffuseness, and color (Kartashova, de Ridder, te Pas, Schoemaker, & Pont, 2015; Kartashova, Sekulovski, de Ridder, te Pas, & Pont, 2016; Koenderink et al., 2007; Morgenstern et al., 2011; Toscani, Gegenfurtner, & Doerschner, 2017; Xia, Pont, & Heynderickx, 2013, 2014). Another approach is to use images of shaded objects as stimuli to investigate the perception of illumination properties, such as direction and diffuseness (Morgenstern, Geisler, & Murray, 2014, 2015; Pont & Koenderink, 2007; Xia et al., 2014), position of the light source (Schütt, Baier, & Fleming, 2016), complex 2-D light fields (van Doorn, Koenderink, & Wagemans, 2011), and complex natural 3-D light fields (Kartashova et al., 2016). Numerous studies implemented variation of illumination for measuring shape or material perception (e.g., Doerschner, Boyaci, & Maloney, 2010; Dror, Willsky, & Adelson, 2004; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2006, 2008; M. Kim, Wilcox, & Murray, 2016; Marlow, Kim, & Anderson, 2012; Motoyoshi & Matoba 2012; Olkkonen & Brainard, 2010; Pont & te Pas, 2006; Wijntjes & Pont, 2010; Zhang et al., 2015) and found out that illumination influenced the judgments of shape and materials. Yet whether or not observers could perceive the changes of illumination remained unknown. In addition, the lightings involved in the abovementioned studies were mostly arbitrary complex natural luminance maps.

Mathematically, a light field can be described by five parameters $\{\theta, \phi, x, y, z\}$ that describe the luminance for all directions and throughout the space (note that we neglect color and time for simplification). For a given position (knowing $\{x, y, z\}$), the local light field can be defined by just two parameters $\{\theta, \phi\}$ that define the directions. Thus, the local light field can be defined as a spherical function and reconstructed by the sum of its spherical harmonics (SH): $f(\theta, \phi) = \sum_{l=0}^{\infty} SH_l$, where l is the order of the angular mode (Mury, Pont, & Koenderink, 2007; Xia, Pont, & Heynderickx, 2016). The zeroth-order SH component (SH_0) is known as the "light density," and the first order SH component (SH_1) is known as the "light vector" (Mury et al., 2007). The diffuseness of a local light field can be calculated by subtracting the ratio of the powers of light vector SH_1 and light density SH_0 from one (Xia's diffuseness metric; see Xia, Pont, & Heynderickx, 2017a, 2017b). It ranges from zero, the most directed light, to one, the most diffuse light. In architectural perception-based lighting design, many designers build up their light plans in three canonical modes (Ganslandt, & Hofmann, 1992; Kelly, 1952), namely ambient, focus, and brilliance light. Phenomenologically, these modes correspond to the zeroth-, first-, and higher (than second) order components of the SH decompositions of the local light fields in physics (Mury, 2009). In this study, we implemented three canonical lighting modes by creating scenes representing the three abovementioned illuminations. The second order of the SH component of the physical light field is known as the "squash tensor," which we did not recreate in our laboratory environment. We ignored this component here because, in lighting architecture, it is not "designed" or addressed explicitly, probably because this component mostly comes from inter-reflections in natural scenes (Mury et al., 2007).

4.1.2. CANONICAL MATERIAL MODES

In material-perception studies, we are trying to understand to what extent and how we are able to recognize what things are made of (material categories, such as fabric, paper, plastic, etc.) or to make subjective judgments about the physical characteristics (material qualities, such as soft, smooth, glossy, etc.) or to attribute concepts to certain materials (material meanings, such as aggressive, nostalgic, industrial, etc.). In the material-perception literature, most often, computer graphic renderings are being used as stimuli, especially for materials within the glossy–matte variation. Computer graphics allows users to manipulate a large number of parameters to vary the geometry and surface reflectance of a 3-D object as well as the illumination to create stimuli sets. Using parametric models, it is calculated how incident light scatters from surfaces, resulting in a certain appearance of the rendered objects. It allows systematic control over the changes in the stimuli and, thus, often gives results that can be easily interpreted, but yet it consumes quite an amount of computational power and sometimes generates images that appear unnatural or unrealistic. Because existing models (Blinn, 1977; Cook & Torrence, 1982; Ward, 1992) simulate glossy materials well, perceived glossiness has been studied intensively (Anderson & Kim, 2009; Fleming et al., 2003; Ho et al., 2006; J. Kim, Marlow, & Anderson, 2011; Marlow et al., 2012; Motoyoshi et al., 2007; Nishida & Shinya, 1998; Pellacini, Ferwerda, & Greenberg, 2000; Vangorp, Laurijssen, & Dutré, 2007). There are also some studies addressing how we perceive other (opaque) material qualities, such as velvetiness (Koenderink & Pont, 2003; Nishida, Sawayama, & Shimokawa, 2015). Other approaches include using real and photographed objects for glossiness perception (Hansmann-Roth, Pont, & Mamassian, 2017; van Assen, Wijntjes, & Pont, 2016), material categorization (Fleming, Wiebel, & Gegenfurtner, 2013; Sharan, Rosenholtz, & Adelson, 2009, 2014), or meaning attribution (Karana, Hekkert, & Kandachar, 2009).

We previously developed a material probe, MatMix 1.0, and found that it provided a perceptually intuitive measuring tool (Zhang, de Ridder, Fleming, & Pont, 2016). It was integrated in an interface for matching tasks, which allowed measurements of material perception in a purely visual and quantitative way. The probe implements optical

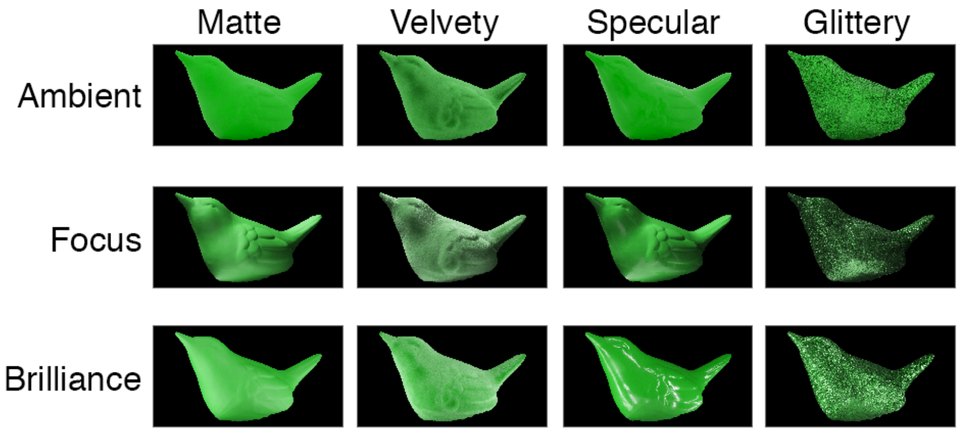

Figure 4.1: The 12 basis images combining three canonical lighting modes and four canonical material modes, i.e., the "bird set". From left to right, each column represents a canonical material mode (matte, velvety, specular, and glittery). From top to bottom, each row represents a canonical lighting mode (ambient, focus, brilliance). In the matching experiments of the previous work, we optically mixed basis images per row such that materials were optically mixed (Zhang et al., 2015). In the current study, we optically mixed the basis images per column, such that lighting was optically mixed in the stimuli and the probe.

mixing of four canonical material modes, namely matte, velvety, specular, and glittery. Each of them represents a very different surface scattering mode, and altogether they span a large part of the bidirectional reflectance distribution function (BRDF) space. In a previous study implementing MatMix 1.0, observers were asked to adjust the material probe and match the material to that of the stimuli, which were optical mixtures of photographs taken under one of three canonical lighting modes (Zhang et al., 2015). Results showed systematic, material-dependent influences of lighting on material perception, which was confirmed in an extra experiment using computer-rendered birds. In the current study, we implemented the same set of photographed basis images, the "bird set" (Figure 4.1), and conducted light-matching experiments by adjusting the probe to allow optical mixing of canonical lighting modes, i.e., by optically mixing the basis images per material instead of per lighting.

To first answer to what extent observers can discount material while matching optically mixed canonical lighting modes, we conducted Experiment 1, in which observers were asked to mix and match the lighting modes of the probe to a mixed illumination in the stimulus. The material modes in the stimulus and the probe could be either the same or different. Observers could only manipulate the illumination of the probe in this task, not its material. In Experiment 2, using a four-category discrimination task and a different group of observers, we tested to what extent observers can simultaneously discriminate materials and lightings. They were shown a pair of basis images selected from the 12 basis images shown in Figure 4.1 and asked to make simultaneous judgments about whether the materials were the same or not and whether the illuminations were the same or not. In Experiment 3, we compared the matching and four-category discrimination tasks for a reduced stimulus set. A third group of observers was asked to first finish a reduced version of the matching experiment and then, after a short break, a reduced version of the four-category discrimination experiment. The reduction concerned removing the brilliance light stimuli and keeping those of the ambient and focus light, i.e., only using the images in the first two rows in Figure 4.1.

4.2. EXPERIMENT 1: CAN PEOPLE DISCOUNT MATERIALS WHILE MATCHING LIGHTING?

4.2.1. METHODS

THE MIXIM 1.0 INTERFACE

In previous work, we found that even inexperienced observers performed well above chance in matching optically mixed materials using our MatMix 1.0 interface (Zhang et al., 2016). In this study, MatMix 1.0 was adjusted to MixIM 1.0 (mix illuminations and materials) to allow light mixing and study whether people can match optically mixed canonical lighting modes for objects that are made of the same material or different ones. In contradistinction to optically mixing materials, mixing canonical lighting modes is actually physically realistic. In the MixIM 1.0 interface (Figure 4.2), three sliders below the right image (probe) represent the three canonical lighting modes, namely ambient, focus, and brilliance light, respectively. How a golf ball appeared under the corresponding light was shown next to each slider to give observers a purely visual reference about what each slider represents. The use of a golf ball as a light probe (Kartashova et al., 2015; Pont & Koenderink, 2007) was chosen because the texture gradients due to the surface structure of the golf ball helps to disambiguate the diffuseness and direction of the light (Xia et al., 2014). In each matching trial, a stimulus image (at left) and the probe image (at right) were presented to observers in corresponding image windows for comparison and matching. The interface was developed using the graphic user interfaces features in MATLAB R2014a (MathWorks, Natick, MA) and presented to the observers on a linearly

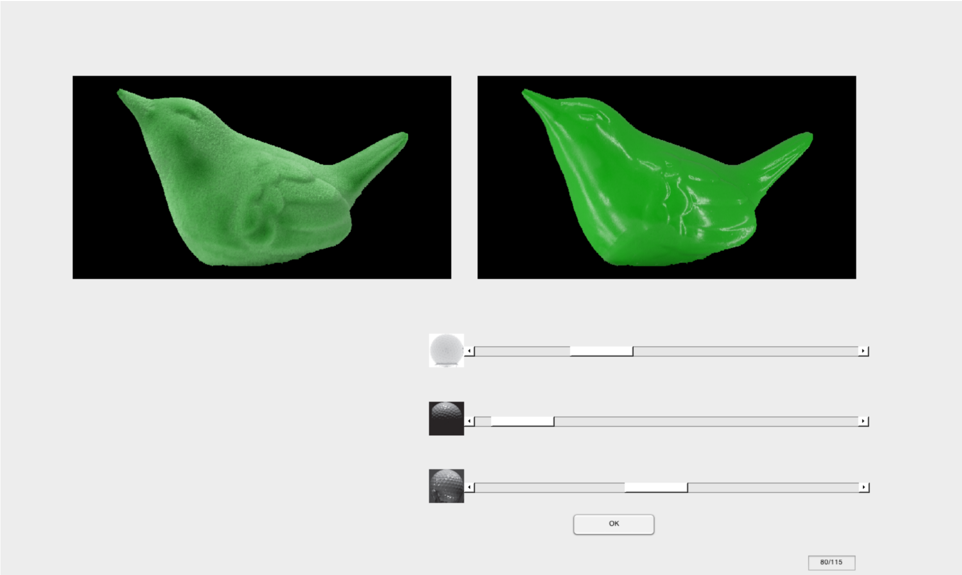

Figure 4.2: The interface of Experiment 1. Left: A stimulus image. Right: The probe image. The material of stimulus and probe could be the same or different (here they are different). The three sliders represent the three canonical lighting modes. The icon next to each slider visualizes the corresponding lighting mode. The position of each slider bar represents a weight value, ranging from zero to 1.2. The task of the observers was to move the sliders to match the illumination of the probe image with that of the stimulus image. In this figure, the illumination of the probe image does not match the illumination of the stimulus image.

calibrated Apple, Inc., 15-in. retina display.

BASIS IMAGES

In our laboratory, we simulated the three canonical lighting modes and took photographs of each canonical material mode under each lighting mode (Zhang et al., 2015) as already shown in Figure 4.1. For the ambient light, we placed both the camera and the object into a white photo tent and then took the photographs for each canonical material mode. For the focus light, we illuminated the object from the left upper side with a halogen spotlight. For the brilliance light, we hung an LED-strip (150 LEDs) surrounding the object. Note that, in order to register the basis images when performing optical mixing, it was important to keep the same relative position between the objects and the camera. This was done by attaching a horizontal, 1-m-long camera slider on a tripod on wheels. The camera was fixed on one side of the camera slider and the object on the other side. The whole setup could then be moved from one scene to another. The photograph was calibrated by adjusting the white balance of the raw images to set the highlights to be white.

No.	$w_{ambient}$	w_{focus}	$w_{brilliance}$
1	1	0	0
2	0	1	0
3	0	0	1
4	0.5	0.5	0
5	0.5	0	0.5
6	0	0.5	0.5
7	0.33	0.33	0.33

Table 4.1: Weight of each canonical lighting mode in the stimuli for Experiment 1.

Then, to avoid color interaction, we set the hue value to 0.33 (green) for all images using MATLAB. The influence of the hue transformation was negligible as the birds were pure green (RAL 6018, except the glittery bird for which the color was matched visually).

STIMULI

For Experiment 1, we designed seven weight combinations of the three lighting modes as shown in Table 4.1. Basis images in each column in Figure 4.1 were linearly superimposed by implementing Equation 4.1 as shown below, per material mode:

$$I_{s.mat.} = w_{ambient} \cdot I_{ambient.mat.} + w_{focus} \cdot I_{focus.mat.} + w_{brilliance} \cdot I_{brilliance.mat.}, \quad (4.1)$$

where $\{w_{ambient}, w_{focus}, w_{brilliance}\}$ are the weights of the lighting modes (Table 4.1) and $\{I_{ambient.mat.}, I_{focus.mat.}, I_{brilliance.mat.}\}$ are the basis images shown in Figure 4.1 with material denoting one of the four canonical material modes: either matte, velvety, specular, or glittery. No linear combinations of materials were used; i.e., the optical mixing of three lighting modes were performed per material. As a result, the linearly mixed stimulus image $\{I_{s.mat.}\}$ presents matte, velvety, specular, or glittery material in a combination of ambient, focus, and brilliance light. In Figure 4.2, the top left image gives an example of stimulus no. 7 for velvety material; i.e., the weights for all basis images of the velvety bird were equal to 0.33.

PROBE

In Experiment 1, observers could manipulate the appearance of the probe image by moving the sliders and, thus, perform the matching accordingly. The probe image was also a linearly superimposed optical mixing result of the basis images per material mode. The mixing process can be illustrated by Equation 4.2:

$$I_{p.mat.} = w_{ambient}' \cdot I_{ambient.mat.}' + w_{focus}' \cdot I_{focus.mat.}' + w_{brilliance}' \cdot I_{brilliance.mat.}', \tag{4.2}$$

where $\{w_{ambient}', w_{focus}', w_{brilliance}'\}$ are the weight values corresponding to the positions of the slider bars in the corresponding sliders (see Figure 4.2: the interface) and $\{I_{ambient.mat.}', I_{focus.mat.}', I_{brilliance.mat.}'\}$ are the basis images shown in Figure 4.1 per material mode, which could be either the same or a different material mode than the material mode used in the stimulus image. No linear combinations of materials were used in the probe either. The linearly mixed probe $\{I_{p.mat.}\}$ allows real-time dynamic and interactive variation of a visual presentation of canonical lighting modes through adjustments of the slider bars.

PROCEDURE

The positions of the slider bars were randomly initialized in each trial. The trials were presented in pseudorandom order. At the start of the experiment, observers were instructed that their task was to move the sliders to adjust the appearance of the bird in the top right window (probe) until it appeared to be in the same illumination as the bird in the top left window (stimulus). They were told that the materials could be the same or different, so the task was not to match the images themselves, but the illumination of the birds. Three trials were performed as practice trials before the first session started. In the practice trials, participants were told that they could move the slider bars by dragging the mouse or pressing the left and right arrow keys on the keyboard. Moving the slider bars by dragging the mouse resulted in bigger steps, and pressing the arrow keys resulted in smaller steps and more gradual changes in the probe. In the actual experiment, four material modes in the probe image were combined with four material modes in the stimulus, resulting in 16 material combinations. Together with seven weight combinations for the stimuli lighting in the optical mixture (Table 4.1) per material combination, there were 112 trials in total for each observer. It took around 60 min to finish the experiment.

OBSERVERS

We recruited four unpaid observers who had participated in at least five psychophysical experiments, and 11 paid inexperienced observers participated in Experiment 1. The four unpaid observers are grouped as “experienced” as they had participated in former experiments working with the experimental interface. All 15 participants had normal or corrected-to-normal vision. Participants read and signed a consent form before the experiments. The experiments were approved by the human research ethics committee of Delft University of Technology and conducted in accordance with the declaration of Helsinki and Dutch law.

4.2.2. ANALYSIS AND RESULTS

LEAST SQUARES FIT

The matching performance using the MixIM 1.0 interface can be evaluated by solving the linear factor matrix **X** of Equation 4.3 using least squares fitting:

$$[\mathbf{P}]_{3\times(112\times N)} = [\mathbf{X}]_{3\times3}\cdot[\mathbf{S}]_{3\times(112\times N)} + [\mathbf{E}]_{3\times(112\times N)}, \tag{4.3}$$

where, $[\mathbf{P}] = \begin{bmatrix} w_{ambient}' \\ w_{focus}' \\ w_{brilliance}' \end{bmatrix}$, $[\mathbf{S}] = \begin{bmatrix} w_{ambient} \\ w_{focus} \\ w_{brilliance} \end{bmatrix}$, and the residual $[\mathbf{E}] = \begin{bmatrix} e_{ambient} \\ e_{focus} \\ e_{brilliance} \end{bmatrix}$.

In Equation 4.3, each row represents a canonical lighting mode, specifically the ambient, focus, and brilliance lighting mode from top to bottom. Per observer, there were 112 trials, and together with the number of participants N, there were in total (112×N) columns in matrix **P**, matrix **s**, and matrix **E**. Each column in matrix **S** represents the weights of the three canonical lighting modes in the stimulus image, and the corresponding column in matrix **P** represents the weights of the three canonical lighting modes in the probe image, i.e., the values represented by the positions of the three sliders set by the observers. The 3×3 linear factor matrix **X** was solved using a least squares fit in MATLAB, and then matrix **E** was the subtraction between **P** and $\mathbf{X}\cdot\mathbf{S}$. If the matching would be veridical, **X** would be a 3×3 identity matrix, and the matrix **E** would be a zero matrix. The ratio r between the sum of the diagonal values in **X** and the sum of **X**, i.e., $r = \sum diag(\mathbf{X})/\sum(\mathbf{X})$, can be used to evaluate the performance, ranging from zero (only possible mathematically) to one (veridical) with 0.33 being the chance level.

OVERALL RESULTS

The overall results of all observers in Experiment 1 is expressed as the linear factor matrix X, solved by least squares fitting, and is shown in Table 4.2 (N=15). In the matrix,

r=0.56	$w_{ambient}'$	w_{focus}'	$w_{brilliance}'$
$w_{ambient}$	0.63	0.22	0.31
w_{focus}	0.29	0.66	0.29
$w_{brilliance}$	0.20	0.22	0.62

Table 4.2: The linear factor matrix **X** in Equation 4.3 solved using the least square method (N = 15).

the diagonal values are 0.63, 0.66, and 0.62 for ambient, focus, and brilliance light, respectively, and the nondiagonal values are all between 0.20 and 0.31, so the matrix is dissimilar from an identity matrix. The ratio r is 0.56, which is far above chance level (r=0.33, see individual performance).

The performance per material combination in stimulus and probe for all observers can be seen in Figure 4.3. The plot shows the ratio r calculated per material combination with the colors of the bars coding the materials of the probe. Each subplot shows results for one material of the stimulus (matte, velvety, specular, and glittery from left to right with labels on the x-axis coded in corresponding colors). When the materials were the same in the stimulus and the probe, the performances were closest to veridical ($r = 1$) in each subplot. When the materials were different in the stimulus and the probe, the performances were still above chance but less close to veridical than when materials were the same. When the velvety material mode was presented, irrespective of whether it was in the probe or in the stimulus, the results were the least veridical. This shows that material differences decreased the performance of matching optically mixed lighting modes. Thus, for our very diverse material and lighting modes, there were strong perceptual interactions between materials and lightings.

INDIVIDUAL PERFORMANCE

The individual matching results (the histogram of the ratios r for all observers) can be seen in Figure 4.4 ($Mean = 0.57, SD = 0.14$). It clearly shows that four out of 15 observers performed just above chance level (0.33), and the other 11 observers performed well above chance; i.e., most of the observers were able to match the optically mixed canonical lighting modes. The four observers who performed just above chance level were all inexperienced observers (colored in blue).

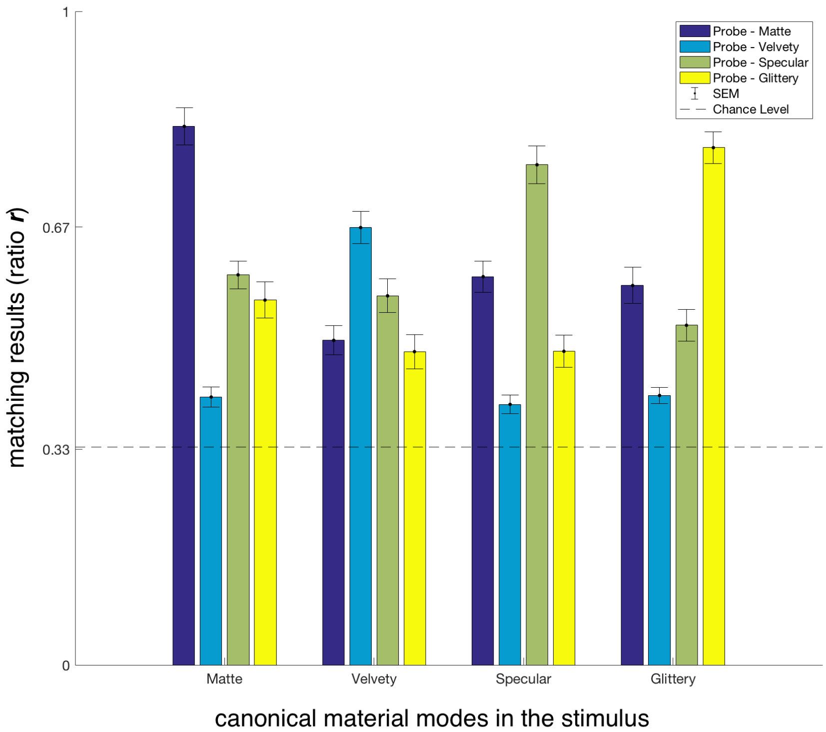

Figure 4.3: Ratio r calculated per material combination of the stimulus and the probe. The four subplots show the results for matte, velvety, specular, or glittery stimuli from left to right, respectively. The material of the probe is color-coded; see legend. The y-axis represents the ratio r. Each ratio was calculated over all data of the 15 observers per material combination. The error bars depict one standard error of the mean.

PERCEPTUAL INTERACTIONS BETWEEN CANONICAL LIGHTING MODES: BIVARIATION PLOT

Another way of interpreting the data from our matching experiment is to visualize the interactions between the basis modes in the mixtures. The interactions between each combination of two lighting modes were visualized by means of ellipses representing one standard deviation values of bivariate normal distributions fitted to the data for all observers for the 16 material combinations (four materials in the stimulus by four materials in the probe). The fitted ellipses are shown per lighting combination in Figure 4.5 and for different groups of observers in Figure 4.6. Every data point represents the settings of two of the three sliders in the probe in one trial. For clarity of presentation, the data points themselves were rendered invisible in the plots. Each subplot contains

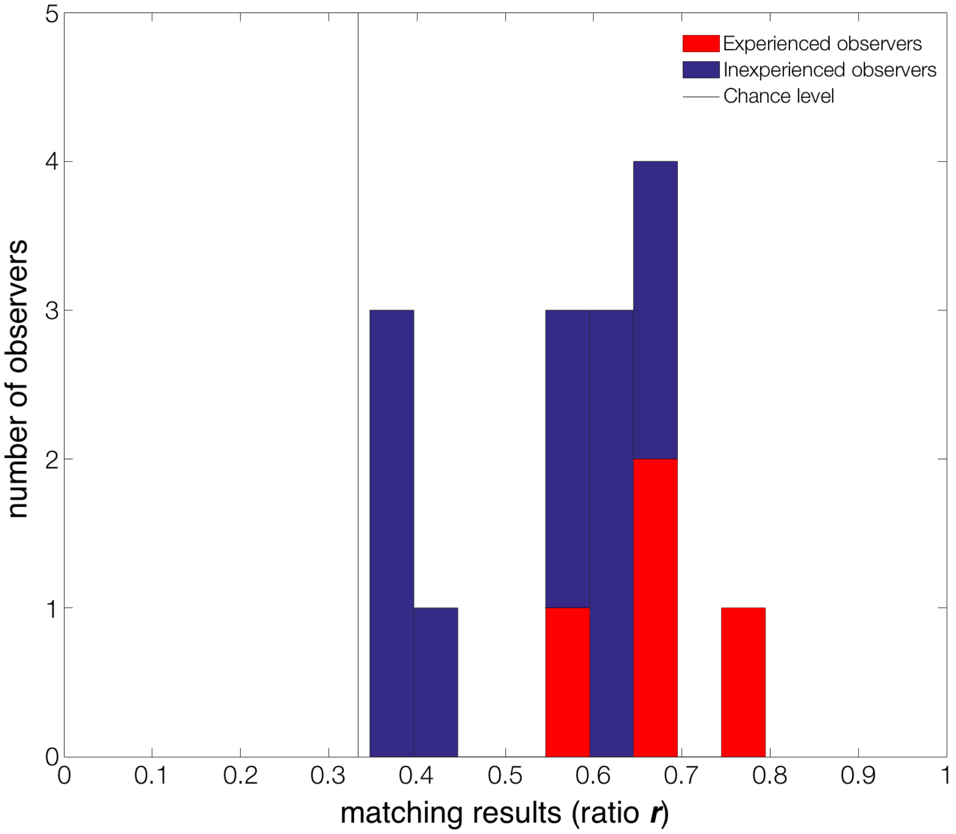

Figure 4.4: Histogram of number of observers for the performance ratio r. The red-colored bars are the results of four observers who are experienced in psychophysical experiments. The blue-colored bars are the results of 11 observers who had no experience in psychophysical experiments at all.

three ellipses, which depict the results for three different weight combinations in the stimuli. The coordinates of the crosses depict the corresponding weight combinations of the stimuli (see Table 4.1). This provides a means to visualize the extent to which participants confuse the lighting modes. In general, if there is less overlap between ellipses, if the ellipses are centered closer to the crosses, and if the ellipses are smaller, then the lighting modes interact less. The general results can be seen in Figure 4.5. In the plots, the red color corresponds to the stimuli in which only ambient light was present, the green color corresponds to the stimuli in which only focus light was present, the blue color corresponds to the stimuli in which only brilliance light was present, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). We find that the ellipses are in the right order but tend to shift toward each other in the center. Blue ellipses shifted away from the blue crosses the most, showing that the responses for mixtures containing the brilliance light were the least veridical.

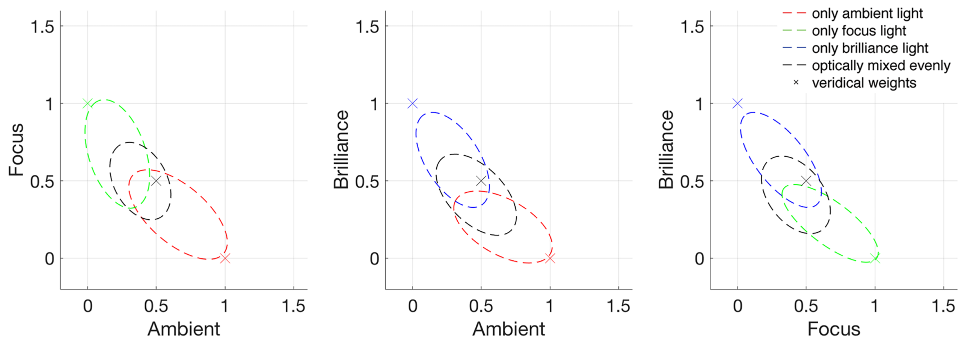

Figure 4.5: Bivariation plots for each combination of two lighting modes for all observers. The three subplots are results for different lighting combinations. Different colors correspond to different lighting-weight combinations in the stimuli, which are depicted by the crosses (the veridical weights). Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

To further analyze the interactions between materials and lightings, we looked into the results per material combination as shown in Supplementary Figure 4.15, Figure 4.16, Figure 4.17, and Figure 4.18 . The rows of Supplementary Figure 4.15 - 4.18, containing three subplots, show the matching results per material combination of the stimuli and the probe under different lightings, corresponding to the results (one of the 16 ratios r) shown in Figure 4.3. For symmetric matching, if the materials in the stimuli and the probe were the same, we found that the crosses (the stimulus centers) fell into the ellipses (one standard deviation of bivariate normal distribution fitting). The only exception happened if velvety was presented in the stimuli and the probe, for which the probing results of the ambient and brilliance light deviated more than one standard deviation. For asymmetric matching, Supplementary Figure 4.15, 4.16, 4.17, and 4.18 shows that when velvety was presented in the probe, the ellipses tended to shift toward the green cross representing focus lighting or to the origin for conditions without focus lighting. This explains why the results were less veridical when the velvety mode was present as shown in Figure 4.3.

N=4

0.64	0.08	0.17
0.24	0.76	0.19
0.13	0.14	0.71

N=7

0.69	0.20	0.30
0.27	0.73	0.25
0.15	0.20	0.67

N=4

0.51	0.38	0.45
0.39	0.44	0.46
0.35	0.33	0.45

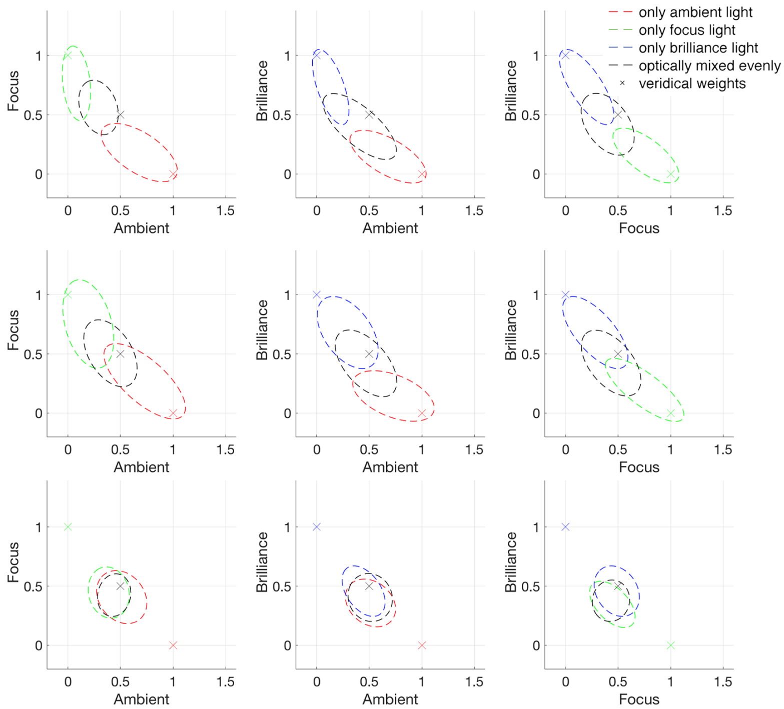

Figure 4.6: Left: Linear factor matrices that were fitted using the least squares method, per group, in the same format as in Table 4.2. Right: Bivariation plots for each combination of two lighting modes (in the columns) for three groups of observers (in the rows). Top: Results of the four experienced observers. Middle: Results of the seven inexperienced observers who performed far above chance. Bottom: Results of the four inexperienced observers who performed just above chance. Different colors correspond to different lighting-weight combinations in the stimuli, which are depicted by the crosses (the veridical weights). Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

To further analyze the individual results, we separated the group of four observers who performed just above chance level from the group of inexperienced observers that performed better, according to both the results from the least squares fitting method (Figure 4.4) and their individual bivariation plots (as shown in Supplementary Figures 4.19, 4.20, and 4.21). In addition, the four experienced observers were separated as one group (colored in red in Figure 4.4). In Figure 4.6, results of the three observer groups can be seen in the rows. In each row, on the left, it shows the 3×3 linear factor matrix $\mathbf{X}$ that was calculated per group (the same format as Table 4.2). On the right, each subplot shows a combination of two lighting modes (in colors). The first row shows the data for the group of the four experienced observers; note that all of them performed well above chance ($r = 0.60, 0.69, 0.69, 0.79$). The second row shows the data for the group of the seven well-performing inexperienced observers ($r = 0.55, 0.55, 0.57, 0.61, 0.62, 0.66, 0.68$). The third row shows the data for the group of the four inexperienced observers that performed just above chance ($r = 0.35, 0.35, 0.39, 0.40$). The ellipses for the experienced observers (the first row) show less overlap than those for the inexperienced observers (the second row) and certainly than those for the just-above-chance performers. The crosses, depicting the veridical settings, were all within the ellipses for the experienced observers, and the blue crosses (brilliance light) were outside the blue ellipses for the well-performing inexperienced observers; i.e., the veridical weights of the brilliance lighting mode differed more from the mean probing results for this group of inexperienced observers. The results of the observers who performed just above chance level according to the least square fitting analysis, as shown in the third row, cluster in the center. Overall, the ellipses tend to shift to the center of the plots. Apparently, the participants always use at least two sliders even when only one slider is required for a perfect match. This is especially obvious with the inexperienced observers, but it is also apparent for the experienced participants.

4.2.3. INTERMEDIATE DISCUSSION

In Experiment 1, we asked observers to match optically mixed lightings in two conditions: symmetric matching (same materials in the stimulus and the probe) and asymmetric matching (different materials in the stimulus and the probe). The goal was to test whether observers could match the mixture of canonical lighting modes while discounting materials. In general, observers were above chance level in the light-matching tasks. Individual differences were found as four out of 15 observers tended to mix all lightings no matter if they were presented in the stimulus, which led to their less-veridical performances. We also found that when velvety was presented in the probe or in the stimulus, the overall performance was significantly less veridical. To conclude, using our optical mixing interface, we found that observers were able to either match lightings while discounting materials (Experiment 1) or match materials while discounting lightings (Zhang et al., 2016). To further investigate the confounds between our canonical material and lighting modes, we designed Experiment 2 to test whether observers could simultaneously discriminate materials and lightings and Experiment 3 to relate the results of the two types of tasks.

4.3. EXPERIMENT 2: CAN PEOPLE SIMULTANEOUSLY DISCRIMINATE MATERIAL AND LIGHTING?

4.3.1. METHODS

This experiment was to test whether observers can discriminate our canonical material and lighting modes simultaneously and to what extent material and lighting perceptions are confounded. The task was similar to a previous study in which observers were asked to judge materials and illuminations separately for a series of spherical objects (te Pas and Pont, 2005). Here, we asked observers to make discrimination judgments for a more systematic set, that is, our canonical material and lighting modes, and observers had to judge materials and lightings simultaneously. In each trial, observers were shown a pair of stimulus images and asked to choose from four response categories—"same materials same lightings," "same materials different lightings," "different materials same lightings," and "different materials different lightings"—based on the appearance of two birds (Figure 4.7). The aim of the experiment was to test whether (and for which modes) observers can judge if differences in appearance are due to material and/or lighting variations for systematically chosen modes that strongly differ optically and together span much of the reflectance and lighting spaces. The interface was developed with the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) in MATLAB and presented to the observers on a linearly calibrated Apple, Inc., 15-in. retina display.

4.3.2. STIMULI

Only the 12 basis images were used as stimuli in Experiment 2 (Figure 4.1); i.e., no optical mixing was performed in Experiment 2.

4.3.3. OBSERVERS

Eight paid inexperienced observers participated in Experiment 2. All participants had normal or corrected-to-normal vision. Participants read and signed the consent form before the experiments. The experiments were approved by the human research ethics committee at Delft University of Technology and conducted in accordance with the declaration of Helsinki and Dutch law.

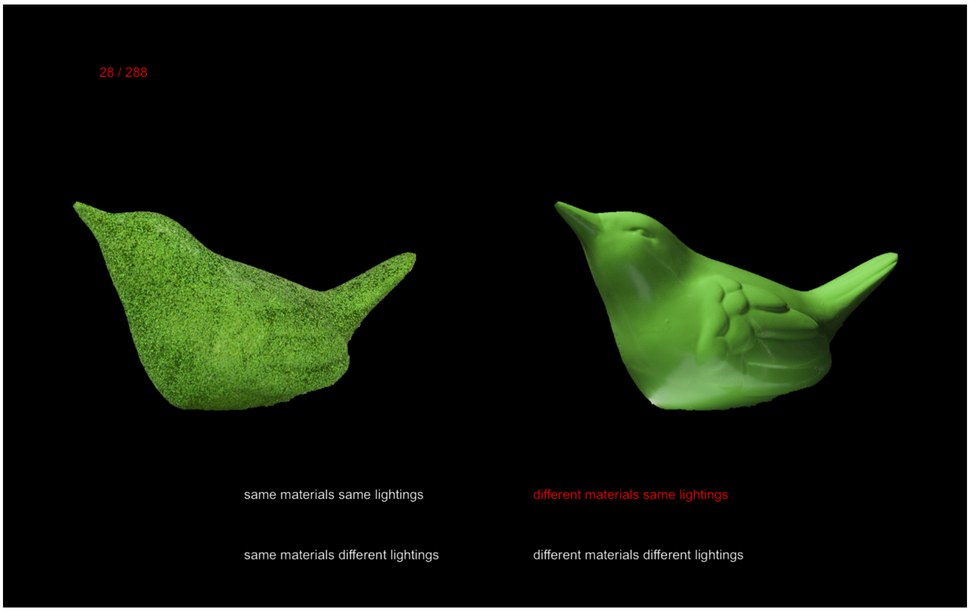

Figure 4.7: The interface of Experiment 2. Left: Glittery material under ambient light. Right: Specular material under focus light. The four response options are listed below the images. The selected option is marked red. The number in the top left corner indicates the progress (number of trials done as a ratio of the total number of trials). Here, the selected option is not correct.

4.3.4. PROCEDURE

Because all observers were inexperienced and did not participate in Experiment 1, they were instructed to browse through all stimulus images in pseudorandom order before the actual experiment started to give them a brief idea of how similar or different the images could be. Each stimulus image was repeated twice and displayed for at least 0.5 s before the observer could click a button to display the next one. They were told that there were four different material types and three lighting types and every image would be one of the four materials in one of the three types of lighting. They were also told that, in the actual experiment, their task would be to compare two of the images and answer whether the materials are the same or different and whether the lightings are the same or different.

With 12 basis images as stimuli, there were 78 possible combinations, 12 of "same materials same lightings," 12 of "same materials different lightings," 18 of "different materials same lightings," and 36 of "different materials different lightings." In order to balance the number of trials for each stimulus category, they were repeated six, six, four, and two times per category, respectively, so that we got 72 trials per stimulus category, i.e., 288 trials per observer. Without time limits for the task, it took around an hour to finish the experiment.

In the actual experiment, a pair of stimuli was displayed and one of the four options was randomly initialized. For the images in each stimuli pair, being left or right was also randomized. Observers were instructed that they could press up, down, left, and right arrow keys on the keyboard to select their answer. The selected one was marked red. Then observers could press the spacebar to finish the current trial and start the next one. The numbers on the top left corner of the interface indicated the progress of the experiment.

4.3.5. RESULTS

OVERALL PERFORMANCE

In Figure 4.8, the fractions of responses per stimulus category are shown. Each square shows the fraction represented as a gray level with the number showing the exact value, calculated by dividing the total counts of the responses by the number of trials per stimulus category (i.e., 72 in this task). Each row represents one stimulus category, and each column represents an answering option. Note that, for each row, the fractions of the four answers add up to one, and the diagonals show the fractions of the correct answers, i.e., the discrimination accuracy. Also note that chance level is 0.25 for this four-category discrimination experiment. As expected, when the materials and lightings were both the same in the stimuli image pair, observers got the highest accuracy (0.97). When the materials were the same and lightings were different, the accuracy somewhat decreased (0.78). But when the materials were different, the accuracy strongly decreased to be just above 0.5 independent of whether the lightings were the same (0.58) or different (0.54). Off-diagonal values are negligible except for two cases (0.27 and 0.33). The responses were found to be significantly associated with the stimulus categories ($\chi^2(9) = 3247.2, p < 0.001$). They also showed that, when materials were different, observers would indeed perceive the materials to be different but then be less accurate about whether the lightings were the same or different. In Supplementary Figure 4.23 , we present the stimulus image pairs that resulted in the least and best performances in Experiment 2 (only for the "different materials different lightings" category). To conclude, both the material and lighting differences caused the accuracy to decrease, but material differences caused the accuracy to decrease more. For different materials, the

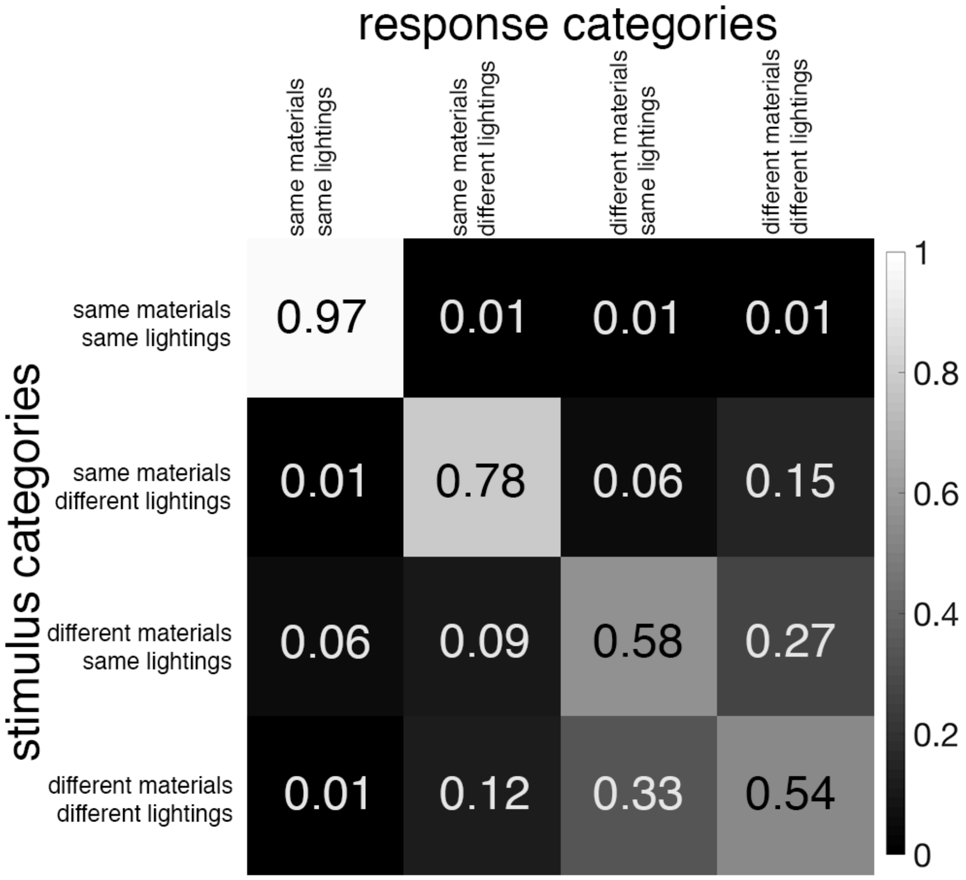

Figure 4.8: The fractions of responses per stimulus category. Each row represents a stimulus category, and each column represents a response category. The squares on the diagonal are the fractions of answering correctly, i.e., the discrimination accuracies.

observers had much difficulty in judging whether the lightings were the same or not but still performed well above chance.

In order to further analyze the results, we implemented signal-detection theory by considering the four-category discrimination task as two yes-or-no questions: (a) "Are the materials the same?" and (b) "are the lightings the same?" Explicitly, when analyzing materials, lighting was not considered and vice versa. For example, stimulus (or response) categories "same materials same lightings" and "same materials different lightings" were combined as one stimulus (or response) category for materials ("the same"). Answering "the same" when the stimuli were the same constitutes a "hit," and answering "the same" when the stimuli were actually different constitutes a "false alarm." The hits and false alarms could be converted to z scores $z(Hit)$ and $z(Fa)$, respectively (Macmillan & Creelman, 2005).

From $z(Hit)$ and $z(Fa)$, one can derive the sensitivity d', where $d' = z(Hit) - z(Fa)$, and the response bias c, where $c = -\frac{z(Hit)+z(Fa)}{2}$. The former refers to the ability to successfully indicate whether two stimuli are the same or different. The latter refers to the tendency to answer "same" independent of the type of stimulus pair (same or different). It turns out that all participants were sensitive to differences in materials as well as in lightings (see Supplementary Table 4.4 presenting the resulting d' and c values per participant). On average, they were significantly more sensitive to the material differences ($d' = 2.36 \pm 0.10$) than to the lighting differences ($d' = 1.82 \pm 0.15$). This was confirmed in a paired t test: $t(7) = 3.86, p = 0.006$. Because we found a significant difference between the averaged hit rates, paired t test, $t(7) = 3.20, p = 0.015$, but not between the averaged false alarms, paired t test, $t(7) = -1.27, p = 0.25$, the higher sensitivities for materials may be attributed to higher hit rates for materials. The average response biases for materials ($c = 0.08 \pm 0.06$) and for lightings ($c = 0.01 \pm 0.12$) were negligible and not significantly different as confirmed in a paired t test: $t(7) = -0.84, p = 0.43$. This is consistent with the observation that the usage of the four types of responses was almost equal: the sums of the columns in Figure 4.8 are 1.05, 1.00, 0.98, and 0.97. Finally, the largest range of individual values happened with $z(Fa)$ for lighting ($SEM = 0.18$; see Supplementary Table 4.4), confirming that there are individual differences comparable to those found for the performance measure in Experiment 1. In Figure 4.9, sensitivity d' and response bias c for materials and lightings in Experiment 2 are plotted in red. It is clear that observers had higher sensitivity for materials than for lightings. Note that, in this figure, we also show the results from Experiment 3 in blue plots.

