Characterising Environmental Illumination and its Effect on Physiology and Behaviour

Samuel Ponting

A thesis submitted for the degree of

Doctor of Philosophy



Pembroke College

Department of Experimental Psychology

University of Oxford

Michaelmas 2025

Abstract

The lighting environments in which we find ourselves on a daily basis have profound effects on a wide range of physiological processes, and in the modern world we experience illuminations that are dynamic in intensity, colour and structure. Daylight illumination in outdoor environments, for example, changes with the diurnal cycle, as well as throughout the year due to the orbital cycles of earth. However, humans spend much of their time indoors relying on artificial lights to provide illumination for interacting with the world, and increasingly task-oriented behaviour draws us to screens that further complicate the nature of light exposure. Often, these sources of illumination interact in complex ways, such as in an office or home with artificial lights, screens and windows that introduce daylight into an otherwise artificially lit scene.

Physiologically, the light in our environment is sensed by the five photoreceptor classes found in the human eye, which are comprised of the Long-, Medium- and Short-wavelength sensitive cone cells, which primarily drive vision in daytime light levels; rod cells, which provide visual function in low-light environments; and intrinsically photosensitive retinal ganglion cells (ipRGCs), which mediate non-visual effects of light such as feelings of wakefulness and changes to the sleep-wake cycle. Since the discovery of ipRGCs in humans, a wealth of research has demonstrated a range of positive effects achieved by modifying exposure to light in the ipRGC sensitivity range and as such there has been significant interest from industry to incorporate knowledge about the way light affects our health into the design of lighting fixtures and indoor spaces in so-called human-centric lighting design. Despite this, many questions about the function of ipRGCs, what stimulation they typically receive, and how this affects health remain unanswered.

This thesis presents a range of new approaches for the measurement, characterisation and analysis of environmental illumination, as well as novel modelling methodologies for predicting the ways in which illumination affects key wellbeing measures such as sleep and alertness. Overall, the focus of this thesis is the quantification of physiologically-relevant light in the built environment