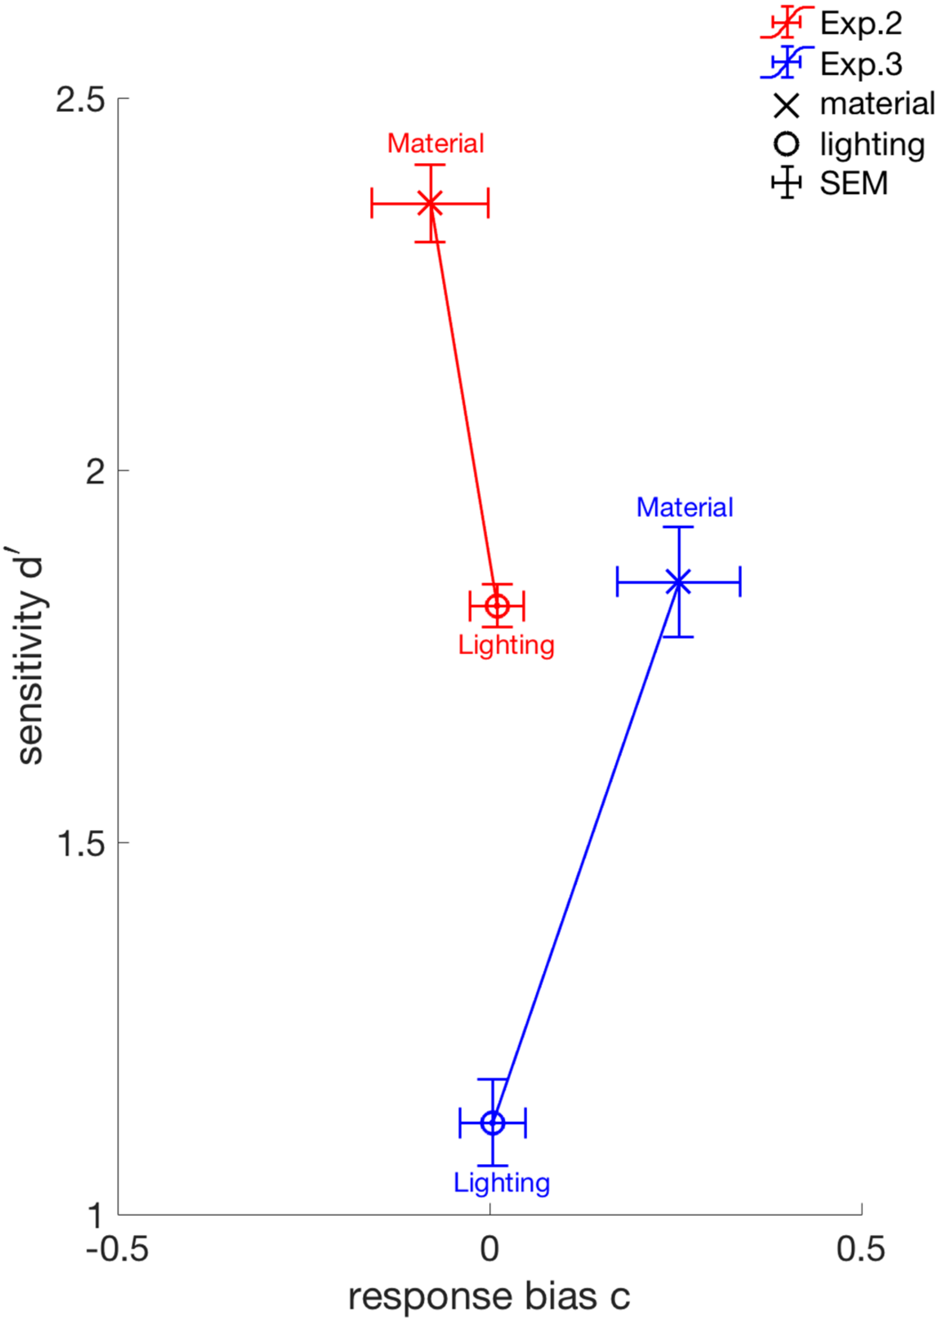

Figure 4.9: Sensitivity d' and response bias c for materials and lighting in Experiments 2 and 3. Red-colored plots show results from Experiment 2, and blue-colored plots show results from Experiment 3. Crosses depict results for materials; circles depict results for lighting. Each error bar depicts the corresponding standard error of the mean for both axes.

4.4. EXPERIMENT 3: ARE MATCHING PERFORMANCES AND DISCRIMINATION ACCURACIES WITHIN OBSERVERS COMPARABLE?

Because we found similar effects and idiosyncratic differences in Experiments 1 and 2, we wanted to further investigate the relationship between the matching and discrimination performances. Thus, we conducted a third experiment consisting of two sessions, one with the matching task and the other with the category discrimination task. A different group of observers was recruited and asked to participate in both sessions in order to be able to directly compare the results of the two tasks. Both tasks were simplified by removing the brilliance lighting mode and keeping ambient and focus lighting modes only in the stimuli.

4.4.1. METHODS

OBSERVERS

Ten inexperienced observers participated in both sessions of Experiment 3. All participants had normal or corrected-to-normal vision. Participants read and signed the consent form before the experiments. The experiments were approved by the human research ethics committee at Delft University of Technology and conducted in accordance with the declaration of Helsinki and Dutch law.

SESSION 1: SIMPLIFIED VERSION OF THE MATCHING TASK

Because the brilliance light was removed from the MixIM 1.0 interface (Figure 4.10), the basis images used in the mixing process are only the top two rows in Figure 4.1. For mixing only ambient and focus light, the mixing process for the stimuli was simply adjusted to Equation 4.4 with the weights as in Table 4.3:

$$I_{s.mat.} = w_{ambient} \cdot I_{ambient.mat.} + w_{focus} \cdot I_{focus.mat.}. \tag{4.4}$$

And similarly, the mixing process for the probe becomes:

$$I_{p.mat.} = w_{ambient}{}' \cdot I_{ambient.mat.}{}' + w_{focus}{}' \cdot I_{focus.mat.}{}'. \tag{4.5}$$

In this session, the four material modes in the stimuli and the four material modes

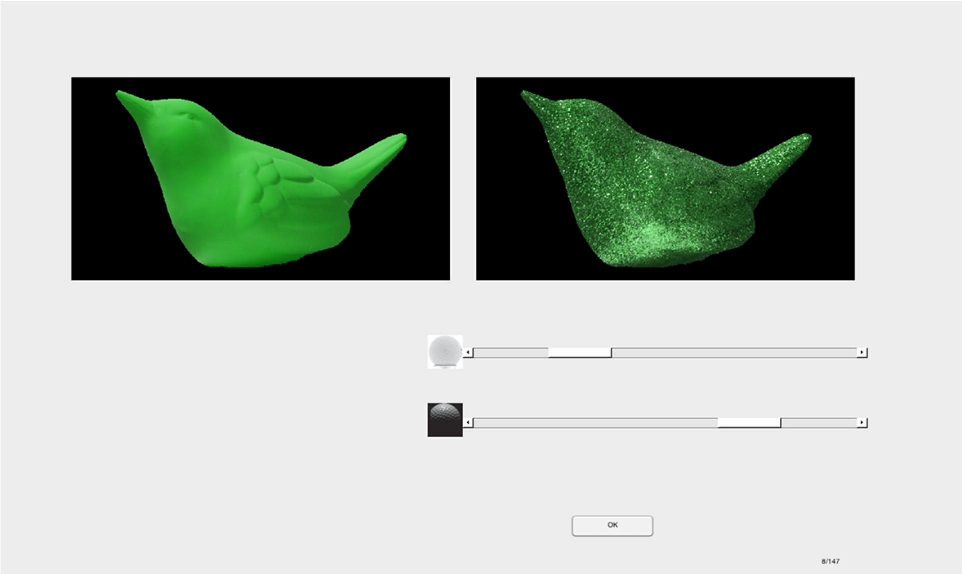

Figure 4.10: The interface for the first session in Experiment 3. Left: A stimulus image, consisting of a mixture of matte material in 50% ambient light and 50% focus light. Right: A probe image (glittery material mode). Top slider represents the contribution of ambient light. Bottom slider represents the contribution of focus light. In this figure, the illumination of the probe image does not match the illumination of the stimulus image.

No.	$w_{ambient}$	w_{focus}
1	1	0
2	0	1
3	0.5	0.5

Table 4.3: Weight of each canonical lighting mode in the stimuli for Experiment 1.

in the probe images were combined with three weight combinations for the light modes, which resulted in 48 trials per run. With three repetitions plus three practice trials, there were 147 trials per observer, which resulted in a session lasting between 30 and 60 min.

Session 2: Simplified version of the four-category discrimination task

After observers finished the first session, they did a second session: the four-category discrimination task using the same interface as in Experiment 2 (Figure 4.7). Unlike in Experiment 2, before the actual experiment started, observers did not browse through all stimuli images. Instead, they were told all stimuli images they were about to see had appeared in the previous session. They were also told that all stimuli images in this session would be images of one of the four material modes in one of the two lighting modes, which they just manipulated by moving the sliders in the first session.

For each observer, with eight basis images as stimuli, there were 36 possible combinations, including eight "same materials same lightings," four "same materials different lightings," 12 "different materials same lightings," and 12 "different materials different lightings." To create the same number of stimuli per category, these combinations were repeated three, six, two, and two times, respectively. The resulting total number of trials was 24 per category, in total 96 trials per observer, which took approximately half an hour to finish.

4.4.2. Results

Matching

Because the brilliance lighting mode was removed in Experiment 3, the linear factor matrix X solved by least square fitting changed accordingly as in Equation 4.6:

$$[\mathbf{P}]_{2\times(96\times N)} = [\mathbf{X}]_{2\times 2} \cdot [\mathbf{S}]_{2\times(96\times N)} + [\mathbf{E}]_{2\times(96\times N)}\,, \tag{4.6}$$

where, $[\mathbf{P}] = \begin{bmatrix} w_{ambient}{}' \\ w_{focus}{}' \end{bmatrix}$, $[\mathbf{S}] = \begin{bmatrix} w_{ambient} \\ w_{focus} \end{bmatrix}$, and the residual $[\mathbf{E}] = \begin{bmatrix} e_{ambient} \\ e_{focus} \end{bmatrix}$.

For N participants, there were in total $(96 \times N)$ columns in matrix **S**, matrix**P**, and matrix **E** when solving the Equation 4.6. The linear factor matrix **X** became 2×2, and the matrix **E** was again the subtraction between **P** and $\mathbf{X} \cdot \mathbf{S}$. If the matching would be veridical, **X** would be a 2×2 identity matrix, and matrix **E** would be a $2 \times (96 \times N)$ zero matrix. The matching performance could be evaluated in the same manner as in Experiment 1, i.e., taking the ratio between the sum of the diagonal values and the sum of **X**, i.e., $r = \sum diag(\mathbf{X}) / \sum(\mathbf{X})$, which ranges from zero (only possible mathematically) to one (veridical). Note that, in Experiment 3, the chance level is 0.5, which is higher than the chance level (0.33) in Experiment 1.

The overall matching results of all observers in Experiment 3 (N = 10) are:

$[\mathbf{X}] = \begin{bmatrix} 0.72 & 0.35 \\ 0.41 & 0.73 \end{bmatrix}$, $r = 0.66$.

Because the chance level is 0.50, the ratio r being 0.66 shows that, overall, observers performed above chance in the matching session in Experiment 3. The bivariation plot of all observers is shown in Figure 4.11 in the same format as in Figures 4.5 and 4.6 for Experiment 1. Each ellipse represents one standard deviation of bivariate normal distribution fitted to 16 data points (rendered invisible for clarity of presentation). The coordinates of the crosses depict the corresponding weight combinations of the stimuli as shown in Table 4.3, corresponding to the color of the ellipses. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the black color corresponds to the stimuli in which both lighting modes were optically mixed 50% each in the mixture. Similar to what can be seen in Figure 4.5, the ellipses show a shift from the veridical values toward the center but are still in the correct order. Check Supplementary Figure S6 for individual results.

Discrimination

The results of the four-category discrimination task are shown in Figure 4.12. Similar to the results of Experiment 2 (Figure 4.8), when the materials and lightings were both the same in the stimuli image pair, observers got the highest accuracy, being 0.89. When the materials were the same and lightings were different, the accuracy decreased to 0.48. The accuracy was 0.43 when the materials were different and the lightings were the same, and 0.57 when both the materials and the lightings were different. Note that here the chance level is 0.25, the same as in Experiment 2. We again found a strong association between the responses and the stimulus categories ($\chi^2(9) = 926.32, p < 0.001$).

Again, we implemented signal-detection theory by considering the four-category discrimination task as two yes-or-no questions: (a) “Are the materials the same?” and (b) “are the lightings the same?” As in Experiment 2, observers were all found to be sensitive to the differences in both the materials and the lightings (the resulting values of sensitivity d' response bias c are listed in Supplementary Table 4.5). They were also significantly more sensitive to material differences ($d' = 1.85 \pm 0.16$) than to lighting differences ($d' = 1.12 \pm 0.16$) , confirmed in a paired t test, $t(9) = 6.832, p < 0.001$. The average response bias for materials ($c = 0.25 \pm 0.07$) was not significantly different from that for lightings ($c = 0.00 \pm 0.09$), confirmed in a paired t test,$t(9) = 2.151, p = 0.06$. Unlike in Experiment 2, there was no significant difference between the averaged hit rates, paired t test, $t(9) = 0.97, p = 0.36$, but now there was one between the averaged false alarms, paired t test, $t(9) = -4.45, p = 0.002$, suggesting that the higher sensitivities for materials may be attributed to lower false alarm rates.

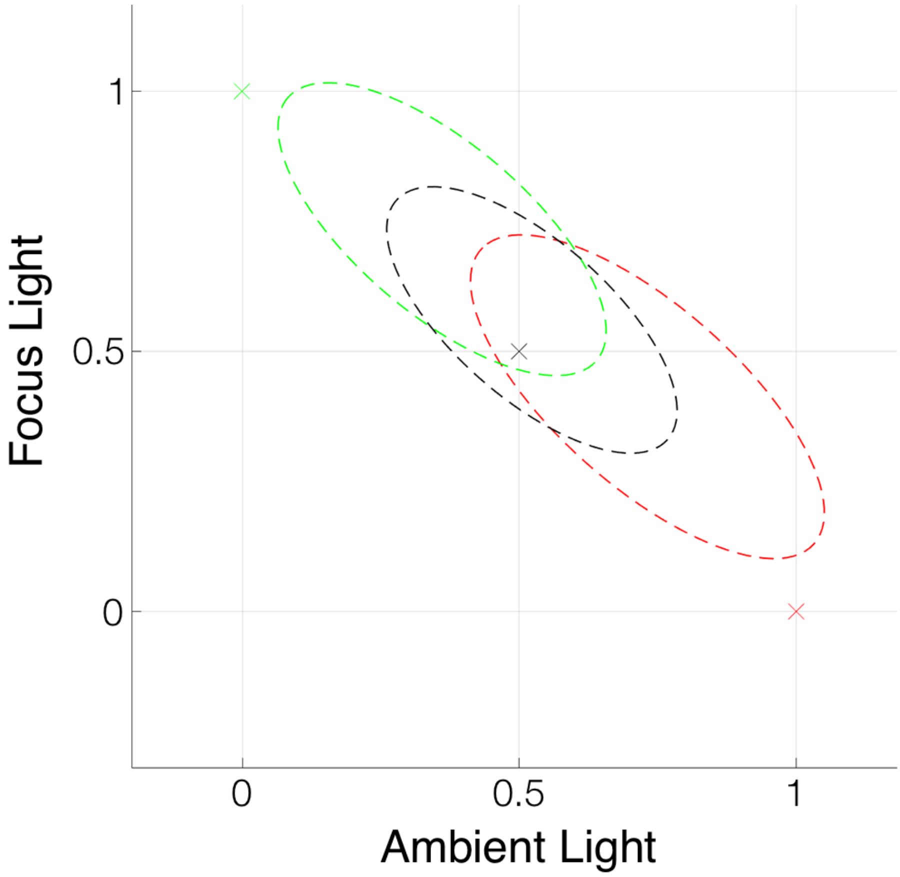

Figure 4.11: The bivariation plot of the overall matching results in Experiment 3 (N=10). Different colors correspond to weight combinations (Table 4.3) of ambient light (x-axis) and focus light (y-axis) in the stimuli, which are depicted by the crosses. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

COMPARISON

To directly compare the performances of the matching task (session 1) and the four-category discrimination task (session 2), we first tested at a global level by correlating the individual light-matching accuracies with the corresponding sensitivities d' and response biases c for both materials and lighting (see Supplementary Table 4.4). The light-matching accuracy was found to be significantly correlated with one variable only, namely

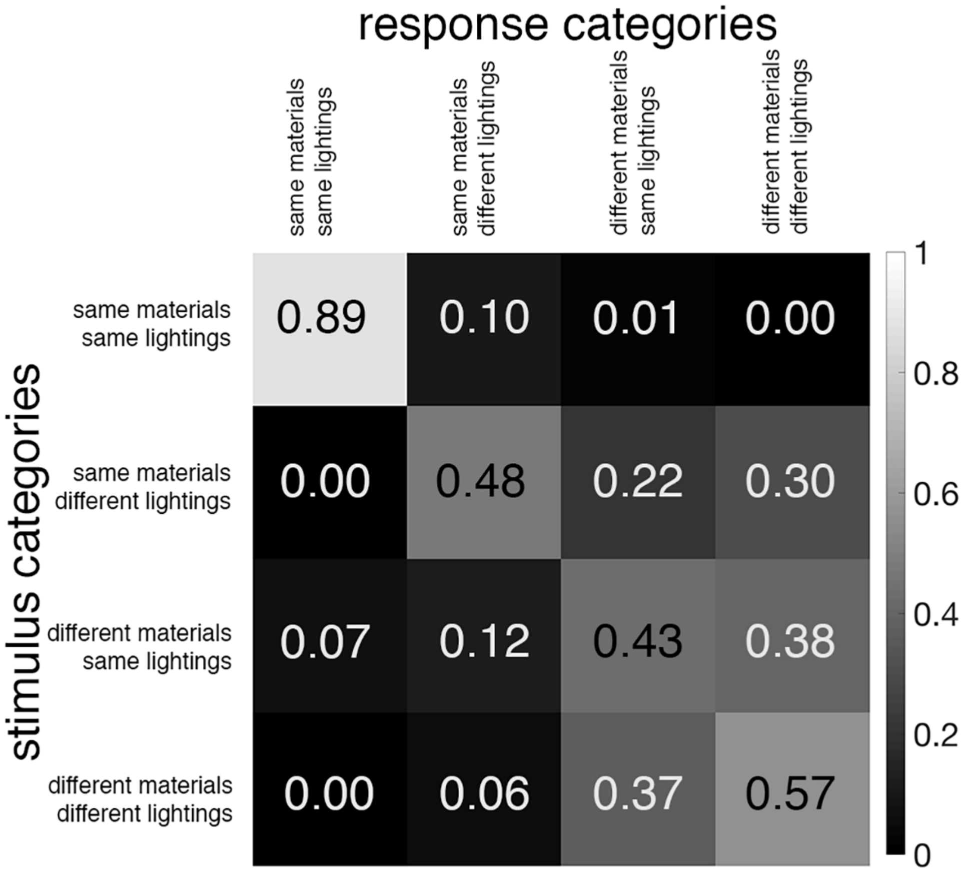

Figure 4.12: The fractions of responses per stimulus category of the four-category discrimination task in Experiment 3. Each row represents a stimulus, and each column represents a response category. The squares on the diagonals are the fractions of answering correctly, i.e., the accuracies.

response bias c for light discrimination (negatively correlated, $r^2 = 0.40, p = 0.049$).

Subsequently, we further tested the correlation between the light-matching accuracy (the ratio r) and the light-discrimination accuracy (the fraction of correctly answering "same lighting") per material combination (Figure 4.13). Overall, a significant correlation between the light-matching and light-discrimination accuracy was found in Experiment 3 ($r^2 = 0.45, p < 0.01$). More specifically, some observations are listed below:

- For the symmetric cluster in which materials were the same (blue data points), we observed that the material combinations including velvet tend to produce lower performances in both the discrimination and matching tasks.

- For the asymmetric cluster in which materials were different (red data points), we observed that when specular material was involved, the discrimination accuracy (0.57 ± 0.04) was significantly higher than when specular material was not involved (0.49 ± 0.03).
- The combinations with specular and glittery resulted in the highest performance, showing that those two modes interacted least of all asymmetric combinations.
- The combinations with velvety and matte gave the lowest performance among all cases, showing that these modes interacted most of all our material modes.

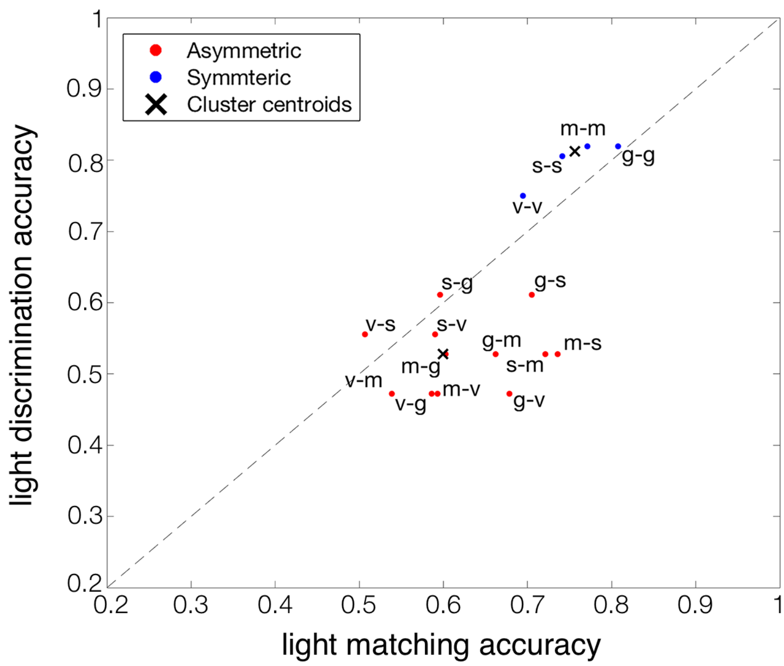

Figure 4.13: Comparison between the lighting-matching and the discrimination results in Experiment 3. The data points depict different material combinations with "m", "v", "s", and "g" denoting matte, velvety, specular, and glittery, respectively (e.g., "m-m" means the materials in the trial were both matte; "s-g" means the materials in the trial were specular and glittery). Colors were assigned using a k-means clustering algorithm for two clusters with the crosses depicting the cluster centroids. The dashed line depicts the identity line.

4.5. DISCUSSION

In this paper, we present three experiments. In Experiment 1, we asked observers to optically mix three canonical lighting modes (ambient, focus, and brilliance) while discounting our canonical material modes (matte, velvety, specular, and glittery) in a matching task. Eleven out of 15 observers' performance levels were well above chance, and the remaining four observers performed just above chance (Figure 4.4). In Experiment 2, we asked observers to simultaneously discriminate materials and lightings in a four-category discrimination task and found that observers were more sensitive in discriminating our material modes than our lighting modes; i.e., they were better at judging the material modes than the lighting modes to be the same or not. In Experiment 3, we implemented a simplified version of both the matching and four-category discrimination tasks by removing the brilliance light and then asked observers to first perform the matching task and then the four-category discrimination task. Results from Experiment 3 showed that the matching and discrimination results were comparable and confirmed the asymmetric perceptual confounds between materials and lightings that we observed in Experiments 1 and 2. Across these experiments, observers were found to be more sensitive to material differences than to lightings differences.

For the matching task, an interface inspired by audio-mixing desks was tested in a previous study (Zhang et al., 2016) and further developed in this study. Here, the number of sliders in the interface was reduced from four for the material mixing in the previous study to three in Experiment 1 and two in Experiment 3 of this study for lighting matching. This actually reduced the level of complexity of manipulating the interface and increased the level of chance performance from one of four to one of three for Experiment 1 and one of two for Experiment 3 (if calculated as the ratio r using least squares fitting). However, the general performance of the light-matching task in this study was lower than the performance of the material-matching task in our former studies. So observers were better at discounting our lightings in matching the optically mixed canonical material modes than discounting our materials in matching the optically mixed canonical lighting modes. This again confirms the asymmetric perceptual confounds we found in Experiments 2 and 3.

One possible cause of this asymmetric perceptual confound might be that we showed the appearance of the objects without a context. In our experiments, observers had to make judgments based purely on the objects' appearances. If observers would have access to other information about the light, such as from the background or the appearance of other objects, it might be easier for them to make more accurate judgments. Indeed, light is usually inferred by looking at the appearance of (the objects in) a scene.

Ecologically, this asymmetric confound makes sense as human beings have to recognize and interact with materials under different illumination in our daily lives. Yet most of us (except for instance lighting professionals) do not normally have the necessity to recognize or interact with different types of lightings. In fact, we may simply be used to changes of illumination in natural environments without realizing it, especially for those changes that occur naturally, which is the case for the variations and modes that we used.

It should be realized, however, that we are comparing apples and oranges (lighting and materials) and that there is no obvious physical basis to compare the magnitudes of the differences between materials and lightings. In this study, we approached this by selecting canonical modes, which are optically very different from each other and altogether span much of the reflectance (BRDF) space and descriptions of natural light fields. The limitations of our conclusions are obviously set by this choice of modes and their representations via the bird photographs. Detailed characteristics of the modes, such as lighting direction, beam width, the statistical characteristics of the brilliance lighting, and microscattering properties of the glittery flakes or velvet hairs, are expected to have an influence on the results. However, considering the coarse characteristics of the modes and especially how wide apart they are in the spaces of possible reflectance and lighting types, we reasoned that the asymmetric confound in this study suggests a more generic phenomenon with an ecologically plausible basis.

This connects to how our visual systems represent materials and lightings. In material-perception studies, instead of the "inverse optics" and the "image statistics" approaches, the "statistical appearance models" approach represents an alternative, for instance, for the study of gloss perception (Fleming, 2014). Similarly, in our studies, we presented "a painterly approach" (Zhang et al., 2016), i.e., optical mixing of canonical material or lighting modes, that allows observers to intuitively manipulate the midlevel image cues in a weighted-mixture manner. From the results of our earlier material-matching experiment, we argued that these key midlevel image features form the triggers for material perceptions, such as the smooth shading along the surface of the matte mode, the bright contours for the velvety mode, the highlights for the specular mode, and the bright speckles all over the surfaces of the glittery mode (Zhang et al., 2016). Here, we argue that midlevel image features could also be the triggers for our lighting perceptions: the overall brightness and lack of gradients for the ambient or mathematical zeroth-order component of the light; the contrast, main highlight, and the shading gradient direction for the focus or first-order component of the light; and the contrast and spatial patterns of the glint for the brilliance or higher order components of the light (Ganslandt & Hofmann, 1992; Kelly, 1952). Close observation of our photographs in Figure 4.1 plus their mixtures and computer-rendered simulations may suggest that these features are, overall, less robust for variations of material than for variations of lighting (Figure 4.14). In Figure 4.14A, we show the top 5% brightest pixels in each basis image by applying thresholding to the red channel of the images. In Figure 4.14B, we show the shading patterns by posterizing the green channel of the basis images from 255 to four levels. The last column of the thresholded images shows that the images of glittery material are clearly dominated by the spread of the dots, i.e., the glints, that result in the glittery appearance regardless of illumination. The images of the matte, velvety, and specular materials show otherwise spatially varied patterns. Specifically, we observed smooth shading gradients for matte mode; smooth shading gradients, bright contours, and fine-grained textures that might trigger the velvetiness in the velvety image; and the specular highlight regions spread along the curvature of the surface for specular mode (except for specular under the ambient lighting, which caused interactions with matte mode). One may argue that, in ambient lighting, the bright contours, which we suggested trigger velvetiness, can be observed in the thresholded images for matte, velvety, and specular material, too. However, by closely looking at the spread of the pixels on those bright contours in velvety images, combining the patterns of their shadings, we could discriminate velvety from matte or specular (not quantitatively though). In natural scenes with arbitrary materials and light, this difference in feature robustness would make it harder to judge the lighting than the material. Similarly, these midlevel image features varied differently for matte, specular, and glittery materials under the canonical lighting modes, being more diagnostic for our canonical materials than our canonical lightings, causing the asymmetric perceptual confounds. Simple image statistics (such as comparing the image histograms, then calculating the difference between each two images, and the correlation between each two images) could not explain the asymmetric confounds. In order to better understand what and how midlevel image features account for material and lighting perception, novel quantitative metrics are required for image analysis, such as separating specific features from object color (Klinker, Shafer, & Kanade 1987).

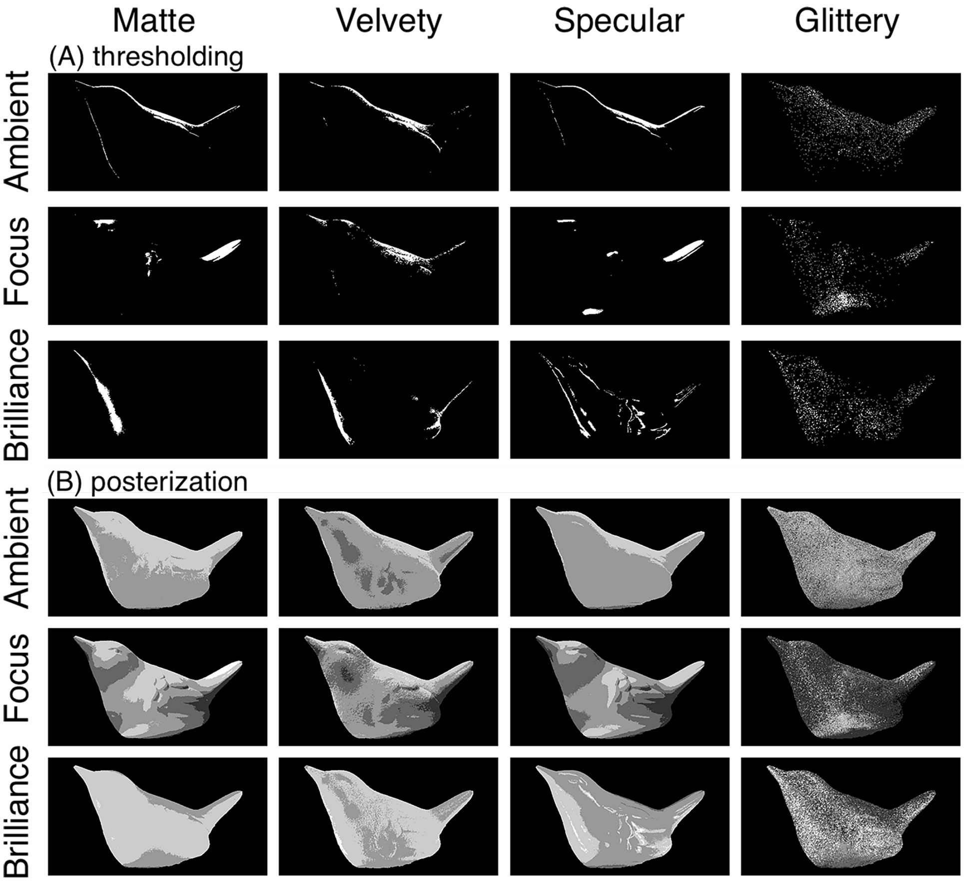

Figure 4.14: Examples of image analyses of the basis images (in the same format as in Figure 4.1). Top: The red-channel thresholding showing the upper 5% percentile of brightest pixels. Note that the thresholding level varies per image. Bottom: The green-channel of the basis images after posterization from 255 to four levels.

4.6. CONCLUSION

In this study, we implemented two types of tasks, namely a light-matching task and a four-category discrimination task for our canonical material and lighting modes. From the results of the light-matching tasks in Experiments 1 and 3, we found that most of our observers could match optically mixed canonical lighting modes while discounting materials although a small portion of the observers tended to use only a narrow range around the center of the possible slider positions. In particular, observers performed better when the materials in the stimulus and the probe were the same than when they were different. From the results of the four-category discrimination tasks in Experiments

2 and 3, we found that observers could discriminate our material modes better than our lighting modes. Their sensitivities for the material discrimination were found to be higher than those for the lighting discrimination. Observers also found it difficult to discriminate lighting modes when the materials were different. Moreover, in Experiment 3, by conducting a simplified version of both matching and discrimination tasks with the same group of observers, we found that the performances of matching and discrimination task were indeed comparable.

To conclude, in all three experiments and across all observers, the sensitivities for judging the differences between our canonical material modes are higher than those for the canonical lighting modes. If materials are different, it is harder to see whether or not the illuminations are different than if materials are the same. If lightings are different, it is almost as easy to see whether the materials are different or not as when the lightings are the same. Our findings suggest that midlevel image features are more robust across different materials than across different lightings and, thus, more diagnostic for our canonical materials than our canonical lightings, causing the asymmetric perceptual confounds.

BIBLIOGRAPHY

Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. *Journal of Vision, 9* (11):10, 1–17, https://doi.org/10. 1167/9.11.10.

Blinn, J. F. (1977). Models of light reflection for computer synthesized pictures. *ACM SIGGRAPH Computer Graphics, 11* (2), 192–198.

Brainard, D. H. (1997) The Psychophysics Toolbox, *Spatial Vision* 10: 433–436.

Cook, R. L., & Torrence, K. E. (1982). A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1(1), 7–24.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. *Journal of Vision, 10* (4):8, 1–9, https://doi.org/10. 1167/10.4.8.

Dror, R. O., Adelson, E. H., & Willsky, A. S. (2001). Estimating surface reflectance properties from images under unknown illumination. In *Human Vision and Electronic Imaging VI* (pp. 231–243). International Society for Optics and Photonics.

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. *Journal of Vision*, 4(9):11, 821–837, https:// doi.org/10.1167/4.9.1.

Fleming, R. W. (2014). Visual perception of materials and their properties. *Vision Research*, 94, 62–75.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, 3(5):3, 347–368, https://doi.org/10.1167/3.5.3.

Fleming, R. W., Wiebel, C., & Gegenfurtner, K. (2013). Perceptual qualities and material classes. *Journal of Vision*, 13(8):9, 1–20, https://doi.org/10.1167/13.8.9.

Ganslandt, R., & Hofmann, H. (1992). *Handbook of lighting design*. Brauschweig/Wiesbaden, Germany: ERCO GmbH.

Hansmann-Roth, S., Pont, S. C., & Mamassian, P. (2017). Contextual effects on real bicolored glossy surfaces. *Journal of Vision*, 17(2):17, 1–13, https:// doi.org/10.1167/17.2.17.

Ho, Y. X., Landy, M. S., & Maloney, L. T. (2006). How direction of illumination affects visually perceived surface roughness. *Journal of Vision*, 6(5): 8, 634–648, https://doi.org/10.1167/6.5.8.

Ho, Y. X., Landy, M. S., & Maloney, L. T. (2008). Conjoint measurement of gloss and surface texture. *Psychological Science*, 19(2), 196–204.

Karana, E., Hekkert, P., & Kandachar, P. (2009). Meanings of materials through sensorial properties and manufacturing processes. *Materials & Design*, 30(7), 2778–2784.

Kartashova, T., de Ridder, H., te Pas, S. F., Schoemaker, M., & Pont, S. C. (2015). The visual light field in paintings of Museum Prinsenhof: Comparing settings in empty space and on objects. In *B. E. Rogowitz, T. N. Pappas, H. de Ridder (Eds.), SPIE/IS&T Electronic Imaging* (pp. 93941M–93941M). International Society for Optics and Photonics.

Kartashova, T., Sekulovski, D., de Ridder, H., te Pas, S. F., & Pont, S. C. (2016). The global structure of the visual light field and its relation to the physical light field. *Journal of Vision*, 16(10):9, 1–16, https:// doi.org/10.1167/16.10.9.

Kelly, R. (1952). Lighting as an integral part of architecture. *College Art Journal*, 12(1), 24–30.

Kim, J., Marlow, P., & Anderson, B. L. (2011). The perception of gloss depends on highlight congruence with surface shading. *Journal of Vision*, 11(9): 4, 1–19, https://doi.org/10.1167/11.9.4.

Kim, M., Wilcox, L. M., & Murray, R. F. (2016). Perceived three-dimensional shape toggles perceived glow. *Current Biology*, 26(9), R350–R351.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1–16.

Klinker, G. J., Shafer, S. A., & Kanade, T. (1987). Using a color reflection model to separate highlights from object color. In *Proceedings - First international conference on computer vision*. Vol. 87 (pp. 145–150). New York, NY: IEEE.

Koenderink, J., & Pont, S. (2003). The secret of velvety skin. *Machine Vision and Applications*, 14(4), 260– 268.

Koenderink, J. J., Pont, S. C., van Doorn, A. J., Kappers, A. M., & Todd, J. T. (2007). The visual light field. *Perception*, 36(11), 1595–1610.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide. (2nd ed.)*. Mahwah, NJ: Lawrence Erlbaum Associates.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology*, 22(20), 1909–1913.

Marr, D. (1982). *Vision*. San Francisco: W. H. H. Freeman.

Morgenstern, Y., Geisler, W. S., & Murray, R. F. (2014). Human vision is attuned to the diffuseness of natural light. *Journal of Vision*, 14(9):15, 1–18, https://doi.org/10.1167/14.9.15.

Morgenstern, Y., Geisler, W. S., & Murray, R. F. (2015, March). The role of natural lighting diffuseness in human visual perception. In *SPIE/ IS&T Electronic Imaging* (pp. 93940P–93940P). International Society for Optics and Photonics, https://doi.org/10.1117/12.2086258.

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

Motoyoshi, I., & Matoba, H. (2012). Variability in constancy of the perceived surface reflectance across different illumination statistics. *Vision Research*, 53(1), 30–39.

Motoyoshi, I., Nishida, S. Y., Sharan, L., & Adelson, E. H. (2007, May 10). Image statistics and the perception of surface qualities. *Nature*, 447(7141), 206–209.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2007). Light field constancy within natural scenes. *Applied Optics*, 46(29), 7308–7316.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2009). Structure of light fields in natural scenes. *Applied Optics*, 48(28), 5386–5395.

Nishida, S. Y., Sawayama, M., & Shimokawa, T. (2015). Material-dependent shape distortion by local intensity order reversal. *Journal of Vision*, 15(12):940, https://doi.org/10.1167/15.12.940.

Nishida, S. Y., & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties. *Journal of the Optical Society of America. A*, Optics, image science, and vision, 15(12), 2951–2965.

Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. *Journal of Vision*, 10(9):5, 1–19, https://doi. org/10.1167/10.9.5.

Pellacini, F., Ferwerda, J. A., & Greenberg, D. P. (2000). Toward a psychophysically based light reflection model for image synthesis. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (pp. 55–64). New York, NY: ACM Press/Addison-Wesley Publishing Co.

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision* 10, 437–442.

Pizlo, Z. (2001). Perception viewed as an inverse problem. *Vision Research*, 41(24), 3145–3161.

Poggio, T., & Koch, C. (1985). Ill-posed problems in early vision: From computational theory to analogue networks. *Proceedings of the Royal Society of London B: Biological Sciences*, 226(1244), 303–323.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature*, 317(6035), 314–319.

Pont, S. C., & Koenderink, J. J. (2007). Matching illumination of solid objects. Attention, *Perception, & Psychophysics*, 69(3), 459–468.

Pont, S. C., & te Pas, S. F. (2006). Material-illumination ambiguities and the perception of solid objects. *Perception*, 35(10), 1331–1350.

Schirillo, J. A. (2013). We infer light in space. *Psychonomic Bulletin & Review*, 20(5), 905–915.

Sharan, L., Rosenholtz, R., & Adelson, E. (2009). Material perception: What can you see in a brief glance? *Journal of Vision*, 9(8): 784, https://doi.org/ 10.1167/9.8.784.

Sharan, L., Rosenholtz, R., & Adelson, E. H. (2014). Accuracy and speed of material categorization in real-world images. *Journal of Vision*, 14(9):12, 1–24, https://doi.org/10.1167/14.9.12.

Schu¨tt, H. H., Baier, F., & Fleming, R. W. (2016). Perception of light source distance from shading patterns. *Journal of Vision*, 16(3):9, 1–20, https:// doi.org/10.1167/16.3.9.

te Pas, S. F., & Pont, S. C. (2005). A comparison of material and illumination discrimination performance for real rough, real smooth and computer generated smooth spheres. In *Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization* (pp. 75–81). New York, NY: ACM.

Toscani, M., Gegenfurtner, K. R., & Doerschner, K. (2017). Differences in illumination estimation in #thedress. *Journal of Vision*, 17(1):22, 1–14, https:// doi.org/10.1167/17.1.22.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2), 265–272.

van Assen, J. J. R., Wijntjes, M. W., & Pont, S. C. (2016). Highlight shapes and perception of gloss for real and photographed objects. *Journal of Vision*, 16(6):6, 1–14, https://doi.org/10.1167/16.6.6.

van Doorn, A. J., Koenderink, J. J., & Wagemans, J. (2011). Light fields and shape from shading. *Journal of Vision*, 11(3):21, 1–21, https://doi.org/10. 1167/11.3.21.

Vangorp, P., Laurijssen, J., & Dutré, P. (2007). The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics*, 26(3), 77.

Wijntjes, M. W., & Pont, S. C. (2010). Illusory gloss on Lambertian surfaces. *Journal of Vision*, 10(9):13, 1– 12, https://doi.org/10.1167/10.9.13.

Xia, L., Pont, S. C., & Heynderickx, I. (2013). Probing light in real scenes using optical mixtures. In *Proceedings of the ACM Symposium on Applied Perception* (pp. 137–137). New York, NY: ACM.

Xia, L., Pont, S. C., & Heynderickx, I. (2014). The visual light field in real scenes. *i-Perception*, 5(7), 613–629.

Xia, L., Pont, S. C., & Heynderickx, I. (2016). Effects of scene content and layout on the perceived light direction in 3D spaces. *Journal of Vision*, 16(10):14, 1–13, https://doi.org/10.1167/16.10.14.

Xia, L., Pont, S. C., & Heynderickx, I. (2017a). Light diffuseness metric part 1: Theory.

Lighting Research & Technology, 49(4), 411–427, https://doi. org/10.1177/1477153516631391.

Xia, L., Pont, S. C., & Heynderickx, I. (2017b). Light diffuseness metric, Part 2: Describing, measuring and visualising the light flow and diffuseness in three-dimensional spaces. *Lighting Research & Technology*, 49(4), 428–445, https://doi.org/10.1177/1477153516631392.

Zhang, F., de Ridder, H., & Pont, S. (2015). The influence of lighting on visual perception of material qualities. In *SPIE/IS&T Electronic Imaging* (pp. 93940Q–93940Q). International Society for Optics and Photonics.

Zhang, F., de Ridder, H., Fleming, R. W., & Pont, S. (2016). MatMix 1.0: Using optical mixing to probe visual material perception. *Journal of Vision*, 16(6): 11, 1–18, https://doi.org/10.1167/16.6.11.

SUPPLEMENTS

Participant	materials				lightings			
	$z(Hit)$	$z(Fa)$	d'	c	$z(Hit)$	$z(Fa)$	d'	c
1	1.82	-1.06	2.87	-0.38	0.89	-0.89	1.77	0.00
2	1.09	-1.23	2.31	0.07	0.95	-0.55	1.49	-0.20
3	1.15	-1.09	2.24	-0.03	0.73	-0.89	1.61	0.08
4	1.12	-1.30	2.41	0.09	0.59	-1.53	2.12	0.47
5	1.49	-1.09	2.57	-0.20	0.95	-0.89	1.83	-0.03
6	1.16	-0.84	1.99	-0.16	1.03	-0.57	1.59	-0.23
7	1.09	-0.91	2.00	-0.09	1.19	-0.25	1.43	-0.47
8	1.18	-1.30	2.48	0.06	0.89	-1.81	2.70	0.46
Mean	1.26	-1.10	**2.36**	-0.08	0.90	-0.92	**1.82**	0.01
SEM	0.09	0.06	0.10	0.06	0.06	0.18	0.15	0.12

Table 4.4: Individual results of the SDT data of Experiment 2 (N = 8).

4

Participant	Discrimination								Matching
	materials				lightings				
	$z(Hit)$	$z(Fa)$	d'	c	$z(Hit)$	$z(Fa)$	d'	c	r
1	0.43	-1.53	1.96	0.55	0.33	-0.97	1.29	0.32	0.51
2	0.61	-0.89	1.50	0.14	0.05	-0.89	0.94	0.42	0.54
3	0.11	-1.15	1.25	0.52	0.32	0.16	0.16	-0.24	0.59
4	0.61	-1.39	1.99	0.39	0.89	-0.49	1.38	-0.2	0.60
5	0.82	-1.26	2.07	0.22	0.06	-0.68	0.73	0.31	0.66
6	0.21	-0.75	0.95	0.27	0.27	-0.55	0.81	0.14	0.66
7	1.54	-1.06	2.59	-0.24	1.06	-1.06	2.11	0.00	0.71
8	1.06	-1.26	2.31	0.10	0.97	-0.55	1.52	-0.21	0.71
9	0.68	-1.38	2.06	0.35	0.68	-0.44	1.11	-0.12	0.73
10	0.67	-1.15	1.82	0.24	0.97	-0.21	1.18	-0.38	0.83
Mean	0.67	-1.18	**1.85**	0.25	0.56	-0.57	**1.12**	0.00	0.65
SEM	0.13	0.08	0.16	0.07	0.13	0.12	0.16	0.09	0.03

Table 4.5: Individual results of the SDT data of Experiment 3 (N = 10).

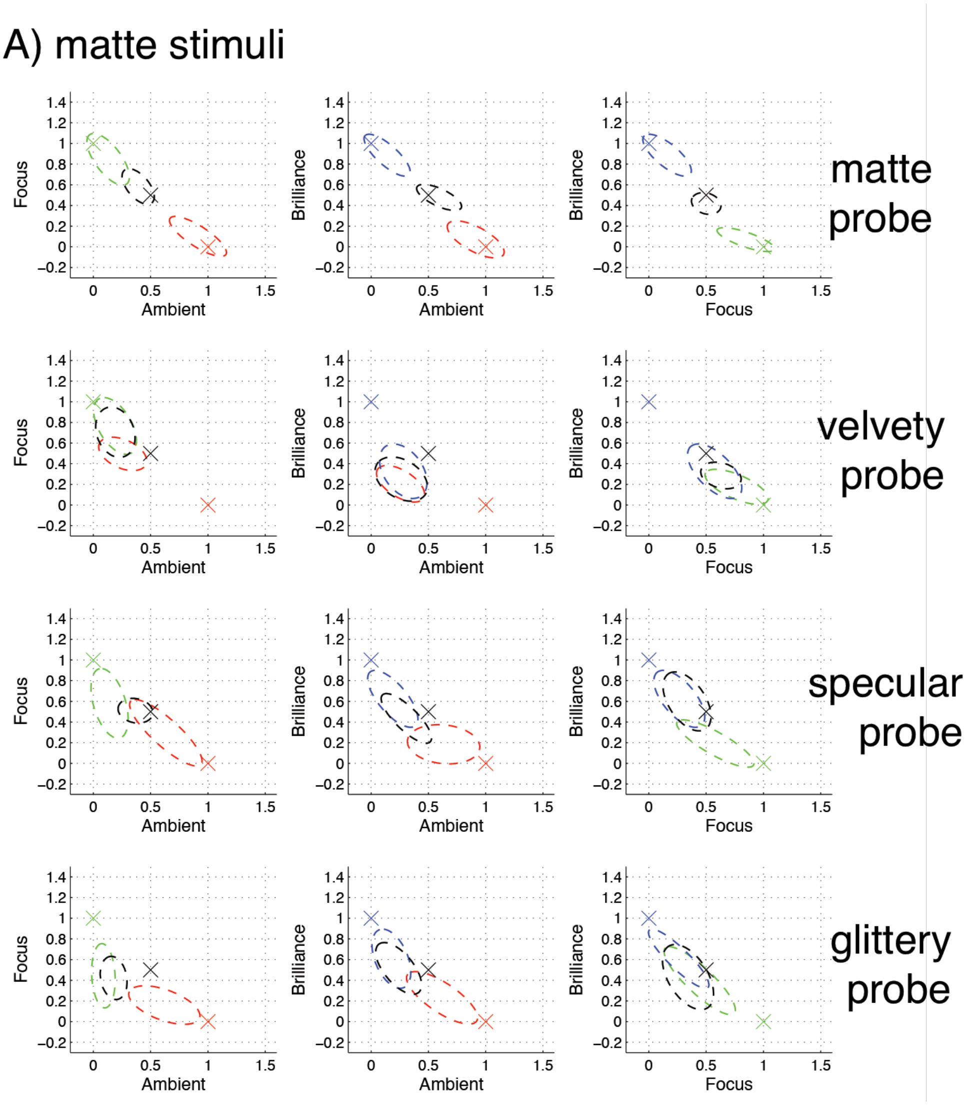

Figure 4.15: The bivariation plots for combinations of canonical material modes in Experiment 1 with A) matte mode in the stimuli. The subplots in each row are results from one of the four material modes in the probe, specifically matte, velvety, specular, and glittery from top to bottom, respectively. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

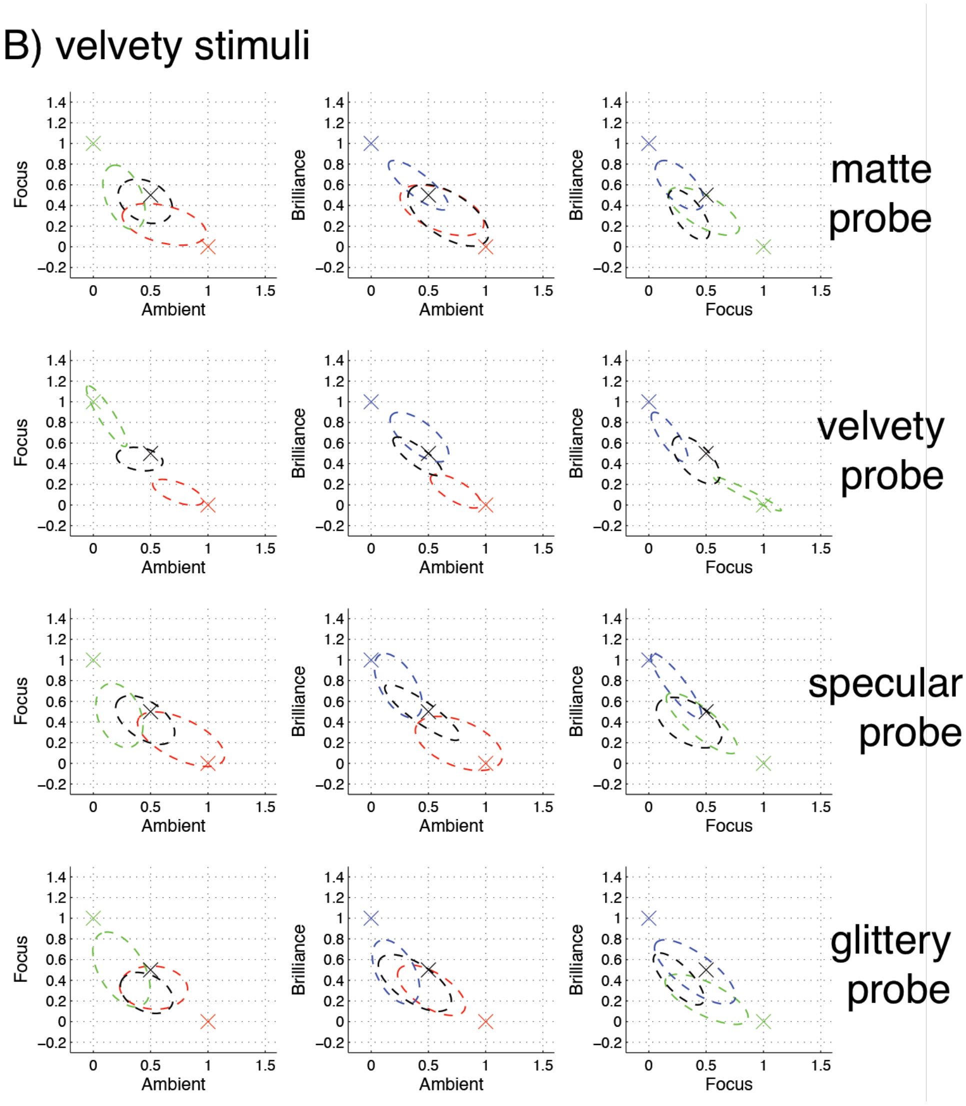

Figure 4.16: The bivariation plots for combinations of canonical material modes in Experiment 1 with B) velvety mode in the stimuli. The subplots in each row are results from one of the four material modes in the probe, specifically matte, velvety, specular, and glittery from top to bottom, respectively. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

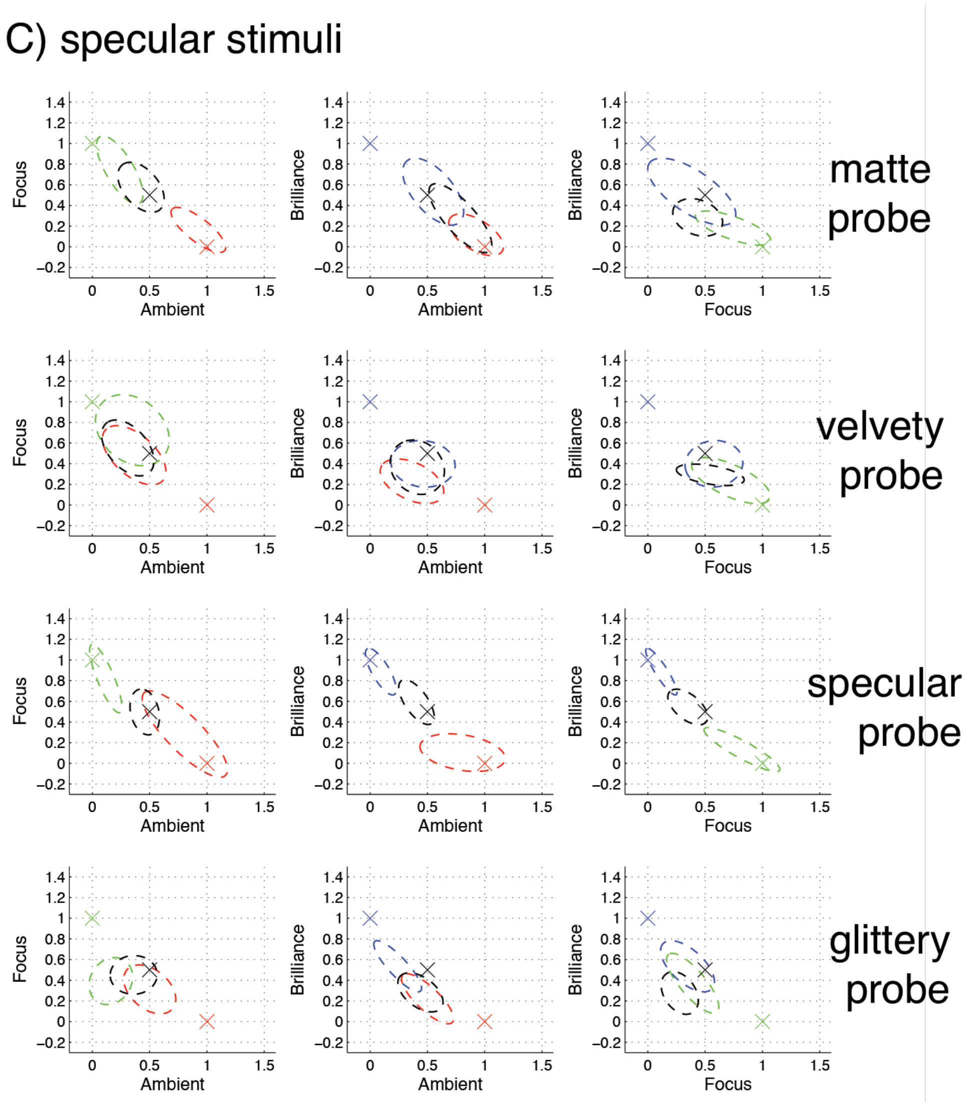

Figure 4.17: The bivariation plots for combinations of canonical material modes in Experiment 1 with C) specular mode in the stimuli. The subplots in each row are results from one of the four material modes in the probe, specifically matte, velvety, specular, and glittery from top to bottom, respectively. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

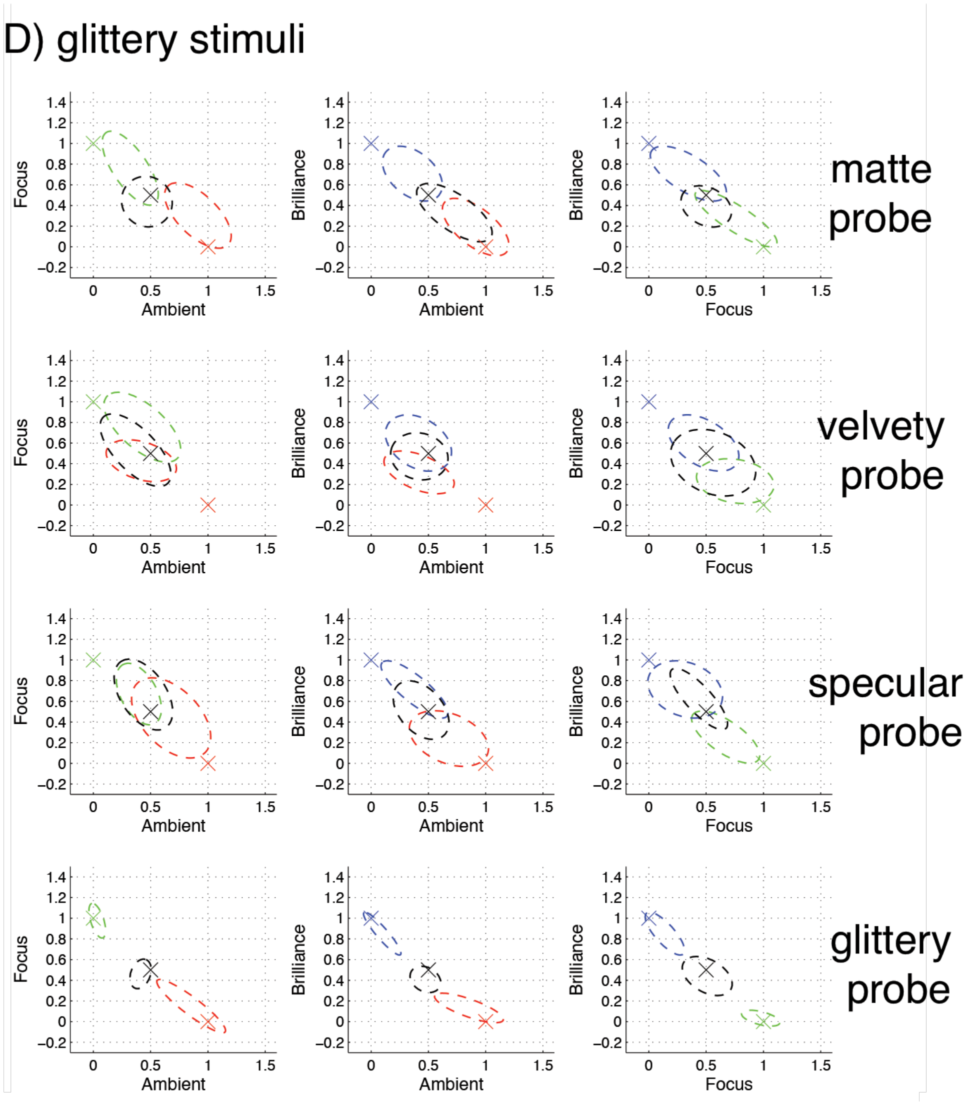

Figure 4.18: The bivariation plots for combinations of canonical material modes in Experiment 1 with D) glittery mode in the stimuli. The subplots in each row are results from one of the four material modes in the probe, specifically matte, velvety, specular, and glittery from top to bottom, respectively. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Specifically, the red color corresponds to the stimuli in which only ambient light was presented, the green color corresponds to the stimuli in which only focus light was presented, the blue color corresponds to the stimuli in which only brilliance light was presented, and the black color corresponds to the stimuli when two lighting modes were optically mixed (each 50% in the mixture). The ellipses represent one standard deviation of bivariate normal distributions fitted to the data.

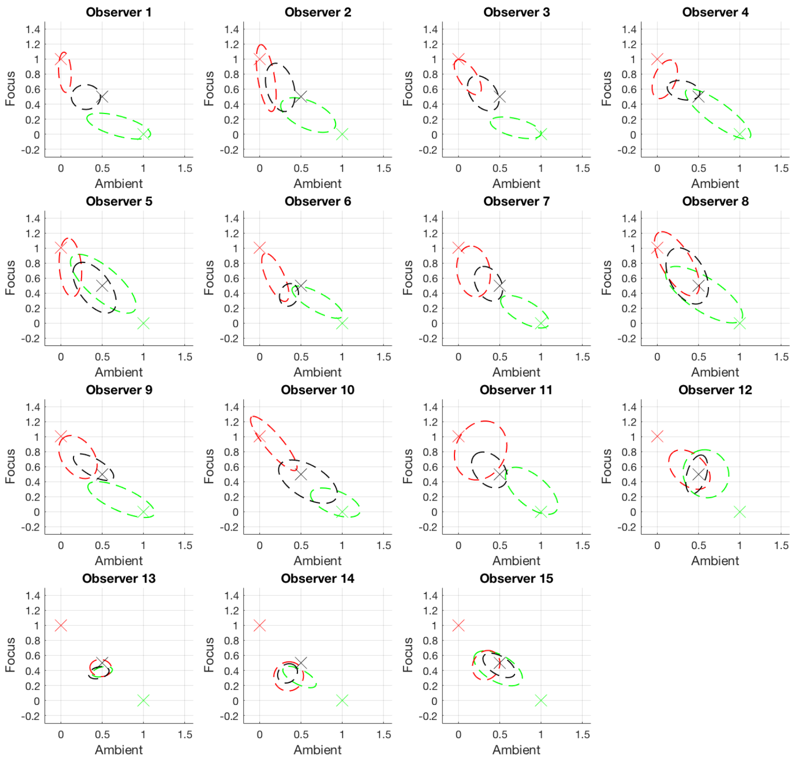

Figure 4.19: The bivariation plots for combination of ambient (green) and focus (red) lighting modes for each observer in Experiment 1. The subplots in each row are results for different lighting combinations. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Observer 1 to 4 are experienced observers. Observer 5 to 15 are inexperienced observers. Except for the last four observers, all performed well above chance if evaluated by the ratio r (Figure 4.4).

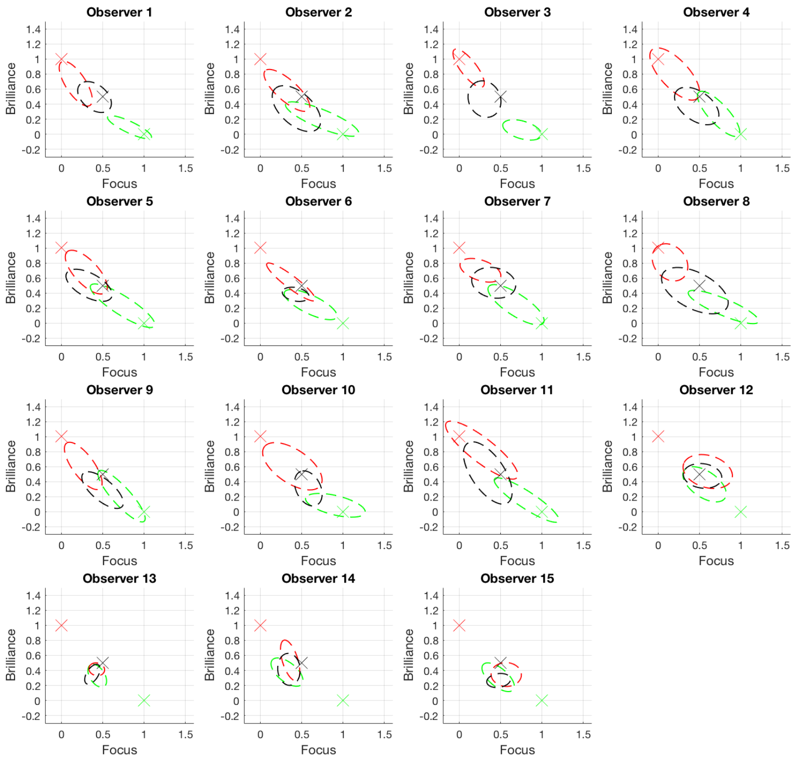

Figure 4.20: The bivariation plots for combination of focus (green) and brilliance (red) lighting modes for each observer in Experiment 1. The subplots in each row are results for different lighting combinations. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Observer 1 to 4 are experienced observers. Observer 5 to 15 are inexperienced observers. Except for the last four observers, all performed well above chance if evaluated by the ratio r (Figure 4.4).

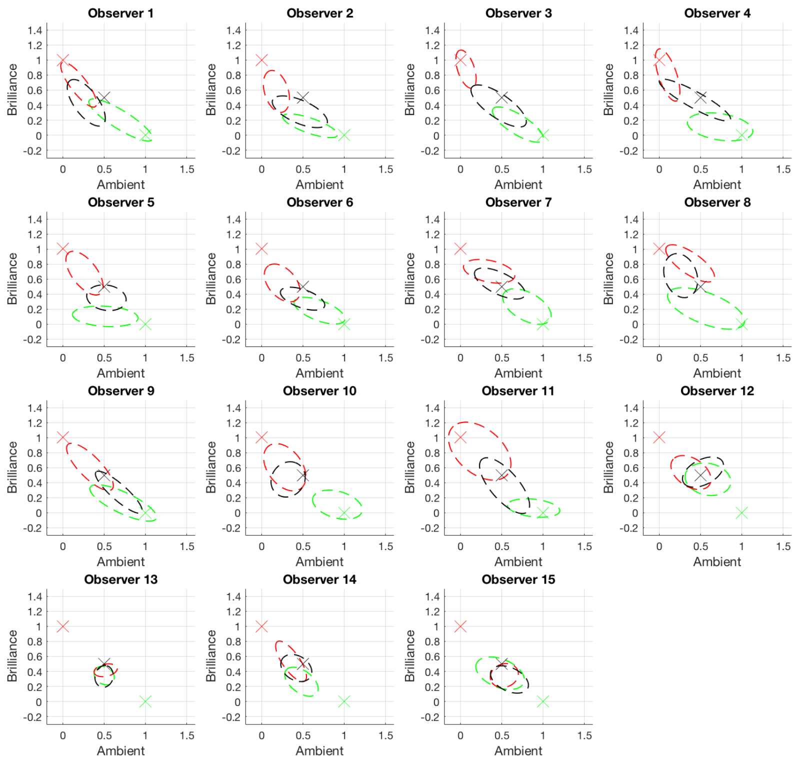

Figure 4.21: The bivariation plots for combination of ambient (green) and brilliance (red) lighting modes for each observer in Experiment 1. The subplots are results for different observer. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. Observer 1 to 4 are experienced observers. Observer 5 to 15 are inexperienced observers. Except for the last four observers, all performed well above chance if evaluated by the ratio r (Figure 4.4).

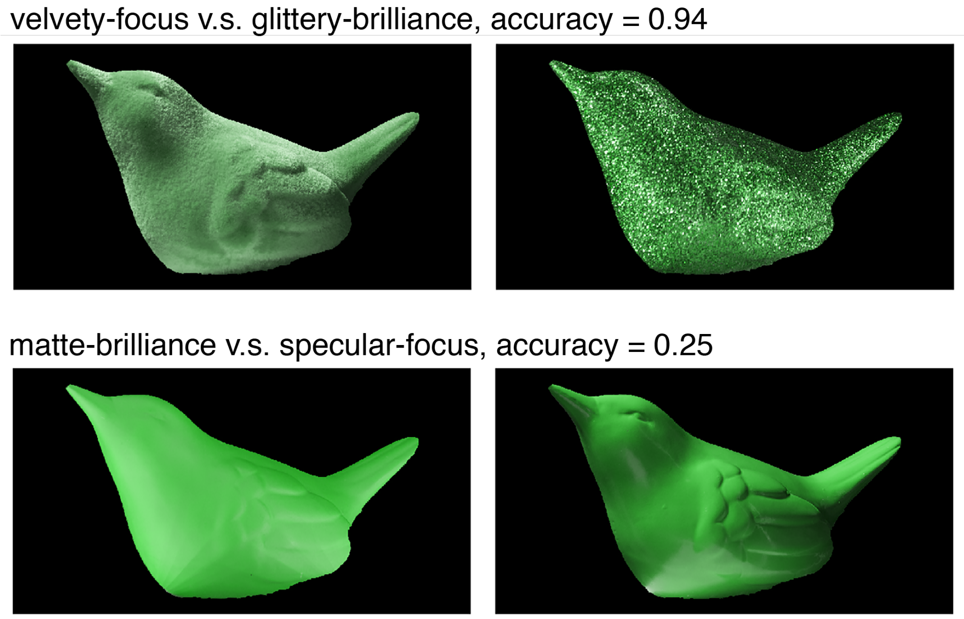

Figure 4.22: The examples of the pairs of stimulus images in the "different materials different lightings" category for which the observers were the most accurate (top) and the least accurate (bottom).

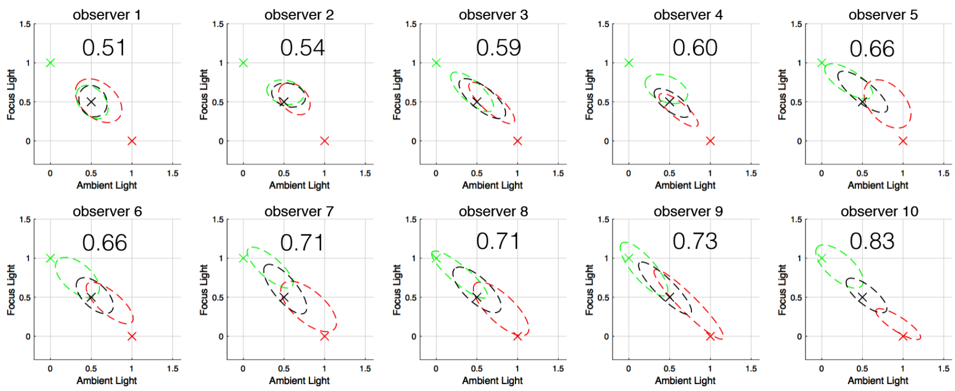

Figure 4.23: The bivariation plots for combination of ambient (red) and focus (green) lighting modes for each observer in Experiment 3, ordered according to performance measure r. The subplots are results for different observers. Different colors correspond to different lighting weight combinations in the stimuli, which are depicted by the crosses. The ellipses represent one standard deviation of bivariate normal distributions fitted to the data. The number in each subplot represents the ratio r for each observer (chance level 0.5).

5

A SYSTEMATIC APPROACH TO TESTING AND PREDICTING LIGHT-MATERIAL INTERACTIONS

Abstract

Photographers and lighting designers set up lighting environments that best depict objects and human figures to convey key aspects of the visual appearance of various materials, following rules drawn from experience. Understanding which lighting environment is best adapted to convey which key aspects of materials is an important question in the field of human vision. The endless range of natural materials and lighting environments poses a major problem in this respect. Here we present a systematic approach to make this problem tractable for lighting–material interactions, using optics-based models composed of canonical lighting and material modes. In two psychophysical experiments, different groups of inexperienced observers judged the material qualities of the objects depicted in the stimulus images. In the first experiment, we took photographs of real objects as stimuli under canonical lightings. In a second experiment, we selected three generic natural lighting environments on the basis of their predicted lighting effects and made computer renderings of the objects. The selected natural lighting environments have characteristics similar to the canonical lightings, as computed using a spherical harmonic analysis. Results from the two experiments correlate strongly, showing (a) how canonical material and lighting modes associate with perceived material qualities; and (b) which lighting is best adapted to evoke perceived material qualities, such as softness, smoothness, and glossiness. Our results demonstrate that a system of canonical modes spanning the natural range of lighting and materials provides a good basis to study lighting–material interactions in their full natural ecology.

Published as: Zhang, F., de Ridder, H., Barla, P., & Pont, S. (2019) A systematic approach to testing and predicting light-material interactions. *Journal of vision*,19(4), 1-22. doi:10.1167/19.4.11

5.1. INTRODUCTION

Lighting is a vital part of several artistic and industrial activities including photography, cinematography, theater, architecture, and design. Taking the example of photography, there are several common practices in setting up lighting environments that permit clear depiction of certain aspects of the shape or material of an object or subject. This is testified by the ever-growing list of tutorials found online, as well as in dedicated books on the topic (Hunter, Biver, & Fuqua, 2015). For example, key light sources may be employed to create highlights or body shadows, bringing out aspects such as glossiness and body shape, and fill lights are used to reduce shad(ow)ing, bringing out colors and delineating shape contours. Material appearance can be modulated by a judicious use of lighting tools; for example using hard-edged spot lights or so-called soft boxes at frontal, oblique, or grazing angles, and using white reflectors and/or black screens. As shown in Figure 5.1, lighting can render certain material qualities visible or invisible. The lighting techniques that permit to achieve such visual effects are usually learned from experience, with specific mechanisms applying to specific cases. The tutorials and books on this topic describe such specific cases, but lack a generic system to predict lighting–material interactions for any material. Here we propose and test an approach to develop such a system.

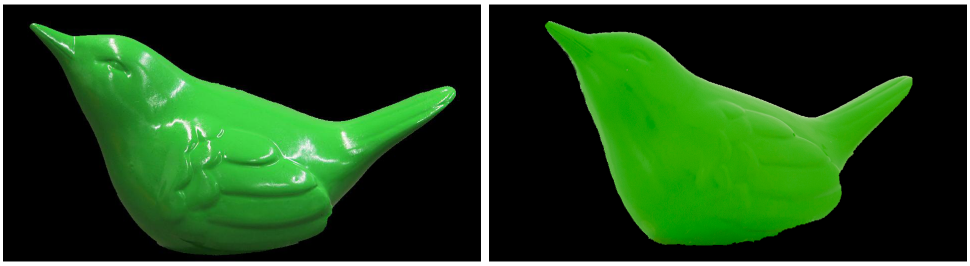

Figure 5.1: Photos of the same glossy bird-shaped object under two different illuminations. Left: a common office lighting environment; Right: a canonical ambient lighting environment. The object appears to be relatively smoother, harder, and glossier on the left.

An important goal in the field of human vision is to understand how lighting may affect appearance in generic scenes. In this paper, the specific aim is to understand and predict how lighting systematically influences material perception. We restrict our study to opaque materials and neglect texture. In the material perception literature, perceptual interactions between materials and lightings were previously reported by several studies, especially for matte versus glossy materials. For example, smooth surfaces appear glossier under collimated light sources than under broad diffuse light sources (Dror, Willsky, & Adelson, 2004; Pont & te Pas, 2006); different natural lighting environments have been shown to affect perceived glossiness to different amounts (Fleming, Dror, & Adelson, 2003; Doerschner, Boyaci, & Maloney, 2010; Olkkonen & Brainard, 2010; Marlow, Kim, & Anderson, 2012). More recently, Motoyoshi and Matoba (2012) found that changes in the contrast and gamma of the illumination affect perceived glossiness, and the orientation of highlights and the shape of highlights (due to differently shaped light sources) was found to affect the perception of gloss (Marlow, Kim, & Anderson, 2011; van Assen, Wijntjes, & Pont, 2016). Human judgments of qualities such as glossy, smooth, or soft have been found to be systematically related to material classes (Fleming, Wiebel, & Gegenfurtner, 2013) and lighting (Barati, Karana, Sekulovski, & Pont, 2017).

The main goal of this paper is to systematically study and predict the effects of a large variety of lighting environments on the perception of several qualities for a large range of materials. However, the range of naturally occurring lighting environments and materials seems to be endless, even if we restrict ourselves to opaque materials and neglect texture. In order to make this problem tractable, we propose an approach on the basis of canonical modes; that is, stereotypical representations of the basic components of naturally occurring light and materials. Here we use three lighting and four material modes, following our previous work on the topic (Zhang, de Ridder, & Pont, 2015; Barati et al., 2017). The modes are based on an optical model describing natural light fields (Mury, Pont, & Koenderink, 2007) and several optical models describing the bidirectional reflectance distribution functions (BRDFs) of opaque materials (Ward, 1992; Koenderink & Pont, 2003; Barati et al., 2017). The modes can occur in isolation but can also be linearly superposed in order to create generic luminous environments and materials, analogous to how semiglossy materials can be rendered using linear combinations of matte and specular reflectance components. Specifically, for the lighting we consider an *ambient light mode* consisting of a spherically diffuse light environment, a *focus light mode* represented by a collimated light source, and a *brilliance light mode* in the form of a large number of small light sources. These modes represent the zero order, first order, and higher order contributions of a spherical harmonic decomposition of the local light field (Mury, Pont, & Koenderink, 2007), have a physical and perceptual meaning (Pont, 2009), and correspond to the basic layers that are used in perception-based lighting design (Kelly, 1952; Ganslandt & Hofmann, 1992; Pont, 2009). For the materials, we covered smooth objects with four different types of finishes: matte paint, glossy paint, a velvet-like (or flocked) layer, and a glittery layer, representing, respectively, a constant

BRDF (diffuse scattering), a peaked BRDF in the forward (mirror) direction, a BRDF that "explodes" along the surface (asperity scattering), and a broadened noisy BRDF (specular multifacet scattering). We denote these material modes by the terms *matte*, *specular*, *velvety*, and *glittery*, respectively.

The hypothesis we want to test is whether some characteristics of lighting are more amenable to convey material qualities than others, and if yes, which ones? In one of our previous studies (Zhang et al., 2015), we found that the brilliance lighting could make certain material appear glossier but less velvety than the focus lighting. In our former research, we also found that for the modes we used, the matching and comparison results were more robust for materials than for lighting. In that study, we did not test which qualities were perceived for our lighting–material combinations. So the first goal of the current study is to test a large range of material qualities and how they are brought out by specific lighting–material combinations. We take on a systematic approach, both with real and synthetic stimuli, to determine whether our canonical modes can help understand and predict the effects.

In Experiment 1, we test how our stimuli map onto the space of perceived material qualities. To this end, we conducted a rating task using real stimuli photographed under three canonical lighting modes. In order to test whether these observations remain valid for natural lighting environments, we conducted Experiment 2 to model and validate the canonical approach. First, we have quantitatively analyzed environment maps (Debevec, 1998) from the University of Southern California (USC) database (http://gl.ict.usc.edu/Data/HighResProbes/; accessed October 23, 2015) and selected three candidate lightings that best represent the canonical ones. The selection was done on the basis of the power of the components of their spherical harmonic decompositions. The canonical material modes were similar to the ones used in an earlier study in which we compared real to rendered materials (Zhang, de Ridder, Fleming, & Pont, 2016). We then rendered objects of the four canonical material reflectance modes in the three selected lighting environments that best approach the canonical lighting modes. In Experiment 2 we then conducted the same rating experiment as in Experiment 1 using rendered stimuli and selected natural lighting environment maps. The main motivation to conduct Experiment 2 was to model and simulate the canonical modes, and see if we could validate the canonical approach for generic illumination environments. This is interesting for both perception and computer graphics studies. After all, it is easier to render stimuli than photograph real materials under controlled lighting environments.

First, we repeated the analysis of Experiment 1 for Experiment 2 to test how the combinations of modes map onto the data space of perceived material qualities. Next, we tested our predictions that the selected generic natural lighting environments would have similar effects depending on the material mode, as in Experiment 1, via detailed comparisons of the results between the two experiments. These detailed comparisons reveal strong correlations, confirming that lighting–material effects are systematic and can be predicted based on a coarse categorization into canonical modes. In the discussion, we address limitations and future challenges of this systematic approach.

5.2. METHOD

5.2.1. REAL STIMULI: EXPERIMENT 1

In Experiment 1, the stimuli were the cropped images processed from the photos taken from multiple copies of the same real bird object under three canonical lighting modes. The objects were covered with matte, velvety, specular, and glittery finishes to represent the four canonical material modes (Figure 5.2). For the ambient light, we put both the birds and the camera in a white photo tent and illuminated the photo tent with fluorescent tube lamps from the ceiling. For the focus light, we illuminated the birds with a halogen lamp from upper left. For the brilliance light, we surrounded the birds with 150 LED lights. The same images were previously used as stimuli in a matching experiment (Zhang et al., 2015) to study the lighting–material interactions. All photographs were taken using a Canon EOS 400D DIGITAL camera (focal length 57 mm) in controlled laboratory environments, then edited to find the shared contour and make the background black (resolution 954 × 512 pixels). The luminance was photometrically calibrated to be linear.

5.2.2. RENDERED STIMULI: EXPERIMENT 2

ILLUMINATION ENVIRONMENTS

In order to validate our canonical modes approach, we extended it to generic natural lighting environments and tested whether this allows for coarse predictions of the lighting effects on material appearance. Thus, we took a small database of natural illumination maps and then aimed to recreate the effects of Experiment 1. From the USC's high resolution recreations of the Debevec's light probe images we selected three maps

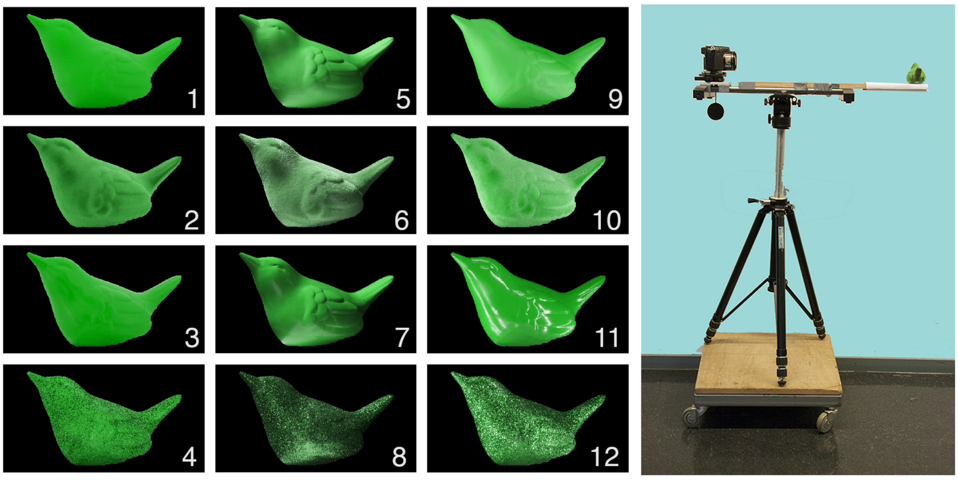

Figure 5.2: Left: The stimuli of Experiment 1. From top to bottom, the four rows represent the matte, velvety, specular, and glittery material modes. From left to right, the three columns represent the ambient, focus, and brilliance lighting modes. Right: the photography setup, which kept the relative position of the object and the camera fixed when switching (riding) between lighting environments. All photographs were taken (using a Canon EOS 400D DIGITAL camera, focal length 57 mm) in controlled laboratory environments, edited to find the shared contour and make the background black (resolution 954×512). The luminance was photometrically calibrated to be linear.

that optically should be the best representatives of the ambient light, focus light, and brilliance light in this database. First, we reconstructed the spherical function of the illumination environment $f(\theta,\phi)$ by the sum of its spherical harmonics (SH):

$$f(\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} C_l^m \cdot Y_l^m(\theta,\phi), \tag{5.1}$$

where C_l^m are the coefficients, Y_l^m are the basis functions, and l is the order of the angular mode varying from zero to infinity. The power of the l-th order, denoted as d_l, physically characterizes the angular distributions of the illumination at order l, which is orientation-invariant and can be calculated as (Stock & Siegel, 1996):

$$d_l = \sqrt{\sum_{m=-l}^{l} (C_l^m)^2} \tag{5.2}$$

Using Xia's diffuseness metric (Xia, Pont, & Heynderickx, 2017) we could calculate the diffuseness of all high-resolution high dynamic range (HDR) maps, via the ratio of the power of the first order $d_{l=1}$ to the power of the zeroth order $d_{l=0}$ of the SH decomposition. By normalization, diffuseness scores range between 0 (zero), which corresponds to extremely directed lighting (i.e., focus lighting) to 1, which corresponds to fully diffuse lighting (i.e., ambient lighting), as in the equation below:

$$(D_{xia})_{normalized} = 1 - d_{l=1}/d_{l=0}/\sqrt{3}. \tag{5.3}$$

As a result, the *Glacier* environment map scored the highest in Xia's diffuseness metric and was being selected as the representative map for the ambient light; the Ennis environment map scored the lowest in Xia's diffuseness metric and thus, was selected as the representative map for the focus light (Figure 5.3, on the left).

5

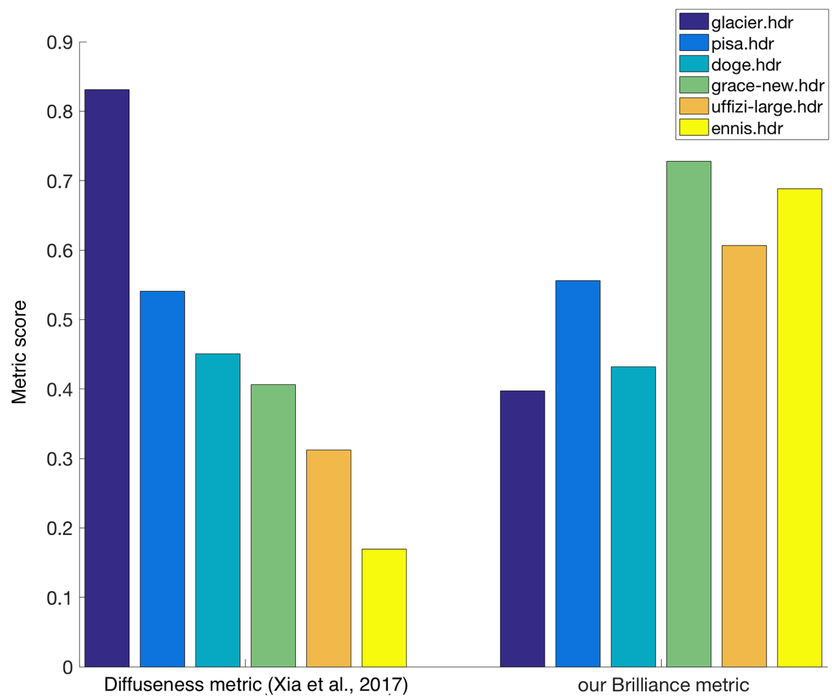

Figure 5.3: Metrics using the spherical harmonics decomposition for the USC high-resolution HDR maps. Left: results of the diffuseness metric (Xia et al., 2017) used to select the Glacier map for the ambient lighting and the Ennis map for the focus lighting. Right: results of the brilliance metric sampled up to the 10th order.

To select the Debevec environment map that could best represent the brilliance lighting mode, we propose a brilliance metric, calculating the ratio between the sum of the higher orders statistics ($l \geq 3$) to the sum of all orders:

$$B = \frac{\sum_{l=3}^{\infty} d_l}{\sum_{l=0}^{\infty} d_l} \tag{5.4}$$

The result of the brilliance metric B will vary from 0 (no brilliance at all) to 1 (pure brilliance). In practice, we implemented finite sampling and found that the brilliance metric gave robust results beyond the 10th order. As shown in Figure 5.3, the *Grace-new* lighting scored the highest in our brilliance metric when calculating up to the 10th order or higher, and thus was selected as the representative map for the brilliance lighting. The selected maps can be seen in the top row of Figure 5.4.

The illumination environments are provided as HDR panoramic images. Since their main lighting directions occur in different locations than our canonical focus lighting, we had to manually adjust their orientations such that their light directions in the first order spherical harmonics were matched. They were adjusted per type of lighting by applying a rotation of θ radians around the horizontal camera axis, followed by a rotation of ϕ radians around the vertical camera axis. In particular, the brilliance lighting had to be rotated around the horizontal camera axis. For the values, see Table 5.1. Each illumination environment was then converted from RGB to gray values (i.e., the relative luminance) using the formula: $0.2126 \times R + 0.7152 \times G + 0.0722 \times B$ (Stokes, Anderson, Chandrasekar, & Motta, 2012).

	θ	ϕ
L1: Glacier (ambient)	-0.170	-1.796
L2: Ennis (focus)	-0.170	1.413
L3: Grace-new (brilliance)	-0.961	-0.615

Table 5.1: The orientation parameters per illumination environment.

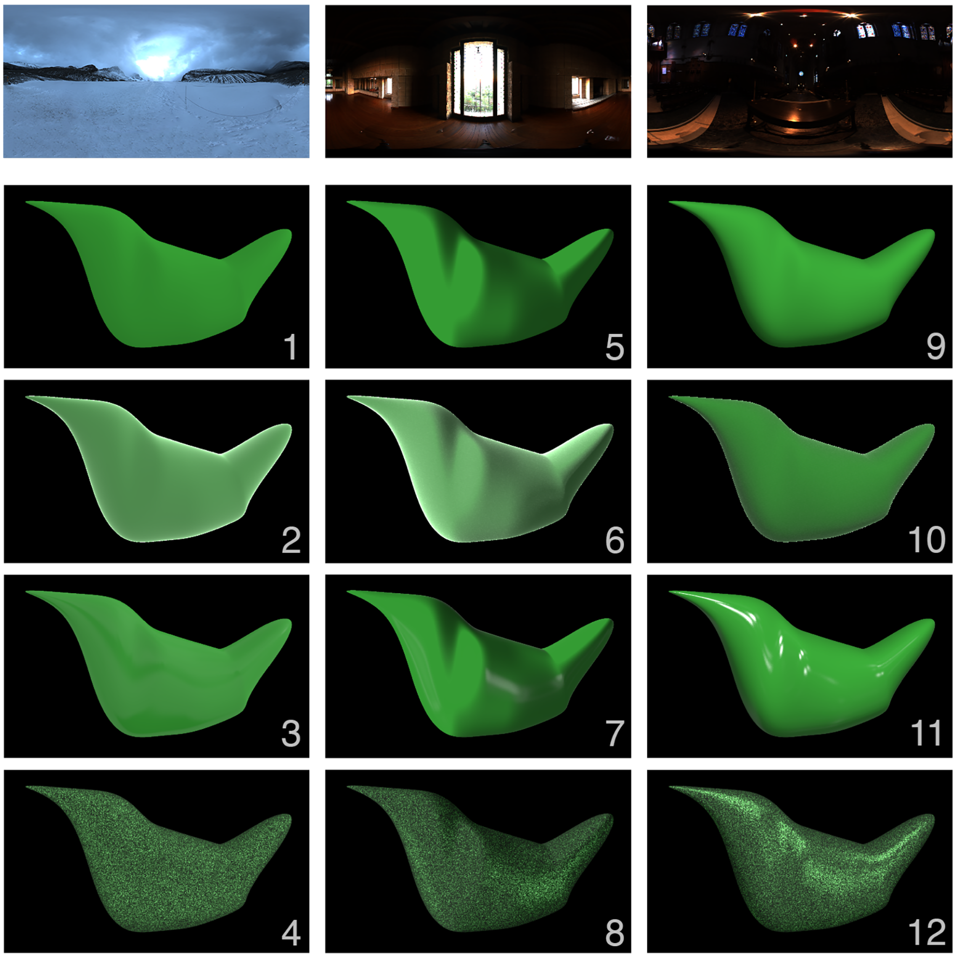

Figure 5.4: The stimuli of Experiment 2. From top to bottom, the first row shows the three lighting maps we selected to be the best representatives for the canonical lighting modes of the high resolution USC database. The second to the last rows represent the rendered matte, velvety, specular, and glittery material modes. From left to right, the three columns represent the materials under the Glacier, the Ennis, and the Grace-new illumination environments, respectively, as visualized in the first row. Compared to the initial Glacier environment, our version has been modified by filling in the black region originally found at the bottom of the image (which is due to the tripod base). The gammas of the images for the glittery mode (last row) were adjusted from 1.0 to 1.8 in order to make the features more visible in the printed version.

MATERIAL MODES

We considered four different material modes: matte, glossy, glittery, and velvety. Each mode may be described by its BRDF $f_r = (\omega_i, \omega_o)$, where ω_i represents an incoming di-

Mode	Parameters		
Matte	$\rho_d = 0.687$		
Glossy	$\rho_d = 0.687$	$\rho_s = 0.067$	$\alpha_s = 0.037$
Glittery	$\rho_d = 0.5$	$\rho_s = 0.9^{\star}$	$\alpha_s = 0.1$
Velvety	$\rho_d = 0.4$	$\rho_v = 0.4$	

Table 5.2: Parameters used in each mode. The $\star$ symbol indicates a surface-varying parameter.

rection vector (e.g., from a light source) and ω_o represents an outgoing direction vector (e.g., toward the viewpoint). Note that both vectors ω_i and ω_o are defined by two angles. In computer graphics, f_r outputs three values, one for each of the RGB color channels. Table 5.2 provides the parameters used in these BRDFs. They are explained in the following list:

1. The *matte* material mode was implemented with a Lambertian BRDF; that is,

$$f_r(\omega_i, \omega_o) = \frac{\rho_d C_d}{\pi} \tag{5.5}$$

where $\rho_d \in [0,1]$ controls the intensity of the Lambertian reflectance, and C_d is a color that we set to $\{24, 253, 22\}$ in RGB to yield a greenish tint resembling the stimuli of Experiment 1. All other material modes exhibit a diffuse component with the same color, but a potentially different intensity as detailed in Table 5.2.

2. For the *velvety* material mode, we used the asperity scattering BRDF model of Koenderink and Pont (2003), to which we also added a diffuse term as before, yielding:

$$f_r(\omega_i, \omega_o) = \frac{\rho_d C_d}{\pi} + \frac{\rho_v}{2\cos\theta_i \cos\theta_o}, \tag{5.6}$$

where ρ_v controls the intensity of the velvety component.

3. The *specular* material mode was implemented with an isotropic Ward BRDF (Ward, 1992):

$$f_r(\omega_i, \omega_o) = \frac{\rho_d C_d}{\pi} + \frac{\rho_s}{4\pi\alpha_s^2\sqrt{\cos\theta_i \cos\theta_o}} e^{-\frac{\tan^2\theta_h}{\alpha_s^2}}, \tag{5.7}$$

where θ_i, θ_o, and θ_h denote angles made by either ω_i, ω_o or $\mathbf{h} = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}$ with the surface normal. The $\rho_s \in [0,1]$ parameter controls the intensity of glossy reflection, while $\alpha_s \in [0, 0.2]$ controls the roughness of the material (higher values meaning blurrier reflections).

4. The *glittery* material mode: we simulated the broadened and noisy forward scattering of the glittery material by mimicking the occurrence of multifacet flakes at the surface of the object. The first step in our mimicking procedure was to use the BRDF of the specular material mode, varying the ρ_s parameter. To this end, we used an isotropic Gabor noise (Lagae, Lefebvre, Drettakis, & Dutré, 2009), which permits the production of an even distribution of flakes as shown in the left part of Figure 5.5. We employed the implementation of Lagae and Drettakis (2011), using a tangent space distribution of 150 impulses with frequency of 0.157, bandwidth of 2.724, and truncated to 0.013. The sparsity of the distribution is then adjusted by remapping noise values $\mathcal{N}$ to $\rho_s = 0.9 \times |(1.1 * \texttt{smoothstep}(\mathcal{N}))^8|_0^1$, where $|\cdot|_0^1$ clamps values between 0 and 1; this is shown in the right part of Figure 5.5. Exponentiation permits selection of the brightest noise values, the 1.1 factor slightly saturates them to obtain apparent flakes of varying size, while the GLSL `smoothstep` function softens the selection.

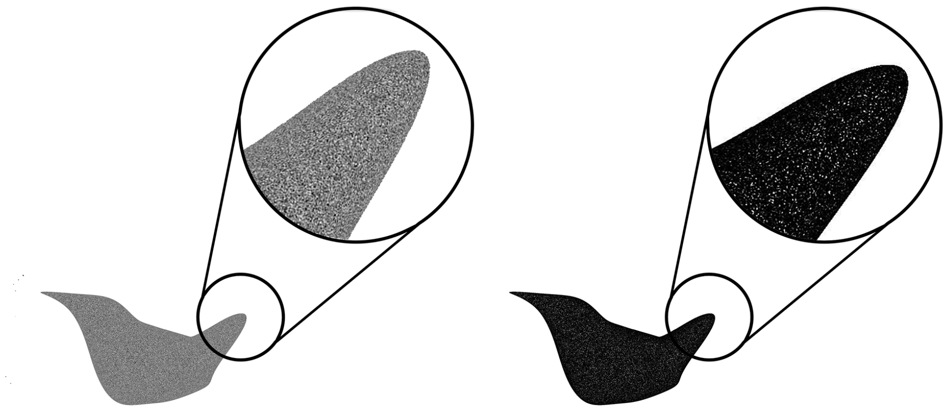

Figure 5.5: The sparsity of the distribution of the flakes for glittery effects, with zoomed view for the tails. Left: output of the Gabor noise function $\mathcal{N}$. Right: remapping of $\mathcal{N}$ values to yield a sparser distribution of apparent flakes.

The second step in the mimicking procedure was to additionally modify the surface normals by a different noise function to take into account the slight variations of flake orientations with respect to the object surface. This time we used a value noise based on the position of each surface point in 3D, with an amplitude of 1.0, a frequency of 0.157

(same as before), a persistence of 0.8, and four octaves. It is used to perturb the direction of the normal around the local geometric normal.

RENDERING PROCESS

The 3D modeling of the bird shape was created in Blender (Blender 2.79b; Blender Foundation, Amsterdam, the Netherlands), as a simplified version of the shape of the real objects, and is the same as the 3D model we used in previous work (Zhang et al., 2016). The black region at the bottom of the Glacier environment map (due to the tripod base) was filled in to make it more ambient and natural (shown in Figure 5.4). The filled-in version of the Glacier environment map can be found in the supplementary materials (Supplementary Figure 5.19). To render the stimuli of Experiment 2 we used Gratin version 0.3 for Apple Mac OS (Vergne & Barla, 2015) with rendering made using OpenGL shading language (GLSL) version 410.

Rendering was performed ignoring shadowing and interreflection effects. The diffuse component may then be equivalently represented using a diffuse-filtered version of the illumination environment, as provided on the USC website (diffuse convolution links). Rendering the diffuse component then simply amounts to look up the diffuse-filtered environment in the direction of the surface normal. We used this pre-filtered approach as it completely removes noise coming from the rendering process of the diffuse component. For the other components, we employed Monte Carlo integration, using importance sampling of the Ward material (i.e., for the specular and glittery modes) to speed up convergence. The rendering results (i.e., the stimuli for Experiment 2) can be found in Figure 5.4.

5.2.3. PROCEDURE

The same procedure was used in both Experiments 1 and 2. For each observer, a list of nine qualities was first explained before the experiment started. The list consisted of the four names of our canonical material modes, namely matte, velvety, specular, and glittery; and five terms that are often used in material perception studies, namely hard, soft, rough, smooth, and glossy (Fleming et al., 2013; Barati et al., 2017). All stimuli were then shown twice in random order to give the observers an idea about how different the stimuli are from each other, so that they could use the whole scale during the following rating procedure. In each trial, the observers had to answer two questions, the first one being a yes or no question: "Is the bird [. . .]?", with the [. . .] displaying one of the nine qualities?"

If they answered Yes, they would then have to answer the second question, "how [...] is the bird on a scale from 1 to 7?" by moving the cursor on a slider. If they answered No to the first question, they would skip the second question and jump to the next trial. The observers were explicitly instructed that they did not have to balance the answers of the yes or no questions. With 12 stimuli, nine material quality terms, and three repetitions, there were 324 trials per observer per experiment. The trials were presented in nine blocks based on the qualities. The nine blocks were presented in a randomized order across observers. The interface was developed with the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007) in MATLAB R2016b (MathWorks, Natic, MA), and presented on a linearly calibrated Apple Inc. 15-in. Retina display, with a resolution of 1440 × 900 pixels, ranging from 0.11 cd/m^2 to 75cd/m^2. The stimulus was presented as 954 × 512 pixels in the middle of the screen. The viewing distance between the observer and the screen was about 0.5 m and kept constant.

5.2.4. OBSERVERS

Fifteen paid observers participated in Experiment 1. A different group of 12 paid observers participated in Experiment 2. All participants had normal or corrected-to-normal vision, and were inexperienced in psychophysical experiments. Participants read and signed a consent form before the experiments were conducted. The study was approved by the Human Research Ethics Committee at Delft University of Technology, and conducted in accordance with the Declaration of Helsinki and Dutch law.

5.2.5. ANALYSIS

To investigate how our canonical modes map onto the space of perceived material qualities, for both Experiment 1(real stimuli) and 2 (rendered stimuli) we analyzed the Yes or No (Y/N) data to see (a) whether the four names (matte, velvety, specular, and glittery) agree with the corresponding material modes, and (b) how combinations of canonical material and lighting modes associate with perceived qualities. To answer the first question, the raw data of the Y/N results for the four names are presented as percentage of answering Yes for each material mode. To answer the second question, we performed a correspondence analysis (CA) on the Y/N data for all qualities to further present the associations between the modes and qualities. Subsequently, we performed a principle component analysis (PCA) on the rating data to further explore the data space of perceived material qualities and how our real and rendered stimuli are positioned in that space. Finally, we compared the raw rating data and the PCA data space of Experiment 1 with those of Experiment 2 to examine how well the renderings correlate with the real objects, and to look into how lightings evoke material dependent effects for material modes.

5.3. EXPERIMENT 1

5.3.1. ANALYSIS AND RESULTS

YES/NO DATA

The first issue we wanted to look into was whether the names matte, velvety, specular, and glittery for the four canonical material modes agrees with the observers' judgments. To answer this, we analyzed the relevant subset of the results from the Y/N questions. In Figure 5.6, the percentage of answering Yes for each quality is shown per material mode (i.e., the fractions of responses for the names matte, velvety, specular, and glittery) averaged across three lightings (for all Y/N data per lighting–material condition, see Figure 5.17), coded as a gray level. Each row represents a quality and each column represents a stimulus material mode. Note that for each row or column, the percentages of the four values do not necessarily add up to 100%. In general, the names were found to be associated with the material modes ($\chi^2(9) = 571, p < 0.01$). The diagonal values show the percentages of the answers for the congruent naming and material modes. It shows that the names velvety, specular, and glittery agree with the corresponding material modes, as those diagonal values are 0.84, 0.76, and 0.99, respectively. Although the name matte was found to agree with the observers' responses to our matte material (0.81), it also applied to our velvety material (0.70), and sometimes also to the specular (0.28) and glittery (0.33) modes. This mostly applied to the stimuli under the ambient light (0.58 for specular mode and 0.53 for glittery mode), sometimes to the stimuli under the focus light (0.27 for specular mode and 0.29 for glittery mode), and less often to the stimuli under the brilliance light (0.18 for the glittery mode and never for the specular mode). Meanwhile, the name specular also sometimes applied to the matte stimuli (0.39), mostly under the focus lighting (0.64), and sometimes to the ambient lighting (0.24) and the brilliance lighting (0.27). These results confirm our previous findings about the interactions between matte and specular material modes (Zhang et al., 2015; Zhang et al., 2016).

We further analyzed the results of the Y/N questions by performing a correspondence

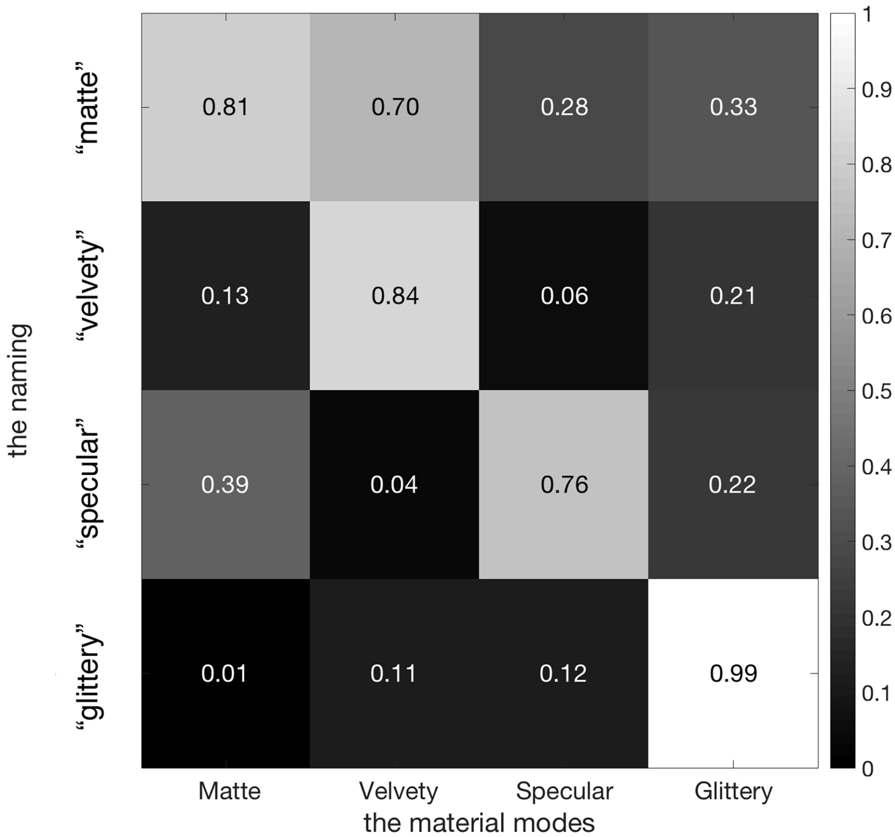

Figure 5.6: The fraction of answering Yes for the four material modes in Experiment 1.

analysis (CA) using all nine qualities. In a 2D correspondence analysis biplot, we visualized the association between the labels based on their proximity (i.e., their distances in that space) and the distinctness based on their distances to the origin. The closer the data points are to each other, the better they associate. The resulting 2D CA biplot of the Y/N questions for Experiment 1 can be found in Figure 5.7. We observe that the name matte is relatively closer to the origin than the other names (velvety, specular, and glittery). The stimuli of the specular and glittery material modes cluster closely around their corresponding names specular and glittery, respectively. Velvety stimuli cluster in the middle of the names matte and velvety, showing that both names could apply to the velvety material mode. Similarly, the matte stimuli cluster in the middle of the names matte and specular, in line with the finding that the name specular also sometimes applied to the matte stimuli (Figure 5.6). These observations suggest that the names velvety, specular, and glittery applied to the corresponding material modes, while for matte, there were larger variations. We also see that the lighting modes have greater influence on the judgments for the specular and glittery materials with brilliance lighting and ambient lighting having the largest and smallest impact, respectively. This effect is virtually absent for the matte and velvety materials. Lastly, note that the space has been rotated 90° anticlockwise and then mirrored around the vertical axis. The resulting y-axis, being the first dimension explaining 53.2% of the variance, shows rough and smooth as opposing qualities while the second dimension (x-axis), explaining 39.3% of the variance, shows soft and hard as opposing qualities. Figure 5.7 also shows that glittery material is primarily associated with rough, specular material with smooth and glossy, velvety material with soft, and matte material with smooth. Hard is close to the origin, meaning that it is associated with almost all stimuli except for velvety material.

RATING DATA

Here we analyze in depth whether certain types of material qualities are actually evoked by certain characteristics of the canonical lighting environments, and if so, to what extent. To answer this, we analyzed the results from the rating sessions by performing a PCA. Since the observers only had to give a rating when they answered Yes, we took all No answers as zero and together with the scale of 1 to 7, we took the medians across repetitions and observers to perform the PCA. The first three principle components (PCs) accounted for 97% of the data variance, with the first and second dimension explaining 52% and 34%, respectively. We therefore believe that a 2D PCA space suffices to visualize the apparent differences between our stimuli. Figure 5.8 shows the resulting data via a 2D PCA biplot, with the data being color-coded for the stimulus material; specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The four corresponding qualities are colored in corresponding colors and the other five qualities are colored in black. The shape marks the lighting mode, specifically ambient light as squares, focus light as circles, and brilliance light as stars. The qualities hard, specular, and glossy load positively on the first principle component PC1, while the qualities soft, matte, and velvety load negatively on PC1. Similarly, the qualities rough and glittery load positively on the second principle component (PC2), while the quality smooth loads negatively on this second component. Furthermore, projecting the stimuli data points onto the quality axes helps us understanding the lighting effects. Per material,

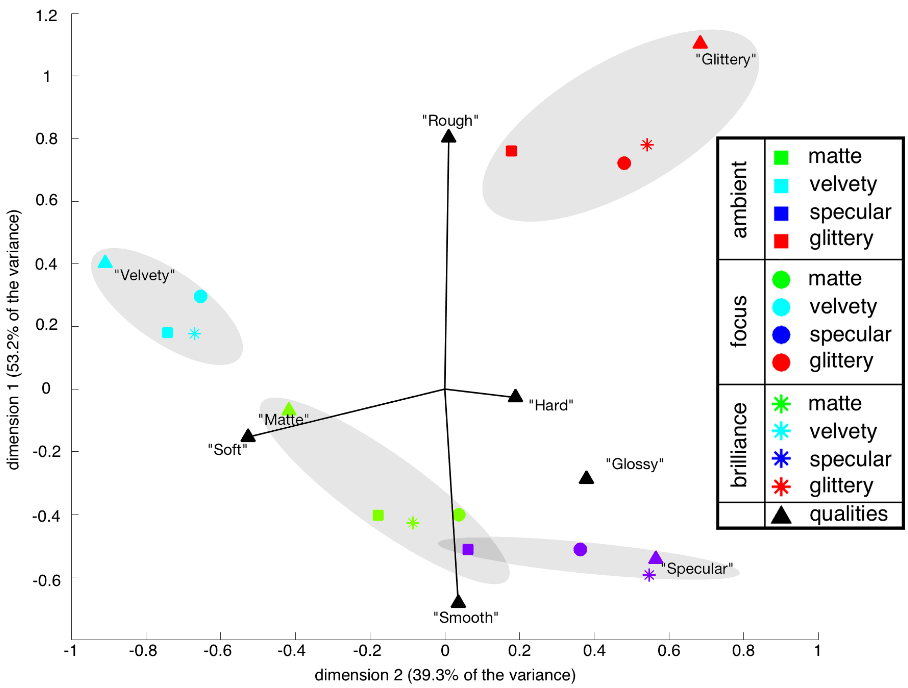

Figure 5.7: Visualization of the correspondence analysis results of the Y/N data for Experiment 1. The first dimension (y-axis) explains 53.2% of the variance. The second dimension (x-axis) explains 39.3% of the variance. The stimuli of the same material mode and their corresponding qualities are colored the same, specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The other five qualities are colored black. The same shape marks the stimuli of the same lighting mode, specifically ambient light as squares, focus light as circles, and brilliance light as stars. The triangles represent the nine qualities that were tested in the Y/N question. The gray ellipses were drawn by hand to show the material clusters. Please note that we connected the origin to soft, hard, rough, and smooth only to show that soft and hard as well as smooth and rough form opposing pairs of qualities and that these two opponent pairs were orthogonal to each other, in line with existing literature (e.g., Fleming et al., 2013). The same was done in Figure 5.10.

the more the projected data points shift along the quality axes, the stronger the change of lighting affects the perception of that quality. For example, the specular stimuli data points (colored in indigo) shift away from the center along the axes of specular, glossy, smooth, and hard as the lighting varies from ambient (square) to focus (circle), and then to brilliance (star). This indicates that the brilliance lighting strongly evoked the quality specular, while ambient light weakened perceived specularity for our specular material mode, confirming earlier findings. Similarly, we found that the ambient light weakened the perception of glittery, hard, and rough for the glittery mode, while it evoked the perception of matte and soft for the matte mode. Other lighting effects were more subtle and will be discussed in detail in the Comparison section hereafter.

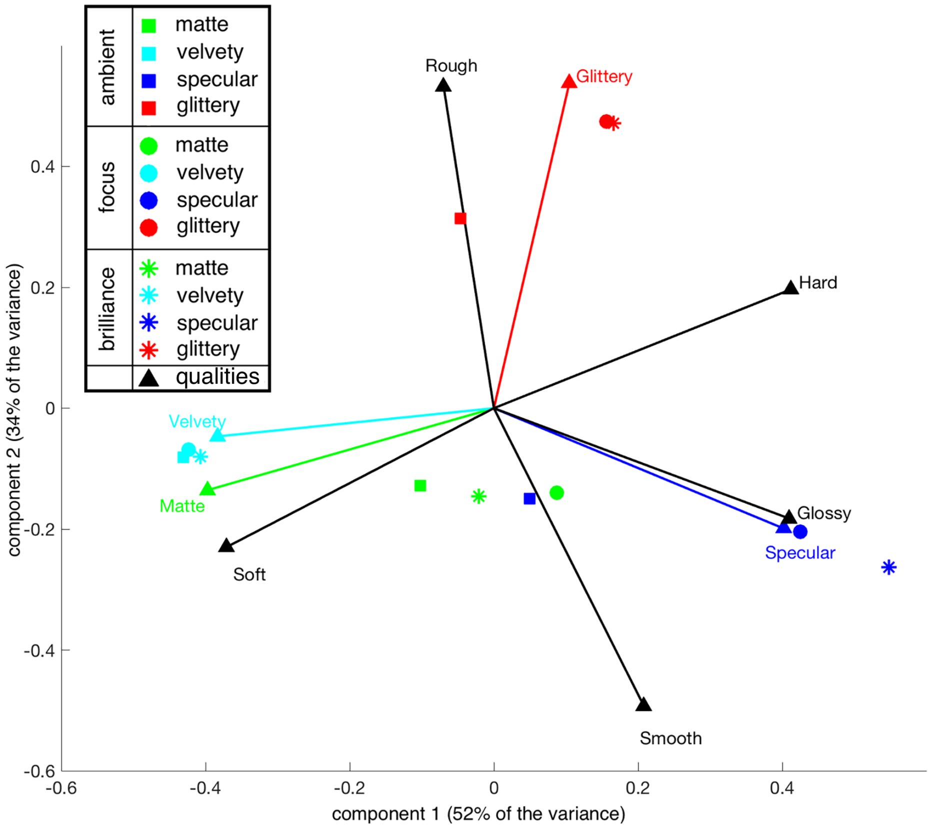

Figure 5.8: Results of the principle component analysis (PCA) for Experiment 1. The PCA was done on the ratings per material quality per stimulus, with the medians determined across all observers and repetitions. The first component (x-axis) explains 52% of the variance. The second component (y-axis) explains 34% of the variance. The stimulus materials and their corresponding qualities are color-coded, specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The other five qualities are colored black. The shapes mark the stimulus lighting, specifically ambient light as squares, focus light as circles, and brilliance light as stars.

5.4. EXPERIMENT 2

5.4.1. ANALYSIS AND RESULTS

The analysis of the data of Experiment 2 was done analogous to and in the same order as that of Experiment 1; that is, we first looked into the Y/N data and then into the rating data.

YES/NO DATA

Figure 5.9 shows the averaged percentage of answering Yes for each material mode quality in the same format that was used in Figure 5.6. The averaging was again done across the three lightings (for all Y/N data per lighting–material condition, see 5.18). Each row represents one quality and each column represents a stimulus material mode, so the diagonal values show the percentages of the answers for the correct naming of the corresponding material modes. Overall, the results are similar to the results of Experiment 1. The names were again found to be associated with the material modes ($\chi^2(9) = 589$, $p < 0.01$). The diagonal values are 0.87, 0.96, 0.76, and 1.00 for matte, velvety, specular, and glittery, respectively. There are two remarkable differences between the two experiments: first, the interaction between matte and velvety increased as the matte material was more often named velvety (0.44) and at the same time, the velvety material was more often named matte (0.91) compared to Experiment 1 (for which we found 0.13 and 0.70, respectively). This confirmed our previous findings about the increased interactions between matte and velvety materials when using computer rendered stimuli (Zhang et al., 2016). Second, the name specular was hardly used for the matte renderings (0.06), in contrast with the frequent naming in Experiment 1 (0.39).

The 2D correspondence analysis results are shown in the biplot in Figure 5.10. Although the lightings were generic natural lighting environments instead of laboratory conditions, we observe quite similar results as in Experiment 1. The first axis, which explains 53% of the variance, shows soft and hard again as opposing qualities. Similarly, the second axis, which explains 40.8% of the variance, shows rough and smooth, also as opposing qualities. Compared to the results from Experiment 1 (Figure 5.7), the only main difference is that the clusters for matte and velvety shifted closer to each other, confirming the increased interactions between the matte and velvety modes.

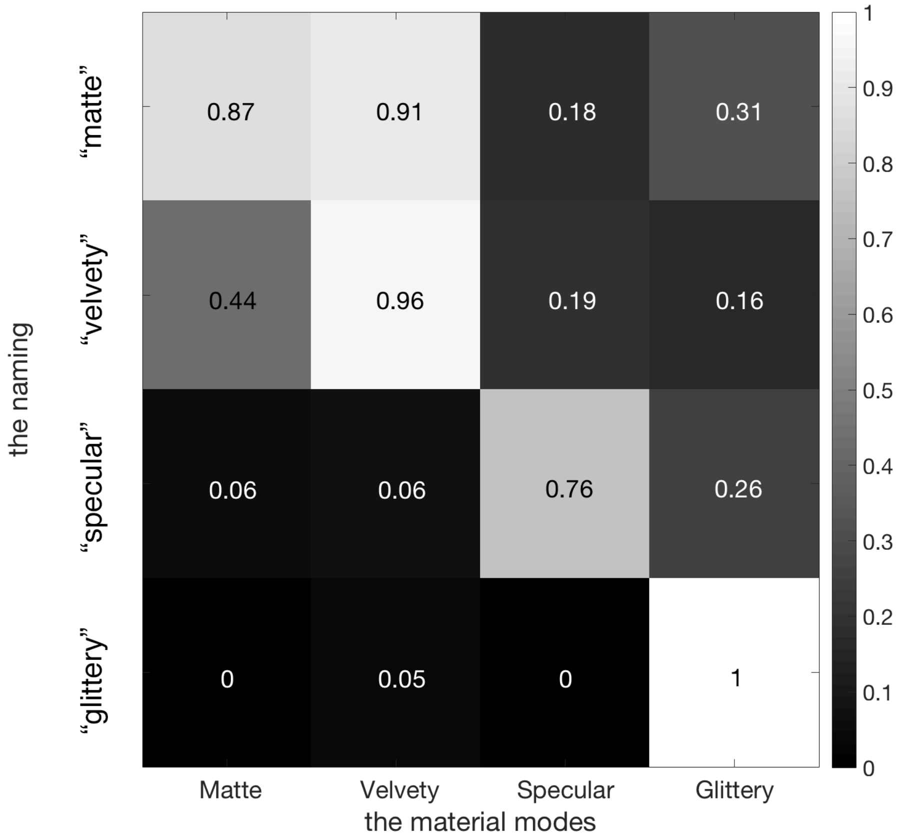

Figure 5.9: The fraction of answering Yes for the four material modes in Experiment 2.

RATING DATA

We also performed a PCA on the medians of the ratings for Experiment 2, in which the medians were calculated across repetitions and observers. The first three PCs accounted for 95% of the data variance, with the first and second components explaining 49% and 35%, respectively. The results are visualized as a 2D biplot in Figure 5.11. The legends are the same as in Experiment 1. Figure 5.11 indicates that, similar to Experiment 1, the qualities hard, specular, and glossy load positively on PC1, while the qualities soft, matte, and velvety load negatively on PC1. The qualities rough and glittery load positively on the PC2, while the qualities smooth, specular, and glossy load negatively on PC2. In line with the observation that the overall space created from the results of Experiment

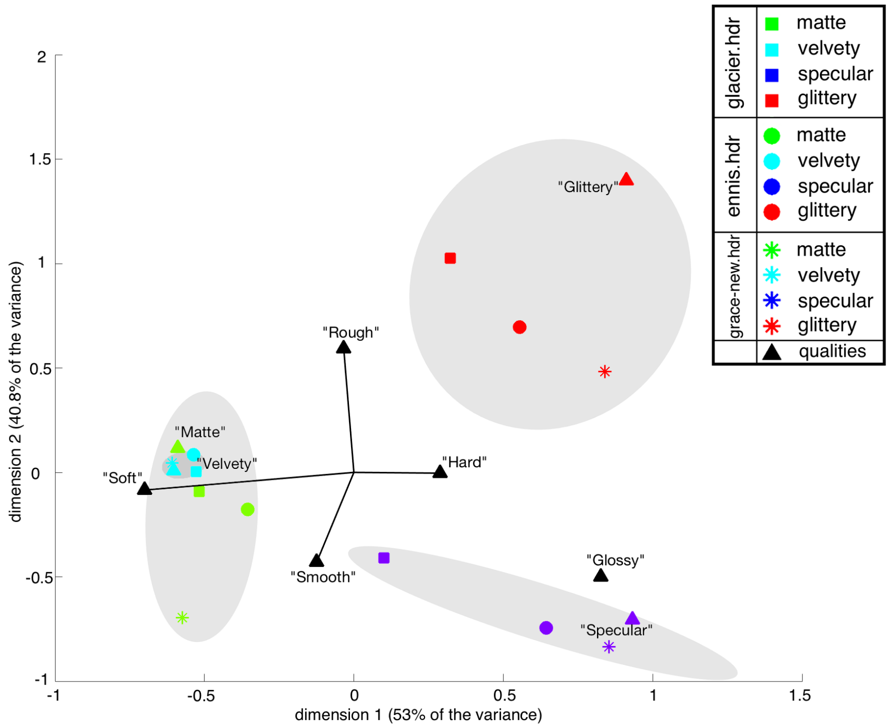

Figure 5.10: The visualization of the correspondence analysis results of the Y/N data for Experiment 2. The first dimension (x-axis) explains 53% of the variance. The second dimension (y-axis) explains 40.8% of the variance. The stimuli of the same material mode and their corresponding qualities are colored the same, specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The same shape marks the stimuli of the same lighting mode, specifically ambient light as squares, focus light as circles, and brilliance light as stars. The triangles represent the nine qualities that were tested in the Y/N question. The gray ellipses were intuitively drawn to show the distance between each stimulus and the quality. As in Figure 5.7, we connected the origin to soft, hard, rough, and smooth only to show that the two opponent pairs were orthogonal to each other.

2 is similar to that of Experiment 1, the specular stimuli again shift away from the center along the specular, glossy, smooth, and hard axes, when the lighting changed from Glacier environment (the most representative light map for the ambient light) to Ennis environment (the most representative light map for the focus light), and then to Grace-new environment (the most representative light map for the brilliance light). The same holds for the glittery stimuli shifting away along the hard axis. Matte and velvety stimuli data points clustered more closely to each other, indicating that they were affected more subtly as the lighting varied. More detailed material-dependent lighting effects will be

discussed in the next section.

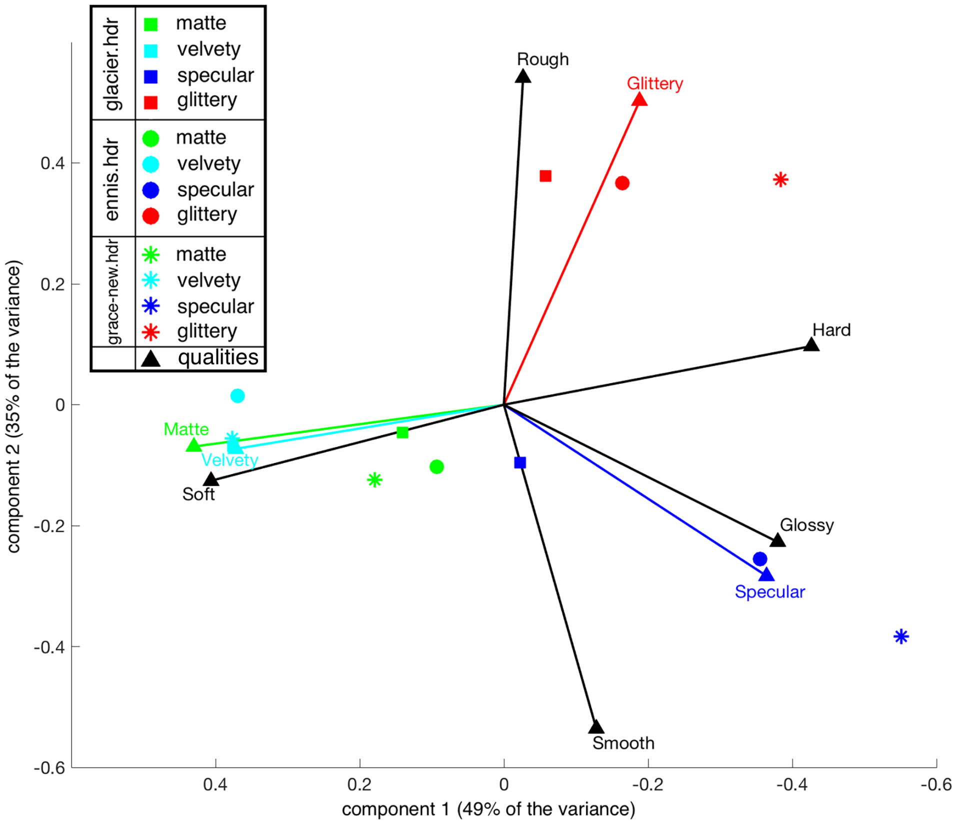

Figure 5.11: Results of the PCA for Experiment 2. The PCA was done on of the ratings per material quality per stimulus, with the medians determined across all observers and repetitions. The first component (x-axis) explains 49% of the variance. The second component (y-axis) explains 35% of the variance. The stimuli of the same material mode and corresponding qualities are colored the same, specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The other five qualities are colored black. The same shape marks the stimuli of the same lighting mode, specifically the Glacier light map as squares, the Ennis light map as circles, and the Grace-new light map as stars.

5.4.2. COMPARISON (REAL STIMULI VS. RENDERINGS)

Here we compare the rating results from the two experiments. Figure 5.12 displays a direct comparison between the medians of the rating results of Experiment 1 (blue plots) and Experiment 2 (red plots) per material (in the columns), per quality term (in the rows), and as a function of type of lighting. In each subplot, L1 denotes the ambient lighting or the Glacier environment, L2 denotes the focus lighting or the Ennis environ-

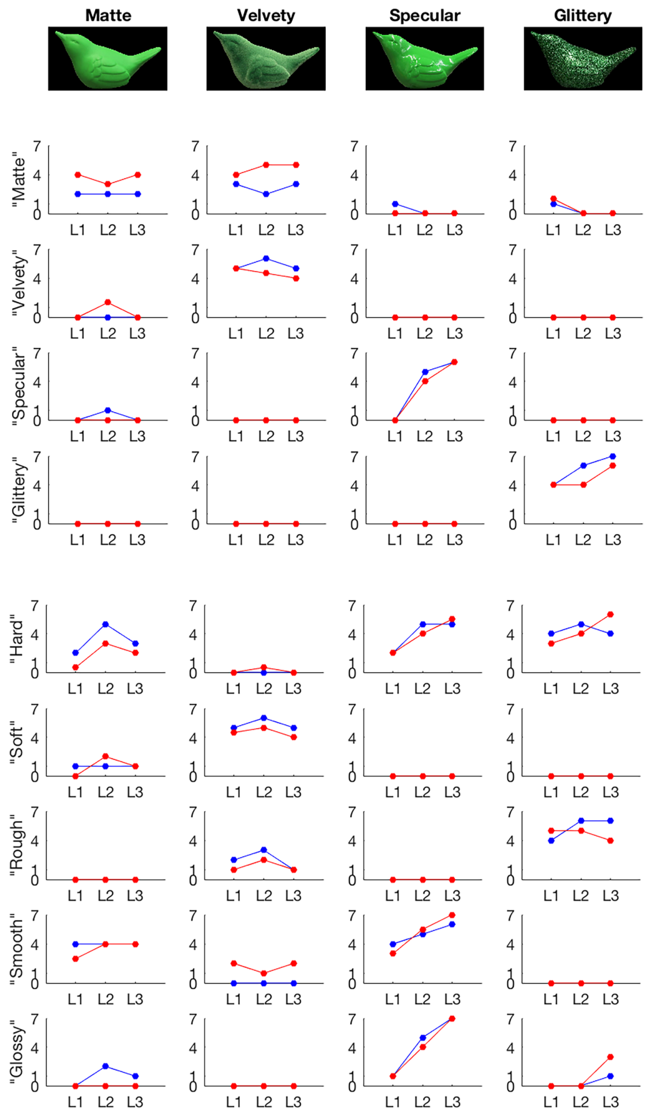

Figure 5.12: The medians of the ratings per material (columns) and quality term (rows), as a function of lighting. Blue plots are results from Experiment 1 (real stimuli); red plots are results from Experiment 2 (rendered stimuli). In each subplot, L1 denotes ambient lighting or the Glacier environment, L2 denotes focus lighting or the Ennis environment, and L3 denotes brilliance lighting or the Grace-new environment.

ment, and L3 denotes the brilliance lighting or the Grace-new environment. With the exception of the matte naming of velvety material, the top four rows of the plots in Figure 5.12 demonstrate how uniquely the material modes have been associated with their corresponding names, again confirming our previous findings. The bottom five rows of the plots show how the material modes have been associated with the other qualities, acknowledging both hard/soft and rough/smooth as mutually excluding qualities. We found that the overall results from the two experiments correlate highly (the correlation coefficient of all medians of the ratings per material and quality: $r = 0.87$, $p < 0.001$), suggesting that the selected generic natural lighting environments (blue plots) had similar effects depending on the material mode and quality as the canonical lighting modes (red plots), as was predicted on the basis of the spherical harmonics calculations. A detailed analysis of the impact of the type of lighting on the quality ratings is provided below in the section on material-dependent lighting effects.

5

We also compared the two PCA spaces (Figures 5.8 and 5.11) by rotating the PCA space of Experiment 2 to match that of Experiment 1 on the basis of the coordinates of the nine qualities. Without translation and scaling, the rotating process can be expressed as:

$$[PCA1] = \begin{pmatrix} \cos\theta \& \sin\theta \\ -\sin\theta \& \cos\theta \end{pmatrix} [PCA2] \tag{5.8}$$

where $PCA1$ and $PCA2$ both are 2×9 matrices containing the coordinates of the nine qualities in the two spaces, and θ is the angle over which we rotate to match the two spaces. Using least squares minimization, we got a small angle $\theta \approx 13°$ with residuals less than 0.09 (negligible). Figure 5.13 denotes the outcome of the 13° anticlockwise rotation. The data points from the two experiments having the same material (same shape) and lighting mode (same color), have been connected by lines for comparison. The small shifts between these corresponding data points confirm and visualize the high correlation between the results of the two experiments. Note that the coordinates of the nine qualities from Experiment 2 overlapped so much with those from Experiment 1 that they are not included for clarity reasons.

5.4.3. MATERIAL-DEPENDENT LIGHTING EFFECTS

A closer look at Figures 5.12 and 5.13 suggests that (a) for matte and velvety materials, some qualities were evoked the most by L2 (i.e., the first order of the SH component) and (b) for specular and glittery materials, some qualities were evoked the most by L3 (i.e.,

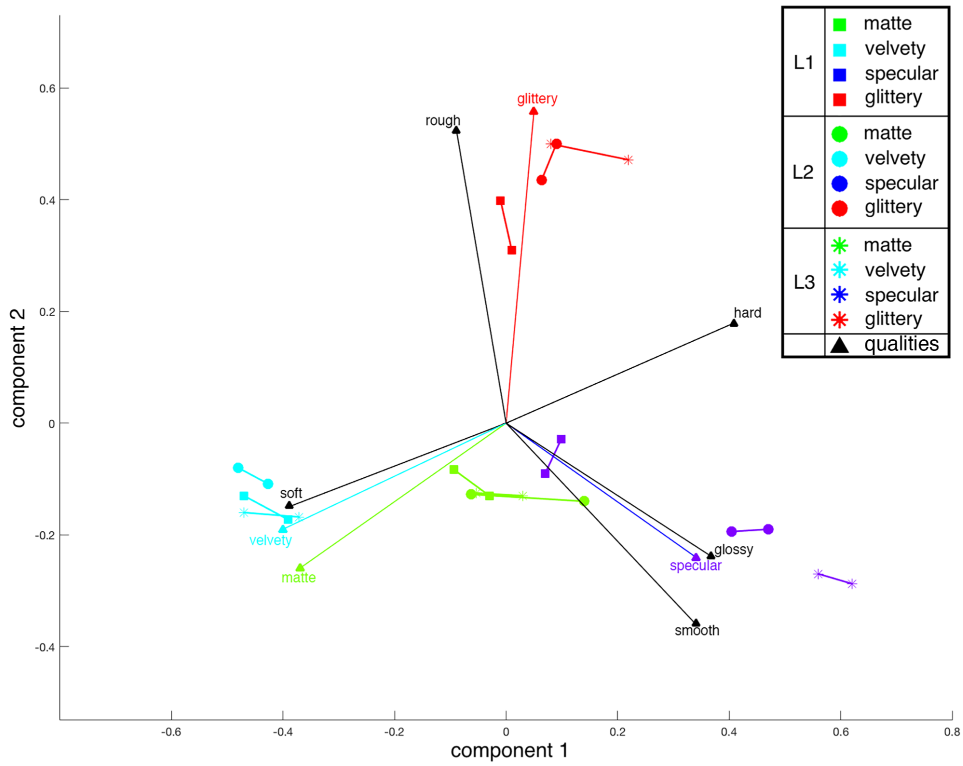

Figure 5.13: Matching the PCA space of Experiment 2 to that of Experiment 1. The four materials and corresponding qualities are colored in the same format as Figures 5.8 and 5.11, specifically matte in green, velvety in light blue, specular in indigo, and glittery in red. The other five qualities are colored black. The same shape marks the stimuli of the same lighting mode, specifically ambient lighting or the Glacier light map as squares (L1), focus lighting or the Ennis light map as circles (L2), and brilliance lighting or the Grace-new light map as stars (L3). The nine qualities from Experiment 1 were marked as triangles. The coordinates of the nine qualities from Experiment 2 overlapped with the corresponding ones from Experiment 1 and thus are not shown for clarity. The data points from the two experiments having the same material (same shape) and lighting mode (same color), have been connected by lines for comparison.

the higher order SH components). To quantitatively validate these possible material-dependent lighting effects, we tested each lighting–material combination for statistical significance by means of two-way ANOVAs, the independent variables experiment (two levels) and lighting (three levels) being the between-subjects and within-subject variable, respectively. The relevant F-values can be found in Table 5.3, 5.4 and 5.5. Overall, the statistics showed no main effect of the type of experiment; that is, no significant differences between real stimuli (Experiment 1) and rendered ones (Experiment 2), nor substantial interactions between type of experiment and type of lighting. However, there were significant material-dependent effects of the type of lighting on the qualities. These have been summarized in Table 5.3, 5.4 and 5.5 . The main findings are: (a) for specular materials, L3 highlights the qualities specular, hard, smooth, and glossy, while L1 reduces these qualities; (b) for glittery materials, L3 highlights and L1 reduces the quality glittery, and to a lesser extent the qualities hard, rough, and glossy; (c) for matte materials, the quality matte is evoked by all lightings, while L2 highlights the quality hard the most; and (d) for velvety materials, the qualities velvety and matte are evoked by all lightings, while L2 highlights to some degree the qualities soft and rough.

5.5. GENERAL DISCUSSION

The Y/N results of both Experiments 1 and 2 have permitted us to assess whether four of the terms we used for material qualities (matte, specular, glittery, and velvety) perceptually correspond to material modes of the same name. From Experiment 1 using real stimuli, we found that the names velvety, specular, and glittery applied to their corresponding material modes, while the name matte applied not only to the matte mode but also to specular and glittery mode under ambient lighting, and to velvety mode under all illuminations. In addition, focus lighting made the matte material look specular (5.17). This is in line with our previous work, where we also observed these interactions between matte, velvety, and specular modes and the canonical illuminations (Zhang et al., 2015, 2016), using the same bird object.

One could argue that the association of matte with other material modes in the present study may be due to our matte material mode not representing a perfect diffuse (i.e., Lambertian) material. However, in Experiment 2, we implemented a computer-rendered matte mode using Lambertian material, and found similar results where the name matte also applied to other materials. This suggests that observers confound materials under certain illuminations, as was shown before (Pont & te Pas, 2006; Zhang, de Ridder, & Pont, 2018). An alternative interpretation would be that the semantic meaning of matte is not unique to purely diffusely scattering materials; in particular, velvety was often judged to be a matte material in both experiments, and even more matte than the matte materials, confirming earlier findings (Zhang et al., 2016). Thus, it is not clear what matte means in terms of perception. In terms of optics it can be defined as the diffuse scattering component of a material's reflectance, which is actually present in most materials, and often determines their body color. In a weighted linear superposition model of glossy materials (as often used in computer renderings) matte and glossy form the op-

Material mode	Evoked qualities	Lighting effects
Matte	matte	Evoked by all lightings,
		independent of Experiment: $F_E(1,25) = 1.17$, $p = 0.29$,
		type of lighting: $F_L(2,50) = 0.71$, $p = 0.50$,
		and no Interaction effect: $F_I(2,50) = 0.86$, $p = 0.43$.
	hard	evoked the most by L2, somewhat less by L3 and L1:
		$F_L(2,50) = 6.10$, $p = 0.004**$
		independent of Experiment: $F_E(1,25) = 2.03$, $p = 0.17$,
		and no Interaction effect: $F_I(2,50) = 0.49$, $p = 0.49$.
	smooth	evoked by all lightings, somewhat less by L1:
		$F_L(2,50) = 9.49$, $p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.24$, $p = 0.63$,
		and no Interaction effect: $F_I(2,50) = 2.42$, $p = 0.10$.
Velvety	matte	evoked by all lightings,
		somewhat less by L2 in Experiment 1 (Interaction effect):
		$F_I(2,50) = 3.88$, $p = 0.03*$
		independent of Experiment: $F_E(1,25) = 2.33$, $p = 0.14$,
		and type of lighting: $F_L(2,50) = 1.17$, $p = 0.32$.
	velvety	Evoked by all lightings,
		independent of Experiment: $F_E(1,25) = 0.003$, $p = 0.96$,
		type of lighting: $F_L(2,50) = 1.28$, $p = 0.29$,
		and no Interaction effect: $F_I(2,50) = 1.46$, $p = 0.24$.
	soft	evoked the most by L2, somewhat less by L1 and L3:
		$F_L(2,50) = 6.39$, $p = 0.003**$
		independent of Experiment: $F_E(1,25) = 0.22$, $p = 0.64$,
		and no Interaction effect: $F_I(2,50) = 1.62$, $p = 0.21$.
	rough	evoked the most by L2, somewhat less by L1 and L3:
		$F_L(2,50) = 14.57$, $p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.29$, $p = 0.60$,
		and no Interaction effect: $F_I(2,50) = 0.98$, $p = 0.38$.

Table 5.3: Summary of the material-dependent lighting effects for matte and velvety material modes. L1 denotes ambient lighting or the Glacier environment, L2 denotes focus lighting or the Ennis environment, L3 denotes brilliance lighting or the Grace-new environment. If a quality is not listed for a specific material mode, it was not evoked by any lighting for that material mode. Significance level: $*p < 0.05$, $**p < 0.01$.

Material mode	Evoked qualities	Lighting effects
Specular	specular	evoked the most by L3, somewhat by L2,
		and the least by L1: $F_L(2,50) = 59.94, p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.72, p = 0.41$,
		and no Interaction effect: $F_I(2,50) = 1.91, p = 0.16$.
	hard	evoked the most by L3 and L2, somewhat by L1,
		$F_L(2,50) = 8.94, p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.00, p = 0.99$,
		and no Interaction effect: $F_I(2,50) = 0.85, p = 0.43$.
	smooth	evoked the most by L3, somewhat by L2,
		and the least by L1: $F_L(2,50) = 28.69, p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.01, p = 0.92$,
		and no Interaction effect: $F_I(2,50) = 2.93, p = 0.06$.
	glossy	evoked the most by L3, somewhat by L2,
		not by L3: $F_L(2,50) = 77.86, p < 0.001$
		independent of Experiment: $F_E(1,25) = 0.47, p = 0.50$,
		and no Interaction effect: $F_I(2,50) = 0.78, p = 0.46$.

Table 5.4: Summary of the material-dependent lighting effects for specular material modes. L1 denotes ambient lighting or the Glacier environment, L2 denotes focus lighting or the Ennis environment, L3 denotes brilliance lighting or the Grace-new environment. If a quality is not listed for a specific material mode, it was not evoked by any lighting for that material mode. Significance level: $*p < 0.05, **p < 0.01$.

Material mode	Evoked qualities	Lighting effects
Glittery	glittery	evoked the most by L2, somewhat less by L1 and L3 in Experiment 1; the most by L3, somewhat less by L2, and the least by L1 in Experiment 2; main effect of Lighting: $F_L(2,50) = 59.94$, $p < 0.001$ interaction effect: $F_I(1,25) = 9.85$, $p < 0.001$, and independent of experiment: $F_E(2,50) = 4.18$, $p = 0.052$.
	hard	evoked the most by L2, somewhat less by L1 and L3 in Experiment 1; the most by L3, somewhat less by L2, and the least by L1 in Experiment 2; main effect of Lighting: $F_L(2,50) = 9.76$, $p < 0.001$ interaction effect: $F_I(1,25) = 6.24$, $p = 0.004**$, and independent of experiment: $F_E(2,50) = 0.004$, $p = 0.95$.
	rough	evoked by all lightings but the least by L1 in Experiment 1, evoked by all lightings but the least by L3 in Experiment 2; main effect of Lighting: $F_L(2,50) = 3.36$, $p = 0.04*$ interaction effect: $F_I(1,25) = 3.19$, $p = 0.05*$, and independent of experiment: $F_E(2,50) = 1.85$, $p = 0.19$.
	glossy	evoked by L3 only, the most in Experiment 2, main effect of Lighting: $F_L(2,50) = 28.95$, $p < 0.001$ interaction effect: $F_I(1,25) = 6.66$, $p = 0.003**$, and independent of experiment: $F_E(2,50) = 0.01$, $p = 0.93$.

Table 5.5: Summary of the material-dependent lighting effects for glittery material modes. L1 denotes ambient lighting or the Glacier environment, L2 denotes focus lighting or the Ennis environment, L3 denotes brilliance lighting or the Grace-new environment. If a quality is not listed for a specific material mode, it was not evoked by any lighting for that material mode. Significance level: $*p < 0.05$, $**p < 0.01$.

posites of the range. In future studies, it might be necessary to investigate whether matte can semantically be considered the opposite of glossy or specular, and define the names of the canonical material modes properly.

One may wonder to which extent the four material modes used in this paper are representative of the space of all possible materials. In Figure 5.14, we superimposed our results with the 10 material classes from the results of Fleming et al. (2013) obtained from a large number of images from MIT-Flickr database (Sharan, Rosenholtz, & Adelson, 2009) as stimuli. Merging the two sets of results was done by mapping their respective main dimensions: soft–hard and rough–smooth. We note that, on the one hand, each canonical mode represents various classes; on the other hand, the appearance of materials within each class can also vary across canonical modes. Even though our four modes cover the space of materials rather well, there remain material classes that are not covered. For example, if we focus on solid opaque materials, the most notable missing class of materials is metal. Depending on their microstructure, metals may appear to be either smooth and specular, or rough and glittery. Todd and Norman (2018) tested how ambient light influences the perception of metals and found that a combination of ambient and focus lighting was optimal for depicting metals in their experiments. Additional experiments may be required to evaluate the influences of generic illuminations containing higher frequency components, for instance, testing to what extent the brilliance lighting mode influences the perception of metals. Moreover, color effects should certainly be taken into account in such experiments, since those might be diagnostic for the difference between dielectrics and metals (simply said, in dielectrics such as plastics, the highlights have the color of the illumination, while in metal they have the color of the metal).

The rating results from Experiment 1 (real stimuli) and Experiment 2 (rendered stimuli) were strongly correlated, showing that the effects of canonical lightings on material perception are reproducible using generic lighting environments having similar spherical harmonic compositions (i.e., the relative power of SH components). In Figure 5.3, we could make an interesting observation that, as the scores of the diffuseness metric descended, the scores of the brilliance metric showed a tendency to ascend. This might indicate that SH compositions of generic natural lighting environments may follow these statistical regularities. Another example of such a regularity was found by Mury, Pont, and Koenderink (2009), namely that the positive component of the second order SH component (the so-called squash tensor) often aligns with the direction of the first order

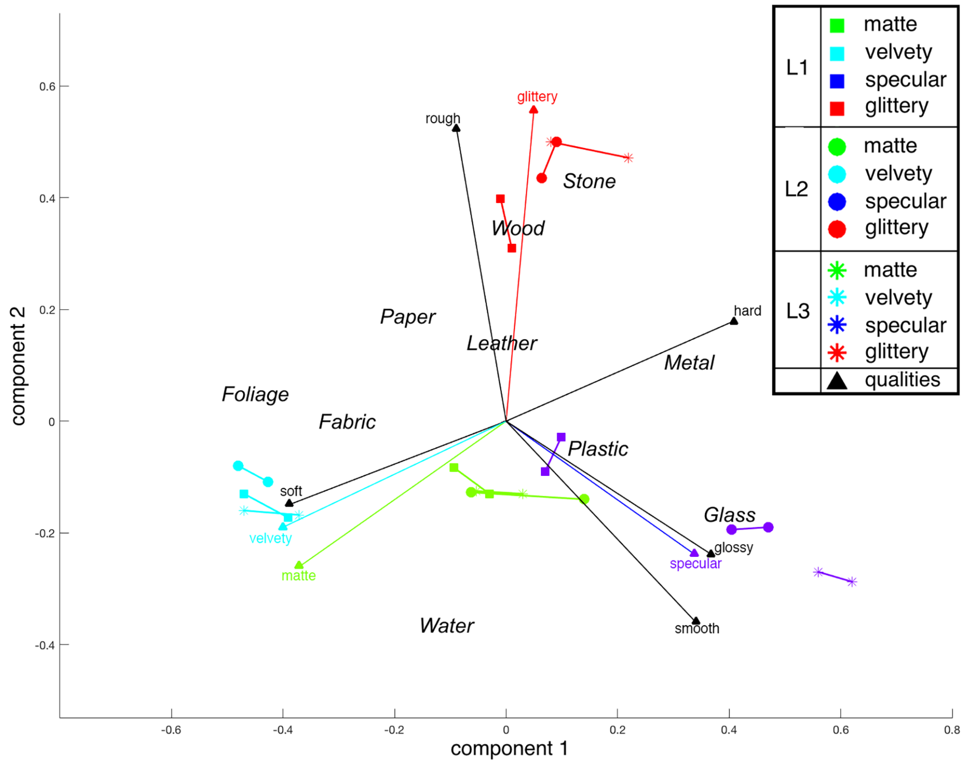

Figure 5.14: The 10 material classes from previous studies (Sharan et al., 2009; Fleming et al., 2013) mapped onto our PCA space taken from Figure 5.13.

SH component (the light vector) of the light field. This effect is due to reflections from the lighted surface, causing a light clamp (Mury et al., 2009). In lighting design, the light clamp does not form a standard component, probably because one gets this effect for free, due to such reflections, and therefore we did not include the squash tensor in our canonical lighting modes and also neglected it in our brilliance metric (Equation 5.4). The optical cause of the regularity found in the present study might be that the more diffuse the light becomes, the more diffusely scattering the environment has to be, with extreme cases such as a white photo tent or integrating sphere.

Overall, the lighting effects are stronger if the BRDF of the material is more peaked, such as our specular and glittery materials, or general material classes such as metal, plastic, and glass (and even stone, wood, or leather if polished). Corresponding material qualities were evoked the most by lightings that were dominated by higher order SH components, then somewhat less by lightings that were dominated by first order SH component, and the least by lightings that were dominated by a zero-order SH component. On the contrary, if the BRDF is less peaked, the lighting effects are subtler, such as our velvety and matte materials, or fabric, foliage, and paper. Meanwhile, some qualities that associated with matte and velvety materials were evoked the most by lightings that were dominated by a first-order SH component (see Figure 5.12 and Table 5.3). The relation between peakedness of the optical functions and the magnitude of the effects might have been expected, since more peaked functions cause larger variations between conditions. The effects for specular and glittery materials might be explained by the observation that the most articulated lighting (brilliance) brings out material properties of materials with peaked BRDFs most, because that combination would lead to the most salient and numerous highlights/glints, being expressions of the most characteristic optic properties of such peaked reflectance. The most bipolar lighting (focus) might be explained to best bring out the characteristic gradients of smoothly varying BRDFs (ambient and brilliance will diffuse out and cause no or fewer strong gradients).

For the specular stimuli in the study of Motoyoshi and Matoba (2012), perceived glossiness changed when varying the contrast and gamma of the illumination. This could be explained in terms of a spherical harmonics decomposition of the illumination. We have applied their contrast and gamma changes to the lighting map they used, Rendering with Natural Light–eucalyptus (RNL-eucalyptus), and recomputed the SH-based diffuseness and brilliance metrics, as shown in Figure 5.15. Interestingly, and in line with what can be seen in Figure 5.3, we found that the scores of the diffuseness metric descended as the scores of the brilliance metric ascend ($r = -0.99$, $p < 0.01$). Here, a lowering of the contrast in the lighting could be interpreted as an increase of the zero-order SH component and a decrease of higher-order SH components, which tends to make the specular surface look more matte. A lowering of the gamma in the lighting could be interpreted as an increase of higher-order SH components and decrease of the zero-order SH component, and thus could evoke the perception of the glossy quality for specular surfaces. One notable advantage of SH-based metrics is that it provides absolute values that permit comparisons between different environments. In future work, it would be interesting to study how manipulating the relative power of SH components could affect material perception in a controllable way, with potential applications in the field of computer graphics. One possibility would be to express both material and lighting in terms of SH coefficients (as done by Ramamoorthi & Hanrahan, 2001, for inverse rendering). However, spherical harmonics characterize a lighting environment in a global fashion: for instance, simply rotating the lighting with respect to the object has the potential to affect material perception but leaves the spherical harmonic coefficients unchanged. This perceptual effect is shown in Figure 5.16, where the perceived material of the specular bird seems to change significantly when it is rotated in the focus canonical lighting. This holds for various materials, as was shown using computer graphics methods (e.g., Bousseau, Chapoulie, Ramamoorthi, & Agrawala, 2011) and conforms with the practice of optimizing the orientation of lighting to maximize heuristic measurements inspired from photographic practices (Hunter et al., 2015). More metrics are needed for detailed analysis of illuminations beyond the first order SH, to quantify their spatial structures (the type of light texture (e.g., dotted or stripy studio lighting). In a subsequent work, we will study how the orientation of a lighting environment or the shape of the object may affect material appearance depending on the choice of material mode.

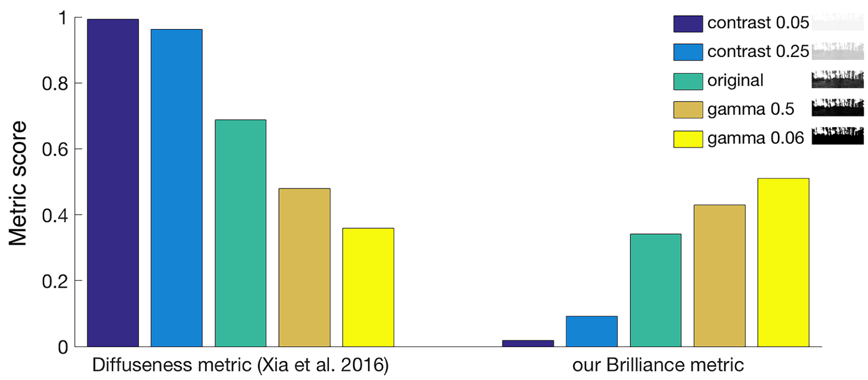

Figure 5.15: The spherical harmonics–based diffuseness and brilliance analysis for the light maps to which a gamma/contrast change was applied. The original light map was used in the study of Motoyoshi and Matoba (2012). The thumbnails of the light maps are shown next to the legend. The scores of the two metrics negatively correlated ($r = -0.99$, $p < 0.01$).

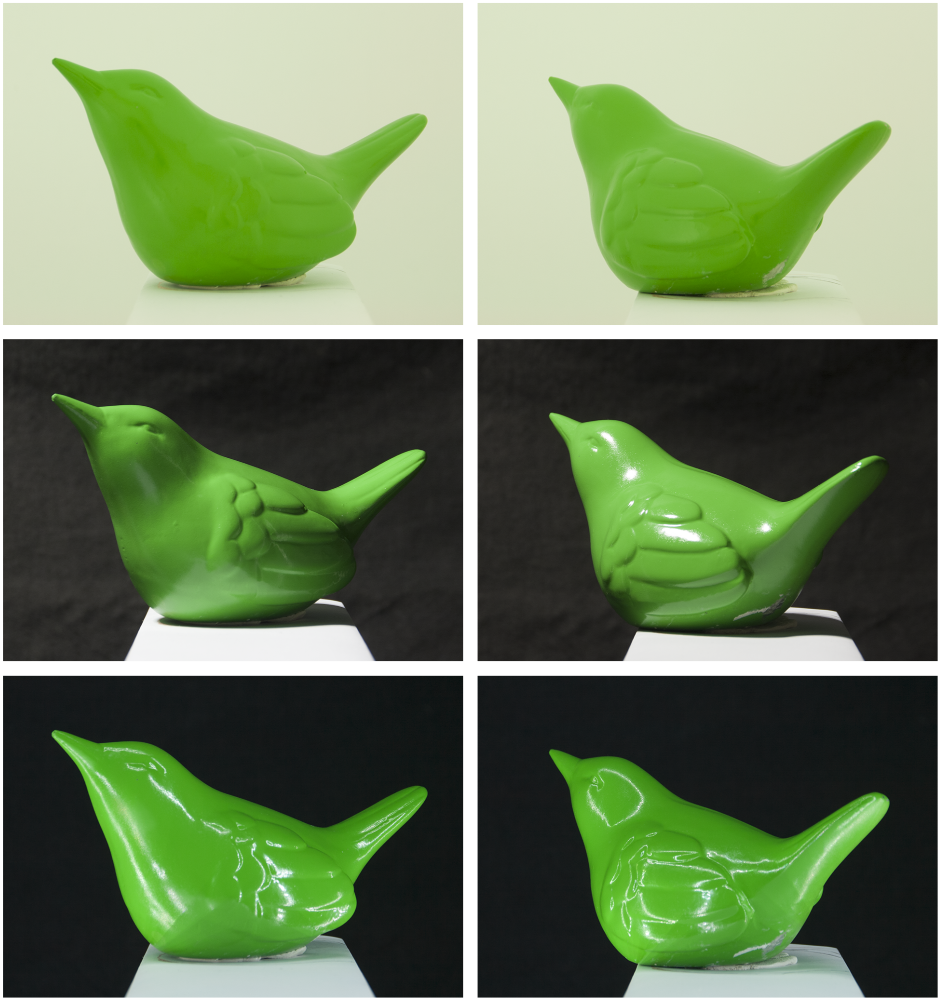

Figure 5.16: With a fixed viewing angle, either rotating the illumination or the illuminated object may result in changes in appearance and thus give a different perception. From top to bottom, each row shows the same specular object in two orientations under ambient, focus, and brilliance light, respectively. Its appearance varies the most across the two orientations when using focus lighting (second row).

5.6. CONCLUSION

We investigated (a) how canonical material modes associate with perceived material qualities and (b) how canonical lighting modes brought out the perception of material qualities for each material. In combination with four canonical material modes (matte, velvety, specular, and glittery) and three canonical lighting modes (ambient, focus, and

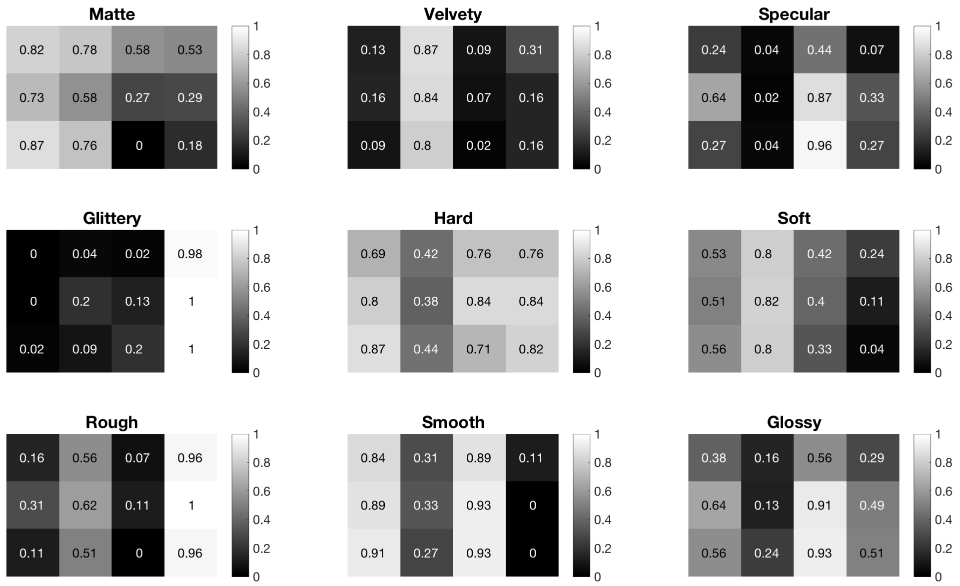

Figure 5.17: The fraction of answering "yes" for the nine qualities in Experiment 1 per illumination-material condition. In each subplot, the title shows the quality being tested. In the three by four matrix it shows the fraction of answering "yes" for that quality. From top to bottom of each subplot, the three rows correspond to the stimuli under ambient light, focus light, and brilliance light, respectively. From left to right of each subplot, the four columns correspond to the stimuli of matte, velvety, specular, and glittery material modes, respectively. The average of the three fractions per column correspond to the results shown in Figure 5.6.

brilliance), 12 stimuli were rated for nine material qualities, namely matte, velvety, specular, glittery, glossy, rough, smooth, hard, and soft. Material-dependent lighting effects were found. Specifically, we performed a pair of experiments in which we presented observers with images of an object made of four materials in three lightings, and asked them to rate each configuration according to nine material qualities. In the first experiment, the stimuli were photographs of real objects lit by canonical lightings, while in the second experiment, the stimuli were rendered using state-of-the-art surface reflectance models for the four material modes lit by three environment maps as illuminations. Three environment maps were selected to represent our canonical lighting modes based on a diffuseness metric and a brilliance metric (Figure 5.3), namely the Glacier, Ennis, and Grace-new environments for the ambient, focus, and brilliance light, respectively. We made predictions of the effects of lighting on material appearance for generic natural lighting environments and validated the predictions: results correlated strongly for the two experiments and reproduced material-dependent lighting effects of former

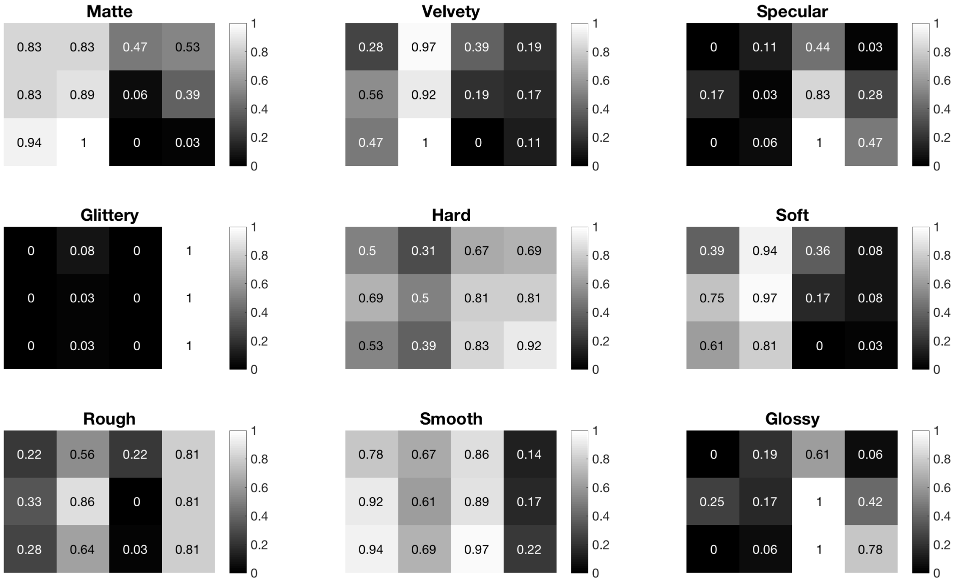

Figure 5.18: The fraction of answering "yes" for the nine qualities in Experiment 2 per illumination-material condition. In each subplot, the title shows the quality being tested. In the three by four matrix it shows the fraction of answering "yes" for that quality. From top to bottom of each subplot, the three rows correspond to the stimuli under ambient light, focus light, and brilliance light, respectively. From left to right of each subplot, the four columns correspond to the stimuli of matte, velvety, specular, and glittery material modes, respectively. The average of the three fractions per column correspond to the results shown in Figure 5.9.

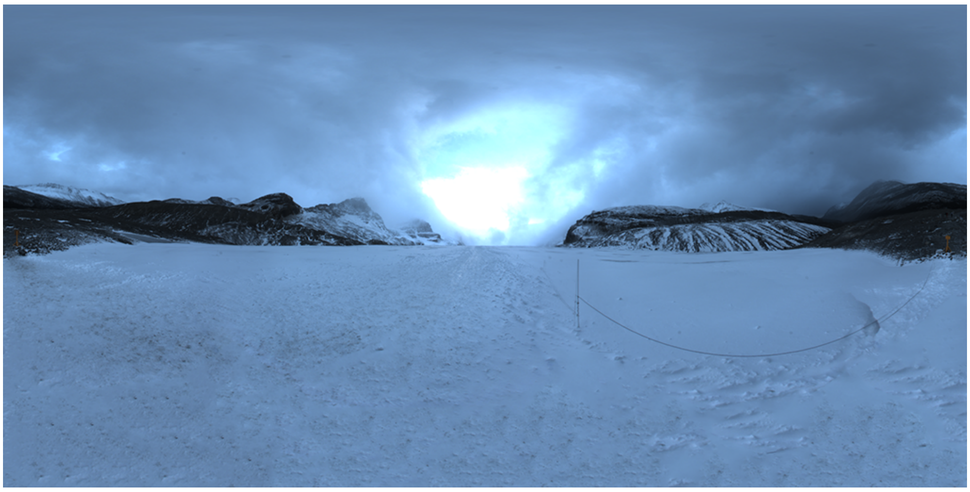

Figure 5.19: The filled-in version of the Glacier environment map.

studies. Our results support the notion that systematically varying the spatial structure of the natural illumination, such as enhancing or attenuating the corresponding modal components of the lighting, can systematically evoke the perception of associated material qualities. The optics-based models span a wide range of natural materials and lighting environments, providing a systematic approach to study the perceptual interactions of materials and lighting.

BIBLIOGRAPHY

Barati, B., Karana, E., Sekulovski, D., & Pont, S. C. (2017). Retail lighting and textiles: Designing a lighting probe set. Lighting Research & Technology, 49 (2), 173–194.

Bousseau, A., Chapoulie, E., Ramamoorthi, R., & Agrawala, M. (2011). Optimizing environment maps for material depiction. Computer Graphics Forum, 30 (4), 1171–1180.

Brainard, D. H. (1997). The psychophysics toolbox. Spatial Vision, 10, 433–436.

Debevec P. E. (1998). Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In Proceedings of SIGGRAPH 98. Computer Graphics Proceedings. Annual Conference Series (pp. 189–198). New York, NY: ACM Press.

Döerschner, K., Boyaci, H., & Maloney, L. T. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. Journal of Vision, 10 (4): 8, 1–9, https://doi.org/10.1167/10.4.8.

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. Journal of Vision, 4 (9): 11, 821–837, https://doi.org/10.1167/4.9.11.

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. Journal of Vision 3 (5): 3, 347–368, https://doi.org/10.1167/3.5.3.

Fleming, R. W., Wiebel, C., & Gegenfurtner, K. (2013). Perceptual qualities and material classes. Journal of Vision, 13 (8): 9, 1–20, https://doi.org/10.1167/13.8.9.

Ganslandt, R., & Hofmann, H. (1992). Handbook of interior lighting. Lüdenscheid, Germany: ERCO.

Hunter, F., Biver, S., & Fuqua, P. (2015). Light science & magic: An introduction to photographic lighting. Boca Raton, FL: CRC Press.

Kelly, R. (1952). Lighting as an integral part of architecture. College Art Journal, 12 (1), 24–30.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. Perception, 36 (14), 1–16.

Koenderink, J., & Pont, S. (2003). The secret of velvety skin. Machine Vision and Applications, 14 (4), 260–268.

Lagae, A., & Drettakis, G. (2011, August). Filtering solid Gabor noise. In ACM Transactions on Graphics (TOG), (Vol. 30, No. 4, p. 51). New York, NY: ACM.

Lagae, A., Lefebvre, S., Drettakis, G., & Dutré, P. (2009, July). Procedural noise using sparse Gabor convolution. In ACM Transactions on Graphics (TOG), (Vol. 28, No. 3, p. 54). New York, NY: ACM.

Marlow, P., Kim, J., & Anderson, B. L. (2011). The role of brightness and orientation congruence in the perception of surface gloss. Journal of Vision, 11 (9): 16, 1–12, https://doi.org/10.1167/11.9.16.

Marlow, P., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. Current Biology, 22 (20), 1909–1913.

Motoyoshi, I., & Matoba, H. (2012). Variability in constancy of the perceived surface reflectance across different illumination statistics. Vision Research, 53 (1), 30–39.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2007). Light field constancy within natural scenes. Applied Optics, 46 (29), 7308–7316.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2009). Structure of light fields in natural scenes. Applied Optics, 48 (28), 5386–5395.

Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. Journal of Vision, 10 (9): 5, 1–19, https://doi.org/10.1167/10.9.5.

Pont, S. C. (2009, February). Ecological optics of natural materials and light fields. In Human Vision and Electronic Imaging XIV (Vol. 7240, p. 724009). Bellingham, WA: International Society for Optics and Photonics. https://doi.org/10.1117/12.817162.

Pont, S. C., & te Pas, S. F. (2006). Material-illumination ambiguities and the perception of solid objects. Perception, 35 (10), 1331–1350.

Ramamoorthi, R., & Hanrahan, P. (2001). A signal-processing framework for inverse rendering. In Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH ’01). New York, NY: ACM Press/Addison-Wesley Publishing Co.

Sharan, L., Rosenholtz, R., & Adelson, E. (2009). Material perception: What can you see in a brief glance? Journal of Vision, 9 (8): 784, https://doi.org/10.1167/9.8.784.

Stock, R. D., & Siegel, M. (1996). Orientation invariant light source parameters. Optical Engineering, 35 (9), 2651–2661.

Stokes, M., Anderson, M., Chandrasekar, S., & Motta, R. (2012). A standard default color space for the Internet—sRGB, 1996. http://www.w3.org/Graphics/Color/sRGB.

Todd, J. T., & Norman, J. F. (2018). The visual perception of metal. Journal of Vision, 18 (3): 9, 1–17, https://doi.org/10.1167/18.3.9.

van Assen, J. J. R., Wijntjes, M. W., & Pont, S. C. (2016). Highlight shapes and perception of gloss for real and photographed objects. Journal of Vision, 16 (6): 6, 1–14, https://doi.org/10.1167/16.6.6.

Vergne, R., & Barla, P. (2015). Designing Gratin, a GPU-tailored node-based system. Journal of Computer Graphics Techniques, 4 (4), 54–71.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. ACM SIGGRAPH Computer Graphics, 26 (2), 265–272.

Xia, L., Pont, S. C., & Heynderickx, I. (2017). Light diffuseness metric part 1: Theory. Lighting Research & Technology, 49 (4), 411–427.

Zhang, F., de Ridder, H., Fleming, R. W., & Pont, S. (2016). MatMix 1.0: Using optical mixing to probe visual material perception. Journal of Vision, 16 (6): 11, 1–18, https://doi.org/10.1167/16.6.11.

Zhang, F., de Ridder, H., & Pont, S. C. (2015, March). The influence of lighting on visual perception of material qualities. In Rogowitz, B. E. Pappas, T. N. & de Ridder H. (Eds.), Proceedings of SPIE 9394, Human Vision and Electronic Imaging XX (p. 93940Q). https://doi.org/10.1117/12.2085021.

Zhang, F., de Ridder, H., & Pont, S. C. (2018). Asymmetric perceptual confounds between canonical lightings and materials. Journal of Vision, 18 (11): 11, 1–19, https://doi.org/10.1167/

6

EFFECTS OF LIGHT MAP ORIENTATION AND SHAPE ON THE VISUAL PERCEPTION OF CANONICAL MATERIALS

Abstract

We previously presented a systematic optics-based canonical approach to test material-lighting interactions in their full natural ecology, combining canonical material (matte, velvety, specular, and glittery) and lighting (ambient, focus, and brilliance) modes (Zhang et al., 2019). Analyzing the power of the spherical harmonics components of the lighting allowed us to predict the lighting effects on material perception for generic natural illumination environments. To further understand how material properties can be brought out or communicated visually, in the current study we tested whether and how light map orientation and shape affect these interactions in a rating experiment: for 36 combinations of four materials, three shapes, and three lightings, we rotated the lighting environment in fifteen different configurations. For the velvety objects, there were main and interaction effects of lighting and light map orientation. The velvety ratings decreased when the main light source was coming from the back of the objects. For the specular objects, there were main and interaction effects of lighting and shape. The specular ratings increased when the environment in the specular reflections were clearly visible in the stimuli. For the glittery objects, there were main and interaction effects of shape and light map orientation. The glittery ratings correlated with the coverage of the glitter reflections as the shape and light map orientation varied. For the matte objects, results were robust across all conditions showing no effects of lighting, light map orientation, and shape. Lastly, we propose combining the canonical modes approach with so-called "importance maps" to analyze the appearance features of the proximal stimulus, the image, in contradistinction to the physical parameters, as an approach for optimisation of material communication.

In review as: Zhang, F., de Ridder, H., Barla, P., & Pont, S. (2019). Effects of light map orientations on the visual perception of canonical materials.

6.1. INTRODUCTION

One of the aims of material perception research is to understand how human beings perceive materials in varying lighting environments. The endless possible combinations of materials and lighting environments pose a difficult challenge on this matter in two important ways, namely A) same material under different lights and belonging to different shapes can have a different appearance, and B) same appearance can be the result of different combinations of lightings, shapes, and materials (image ambiguities). The appearance of materials varies enormously depending on the lighting and shape (Olkkonen & Brainard, 2011), and human observers were found not to be “material-constant” if the shape (Nishida & Shinya, 1998; Vangorp, Laurijssen, & Dutré, 2007) or the lighting varies (Dror, Willsky, & Adelson, 2004; Pont & te Pas, 2006). A well-known lighting effect for glossy surfaces that has been found in many studies is that glossy surfaces are perceived as rather matte under very diffuse lighting, and glossier under directed lighting (Dror, Willsky, & Adelson, 2004; Pont & te Pas, 2006; Zhang, de Ridder, & Pont, 2015, 2018; Zhang, de Ridder, Fleming, & Pont, 2016), or perceived to have different levels of glossiness under different artificial or natural lighting environments (Fleming, Dror & Adelson, 2003; Doerschner, Boyaci & Maloney, 2010; Olkkonen & Brainard, 2010; Motoyoshi & Matoba, 2012; Wendt & Faul, 2018; Adams et al., 2019; Zhang, de Ridder, Barla, & Pont, 2019). In a recent study on textiles, the textiles were analyzed optically and categorized into canonical modes, then combined with two canonical lightings (diffuse lighting and collimated lighting) and in a perception experiment found to have a systematic influence on the perception of six material qualities, namely “textured”, "metallic", "silky", "shiny", "glittery" and “soft” (Barati et al., 2017).

In one of our former studies (Zhang, de Ridder, Barla, & Pont, 2019), we asked observers to judge material qualities for a variety of material-lighting combinations. To this end, a system was developed using optics-based models of canonical material and lighting modes that span a wide range of natural materials and lighting. The four material modes employed (matte, velvety, specular, and glittery), were based on optical models that describe the bidirectional reflectance distribution functions (BRDFs) of opaque materials (Ward, 1992; Koenderink and Pont, 2003; Barati et al., 2017), representing, respectively, diffuse scattering, asperity scattering, forward scattering, and meso-facet scattering modes (Zhang, de Ridder, Fleming, & Pont, 2016), spanning a large part of the BRDF space. The three canonical lighting modes employed (ambient, focus, and brilliance light), were based on a mathematical description of the local light field, representing,

6

respectively, the mathematical zeroth, first, and higher order contributions to the spherical harmonic (SH) decomposition of the local light field (Mury, Pont, & Koenderink, 2007). The mathematical basis of this 3-component framework for light descriptions, which we use as canonical modes, has a physical meaning as the three components correspond to fully diffuse light (the ambient or zeroth order SH component, a monopole), directed light from a single direction (the focus or first order SH component, a dipole), and the fine structure or texture of the light field (the brilliance or sum of the third and higher order SH components), respectively. These modes represent properties of light that human observers can distinguish (Doerschner, Boyaci, & Maloney, 2007; Schirillo, 2013; Morgenstern, Geisler, & Murray, 2014; Kartashova et al., 2016; Xia, Pont, & Heynderickx, 2017). Moreover, they are known in perception-based lighting design as the basic components of an integral lighting plan (Kelly, 1952; Ganslandt, & Hofmann, 1992; Pont, 2009, 2013; Pont, & de Ridder, 2018). They thus span the space of natural light conditions.

In one experimental condition of the above-mentioned study (Zhang et al., 2019), we employed computer renderings of three generic natural lighting environments approximating the ambient, focus and brilliance lighting modes. On the basis of quantitative metrics of the relative power of their SH decomposition components, the following three lighting maps from the high resolution USC database (http://gl.ict.usc.edu/Data/HighResProbes/) were selected to represent the three canonical lightings the best: “Glacier” for ambient lighting (dominated by the zeroth order SH component), “Ennis” for focus lighting (dominated by the first order SH component), and “Grace-new” for brilliance lighting (dominated by the sum of the higher orders SH components). Rendering the four canonical material modes under these three selected light maps resulted in a “rendered stimuli” set, comparable to the “real stimuli” set created by illuminating the four canonical materials (real objects) under the three canonical lightings. We evaluated the perception of a range of material qualities for both stimuli sets and found the results to be mutually consistent. Thus, using an optics-based canonical approach, we showed that material perception could be varied in a systematic and predictable manner. For example, brilliance lighting (the controlled real light condition and also its virtual metrics-based best match of the natural luminance maps) evoked perceived glossiness, hardness, and smoothness for specular material the most, while focus lighting (again, the real and also the metrics-based best match) evoked perceived roughness and softness for velvety material the most.

So far, our research into light-material interactions has been confined to one shape ("bird") illuminated under a fixed light map orientation per lighting (Zhang, de Ridder, & Pont, 2015, 2018; Zhang, et al., 2016; Zhang et al., 2019). In everyday experience, however, one occasionally observes subtler lighting effects on material appearance, even when the lighting environment remains unchanged. For example, a drop of water on a desk may not be visible until one moves to a certain location with respect to the direction of the main light source in the room. In this case, changing the viewing angle does not change the global illumination environment, yet the image projected to our eyes becomes different due to direction-dependent forward scattering (i.e., specular reflection) of the material, and thus triggers a different material perception. Changing viewing or illumination direction to trigger a different percept is often used by lighting designers and photographers to make certain features prominent in the same environment. Marlow and Anderson (2013) found that by varying the light directions of a quite directed lighting, perceived glossiness for specular bumpy objects and surfaces change significantly. They explained the results using image features, such as contrast, coverage, sharpness of the highlights, etc. In addition, changing the shape of an object while keeping the material and the lighting environment (illumination map) the same, can also influence material perception. For example, it was found that shape can affect material perception (Vangorp, Laurijssen, & Dutré, 2007). Specifically, using a blob-shaped object resulted in a more veridical judgment of glossiness than the usual spherical object. So, although a single glossy sphere can be modelled and rendered easily (without the need for complex self-shadowing and interreflection computations), and it conveniently represents all possible visible surface orientations in one visualisation, its global convexity was shown to eliminate certain features that are important triggers for perceptual qualities. A blob-shaped object, for instance, might thus be a visually more informative shape for a material probe, in other words, a more visually intelligent shape for material communication.

In the current study, we applied our optics-based canonical approach for the four materials and three lighting modes (Zhang et al., 2019) and looked into the effects of light map orientation and shape in order to further investigate how to how to bring out the physical material best - which is an important issue for disciplines and applications involving lighting design and material communication (f.i. computer graphics, design visualisations, webshops, and material selection interfaces). Specifically, we tested the visual perception of our four canonical materials (matte, velvety, specular, glittery) under three metrics-matched natural lighting maps best representing our canonical lighting modes (ambient, focus and brilliance), namely the lighting maps “Glacier”, “Ellis” and “Grace-new” (from the USC database http://gl.ict.usc.edu/data/HighResProbes/). This test was confined to one perceptual quality per material, namely, the corresponding material quality (matte, velvety, specular or glittery) for each material mode. The lighting modes are expected to give main effects that are material dependent (Zhang et al., 2019). The variation of lighting direction is expected to result in no or minor effects for the Glacier illumination, since that is the best match to ambient illumination, which in its purest form is fully diffuse and non-directional. The Ennis lighting or best match to focus lighting has one clear average direction and thus is expected to affect material perception, based on the literature. The Grace-new as the best match to brilliance light is expected to result in medium effects, since it is more directed and structured than ambient, but less directed than focus. These results are expected to be material-dependent. With respect to lighting directions, the effects were expected to be significant for specular and glittery material - since for those materials the image features (highlights) are strongly dependent on the directions of illumination and viewing, due to the steep variations of their bidirectional reflectance distribution functions (BRDFs, see Nicodemus et al., 1992). In contradistinction, we expected the lighting direction effects to be much subtler for the matte material than for specular and glittery materials, as its BRDF is rather flat (in the ideal case constant) and the shading gradients smooth. In the case of velvet, the key feature concerns its bright contour, which "sticks" to the silhouette, such that the velvety appearance is also expected to be more robust. To investigate the effect of shape, we implemented, next to our “bird” shape (Zhang et al., 2016), a “blob” shape (Vangorp, Laurijssen, & Dutré, 2007), and a “sphere” shape. To vary only the light map orientations for each combination of shape, illumination, and material, the position of the object was kept fixed relative to the camera during the rendering process.

6.2. METHOD

6.2.1. STIMULI

LIGHTING ENVIRONMENTS AND LIGHT MAP ORIENTATIONS

From the USC’s high resolution recreations of the Debevec’s light probe images (http://gl.ict.usc.edu/ we selected three light maps (“Glacier”, “Ennis”, and “Grace-new”) that best represent our canonical lighting modes (ambient, focus, and brilliance, respectively). The selection was made by using a combination of a diffuseness metric (Xia, Pont, & Heynderickx,

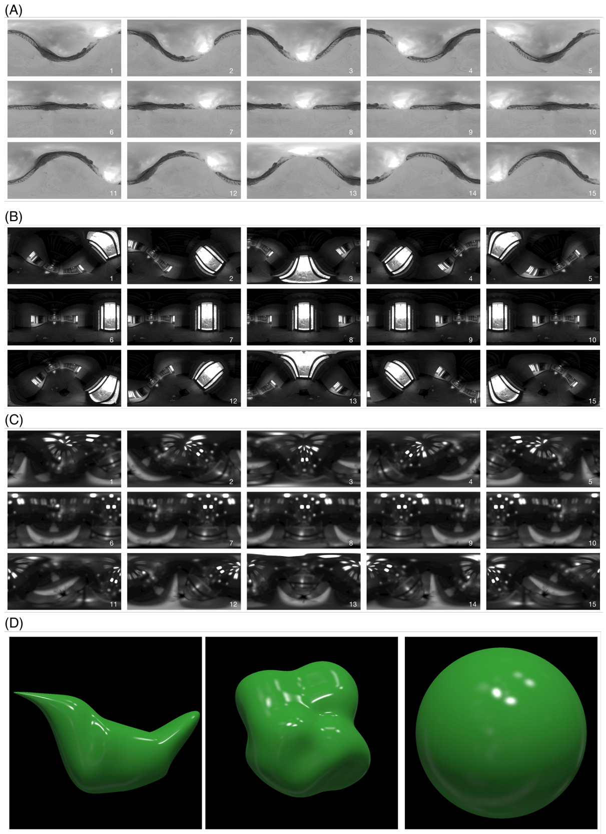

Figure 6.1: (A) – (C) Three light maps rotated vertically for three levels and horizontally for five levels. (A): the “Glacier” map, (B): the “Ennis” map, (C): the “Grace-new” map, representing "ambient", "focus" and "brilliance" lighting. Note that in this paper we label the vertical levels as elevations and the horizontal levels as azimuths, though they do not represent the direction of the main light source in one light map. From left to right: the azimuths are $-2\pi/5, -\pi/5, 0, \pi/5, 2\pi/5$. From top to bottom: the elevations are $-\pi/4, 0, \pi/4$. i.e. the number 8 of each light map was the one with no rotation. These parameters were arbitrarily selected. Also note that the “Grace-new” environment was blurred to reduce noise issues in rendering, due to the presence of very small light sources of very high intensity. After blurring, the map still contains many light sources, but noise in rendering is greatly decreased. (D): Examples of the three shapes used in the experiment, rendered with the specular material and the “Grace-new” light map. From left to right: the “bird”, the “blob”, and the “sphere”.

2017) and a brilliance metric (Zhang et al., 2019), both based on the relative power of their SH decomposition components. Specifically, the “Glacier” map represents the ambient lighting the best, as it scores highest for Xia’s diffuseness metric ($_{DXia} = 0.83; B = 0.42$); the “Ennis” map represents the focus lighting the best, as it scores lowest on Xia’s diffuseness metric ($D_{Xia} = 0.17; B = 0.71$); the "Grace-new" map represents the brilliance lighting the best, as it scores highest on our brilliance metric ($D_{Xia} = 0.40; B = 0.79$). Each light map was rotated vertically and/or horizontally such that the light map orientations varied over three vertical levels (original and $\pm\pi/4$) and five horizontal levels (original, $\pm\pi/5$ and $\pm 2\pi/5$), see Figure 6.1. Note that in this paper we label the vertical levels as elevations and the horizontal levels as azimuths, though they do not represent the direction of the main light source in one light map. At the bottom of the original USC's "Glacier" map there was a large area of black pattern due to the occlusion of a tripod base. To make it more ambient and natural, we removed the occlusion with Photoshop's "content-aware fill" tool (see Figure 6.1A). And the "Grace-new" environment was blurred (see Figure 6.1C) to reduce noise issues in rendering, due to the presence of very small light sources of very high intensity. After blurring, the map still contains many light sources, but noise in rendering is greatly decreased.

RENDERING PROCESS: SHAPES AND MATERIAL MODES

The 3D model of the "bird" shape was created in Blender, and is the same as the 3D model we used in former studies (Zhang et al., 2016). The "blob" shape was taken from Vangorp et al. (2007). The usual "sphere" shape was added as a comparison to the "bird" and the "blob" shapes, as shown in Figure 6.1D. A sphere is the most simple convex object in the sense that it has a constant shape index and curvedness. The blob is smoothly curved with a varying shape index and curvedness. The bird also contains sharp edges.

The matte material mode was simulated to resemble a (hypothetical) material with a Lambertian BRDF. The velvety material mode was implemented with the asperity scattering BRDF model of Koenderink and Pont (2003). The specular material mode was implemented with an isotropic Ward BRDF (Ward, 1992). The glittery material mode was implemented by mimicking the occurrence of multi-facet flakes at the surface of the object (Zhang et al., 2019). Rendering was performed in Gratin version 0.3 for Apple Mac OS (Vergne, & Barla, 2015) to code and compile the computer rendering program in OpenGL shading language (GLSL) version 410 (Figure 6.2).

Irrespective of the choice of material, rendering was performed in RGB (float 32 bits precision), with an orthographic camera, without tone mapping. Rendering was done through ray tracing at 2000 samples per pixel (spp) unless specified otherwise, and considered only direct lighting, a reasonable approximation for the object shapes we consider. We used the same material models as in our previous work (Zhang et al., 2019); we

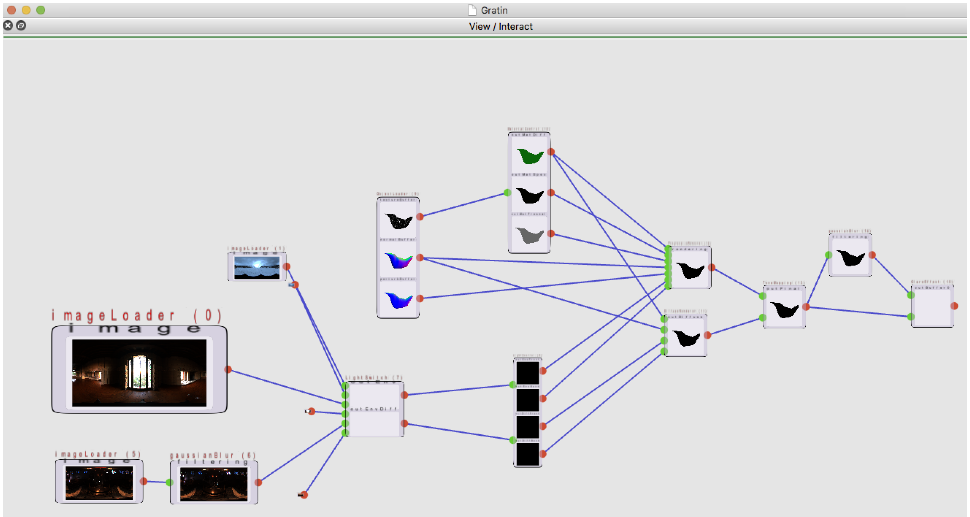

Figure 6.2: The node setups used in Gratin to implement the different material modes.

thus refer the reader to our previous paper for a detailed description.

In order to simulate the matte material, an environment pre-filtering approach was implemented as it completely removes noise coming from the rendering process of the diffuse component. Specifically, assuming that the shadowing and inter-reflection effects can be neglected, the diffuse component of the materials may be represented equivalently using a diffuse-filtered version of the illumination environment, as provided on the USC website
(http://gl.ict.usc.edu/Data/HighResProbes/). Rendering the diffuse component then simply consists in evaluating the diffuse-filtered environment in the direction of the surface normal, and multiplying the result by the colored diffuse albedo.

For the specular and glittery modes, we employed Monte-Carlo integration using importance sampling of the Ward model to speed up convergence (Ward, 2005). We have also included a Fresnel term in the model to improve physical plausibility, using an index of refraction of 1.5, which yields a reflectivity of 4% at normal incidence (typical of dielectrics). For the glittery mode, we used four times more sampes (8000 spp) to capture the fine spatial variations in the flake texture. For the velvety mode, we relied on standard cosine-weighted importance sampling to evaluate the asperity scattering model of Koenderink and Pont, which required longer rendering times. In the supplements Figure S1 – Figure S4, we show all stimuli per material, lighting, and shape. The numbers 1 – 15 correspond to the oriented light maps in Figure 6.1.

6.2.2. PROCEDURE

At the beginning of the experiment, observers were first shown all of the stimuli (540 computer rendered images in total, see Supplementary Materials) twice in a randomized order, to give them an idea about the range of the stimuli and their scale for the rating. They were instructed that in each trial a question "rate how [..] is the object?" was shown on top of the screen, with [..] displaying one of the four names, namely matte, velvety, specular or glittery. In each trial fifteen stimulus panels were shown to the observers below the question, with a slider next to each image (Figure 6.3). The fifteen stimulus panels would have the same material, shape and light map, but only differ in light map orientations (three elevations and five azimuths). The observers were explicitly instructed that the task was to rate the same material using the same range in different trials, instead of using the full scale in each trial per fifteen stimulus panels. This was done to allow analysis per material instead of only per fifteen stimuli. Panels were randomly positioned, and all slider bars were initially set at the bottom. The task of the observers was to rate each image by moving the slider bar, representing "not [..] at all" (or "0" within a "0" to "1" range) at the bottom of the slider to "extremely [..]" (or "1" within a "0" to "1" range) at the top of the slider. When clicking and mouse button within the panel of a stimulus image, a horizontal bar would appear super-positioned on the stimulus, which was slightly thinner than the slider bar attached to the right. When dragging the mouse cursor within the panel, both bars move vertically together. When releasing the mouse button, only the slider bar was shown to indicate the rating, while the thinner bar would disappear. Observers could freely go back-and-forth to rate any one of the fifteen images until pressing the "Enter" key to go to the next trial.

For each canonical material mode, only the corresponding material term [..] was tested, i.e. we asked "rate how matte is the object?" for the stimuli rendered using the matte material mode only. So, for each observer, the experiment contained altogether 4 materials × 3 illuminations × 3 shapes = 36 trials of 15 stimulus panels, in total 36 × 15 = 540 ratings. The experiment took between 40 minutes and an hour per observer. The interface was developed with the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997; Kleiner et al, 2007) in MATLAB R2016b, and presented on a linearly calibrated EIZO ColorEdge CG277 (27-inch class calibration color LCD) display. The viewing distance was around 30 cm.

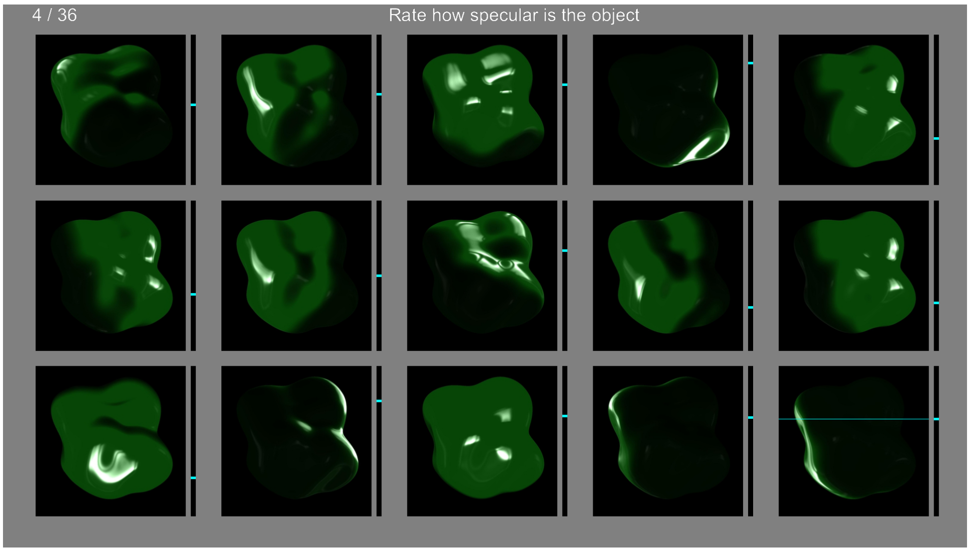

Figure 6.3: The user interface of the experiment developed with the Psychophysics Toolbox. In this screenshot of an example trial, the stimulus is the same specular blob-shaped object rendered using the "Ennis" map for fifteen light map orientations. The resulting fifteen stimulus panels were randomly positioned in each trial. The question "rate how [..] is the object" was positioned above the stimuli, with [..] displaying one of the four material names, in this case "specular". A slider bar was positioned on the right-hand side of each stimulus panel, initially set at the bottom of the slider. Observers were told before the experiment that the vertical slider scales from "not [..] at all" at the bottom to "extremely [..]" at the top. If a slider was moved, an additional horizontal bar was superimposed on the stimulus image to display the current rating value, as demonstrated in the stimulus at the bottom-right corner. The slider bar next to the stimulus image was attached to the thinner horizontal bar, moving vertically to indicate the ratings when dragging the mouse cursor. When releasing the mouse button, only the slider bar was shown, while the thinner bar would disappear. Observers could freely go back and forth to rate any one of the fifteen images until pressing the "Enter" key to go to the next trial. The number of the current trial and the total number of trials were shown in the top-left corner of the screen. Note that the settings shown in this figure are not from any of the observers, but are generated for demonstration purposes only, and the stimuli appeared different as shown in supplements due to the process of taking the screenshot.

6.2.3. OBSERVERS

Twelve paid observers participated in the Experiment. All participants had normal or corrected-to-normal vision, and were inexperienced in psychophysical experiments. Participants read and signed the consent form before the experiments. The study was approved by the Human Research Ethics Committee at Delft University of Technology and conducted in accordance with the declaration of Helsinki and Dutch law.

6.3. RESULTS

Here we present the general results per material mode and, thus, per material quality. We analyzed the rating data per material using a four-way repeated-measures ANOVA, with three lightings, three elevations, five azimuths, and three shapes being the independent variables. We assumed that each observer used one constant perceptual scale for the same material despite the changes of lightings and shapes across all trials (presented in randomized order), and rescaled the 135 data points (3 shapes x 3 lightings x 15 directions) per material per observer such that the ratings range from 0 to 1. For the interpretation of the results, the analysis was confined to the main effects and the first-order interaction effects. We do not present higher order interactions, because we consider those too complicated to be meaningful. In the Supplementary Materials, we plotted the ratings next to the corresponding stimuli, allowing visual inspection of the stimuli and data (Figure S1-S4). Note that since light sources within the three light maps may be located anywhere (f.i. on the side or at the top), it is only meaningful to directly compare the effects of azimuth and elevation within, but not across light maps.

6.3.1. MATTE

We did not find any significant main effect (lighting:$F(2,22) = 1.38, p = 0.27$; shape: $F(2,22) = 0.39, p = 0.69$ azimuth: $F(1.58,17.37) = 1.36, p = 0.28$; elevation: $F(1.27,14.00) = 1.10, p = 0.33$) or any first-order interaction effect for the "matte" ratings of the matte material ($Mean = 0.46, SEM = 0.01$ for all ratings). This suggests that perceived "matteness" was independent of the light maps, lighting directions and shape of the object. This is confirmed in Figure 6.4, showing that the averaged ratings for the matte material mode were robust across all conditions.

6.3.2. VELVETY

Figure 6.5 presents the averaged "velvety" ratings for the velvety material mode. A number of trends can be discerned. First, the overall mean of the "velvety" ratings for "Glacier" map (0.57±0.02) is higher than that for the "Ennis" map (0.44±0.02) and "Grace-new" map (0.44±0.02), a difference that is substantial as confirmed by a significant main effect for lighting environment ($F(2,22) = 5.48, p = 0.012$) Second, there appears to be an effect of azimuth, in particular for "Ennis", in that the means of the ratings for azimuth1 and azimuth5 were lower than for azimuth2, azimuth3, and azimuth4. This is in line with the observation that there is a significant main effect of azimuth ($F(4,44) = 9.13, p < 0.001$)

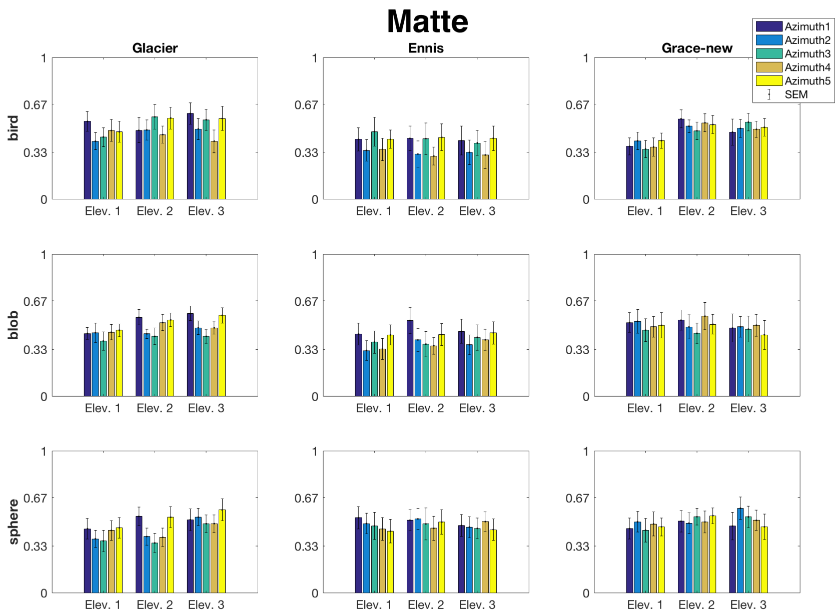

Figure 6.4: The averaged "matteness" ratings of twelve observers per shape (subplot rows), illumination (subplot columns), elevation (x-axis in each subplot), and azimuth (bars for each elevation in each subplot). The error bars indicate ± 1 SEM.

6

combined with a significant interaction effect between lighting environment and azimuth ($F(8,88) = 11.80, p < 0.001$). Figure 6.6A confirms the above-mentioned trend by showing that it can be seen for "Ennis" while the azimuth has hardly any impact on the ratings for "Glacier" and "Grace-new". Note that the ratings for "Grace-new" were lower than for "Glacier". Third, for "Grace-new" the averaged ratings appear to decrease systematically from elevation1 via elevation2 to elevation3, which cannot be seen for "Glacier" and "Ennis". The statistical analysis indicated a significant main effect for Elevation ($F(1.21, 13.34) = 6.63, p = 0.019$); the assumption of sphericity had been violated: $\chi^2(2) = 10.482, p = 0.005$, hence the degrees of freedom were corrected using the Greenhouse-Geisser estimate of sphericity ($\epsilon = 0.61$) In addition, there was a significant interaction effect between the lighting and the elevation ($F(4,44) = 11.72, p < 0.001$) in that, going from elevation1 to elevation3 the ratings drop for "Grace-new" only (Figure 6.6B). Finally, we did not find any significant differentiating effect of shape on the "velvety" judgments for the velvety material (main effect of shape: $F(2,22) = 1.14, p = 0.34$). To summarize, the perceived "velvetiness" of velvety material was affected by the lighting environments (light maps) and light map orientations, with the "Ennis" and "Grace-new" maps reduced perceived "velvetiness" the most, while perceived "velvetiness" was found to be robust for the shape of the object.

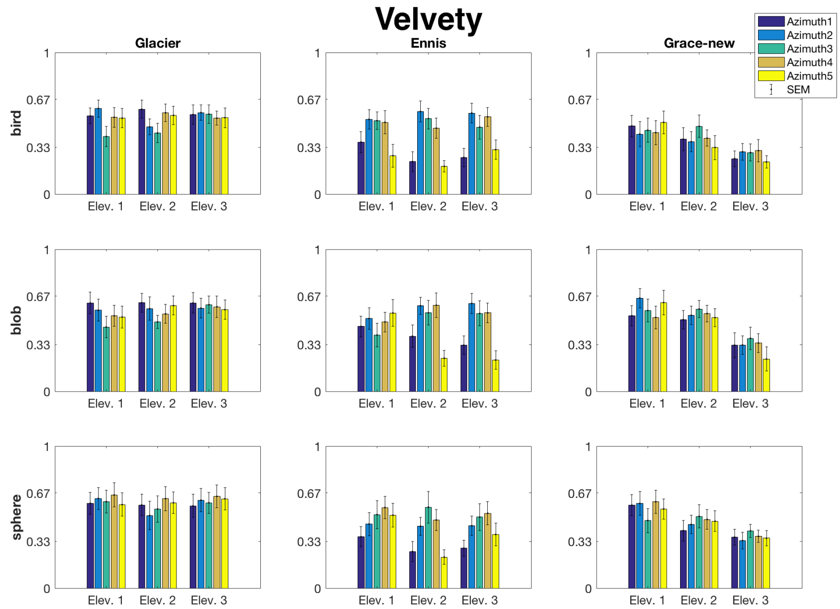

Figure 6.5: The averaged "velvetiness" ratings of twelve observers per shape (subplot rows), illumination (subplot columns), elevation (x-axis in each subplot), and azimuth (bars for each elevation in each subplot). The error bars indicate ± 1 SEM.

6.3.3. SPECULAR

Figure 6.7 presents the averaged "specularity" ratings for the specular material mode. When comparing the columns in Figure 6.7, we found that the averaged ratings of specularity were relatively lower for the "Glacier" map (0.36±0.02) than for the "Ennis" (0.56±0.02) and "Grace-new" maps (0.58±0.02). This was confirmed by a significant main effect for light environment ($F(2,22) = 14.05, p < 0.001$). When comparing the rows in Figure 6.7, we found that the averaged ratings of specularity were relatively lower for the "bird" shape (0.43±0.02) than for the "blob" (0.54±0.02) and the "sphere" (0.54±0.02). This was also confirmed by a significant main effect for shape ($F(2,22) = 8.50, p = 0.002$). We did not find any significant differentiating effect of light map orientations on the "specularity" judgments for the specular material (main effect of azimuth: $F(1.25, 13.74) =$

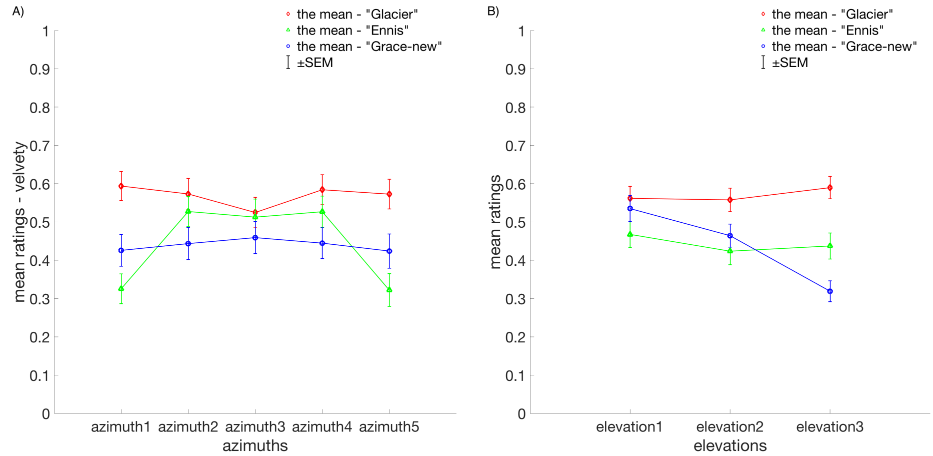

Figure 6.6: A) The mean "velvety" ratings, averaged across observers, shapes and elevations, as a function of azimuth and per light map; B) The mean "velvety" ratings, averaged across observers, shapes and azimuths, as a function of elevation and per light map. The error bars represent ± 1 SEM.

2.70, $p = 0.12$; main effect of elevation: $F(1.10, 12.13) = 3.35, p = 0.089$). The significant interaction effect between light maps and shapes (F(4,44)=2.82,p=0.036) was mainly due to the ratings for the "Ennis" map systematically increasing from the "bird" via the "blob" to the "sphere", while for the "Glacier" and "Grace-new" maps the ratings were rather flat with a small peak for the "blob" shape (Figure 6.8). To summarize, the perceived "specularity" of specular material was affected by the lighting environments (light maps) but not by the light map orientations, with the "Glacier" map reducing perceived "specularity" the most. Similarly, perceived "specularity" depended somewhat on the shape of the object with "bird" under the "Glacier" map reducing "specularity" the most and "sphere" under the "Ennis" map highlighting "specularity" the most.

6.3.4. GLITTERY

Figure 6.9 presents the averaged "glittery" ratings for the glittery material mode. Two trends can be discerned. First, the overall mean of the "glittery" ratings for the "bird" shape (0.31±0.02) is lower than that for the "blob" (0.52±0.02) and the "sphere" (0.52±0.02) shapes, a difference that is substantial as confirmed by a significant main effect for shape ($F(2, 22) = 28.42, p < 0.001$). Second, unlike the light map itself, light map orientations play a role in perceiving "glitteriness" as confirmed by significant main effects for azimuth ($F(4, 44) = 48.81, p < 0.001$) and elevation ($F(2, 22) = 21.62, p < 0.001$) and a non-

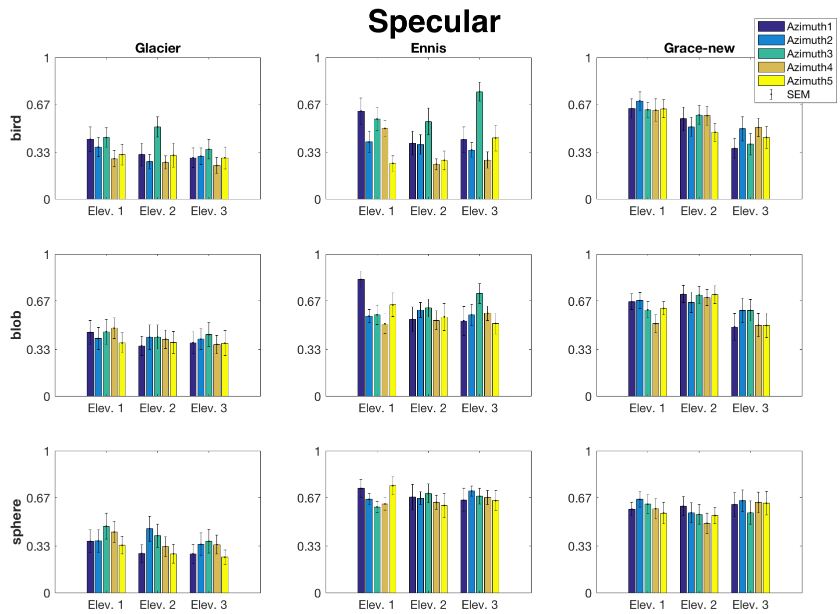

Figure 6.7: The averaged “specularity” ratings of twelve observers per shape (subplot rows), illumination (subplot columns), elevation (x-axis in each subplot), and azimuth (bars for each elevation in each subplot). The error bars indicate ± 1 SEM.

significant main effect for light map ($F(2,22) = 2.06, p = 0.15$). Specifically, the ratings of “azimuth 2” (0.51±0.02) and “azimuth 3” (0.51±0.02) were significantly higher than those for “azimuth 1” (0.41±0.02), “azimuth 4” (0.44±0.02), and “azimuth 5” (0.38±0.02) which was mainly due to the ratings for the “bird” shape (Figure 6.10). The latter is confirmed by a significant interaction effect between the shapes and azimuths ($F(8,88) = 9.63, p < 0.001$). To summarize, the perceived “glitteriness” was affected by the shape of the objects, with the “bird” shape reducing perceived “glitteriness” the most, as well as the light map orientations, in particular the azimuths.

6.4. PERCEPTUAL EFFECTS, QUALITY RATINGS AND IMAGE FEATURES

In the supplements Figure S1 – Figure S4, we show the rating data per material as a function of the azimuth next to the stimuli images per material, from (A) to (C): under the

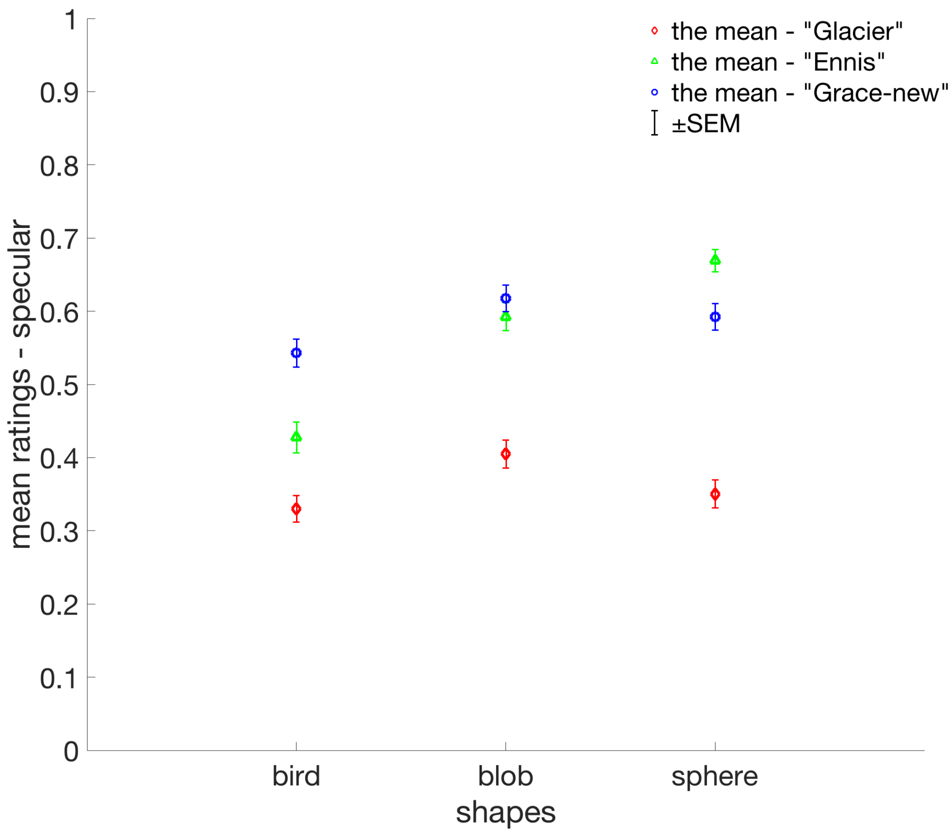

Figure 6.8: The mean "specularity" ratings, averaged across observers, azimuths and elevations, a function of per shape and per light map. The error bars represent ± 1 SEM.

"Glacier" map; (D) – (F): under the "Ennis" map; (G) – (I): under the "Grace-new" map. At left, we show the corresponding ratings. The numbers 1 – 15 on the x-axis correspond to the oriented light maps as shown in Figure 6.1. The stimuli images are shown on the right, with the numbers on the bottom-right corner of each stimulus image corresponding to the oriented light maps. The rows represent the three elevations. In these figures, we could make observations about which image features might have triggered the perceptual effects that triggered the quality assessments. In the following sections, we will describe our observations in detail per material mode and connect the observations to our results and previous findings in literature.

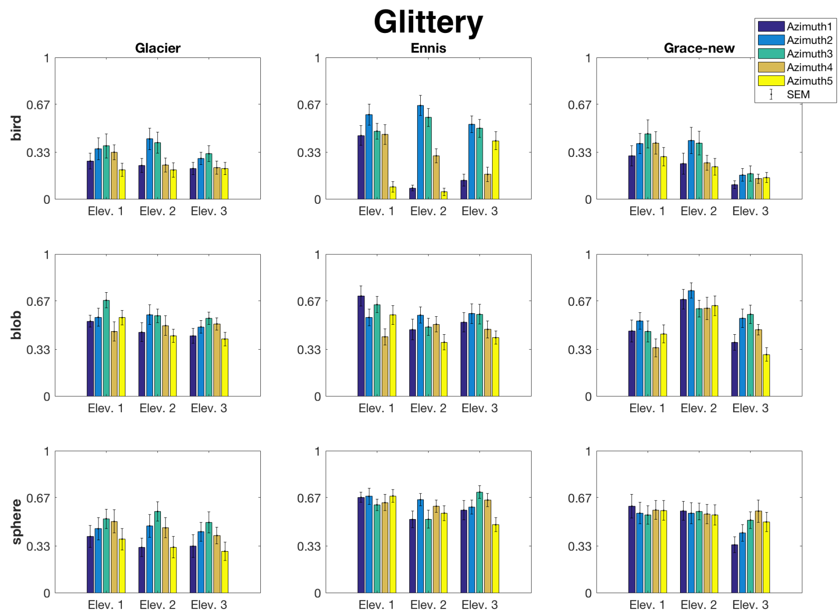

Figure 6.9: The averaged "glitteriness" ratings of twelve observers per shape (subplot rows), illumination (subplot columns), elevation (x-axis in each subplot), and azimuth (bars for each elevation in each subplot). The error bars indicate ± 1 SEM.

6.4.1. MATTE

For matte materials, the diffuse shading gradients vary smoothly and do not show sudden (dis)appearances of highlights or other salient features, which can explain that the perception of "matteness" is quite constant (Figure S1). This kind of robustness implies that its invariants can be used to infer shapes based on the shading patterns (Koenderink & van Doorn, 1980; Belhumeur, Kriegman, & Yuille, 1999; Narasimhan, Ramesh, & Nayar, 2003; Kunsberg & Zucker, 2013, 2018). Simultaneously, knowing the shape can help observers with judging the characteristics of the local light field and material (Koenderink et al., 2007; Xia, Pont, & Heynderickx, 2014, 2016; Kartashova et al., 2016). As an exception, illusory gloss effects were found for matte materials on bumpy surfaces under collimated lighting (Wijntjes and Pont, 2010), where second order shading effects were confused with specular highlights. Such effects only appear for quite non-generic lighting and only in specific cases form a real problem, for instance endoscope lighting, for which the viewing (camera) and lighting directions coincide (Wu et al., 2010)

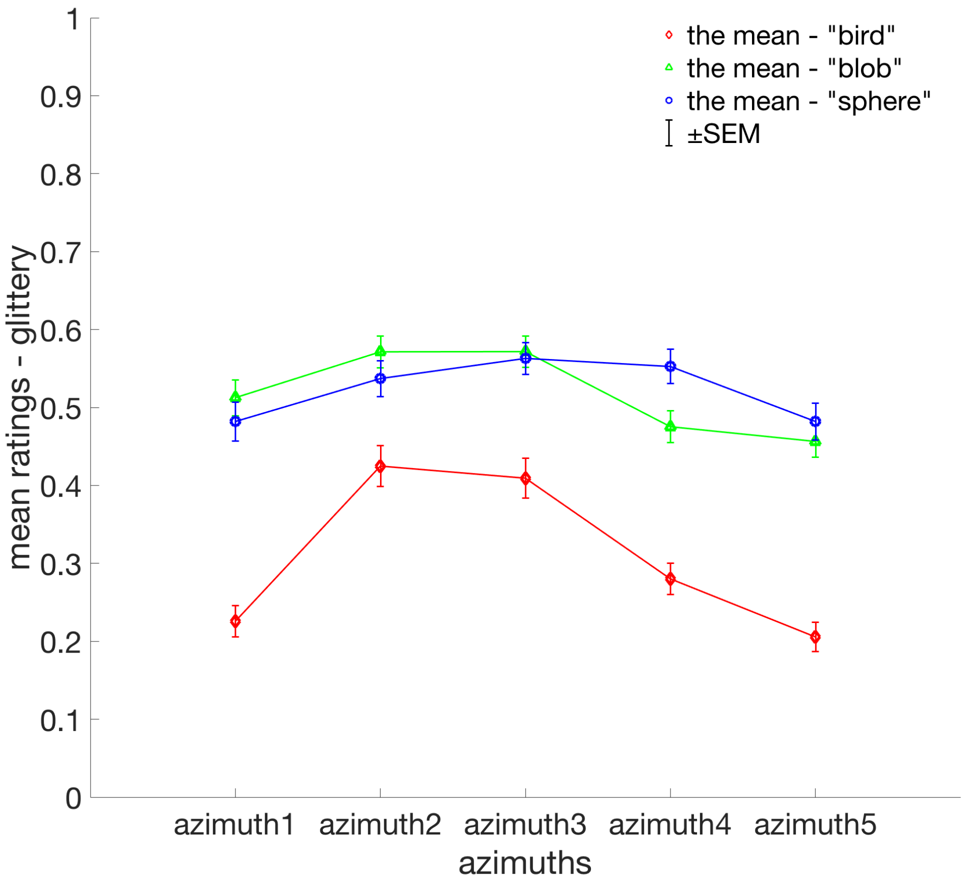

Figure 6.10: The mean “glittery” ratings, averaged across observers, light maps and elevations, as a function of azimuth and per shape. Note that the azimuth variation is confounded with a change of the lighting elevation. The error bars represent ± 1 SEM.

6.4.2. VELVETY

The main effect we found for the velvety materials was that “velvetiness” was sometimes rated low for the “Ennis” and “Grace-new” lightings, and especially when the main light source was coming from the back of the objects. The chief visual cue of velvet appearance is a thin but very steep luminance gradient at the silhouette, i.e. the bright contours along the surface due to surface scattering by the asperities (Koenderink & Pont, 2003). The cue is invariant to lighting directions when using the “Glacier” environment map (shown in Figure S2A - C), as are the ratings, confirming our expectations. Figure S2 reveals that, for the other two lightings, all stimuli with relatively low ratings have other image features in common. When using the “Ennis” environment map, we observe relatively low ratings for stimuli No. 1, 5, 6, 10, 11, 15 of the “bird” shape (Figure S2D), No. 6 10, 11, 15 of the “blob” shape (Figure S2E) and the “sphere” shape (Figure S2F). This corresponds to the results shown in Figure 6.6A, which might be due to an ambiguity caused by the directed light source: when the directed light source is behind the object, the asperity scattering mode’s luminance gradient might be confounded with the diffuse or specular scattering modes gradients. In other words, the bright contour due to asperity scattering in isolation (without diffuse shading over the body) cannot be distinguished from the bright rim that occurs for the combination of backlighting and diffuse and / or specular scattering. The material may then be perceived as matte or even as somewhat specular, rendered using rim lighting. This could also explain the results using the “Grace-new” environment map under elevation 3 (Figure S2G – I, and also see Figure 6.6B). The significant drops in the ratings as a function of elevation correspond to back-lighting configurations.

To conclude, velvetiness seems to require not only bright contours due to the surface scattering, but in addition the co-occurrence of diffusely scattered luminance or smooth gradients over the body. Simply presenting only a “bright contour” on an otherwise dark object will not trigger the perception of velvetiness - but instead may trigger the perception of “matteness”. This corresponds to results from our above-mentioned former work (Zhang et al., 2019), in which we found strong interactions between the matte and velvety material modes.

6.4.3. SPECULAR

The main visual cue for specular materials are the specular highlights. When using the “Glacier” environment map, the overall ratings for the specular materials were relatively lower (Figure S3A – C), corresponding to the results shown in Figure 6.7 and 7. This was within our expectations as it confirmed previous findings in glossiness perception literature indicating that perceived glossiness reduces under diffuse lighting (Dror, Willsky, & Adelson, 2004; Pont & te Pas, 2006; Zhang, de Ridder, & Pont, 2015; Zhang et al., 2016, 2019), because the highlights will be diffused. Meanwhile, perceived glossiness is also affected by negative contrast of reflections caused by dark parts of the environment generating dark specular reflections or lowlights (Kim, Marlow, & Anderson, 2012). The combination of fine-structured bright highlights and dark lowlights might explain the perceived glossiness in Figure S3 A - C.

When using the "Ennis" environment (Figure S3D – F), the most notable image cues are contrast and coverage of specular highlights. The interaction effects show increased ratings for the "sphere" under "Ennis" lighting, i.e. the averaged ratings for "specularity" were highest when combining the "sphere" shape and "Ennis" lighting (Figure 6.8). This might be due to a clear reflection of the illumination. The window-shaped specular highlight patterns are particularly clearly reflected on the "sphere" (Figure S3F). As a comparison, highlights on the "bird" (Figure S3D) and the "blob" (Figure S3E) deformed in a more complex manner. This confirmed previous findings on glossiness perception, namely, that the shape of highlights may influence glossiness perception (van Assen, Wijntjes, & Pont, 2016). It also shows that when highlights reveal real-world illumination properties, they are less likely to be misperceived as texture and thus could increase perceiving glossiness (Fleming, Dror, & Adelson, 2003).

When using the "Grace-new" environment map (Figure S3G – I), the ratings in general were relatively high as the highlights are quite visible in most of the stimuli, as expected, and reflect the fine structure of the brilliance lighting. Coverage and contrast of the highlights were mainly varying as a function of the elevation and the effects of azimuth are not as salient as for the other two light maps, due to the angular structure of the brilliance lighting in the "Grace-new" environment (primarily many tiny hotspots from above).

Unexpectedly, we did not find significant effects of light map orientations for the perception of specularity in this study (unlike for example, in Marlow, Kim & Anderson, 2012). However, we did find higher order interaction effects between light map orientation, shape, and the choice of lighting environment. This was not mentioned in the results since the interpretation of these higher order effects is usually very complex. In the stimuli images we could observe subtle variations in the "specular" ratings as the light map orientations varied. Some of these variations might be due to the changes of the contrast and coverage of the high/lowlights with respect to the diffuse shading.

6.4.4. GLITTERY

The main visual cues for "glitteriness" also seemed to be the features of the highlights on the glitters (see Figure S4). With simple image processing, namely thresholding the top 1% brightest pixels (as the glitters) in each stimulus, we could count the number of glitters as a coarse evaluation of the coverage of the highest intensity glitters. As shown in Table 1, significant correlations were found between the numbers of glitters in the stim-

	Glacier (ambient)	Ennis (focus)	Grace-new (brilliance)
bird	$R^2 = 0.67, p < 0.001$	$R^2 = 0.40, p = 0.01$	$R^2 = 0.67, p < 0.001$
blob	$R^2 = 0.64, p < 0.001$	$R^2 = 0.67, p = 0.22$	$R^2 = 0.01, p = 0.80$
sphere	$R^2 = 0.91, p < 0.001$	$R^2 = 0.26, p = 0.05$	$R^2 = 0.37, p = 0.02$

Table 6.1: The correlations between the numbers of glitters (top 1% brightest pixels in the stimuli) and the perceptual ratings for gltteriness per shape and lighting.

uli and the glittery ratings per shape and lighting, except for the blob under the "Ennis" and "Grace-new" lighting.

6.5. GENERAL DISCUSSION

The main question we pose in this paper is how light map orientation and object shape influence the perception of materials in addition to the material reflectance itself and the main modes of the lighting environment. To answer this question, an experiment was set up in which we combined four canonical material modes, three shapes and three illumination environments, and then oriented the illumination environments in fifteen different directions (varying across three elevations vertically and five azimuths horizontally). In our rating experiment, we found the following main results:

1. For matte materials, perceived "matteness" was robust and constant across all variations, i.e. no effect was found of either light map orientation, shape of the object, or lighting mode.
2. For the perceived "velvetiness" of velvety materials, there were significant effects of light map orientation, which were lighting-dependent but shape-independent. Such effects were evoked the most under the "Ennis" light map and hardly under the "Glacier" and "Grace-new" light maps. The "Glacier" light map highlighted "velvetiness" the most.
3. For specular materials, we found no significant effect of light map orientation (for both elevation and azimuth). The perception of "specularity" was influenced by light mode and shape, with the "Glacier" light map as well as the "bird" shape reducing perceived "specularity" the most.
4. For the perceived "glitteriness" of glittery materials, the effects of direction and shape were significant, with the "bird" shape reducing "glitteriness" the most. Light-

ing mode had only an interaction effect, reducing perceived “glitteriness” the most for the “Glacier” light map for all elevations and “Grace-new” for elevation3.

In a former work (Zhang et al., 2019), we investigated the interaction between material and light modes for one shape only, namely the “bird” shape, and for one lighting direction. In that study, we combined the four canonical material modes (matte, velvety, specular, and glittery) with three canonical lighting modes (ambient/ “Glacier”, focus/”Ennis”, and brilliance/”Grace-new”), and found material dependent lighting effects for nine qualities (“matte”, “velvety”, “specular”, “glittery”, “glossy”, “rough”, “smooth”, “hard”, and “soft”), which were similar to the main lighting effects found in the current study. In particular, the impact of the “Glacier” light map with respect to the other two light maps was similar: reducing perceived “specularity” and “glitteriness” for the specular and glittery materials, respectively, and highlighting “velvetiness” and, to a lesser extent, “matteness” for velvety and matte material, respectively.

We again found a difference between the matte and velvety material modes, on the one hand, and the specular and glittery materials, on the other hand, when considering the effects of shape on the perceived qualities. No systematic effects were found for the matte and velvety materials, whereas both specular and glittery materials showed a reduction in their corresponding perceived qualities for the “bird” shape with respect to the other two shapes. An explanation for this systematic finding is probably that the BRDFs of specular and glittery material are more peaked than those of velvety and matte materials. If the lightings or shapes vary, the appearances of specular and glittery materials then will change more than those of matte and velvety materials.

Interestingly, another differentiation was found in the current study, namely, between matte and specular materials, on the one hand, and velvety and glittery materials, on the other hand. This was based on the (in)sensitivity for lighting direction where velvety and glittery materials showed systematic changes in the quality ratings as a function of azimuth and/or elevation, effects which were absent for the other materials. In Zhang et al. (2019), such a differentiation could also be observed but there it was based on the judgments of roughness and smoothness. In that study, the matte and specular material modes were assessed to be more smooth and velvety and glittery material modes to be more rough. Since image texture (gradients) due to 3D surface corrugations are extremely sensitive to lighting variations (Pont & Koenderink, 2008) these findings might well be related.

The main question we addressed in this paper is how different material and shape and lighting combinations affect perceived material appearance. Endless combinations of materials, shapes, and illuminations may cause a similar appearance while small variations of one of those factors can sometimes cause large variations in appearance. A major challenge is to find a way to predict the appearance within this endless space of possibilities. In order to do so we need to get a grip onto the proximal stimulus, the image, and its features, in contradistinction to the basic physical parameters that determine them (and in the end we also will understand the relationships between the physics and image features). To this end, we want to bring up the notion of the importance map, which characterizes the contribution of different lighting directions depending on surface reflectance (material) and geometry (shape). When we trace the light rays from the viewpoint to the surface back to the environment, we find that more light rays accumulate in some directions than in other directions. An importance map records this accumulation: brighter points in the map correspond to directions where more accumulation has occurred.

Specifically, we show in Figure 6.11 how multiple light rays are traced back to the same direction of an importance map from the surface of a specular bird (not only from the highlights but also the rest of the surface), appearing as a bright spot in the importance map due to the accumulation. It suggests that that spot in the light map has a strong contribution to the final image, which we call “importance”; whereas darker regions in the importance map correspond to directions in the light map that will hardly be reflected to the viewer, i.e. they are less “important”. Since the importance map only depends on shape, material and viewing direction, it is independent of the light map (and its variation after rotation). Hence, if we rotate the light map such that the light sources match the brightest regions in the importance map, the imaged object surface brightens at corresponding locations, depending on its shape and material. Rotating the object (which we did not do in the current study) would also impart a change in the importance map, as the surface shape visible from the viewpoint would change as well.

The twelve importance maps corresponding to the four canonical materials and the three shapes we implemented are shown in Figure S5 in the supplements. We can immediately see that the importance maps of matte objects are robustly diffused and quite symmetric in all shapes, which can explain why perceived “matteness” was found to be constant across all lightings and light map orientations. The importance maps of the velvety material are similar to those of the matte material as they are quite diffused too,

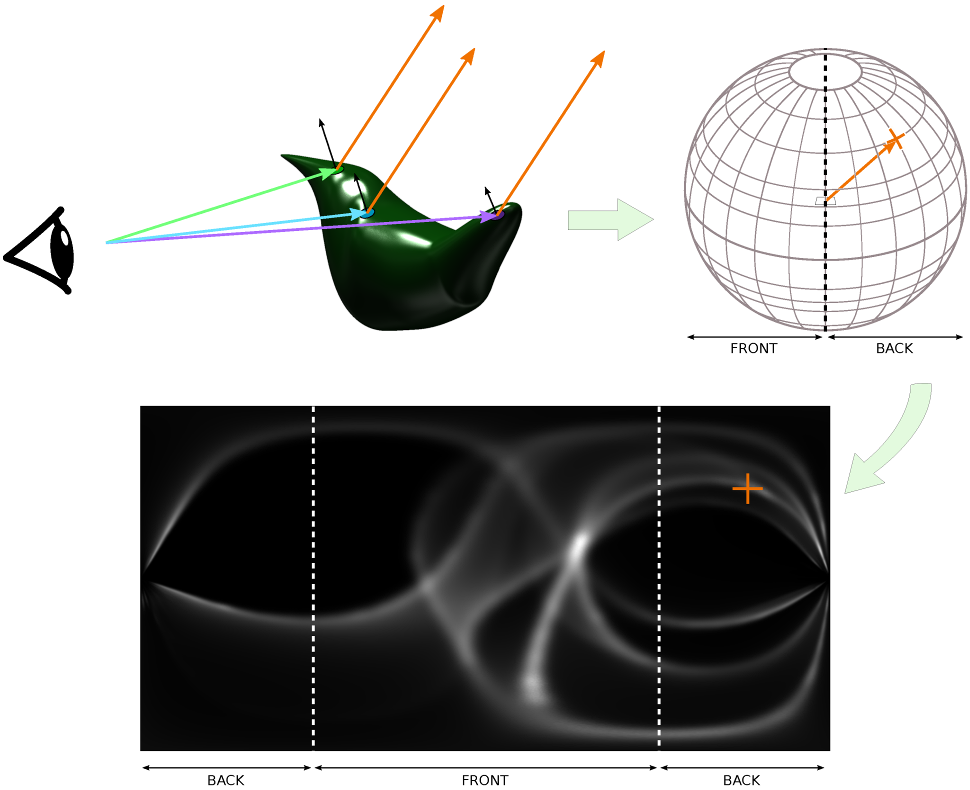

Figure 6.11: A schematic demonstration of importance map construction, taking a specular bird and its importance map as an example. Top left: rays are traced from the viewpoint toward the object, hitting the object surface at different locations; depending on the material's BRDF, new outgoing rays are emitted from these locations. Please note that the way we trace light rays is opposite to how light rays transmit from the sources in lighting environments to object and form an image. Top right: multiple outgoing rays from different positions on the object may have the same direction (in orange); they are then accumulated in the same direction in the distant spherical environment (that is, a single point on the spherical map). Bottom: performing this accumulation for all outgoing rays results in an importance map that we store using a latitude-longitude projection. The central portion of the importance map corresponds to rays that have been projected toward the frontal half of the spherical environment (front), while the sides (where the orange cross is located in our example) corresponds to rays projected towards the rear half of the spherical environment (back). Note that the frontal half of the environment is actually located behind the observer/camera. Brighter regions of the importance map correspond to directions for which more rays have accumulated, due to shape, material or both.

but different in that they show some fine structures for the "blob" and "bird". On the contrary, the importance maps of specular and glittery materials are clearly different for each shape, and show mutually similar asymmetric structures. The importance maps for glittery are more diffuse than those for specular, due to the broadening of the specular peak caused by the distribution of flakes that compose glitter. The (lack of) variations of the importance maps for the shapes explain some of the main effects found in our experiment, namely that perceived "specularity" and "glitteriness" were influenced by shape, while perceived "matteness" and "velvetiness" were not. It also directly shows that the importance maps for specular and glittery materials varied in a more fine-grained way than those of velvety or matte materials, which corresponds to the fact that effects were stronger for materials with peaked BRDFs.

In future work, we would like to explore the use of importance maps to predict how lighting affects image features and thus permits to solve problems such as optimizing lighting for material and shape perception. For example, in combination with metrics they could be used for predicting the strength of image cues such as the sharpness, the contrast and the coverage of the highlights that trigger glossiness perception (Marlow, Kim & Anderson, 2012). In supplementary materials, we illustrate the potential of this approach by showing the product between light maps and importance maps. An example is given in Figure 6.12 where we show the product of the importance map of the "specular bird" with the "Ennis" lighting environment, for two orientations. When the main light source is oriented such that it matches the brightest spot in the importance map (left column in Figure 6.12), large and bright specular highlights appear in the rendered image. When the main light source is oriented such that it matches a less bright spot on the back side of the importance map (right column in Figure 6.12), we observe a small specular highlight on the silhouette of the object.

6.6. CONCLUSION

In this study we primarily investigated how light map orientation and shape influence the visual perception of four canonical materials (matte, velvety, specular, and glittery). Specifically, we performed a rating experiment in which, in each trial, we presented observers fifteen stimuli images that differed in three elevations and five azimuths of the lighting environment, while having the same material, shape, and lighting environment (lighting mode) and instructed them to evaluate the corresponding material quality. Effects of light map orientation were found for velvety and glittery materials, but not for matte and specular materials. Effects of shape were found for specular and glittery materials, but not for matte and velvety materials. Effects of lighting mode were found for velvety and specular materials, but not for matte and glittery materials. Hence, the perception of "matte" for matte materials was found to be the only material quality that is

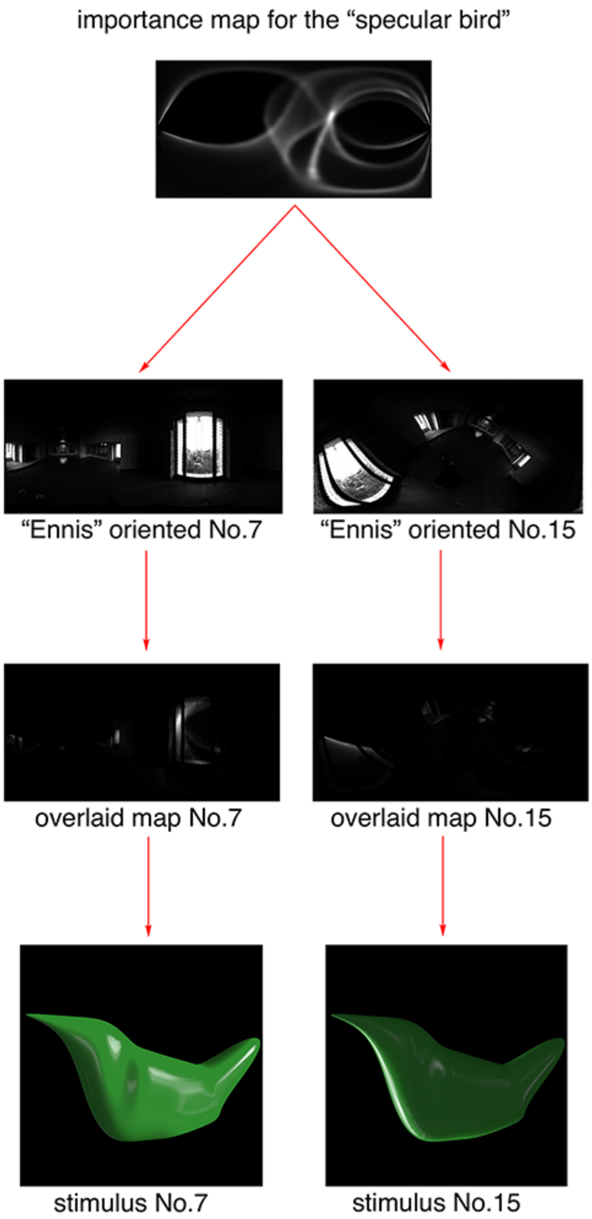

Figure 6.12: An example of the predictive power of the product between an importance map and a light map. From top to bottom, we show the importance map for the "specular bird", the "Ennis" environment map for two orientations, the product between light and importance maps, and the corresponding renderings.

robust across all manipulations of lighting and shape. The results confirmed key image features triggering perceived specularity, glitteriness, velvetiness and matteness.

Bibliography

Adams, W. J., Kucukoglu, G., Landy, M. S., & Mantiuk, R. K. (2018). Naturally glossy: Gloss perception, illumination statistics, and tone mapping. Journal of Vision, 18(13), 4. https://doi.org/10.1167/18.13.4

Barati, B., Karana, E., Sekulovski, D., & Pont, S. C. (2017). Retail lighting and textiles: Designing a lighting probe set. Lighting Research & Technology, 49(2), 173-194.

Basri, R., & Jacobs, D. (2001). Lambertian reflectance and linear subspaces. In Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on (Vol. 2, pp. 383-390). IEEE.

Bousseau, A., Chapoulie, E., Ramamoorthi, R., & Agrawala, M. (2011, June). Optimizing environment maps for material depiction. In Computer graphics forum (Vol. 30, No. 4, pp. 1171-1180). Blackwell Publishing Ltd.

Brainard, D. H. (1997). The psychophysics toolbox. Spatial vision, 10, 433-436.

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

Doerschner, K., Boyaci, H., & Maloney, L. T. (2007). Testing limits on matte surface color perception in three-dimensional scenes with complex light fields. Vision Research, 47(28), 3409-3423.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2010). Estimating the glossiness transfer function induced by illumination change and testing its transitivity. Journal of Vision, 10(4), 8-8.

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

Fleming, R.W., Dror, R.O., & Adelson, E.H. (2003). Real-world illumination and the perception of surface reflectance properties. Journal of vision 3 (5), 3-3

Ganslandt, R., & Hofmann, H. (1992). Handbook Of Interior Lighting. Erco Edition. Germany.

Kartashova, T., Sekulovski, D., de Ridder, H., te Pas, S. F., & Pont, S. C. (2016). The global structure of the visual light field and its relation to the physical light field. Journal of Vision, 16(10), 9-9. doi:10.1167/16.10.9

Kelly, R. (1952). Lighting as an integral part of architecture. College Art Journal, 12(1), 24-30.

Kim, J., Marlow, P. J., & Anderson, B. L. (2012). The dark side of gloss. Nature neuroscience, 15(11), 1590.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. Perception, 36(14), 1.

Koenderink, J. J., & van Doorn, A. J. (1980). Photometric invariants related to solid shape. Optica Acta: International Journal of Optics, 27(7), 981-996.

Koenderink, J., & Pont, S. (2003). The secret of velvety skin. Machine vision and applications, 14(4), 260-268.

Kunsberg, B., & Zucker, S. W. (2013). Characterizing ambiguity in light source invariant shape from shading. arXiv preprint arXiv:1306.5480.

Kunsberg, B., & Zucker, S. W. (2018). Critical Contours: An Invariant Linking Image Flow with Salient Surface Organization. SIAM Journal on Imaging Sciences, 11(3), 1849-1877.

Marlow, P., Kim, J., & Anderson, B. L. (2011). The role of brightness and orientation congruence in the perception of surface gloss. Journal of Vision, 11(9), 16-16.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. Current Biology, 22(20), 1909-1913. https://doi.org/10.1016/j.cub.2012.08.00

Marlow, P. J., & Anderson, B. L. (2013). Generative constraints on image cues for perceived gloss. Journal of vision, 13(14), 2-2.

Marlow, P.J., & Anderson, B.L. (2016) Motion and texture shape cues modulate perceived material properties. Journal of Vision,16(1):5, 1-14. doi:10.1167/16.1.5

Motoyoshi, I., & Matoba, H. (2012). Variability in constancy of the perceived surface reflectance across different illumination statistics. Vision Research, 53(1), 30-39.

Mury, A. A., Pont, S. C., & Koenderink, J. J. (2007). Light field constancy within natural scenes. Applied Optics, 46(29), 7308-7316.

Narasimhan, S. G., Ramesh, V., & Nayar, S. K. (2003). A class of photometric invari-

6

ants: Separating material from shape and illumination. In Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on (pp. 1387-1394). IEEE.

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., & Limperis, T. (1992, January). Geometrical considerations and nomenclature for reflectance. In Radiometry (pp. 94-145). Jones and Bartlett Publishers, Inc..

Nishida, S. Y., & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties. JOSA A, 15(12), 2951-2965.

Olkkonen, M., & Brainard, D. H. (2010). Perceived glossiness and lightness under real-world illumination. Journal of vision, 10(9), 5-5.

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, Spatial Vision 10:437-442.

Pont, S. C. (2009, February). Ecological optics of natural materials and light fields. In Human Vision and Electronic Imaging XIV (Vol. 7240, p. 724009). International Society for Optics and Photonics. doi: 10.1117/12.817162

Pont, S. C. (2013). Spatial and form-giving qualities of light. In: The Wiley-Blackwell Handbook of Experimental Phenomenology. Visual Perception of Shape, Space and Appearance (L. Albertazzi Ed.). Wiley-Blackwell, Ch.8, 205-222

Pont, S. C., & de Ridder, H. (2018). Lighting perceptual intelligence. In Human Vision and Electronic Imaging 2018

Pont, S. C., & Koenderink, J. J. (2008). Shape, Surface Roughness and Human Perception. In Handbook of texture analysis (pp. 197-222).

Pont, S. C., Koenderink, J. J., van Doorn, A. J., Wijntjes, M. W., & te Pas, S. F. (2012, February). Mixing material modes. In Human Vision and Electronic Imaging XVII (Vol. 8291, p. 82910D). International Society for Optics and Photonics.

Pont, S. C., & te Pas, S. F. (2006). Material—Illumination ambiguities and the perception of solid objects. Perception, 35(10), 1331-1350.

Schirillo, J. A. (2013). We infer light in space. Psychonomic bulletin & review, 20(5), 905-915.

van Assen, J. J. R., Wijntjes, M. W., & Pont, S. C. (2016). Highlight shapes and perception of gloss for real and photographed objects. Journal of vision, 16(6), 6-6.

Vangorp, P., Laurijssen, J., & Dutré, P. (2007, August). The influence of shape on the perception of material reflectance. In ACM Transactions on Graphics (TOG) (Vol. 26, No. 3, p. 77). ACM.

Vergne, R., & Barla, P. (2015). Designing gratin, a gpu-tailored node-based system. Journal of Computer Graphics Techniques, 4(4), 54-71.

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. ACM SIGGRAPH Computer Graphics, 26(2), 265-272.

Walter, B. (2005). Notes on the Ward BRDF. Program of Computer Graphics, Cornell University, Technical report PCG-05, 6.

Wendt, G., & Faul, F. (2017). Increasing the complexity of the illumination may reduce gloss constancy. i-Perception, 8(6), 204166951774036. https://doi.org/10.1177/2041669517740369

Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., & Ma, Y. (2010). Robust photometric stereo via low-rank matrix completion and recovery. In Asian Conference on Computer Vision (pp. 703-717). Springer, Berlin, Heidelberg.

Xia, L., Pont, S. C., & Heynderickx, I. (2014). The visual light field in real scenes. i-Perception, 5(7), 613-629.

Xia, L., Pont, S. C., & Heynderickx, I. (2016). Effects of scene content and layout on the perceived light map orientation in 3D spaces. Journal of vision, 16(10), 14-14.

Xia, L., Pont, S. C., & Heynderickx, I. (2017). Light diffuseness metric part 1: Theory. Lighting Research & Technology, 49(4), 411-427.

Zhang, F., de Ridder, H., Fleming, R. W., & Pont, S C.. (2016). MatMix 1.0: Using optical mixing to probe visual material perception. Journal of vision, 16(6), 11-11. doi:10.1167/16.6.11

Zhang, F., de Ridder, H., & Pont, S C.. (2015, March). The influence of lighting on visual perception of material qualities. In Human Vision and Electronic Imaging (p. 93940Q). doi: 10.1117/12.2085021

Zhang, F., de Ridder, H., & Pont, S C..(2018). Asymmetric perceptual confounds between canonical lightings and materials. Journal of vision, 18(11), 11-11. doi:10.1167/18.11.11

Zhang, F., de Ridder, H., Barla, P., & Pont, S C. (2019). A systematic approach to testing and predicting light-material interactions. Journal of vision, 19(4), 11-11. doi:10.1167/19.4.11

Figure S1. The influence of azimuths on the rating results for matte material. The first column subplots on the left show the rating per azimuth. The stimuli are shown on the right. The number 1 – 15 on the x-axis represent the orientation, corresponding to the number shown in Figure 1. (A) – (C): under the "Glacier" map. (D) – (F): under the "Ennis" map. (G) – (I): under the "Grace-new" map.

(A)

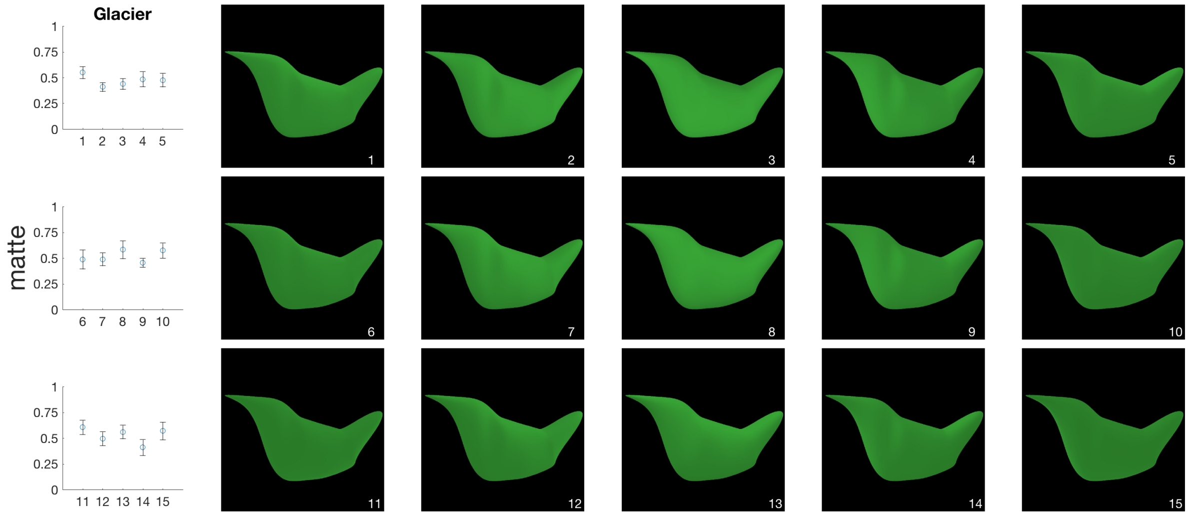

(B)

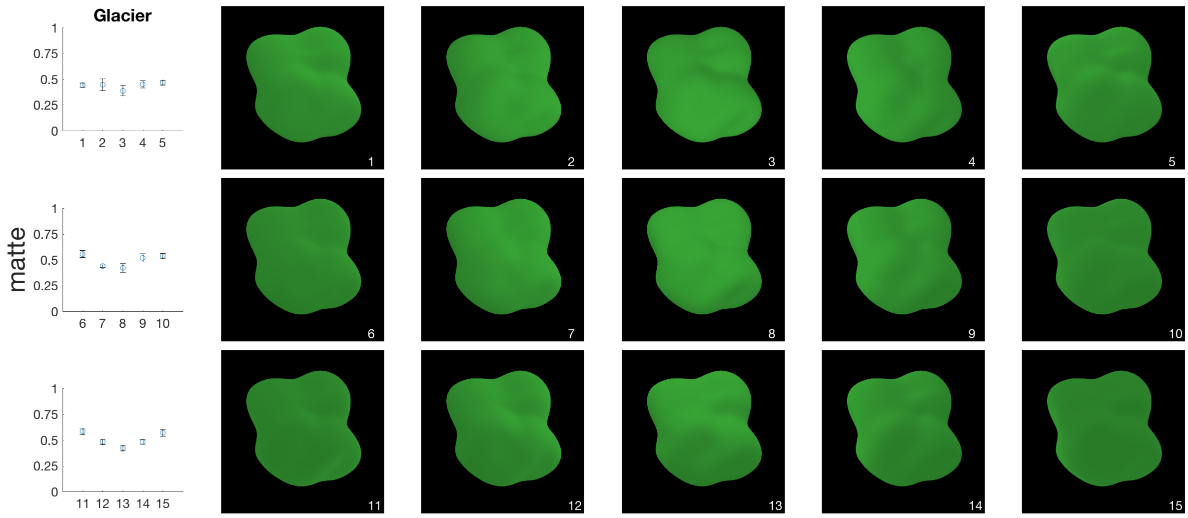

(C)

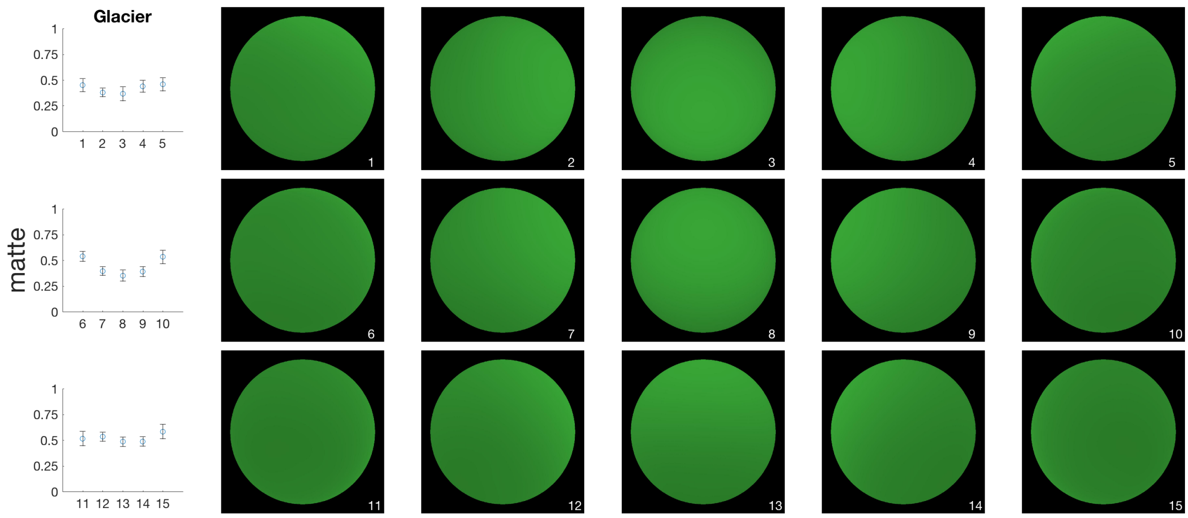

(D)

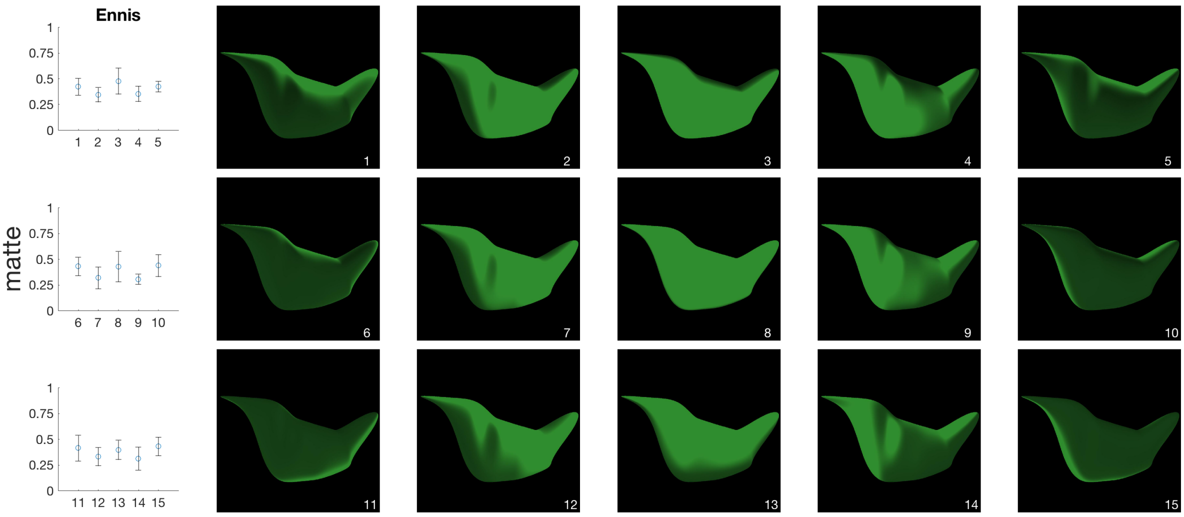
Ennis
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(E)

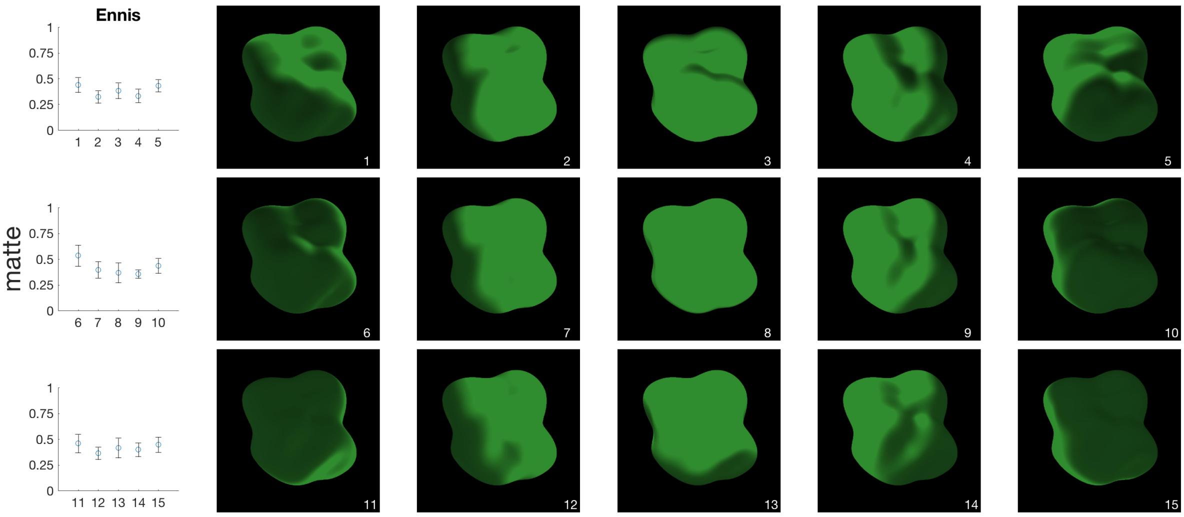
Ennis
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(F)

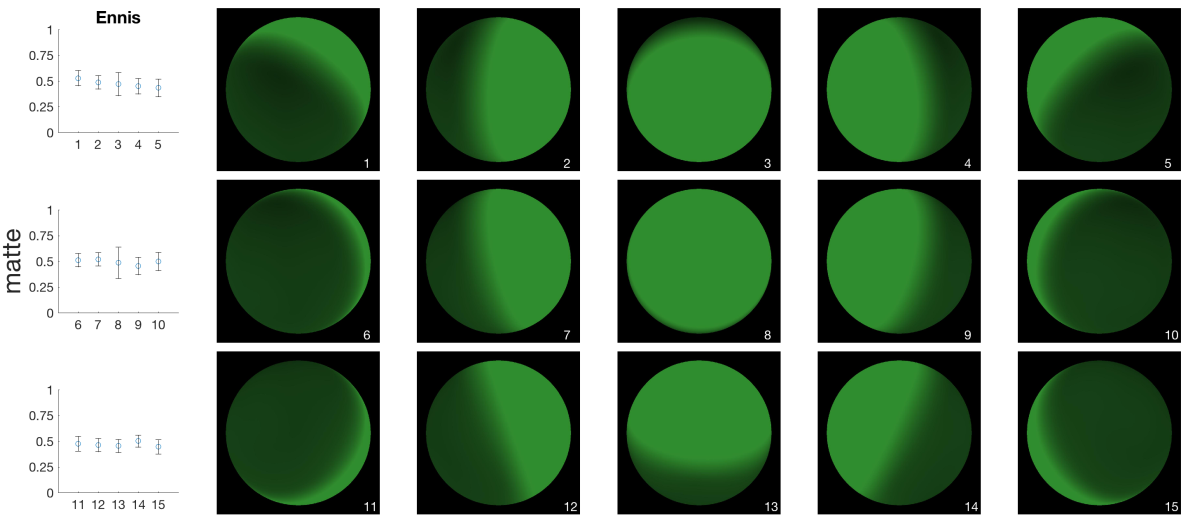
Ennis
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(G)

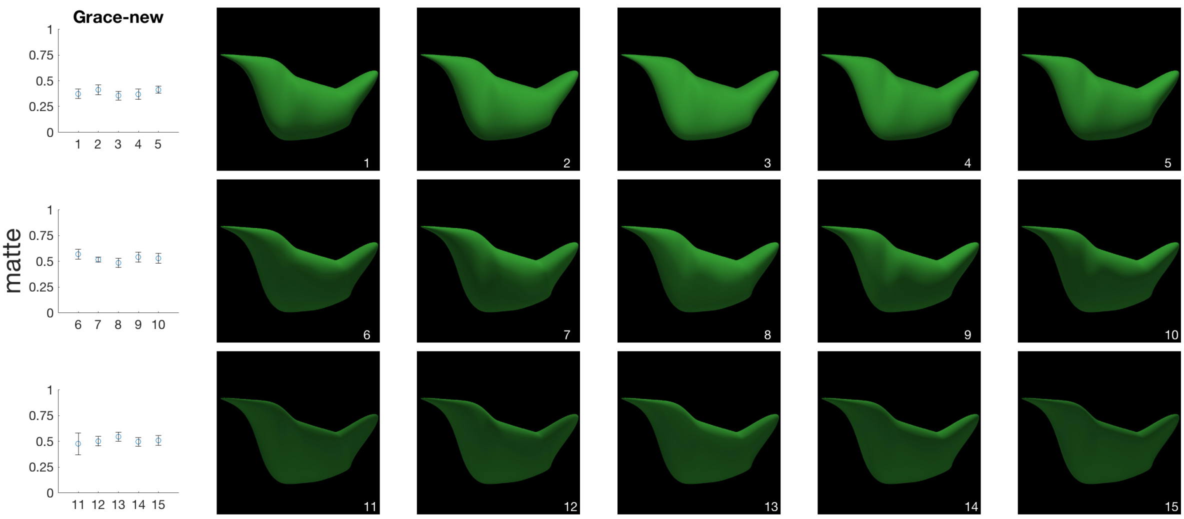
Grace-new
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(H)

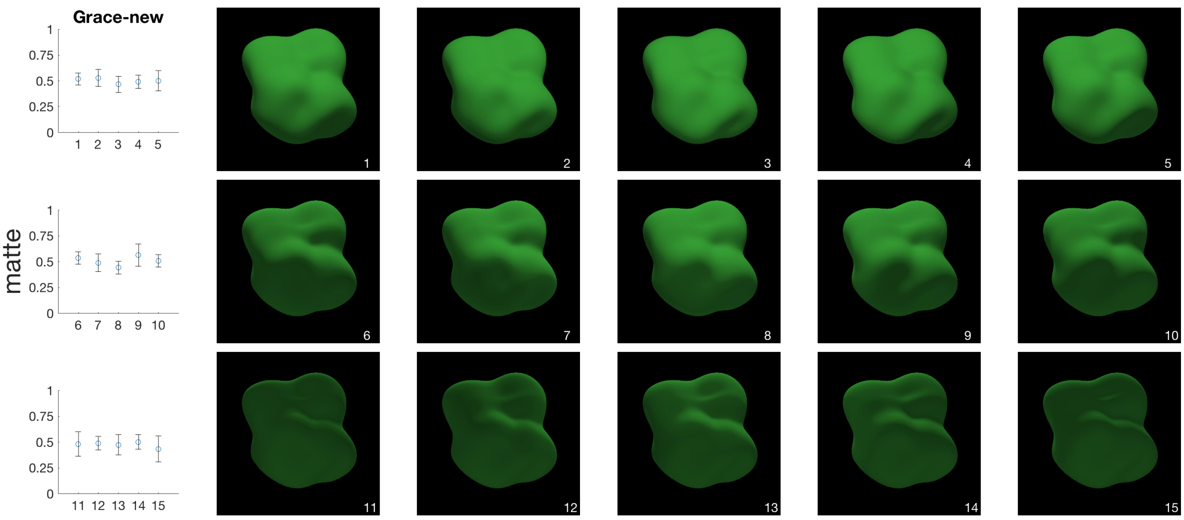
Grace-new
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(I)

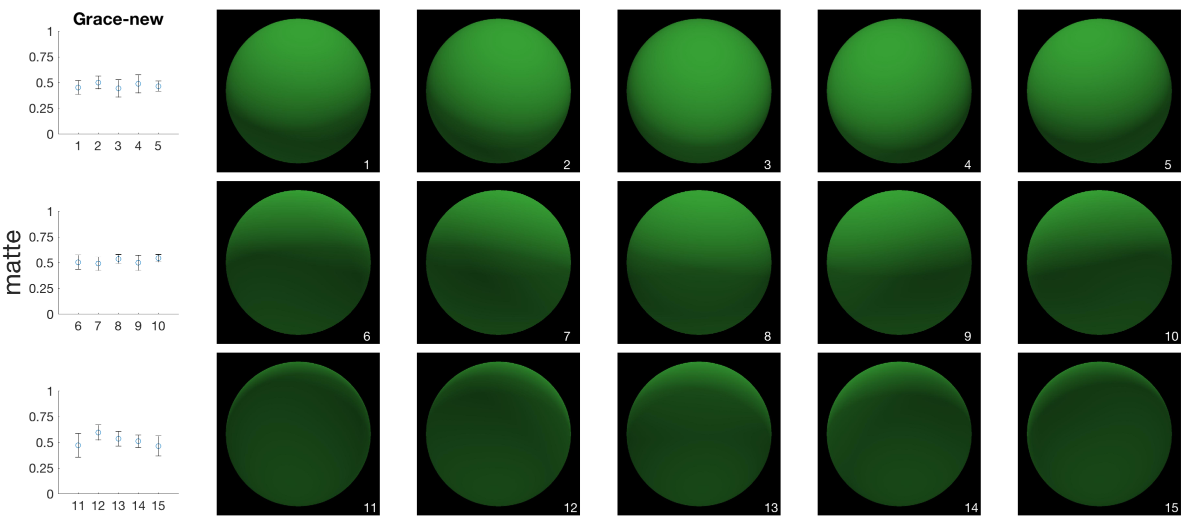
Grace-new
matte
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

Figure S2. The influence of azimuths on the rating results for velvety material. The first column subplots on the left show the rating per azimuth. The stimuli are shown on the right. The number 1 – 15 on the x-axis represent the orientation, corresponding to the number shown in Figure 1. (A) – (C): under the “Glacier” map. (D) – (F): under the “Ennis” map. (G) – (I): under the “Grace-new” map.

(A)

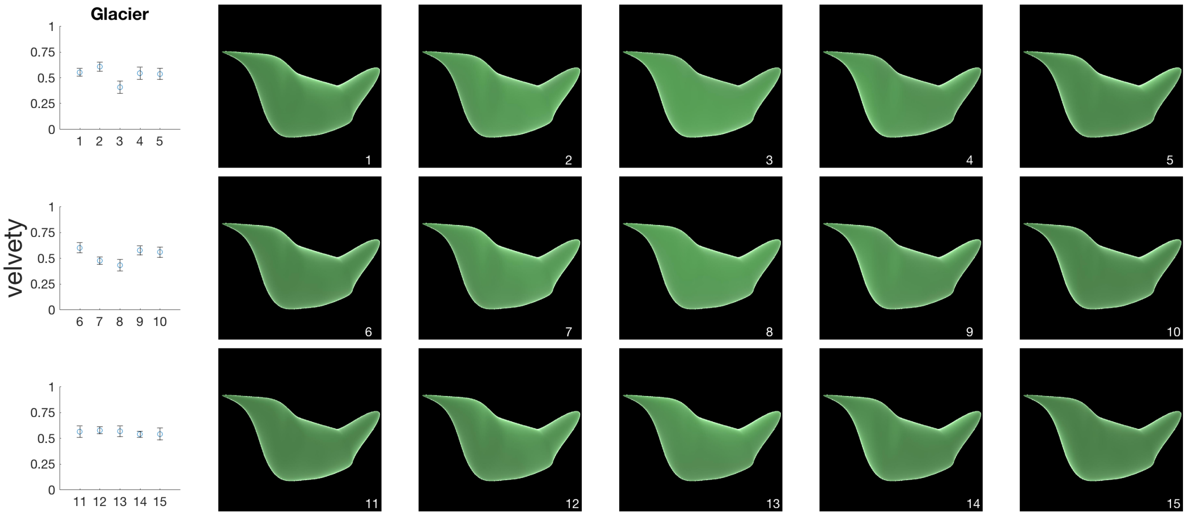

(B)

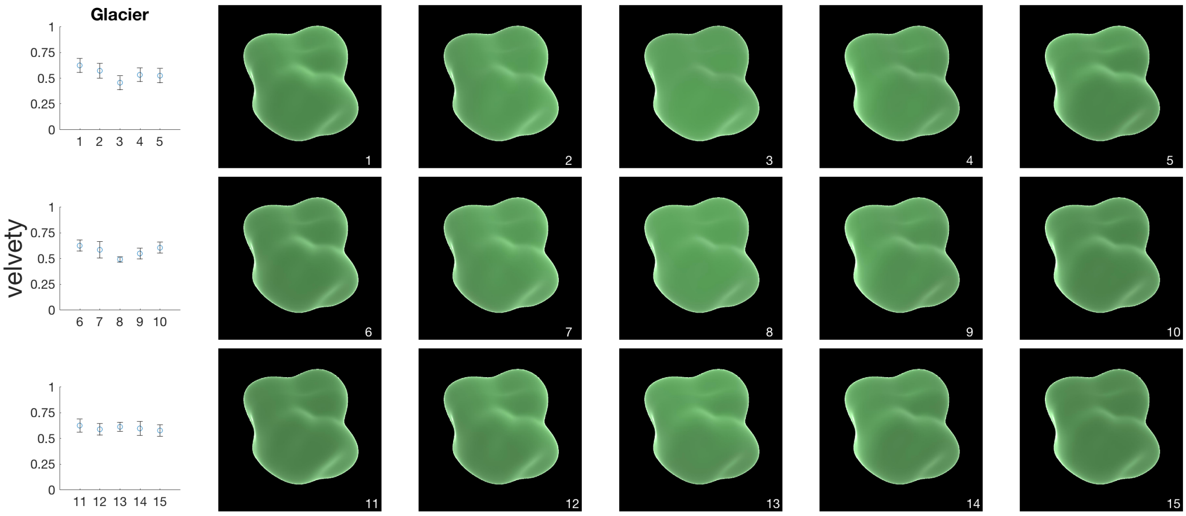

(C)

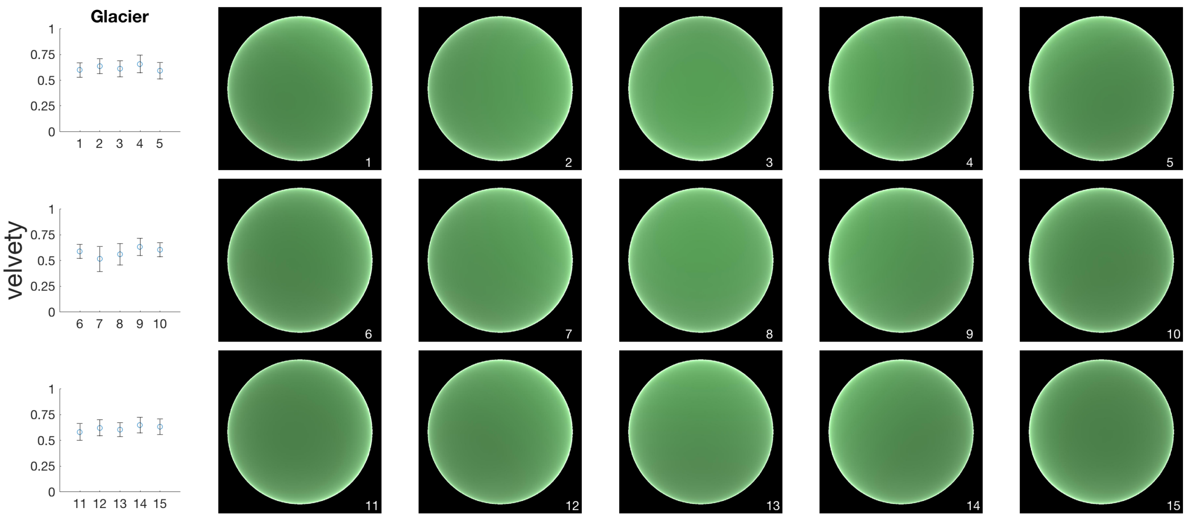

(D)

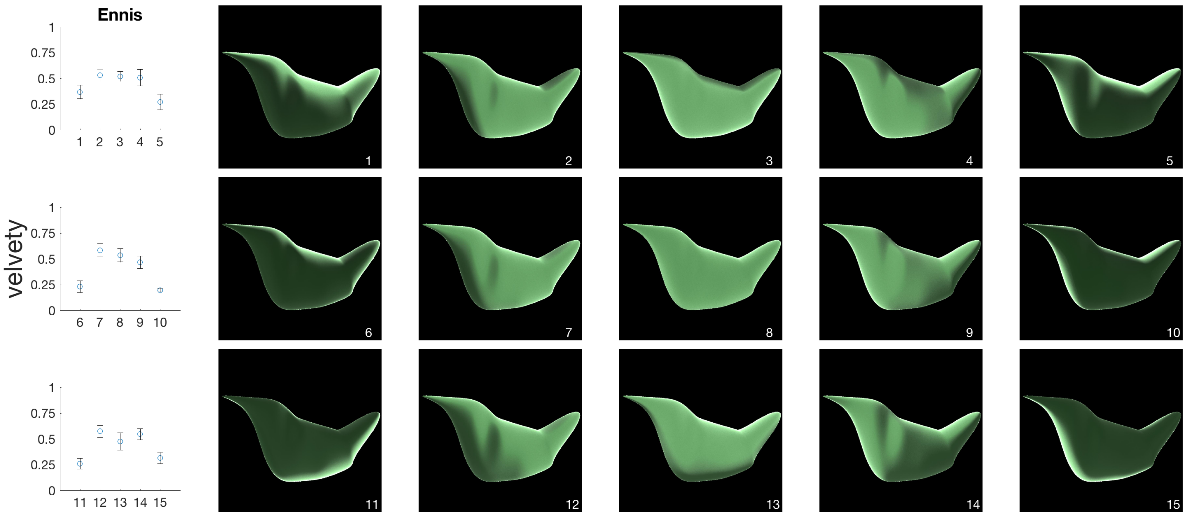

(E)

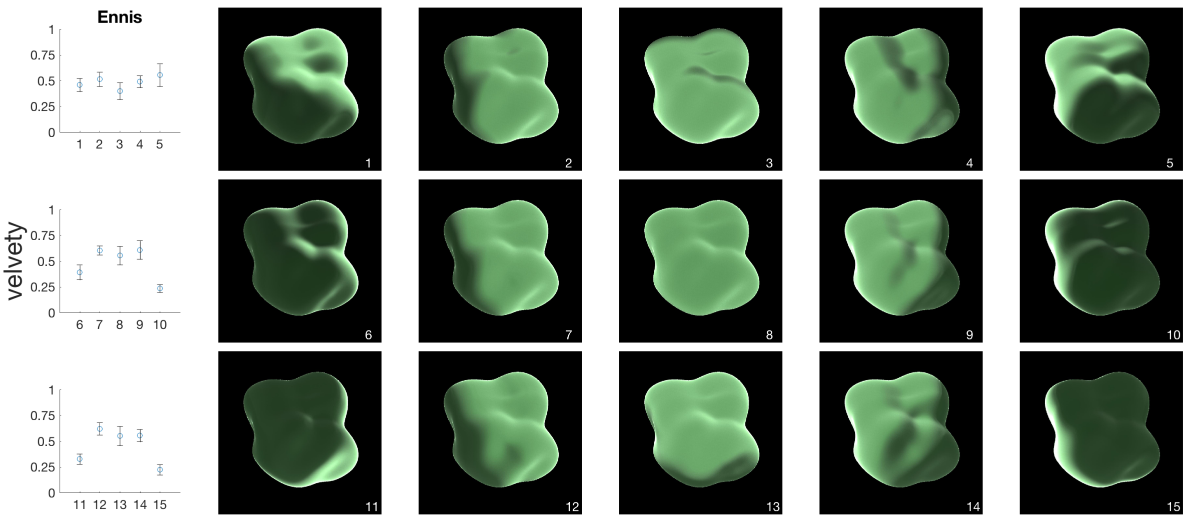

(F)

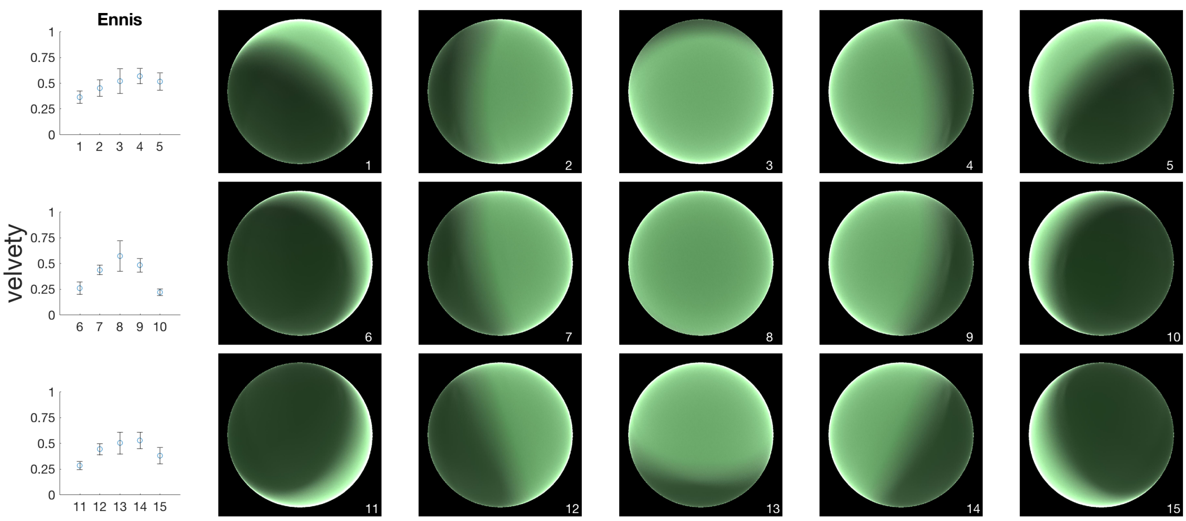

(G)

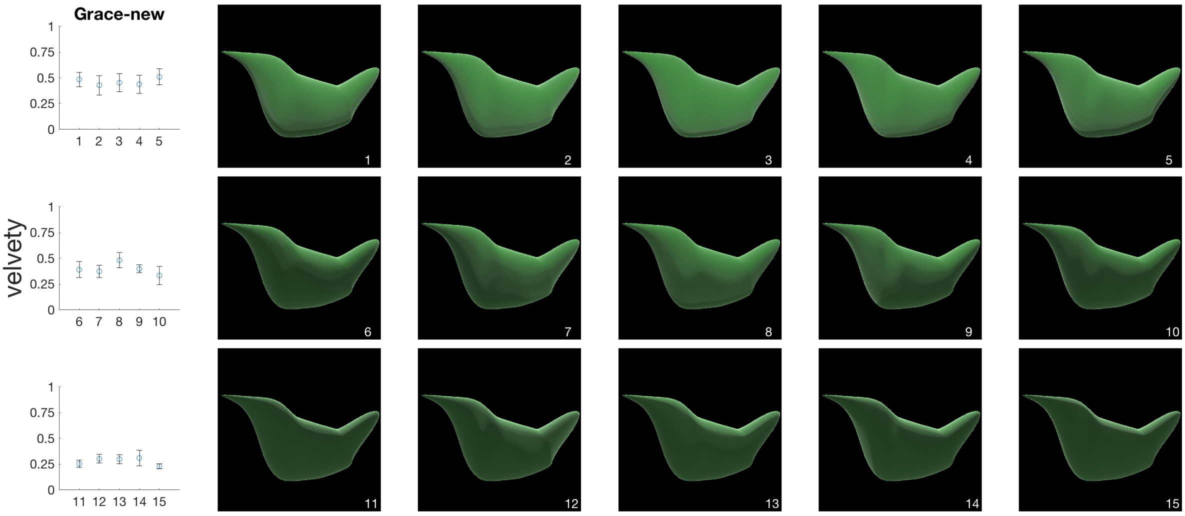
Grace-new
velvety
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(H)

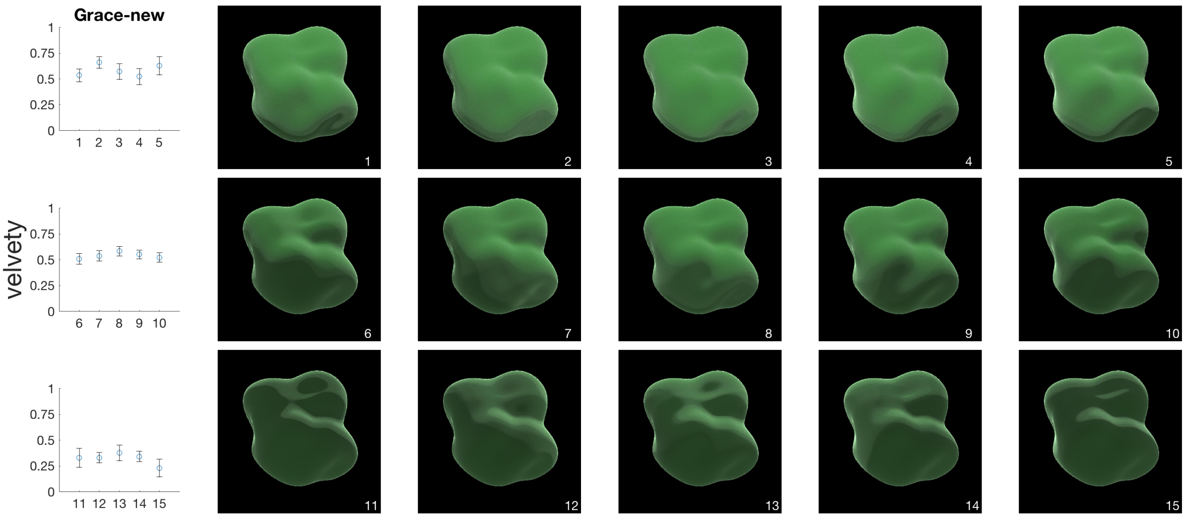
Grace-new
velvety
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(I)

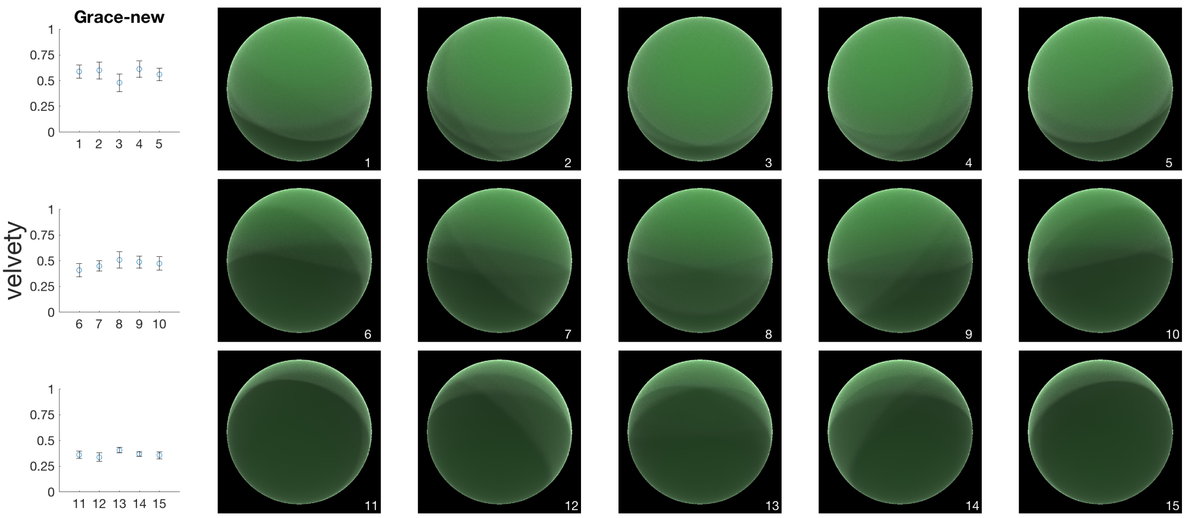
Grace-new
velvety
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

Figure S3. The influence of azimuths on the rating results for specular material. The first column subplots on the left show the rating per azimuth. The stimuli are shown on the right. The number 1 – 15 on the x-axis represent the orientation, corresponding to the number shown in Figure 1. (A) – (C): under the “Glacier” map. (D) – (F): under the “Ennis” map. (G) – (I): under the “Grace-new” map.

(A)

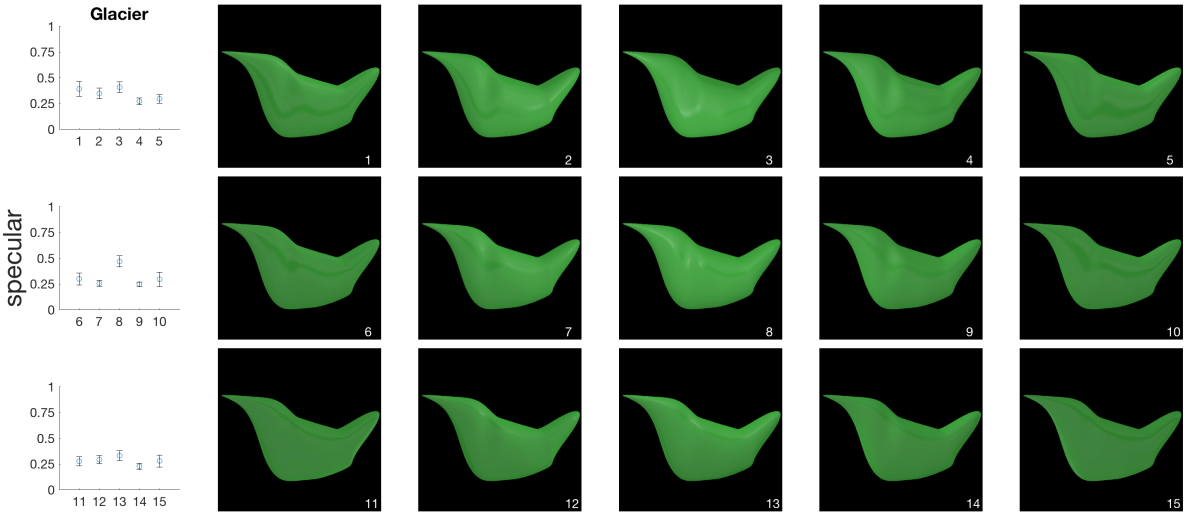

(B)

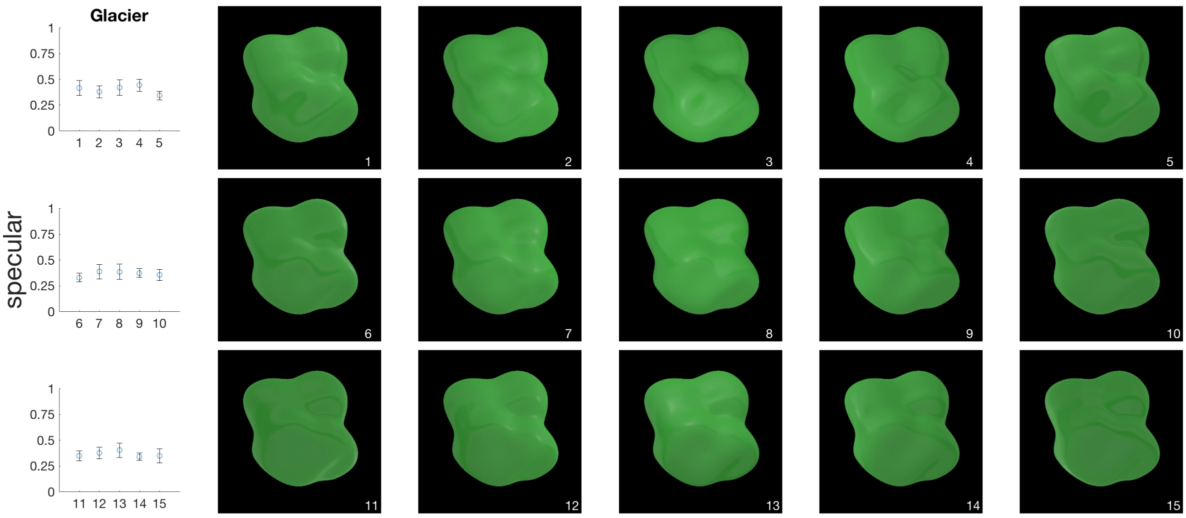

(C)

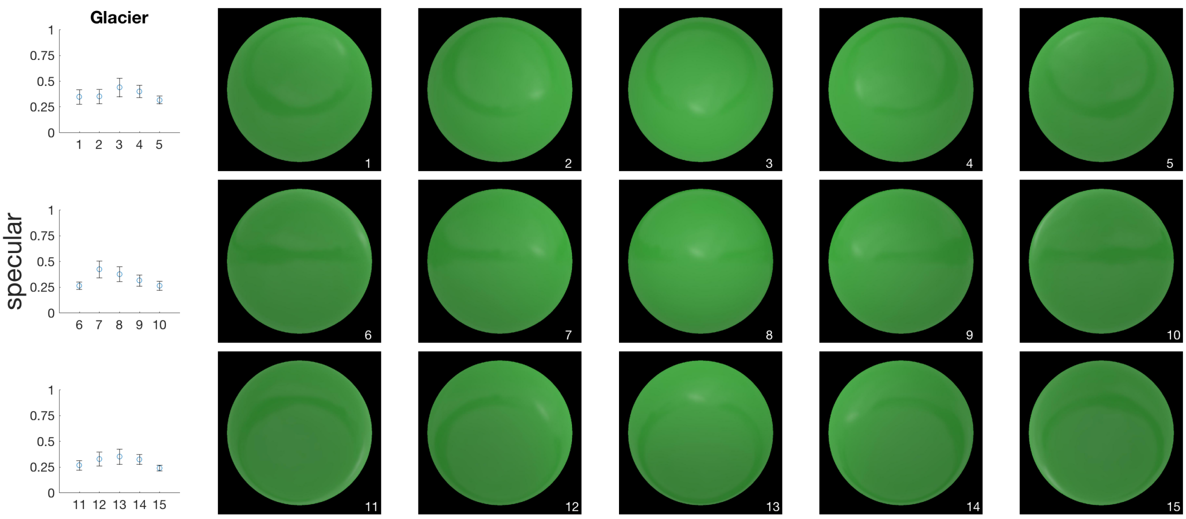

(D)

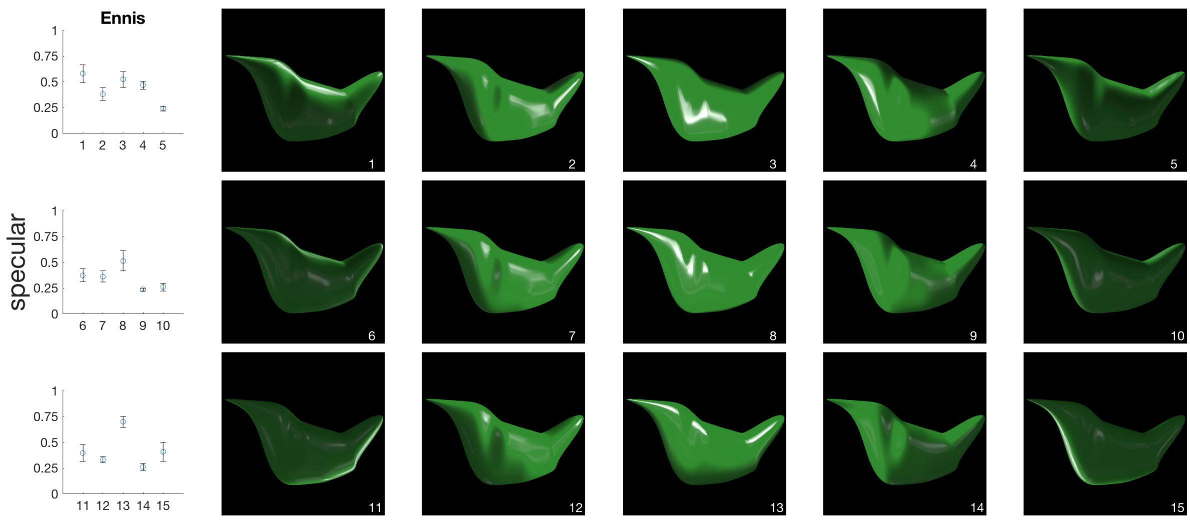
Ennis
specular
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(E)

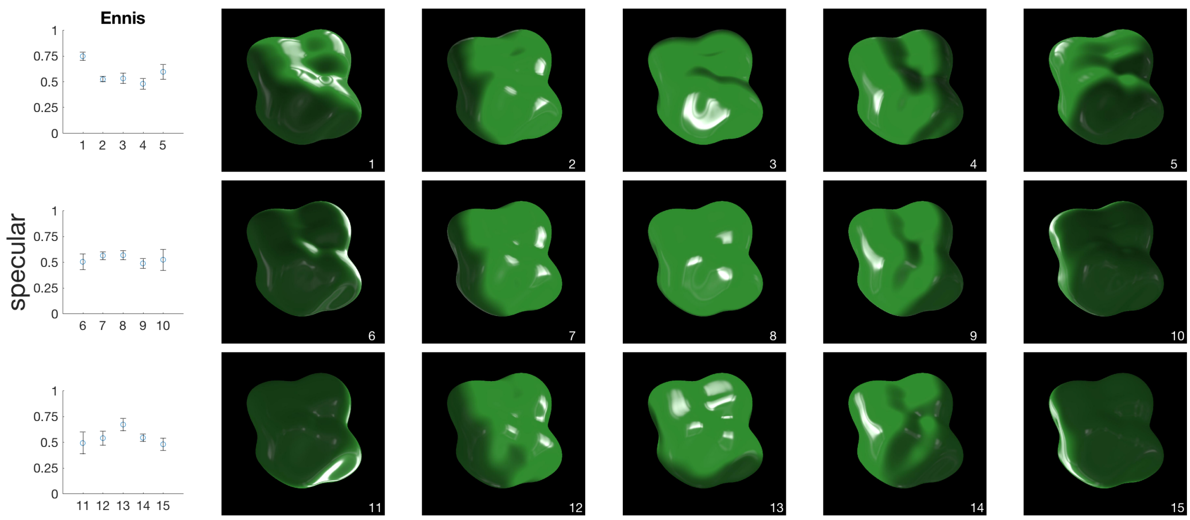
Ennis
specular
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(F)

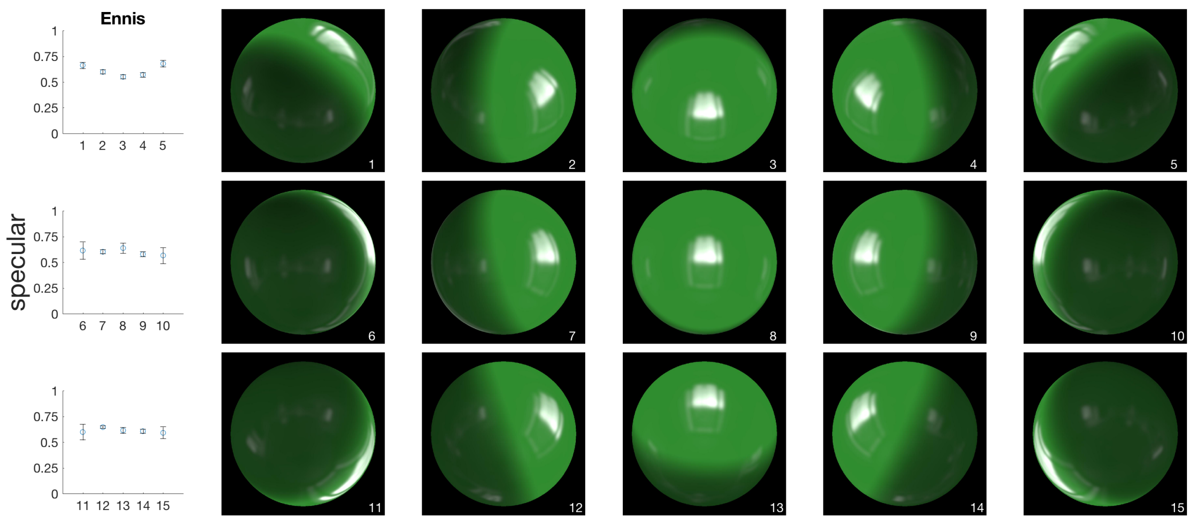
Ennis
specular
1
0.75
0.5
0.25
0
1 2 3 4 5
6 7 8 9 10
11 12 13 14 15

(G)

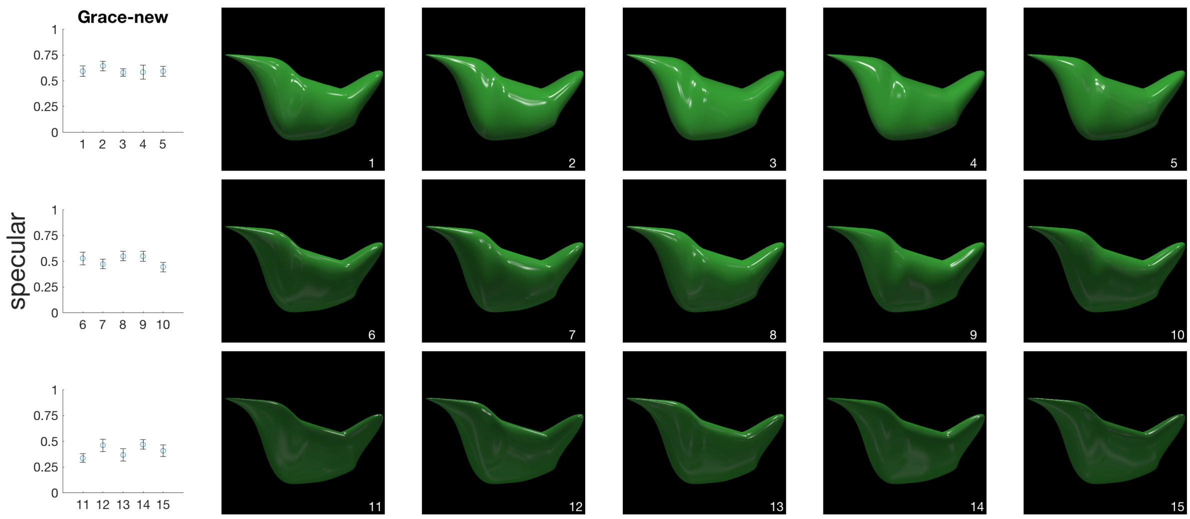

(H)

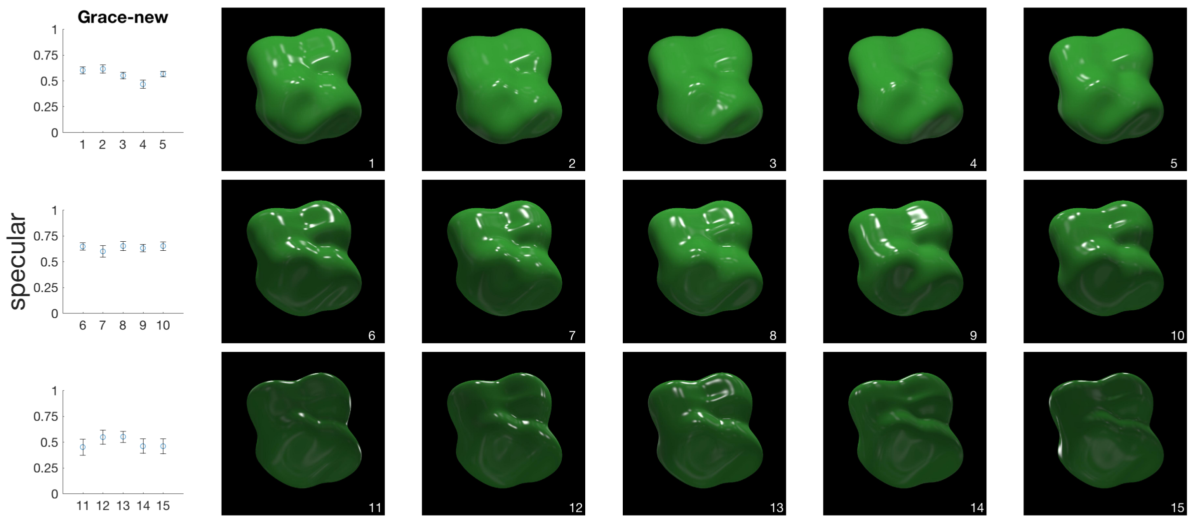

(I)

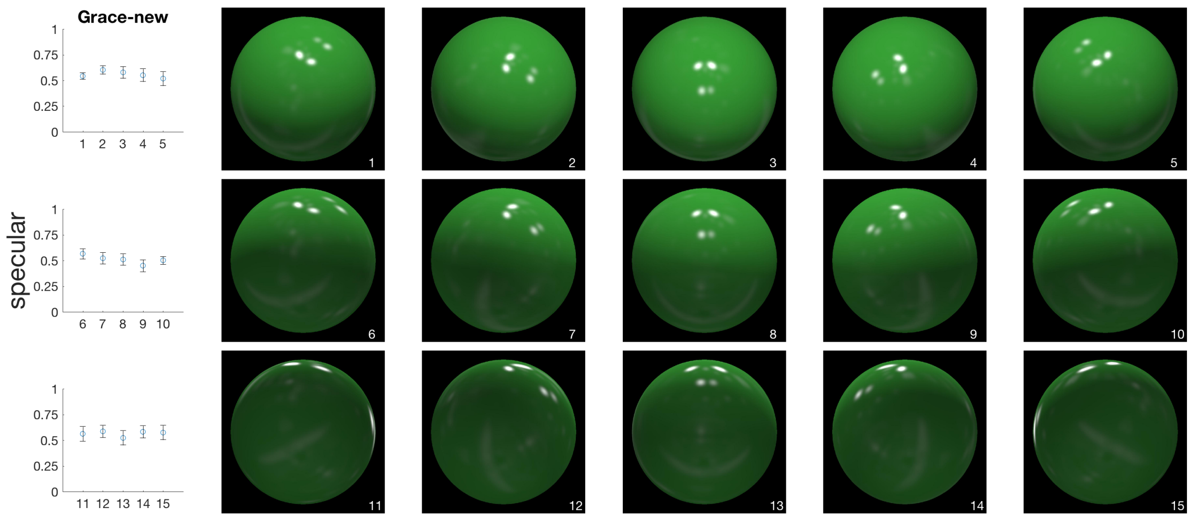

Figure S4. The influence of azimuths on the rating results for glittery material. The first column subplots on the left show the rating per azimuth. The stimuli are shown on the right. The number 1 – 15 on the x-axis represent the orientation, corresponding to the number shown in Figure 1. (A) – (C): under the “Glacier” map. (D) – (F): under the “Ennis” map. (G) – (I): under the “Grace-new” map.

(A)

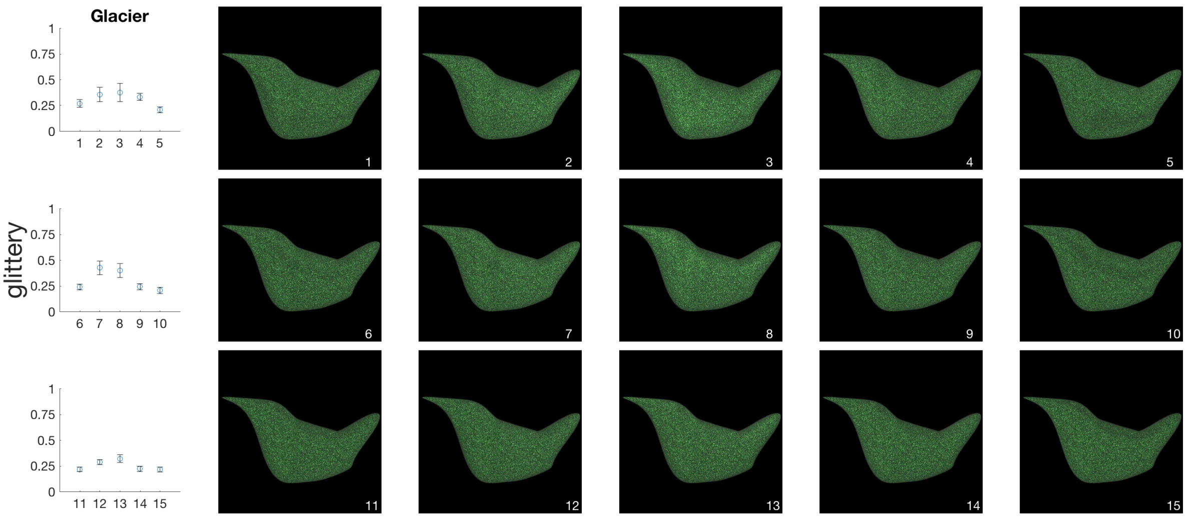

(B)

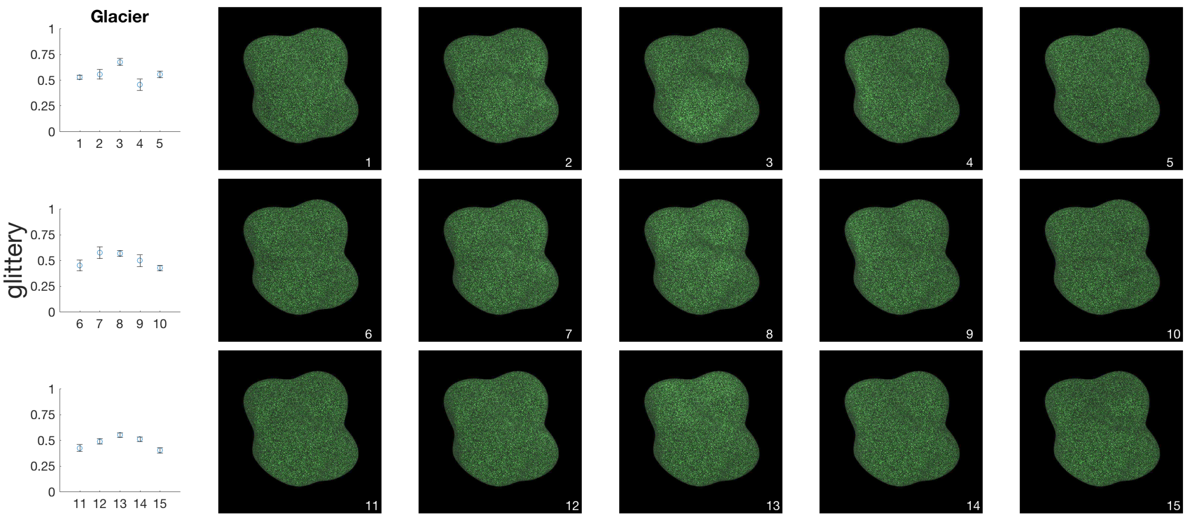

(C)

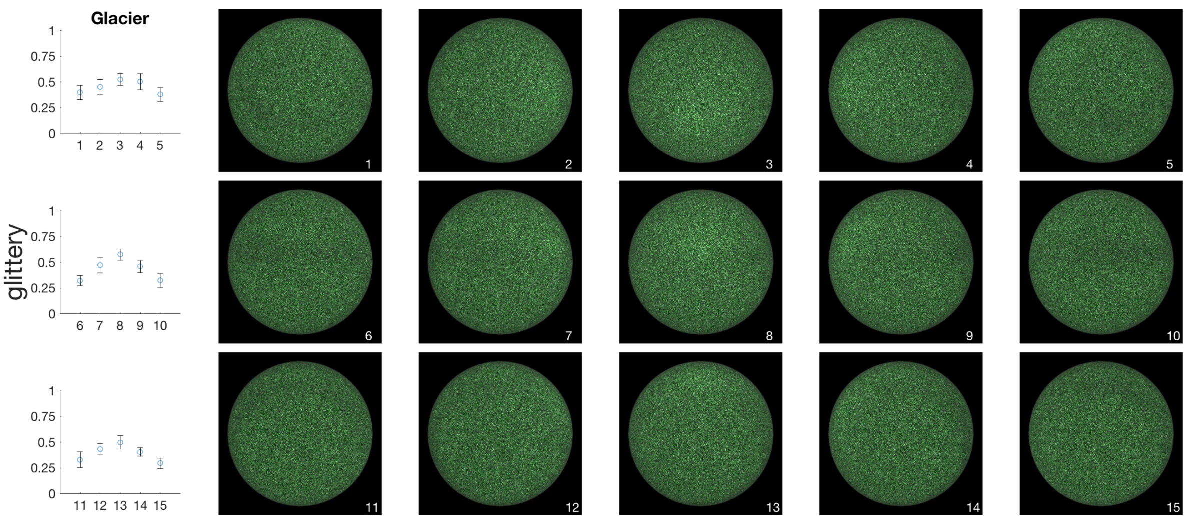

(D)

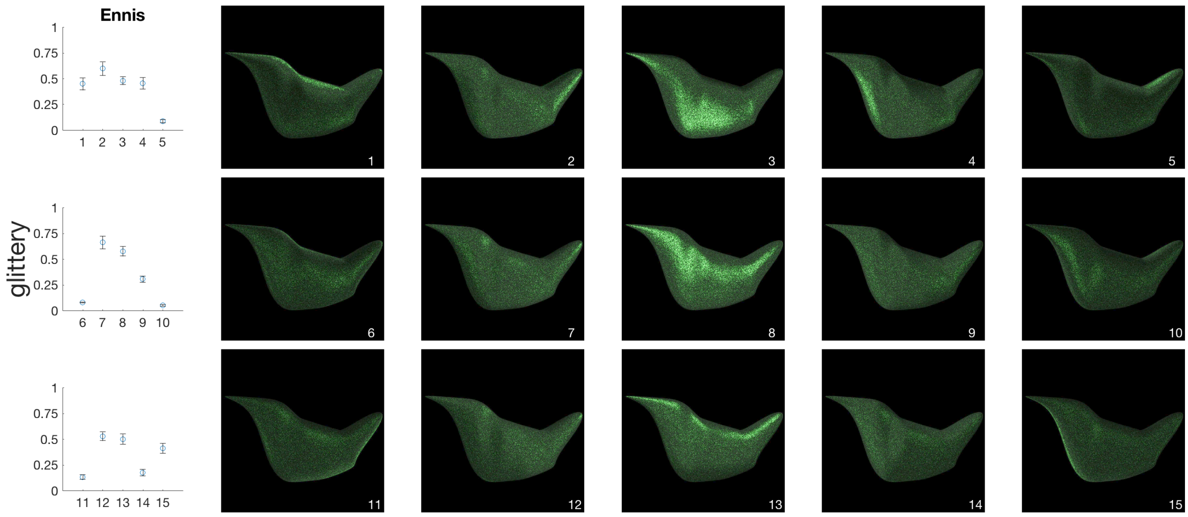

(E)

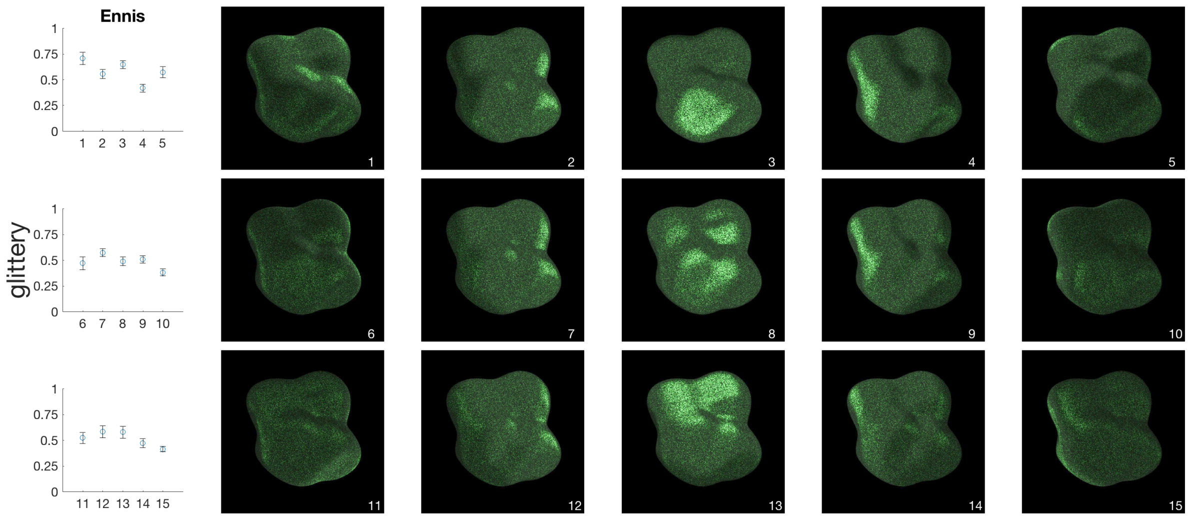

(F)

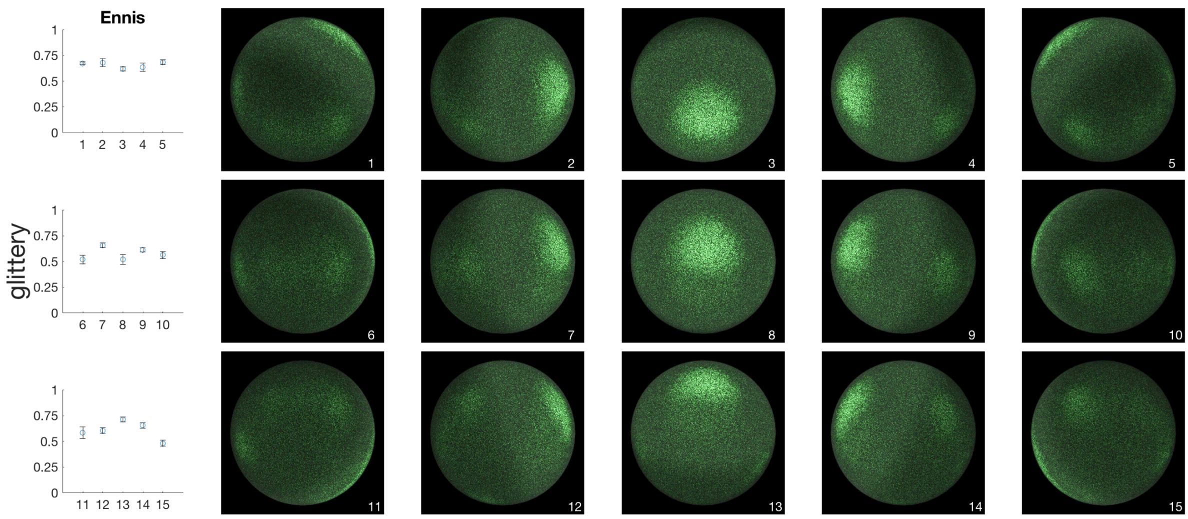

(G)

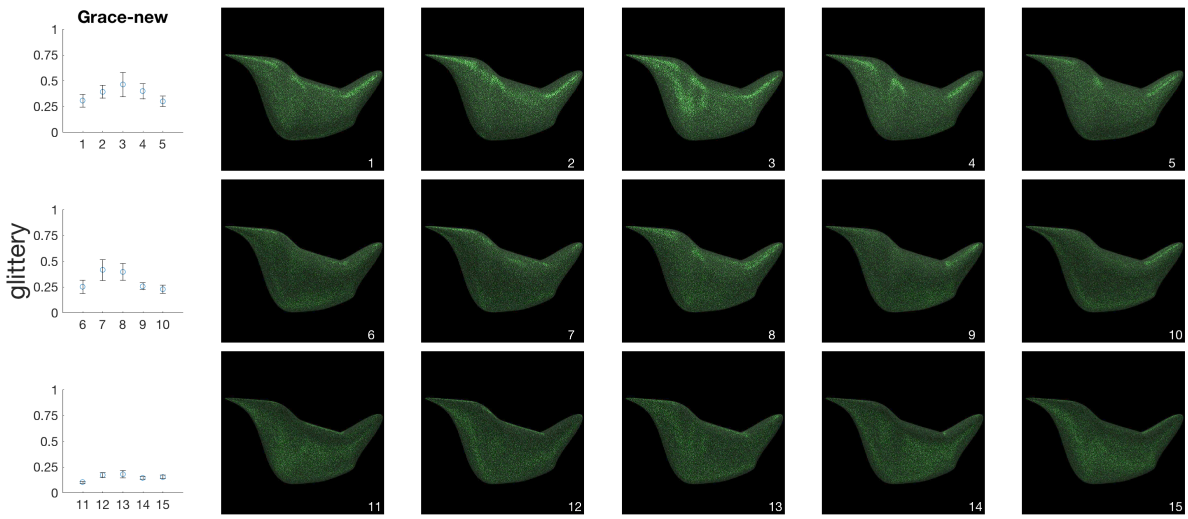

(H)

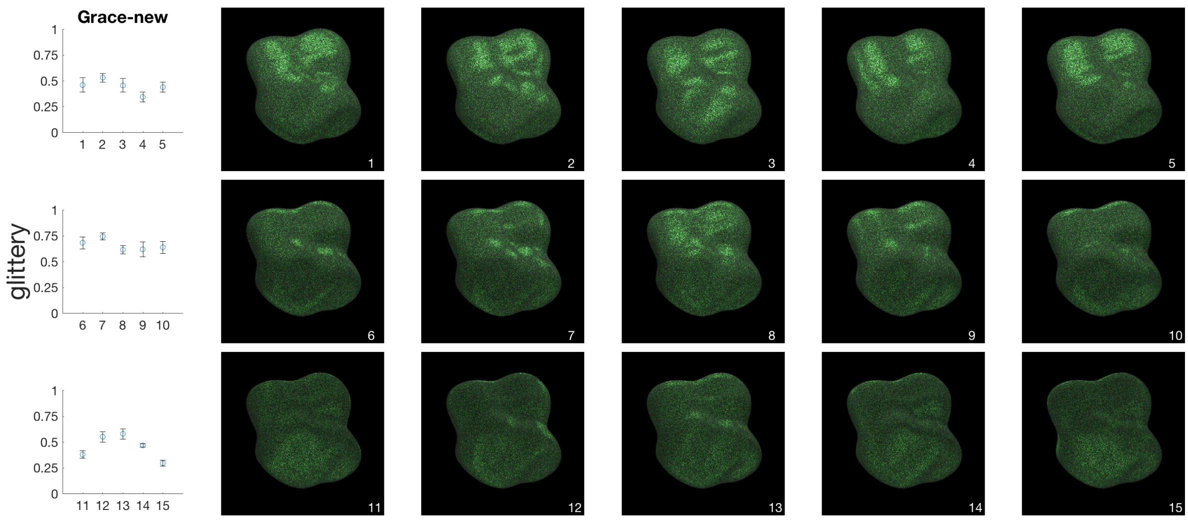

(I)

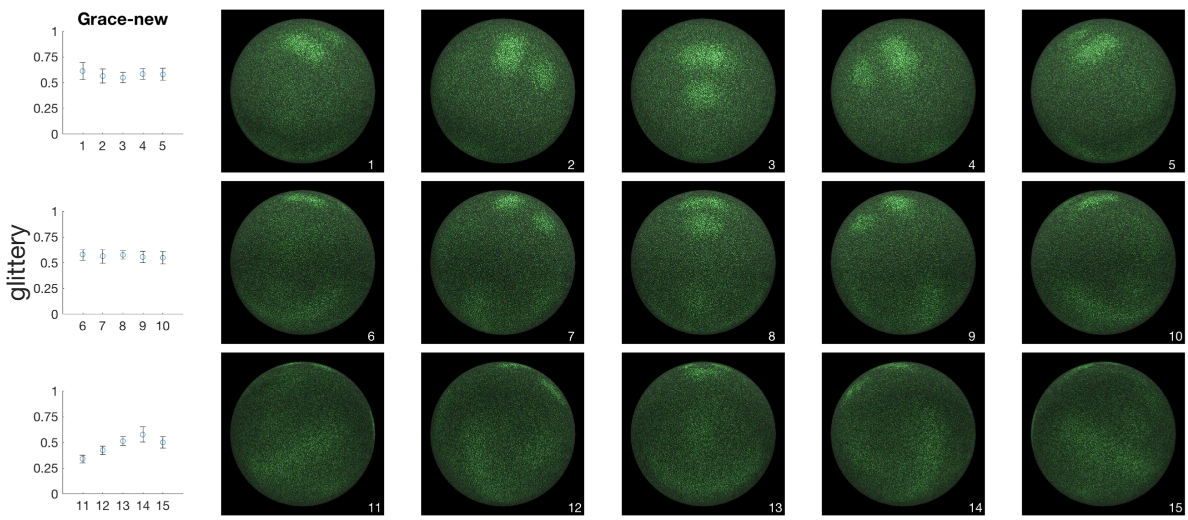

Figure S5. Twelve importance maps combining four canonical material modes (from top to bottom: matte, velvety, specular, and glittery) and three shapes (from left to right: “bird”, “blob”, “sphere”). As in Figure 10, the grayscale represents the amount of light ray accumulation from the environment that illuminate the object. The brighter the pixels are, the more light rays accumulated in the corresponding illumination direction. Some artifacts are present in the back region of the importance map for the specular sphere; these are due to numerical issues in the way we compute the importance map.

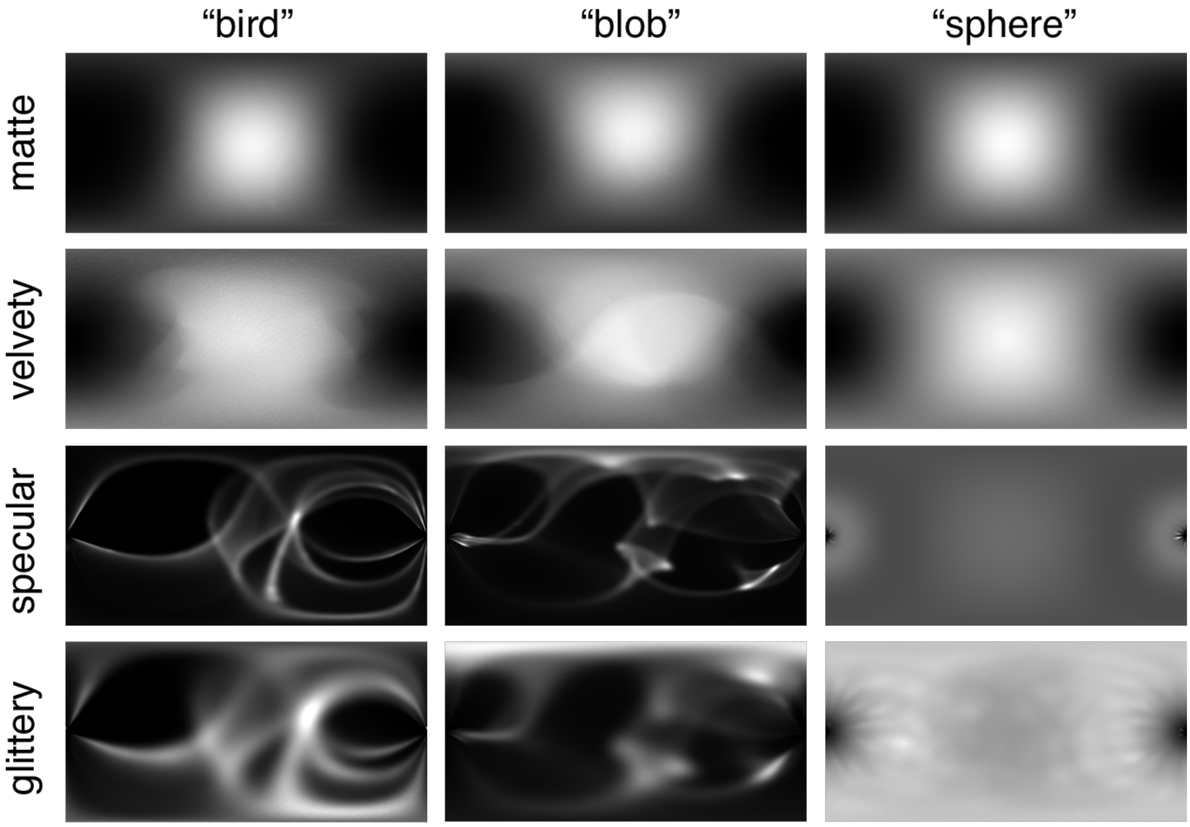

7

Conclusion

7.1. MAIN FINDINGS AND CONTRIBUTIONS

Material perception is multimodal, involving all senses in our daily routine. The physical world consists of materials that can be seen, touched, or interacted with in all kinds of ways. Among all sensory actions, vision plays an important role in material perception. The visual system is able to recover the intrinsic information about the physical properties of the objects and materials in a relatively consistent manner, regardless of the perceptual interactions between lighting, shapes, and materials (Adelson, 2001; Anderson, 2011; Fleming, 2014, 2017). The scientific aim of this thesis was to systematically measure human visual perception of materials and test the influence of lighting and shape on material perception. We developed a novel quantitative probing method that gives consistent results while providing purely visual information, and a canonical approach that allows us to test and predict light-material interactions in a systematic way. Moreover, combined with the optical mixing method the approach provides a way to systematically and smoothly vary the image features, instead of the physical parameters.

Using multiple psychophysical experimental methods, our research was set out to answer the main research questions. They have been answered in Chapters 2 - 6. Below, the main findings and contributions of this thesis are mentioned one by one in the order of the research questions.

- **Q1: Can material perception be measured in a purely visual way?** We introduced a new type of (non-spherical) adjustable probe to measure material perception of opaque materials. The adjustable probe implements an image combination process ("optical mixing"). Beyond just matte and glossy material variations, we included in total four canonical material modes to account for a wide range of materials, namely, diffuse, asperity, forward, and mesofacet scattering for "matte", "velvety", "specular", and "glittery" material modes, respectively. Additionally, we developed an interactive interface that integrated the probe for a matching task, where observers adjusted sliders to vary the weight of each material mode's appearance or image features in the probe. The interface was found to be intuitive even for inexperienced users, allowing quantitative measurements via purely visual probing. Performances were generally well above chance and robust across experiments and observers.
- **Q2: Can observers match optically mixed canonical materials while discounting canonical lightings?** Following the results shown in Chapter 2, we included

three canonical lighting modes, namely "ambient", "focus", and "brilliance" lighting, and created a set of twelve stimuli combining canonical material and lighting modes for the bird-shaped objects. In a matching experiment using the material probe, we tested the effects of canonical lighting on our canonical materials. Again, performances were found to be above chance and robust. Material-dependent lighting effects were found for the optically mixed materials.

- **Q3a: Can observers match optically mixed canonical lightings while discounting canonical materials?**

- **Q3b: Can observers simultaneously discriminate materials and lightings?** We further developed the material probe and expanded it to allow optical mixing of canonical lighting modes. We performed both light matching tasks and 4-category discrimination tasks to investigate the material-light interactions. While observers were far above chance performance, asymmetric perceptual confounds between judgments of material and lighting were found. Specifically, observers were found to match our lightings less well than materials. Our analysis suggests that midlevel image features are more robust across materials than lightings, suggesting they might be more diagnostic for our canonical materials than our canonical lightings.

- **Q4: Can we predict material quality effects of light-material interactions?** We introduced a canonical approach using optics-based models of material and lighting modes. We were able to test and predict light-material interactions in two experiments using both the photographs of the real objects and computer rendered images. To this aim a novel spherical harmonics based metric was introduced for quantifying the "brilliance". Results from the two experiments correlated strongly, showing A) how canonical material and lighting modes associated with perceived material qualities and that we can predict the perceptual effects; and B) which lighting was best adapted to evoke perceived material qualities, such as softness, smoothness, glossiness, etcetera. Our results demonstrate that a system of canonical modes spanning the natural range of lighting and materials provides a good basis to study lighting-material interactions in a wide natural ecology.

- **Q5: Do light map orientations and the shape of the objects influence the perception of the associated material qualities?** Following the work of Chapter 5, using the same canonical approach, we further tested the effects of lighting, light map orientation, and shape on material perception for corresponding canonical material modes. We found that material perception depends on all these factors in a material-dependent manner. We proposed so-called "importance maps" as a manner to analyse the appearance features of the proximal stimulus, the image, in contradistinction to the physical parameters.

7.2. LIMITATIONS AND FUTURE DIRECTIONS

7.2.1. LIMITATIONS

In our study we implemented only four canonical material modes. Yet, there are other modes that can be added to represent more materials, such as the back-scattering mode, a Fresnel term for the forward scattering for sheen effects at grazing angles, or an anisotropy forward scattering mode for a brushed metallic effect. A major question is how to span the perceptually-relevant dimensions of the BRDF space.

Our real stimuli, “the Bird set”, was limited due to the manufacturing method. The surfaces of real objects cannot be perfectly matte or specular. Making photographs of real objects in a laboratory environment limited the choice of illuminations. Also, aligning the photography of real objects for optical mixing was quite difficult in practice.

On the other hand, the use of computer rendered objects was limited due to simplified reflectance models. For example, we did not include inter-reflections and textures. In this way, we had a simple rendering solution that made analysis of relationships with importance maps possible. Inter-reflections and textures are easier to render nowadays and including them might have prevented some of the interaction effects we found. For example, the interactions between matte and velvety were significantly stronger for rendered stimuli than for real stimuli in multiple experiments.

7.2.2. RECOMMENDATIONS FOR FUTURE STUDIES

- **Novel metrics for key image features** We found that key image features trigger the perception of material qualities, such as specular highlights for specular, sparkly glints for glittery, bright contours for velvety, etc. Novel metrics are required to quantitatively evaluate these image features and relate them to perceptual results (Chapter 5).

- **The perceptual qualities of appearance** Connected to the previous point, I propose to develop novel interfaces for rendering materials and things that vary their appearance via the variation of the key image features. Our canonical modes mixer is an example of such a perception-based interface. It was found to be easy to understand and handle. This should in general be the case, since the "knobs" of such an interface will directly be related to the perceived material qualities, in contradiction to rendering interfaces in which basic physical parameters can be varied. A key element in such studies should be that material appearance will be treated as the perceived qualities of stuff on an illuminated thing, and not as the physical specifications of a raw building material. The latter specs will obviously be important in the making, but do not form a perceptually intelligent basis in visual design. However, it remains to be studied which features would be optimal and how to combine them in a user-friendly manner. Thus, it should also be investigated which elements such interfaces should contain and how to design the user interactions.

- **The influence of context on perceived appearance** It goes without doubt that the context in which a thing is observed might have an influence on how its stuff is perceived. Some of the effects found in this thesis for objects shown without an articulated context might be influenced by showing context. Many effects we predict to be rather stable though, such as for instance glossy objects looking matte in diffuse light due to the absence of highlights.

7.2.3. EXAMPLES OF POSSIBLE APPLICATIONS

Knowledge of how people visually perceive materials under the influence of lighting and shape could have great value in both commercial and noncommercial activities. Below we propose a few possible applications that relate to our findings and could be interesting for different fields.

- **An appearance-based material selection interface** Sliders are commonly seen in computer-graphics software for changing parameters of materials. In such computer rendering software, users often explore the effect of combining endless number of different parameters that do not necessarily have an intuitive perceptual meaning. In Chapter 2, we argued that our novel interactive interface is more intuitive than the interface of the traditional computer-graphics software, in a proximal-image feature-oriented manner or painterly approach. Despite the use of sliders for materials, by combining photographs of real materials, optically mixed canonical material modes provide an intuitive way of getting the desired material appearance, instead of an overwhelming combination of different parameters. Further development might allow the creation of an intuitive material picker for computer graphics and prototyping in design processes.

- **Light probes** A probing method for measuring perception of light has been recently developed to test the perception of the light field in empty space (Koenderink et al., 2007; Xia et al., 2014, 2017; Kartashova et al., 2015, 2018). Positioning the probe, namely a white lambertian sphere, in empty space allowed observers to make judgments about the light qualities based on the appearance of the probe. Because the probe was matte, this only works for estimating the qualities of the lower orders spherical harmonics components, such as the intensity, direction, and diffuseness of the local light field (Ramamoorthi & Hanrahan, 2001). Our findings show that matte surfaces are perceived robustly across all lighting and shape conditions (Chapter 6) and thus a matte sphere forms a robust probe in the sense that any change in its appearance will be attributed to the light, not the material. A different material and shape are needed for probing the "brilliance" of the light, i.e. the light qualities of the higher orders of the SH components. The results in this thesis suggest that preferred materials should have more peaked BRDFs such as the canonical specular and glittery modes (Chapter 5 & 6).

It might also be possible to create an interactive lighting environment that allows manual adjustments of each light mode in a real laboratory environment. For example, one could create an interactive lighting environment using high dynamic range screens. In that way, objects surrounded by the screens could be illuminated by systematically controlled illumination, i.e. rendering certain material appearances in a real and controlled lighting environment. This kind of application could be helpful for creating real stimuli in psychophysical experiments concerned with visual perception. For product photography, such an interactive lighting environment would allow systematical manipulations of canonical lighting modes and thereby of their appearance, especially for e-commerce and web-shops. Furthermore, tinkering with light-material-shape interactions in a real environment could support education of students in form-giving and visual communication.

BIBLIOGRAPHY

Adelson, E. H. (2001, June). On seeing stuff: the perception of materials by humans and machines. In Human vision and electronic imaging VI (Vol. 4299, pp. 1-13). International Society for Optics and Photonics.

Anderson, B. L. (2011). Visual perception of materials and surfaces. Current Biology, 21(24), R978-R983.

Kartashova, T., de Ridder, H., te Pas, S. F., Schoemaker, M., & Pont, S. C. (2015). The visual light field in paintings of museum Prinsenhof: Comparing settings in empty space and on objects. Human vision and electronic imaging XX, 9394, 93941M.

Kartashova, T., de Ridder, H., te Pas, S. F., & Pont, S. C. (2018). Visual Light Zones. i-Perception, 9(3), 2041669518781381.

Koenderink, J. J., Pont, S. C., Van Doorn, A. J., Kappers, A. M., & Todd, J. T. (2007). The visual light field. Perception, 36(11), 1595-1610.

Ramamoorthi, R., & Hanrahan, P. (2001). On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object. JOSA A, 18(10), 2448-2459.

Xia, L., Pont, S. C., & Heynderickx, I. (2014). The visual light field in real scenes. i-Perception, 5(7), 613-629.

Xia, L., Pont, S. C., & Heynderick, I. (2017). Separate and simultaneous adjustment of light qualities in a real scene. i-Perception, 8(1), 2041669516686089.

SUMMARY

Materials are omnipresent. Recognizing materials helps us with inferring their physical and chemical properties, for instance if they are compressible, slippery, sweet and juicy. Yet in literature, much less attention has been paid to material perception than to object perception. This dissertation presents studies on a method to systematically measure human visual perception of opaque materials and test the influence of lighting and shape on material perception.

In our studies, we applied multiple psychophysical methods such as matching, discriminating, and perceptual scaling to test the visual perception of materials for human observers. Beyond just matte and glossy material variations that were commonly tested in material perception literature, we included in total four canonical material modes to account for a wide range of materials, namely diffuse, asperity, forward, and mesofacet scattering for "matte", "velvety", "specular", and "glittery" material modes, respectively. For the lightings, we included three canonical lighting modes within a spherical harmonics and perception based framework, namely "ambient", "focus", and "brilliance" lighting. Based on the spherical harmonics analysis of the global lighting environment, we were able to quantify the “diffuseness” and “brilliance” of the light maps by using Xia’s diffuseness metric and a novel brilliance metric we proposed.

Combining the four material modes and three lighting modes, we presented a canonical set that in combination with optical mixing supports a painterly approach in which key image features could be varied directly. With this method we were able to test and predict light-material interactions using both photographs of the real objects and computer rendered stimuli.

We first introduced a new type of non-spherical appearance probe, implementing the painterly approach. Moreover, we developed an interactive interface that integrated the probe for an asymmetric matching task, where observers adjusted sliders to vary each material mode in the probe. The interface was found to be intuitive for inexperienced users and allowed purely visual quantitative measurements. Performances were

generally well above chance and robust across experiments and observers, validating the approach.

We further developed the material probe and expanded it to allow optical mixing of canonical lighting modes. In a light matching experiment and a 4-category discrimination experiment we found asymmetric perceptual confounds between judgments of material and lighting. Specifically, observers were found to be less sensitive to light changes than to material changes. Moreover, using this canonical approach, we were able to test and predict light-material interactions in two perceptual scaling experiments. To this aim a novel spherical harmonics based metric was introduced for quantifying the "brilliance".

Lastly, we compared results from our probing method and results from other psychophysical experimental methods, namely perceptual scaling and discrimination, in which semantic information (for material attributes) was involved. Robust effects of light, shape, and light map orientation were found, in a material dependent way.

To conclude, our research mainly contributed to 1) the development of a novel probing method that mixes image features of the proximal stimulus in a fluent manner instead of varying the distal physical properties of the stimuli, plus a validation that it works and that it allows quantitative measurements of material perception and material-lighting interactions; 2) understanding of visual perception of opaque materials and material-light-interactions in a wide ecological variety; 3) a validated model for predicting the material dependent lighting effects for matte, specular, velvet and glittery materials; and 4) the interpretations of the material perception results in a manner relating to shape and light. Our findings can be further applied to many subjects, such as industrial design, education, e-commerce, computer graphics, and future psychophysical studies.

ACKNOWLEDGEMENTS

As an average male, I do not express personal feelings very often, not even in written form. Now, as one of the perks of finishing this Thesis, I get the chance to write a few words that I would not be able to say in person to many people.

First and foremost, I wish to thank my promotor, professor Sylvia Pont, for "raising" me to be an independent researcher in such an interesting field. Sylvia, you taught me how to "see" the world. Thank you for your constant patience and trust in me. You will always be my role model in both science and life.

I wish to thank another brilliant promotor of mine, professor Huib de Ridder. I will always remember the jokes you made during my Skype interview. More importantly, the way you played with my data means much more than simply having new interpretations. I learned so much from you as well.

Sylvia and Huib, I feel lucky and privileged to have both of you as my promotor. You two together turned a "robotic" guy into a "human" researcher, and guided me how to transform. From now on I am ready to steer the future. Again, thank you for spending all those hours, especially those off-hours, on thinking about how to improve the quality of my work, and for saying "I like it" when I was not confident about it. Thank you for sharing stories of yours and listening to mine. Neither this Thesis nor I myself would be as good without your encouraging, inspiration and guidance.

The next person I wish to thank is Dr. Pascal Barla, with who I have been collaborating since 2016. Pascal, I couldn't thank you more for your contributions to our work. You were always there when I needed a discussion or a green shiny bird. Thank you for being inspiring, fun and smart.

As I always mentioned in the last slide of my presentation, I would like to thank my lab members for all the useful discussions during our group meetings, and so much more. Maarten, to be honest I somehow found you scary at first, because your comments were all very sharp and critical. Only after a while (maybe a few years) I started to be able to follow your thoughts or dare to make jokes with you. Please do not scare your Chinese PhD away in the future. Tatiana, I still remember the time when I was just offered the position in the summer of 2013, you reached me via Facebook, introduced yourself, and gossiped all about me. I was thinking something like you are so curious that you must be really good at doing research with human subjects. I guess I was half right. Anyway, sorry that I did not make it when you finished, and thank you for being an important part of my PhD. Francesca, the “wonder girl”, I feel really honored to be part of the group in your interview, to witness how brilliant you are since then. Thank you for being willing to defend my Thesis in case I pass out, and thank you for teaching me the true Italian, I will keep practicing it. Ling “Shi Jie”, it is my pleasure to be your “Shi Di”. You are also my role model in both science and life. It was really nice to have you and Qian in the office to guide me through my first few months in the Netherlands. I remember the time when Chinese New Year approached right after I arrived in Delft, you really made me feel like being home by having me over and celebrating with your friends, thank you. Mitchell, bro, thanks for bringing masculinity and intelligence to the office. Being the only male in the office was nice, but having you around is better. Jess, it is really great to have you in the lab. You are like a big sister to everyone. Every time I chatted with you, I felt like I learned a little something new. I wish we could chat more in the future. Willemijn, you are always enthusiastic about your work. You may not know it, but I really learned a lot from you. Qian and Yi, you are such a lovely couple. I enjoyed the trip we had together a lot, and I am really happy for you as Oscar is getting a younger brother or sister. I wish all the best for all of you in Nanjing. Jie, thank you for showing me one can achieve so much in so many different things at the same time. You are simply amazing. Cristina, Karina, and Cehao, thank you for bringing new energies and inspirations to the lab in the final stage of my PhD. Hanging out together with you guys helped a lot with releasing my pressure.

I am really proud of being one of the PRISMers. This is one of the tags that I will keep putting on myself in the future. Sabrina, JJ, Lukas, Dicle, Carlos, Vasileios, Thomas, Zarko, Irene, and Tatiana, it is really nice to know you all. You added so much more fun to the conferences and our network meetings and focus groups. Meanwhile, I wish to thank all the PRISM PIs, especially professor Roland Fleming, for initiating this wonderful project that allowed me to pursue science as an ESR, and for showing me that one can be both rigorous and fun when talking about science.

I wish to express my sincere gratitude to the department secretary office, Amanda, Daphne, Denise, Charleyne and Joost, for helping me solving all kinds of problems throughout these years. Meanwhile, I wish to thank the PhDs, postdocs, and professors in Giessen, for all the help during my 3-months visits and all the fun during conferences and meetings. JJ, Vivian, Steven, Patrick, Robert, Ross, Matteo, Matteo, Jing, Jing, Anouk, Lorilei, Filipp and Yaniv, thank you.

A special "thanks" goes to Stefano, for taking care of a house of Chinese students. It is my pleasure to team up with you before Qiong came to the Netherlands. It has always been fun to spend time with you and Tatiana.

I wish to thank my friends from Delft Chinese Football team, Yaxun, Jiakun, KeLu, Vincent Gong, Xiaoqian, Andi, MengGe, HaoGe, KuGe, LaoChen, and Zonghan, your supports on/off the pitch means a lot to me.

In the end I wish to thank my family for their understandings, supports, and company. Dear mom, dad, and my beloved wife, thank you for making it all worthwhile.

Curriculum Vitæ

Fan Zhang

13-06-1989	Born in Beijing, China.
2001–2007	Chaoyang Foreign Language School, Beijing, China
2007–2011	BEng in Mechanical Engineering and Automation Shanghai Jiao Tong Univeristy, Shanghai, China
2012–2013	MSc in Robotics King's College London, London, UK
2013-2019	PhD. Delft University of Technology, Delft, The Netherlands
2018-2019	Postdoc. Delft University of Technology, Delft, The Netherlands

Thesis:	On Probing Appearance: Testing material-lighting interactions in an image-based canonical approach
Promotors:	Prof. dr. S. C. Pont, Prof. dr. H. de Ridder,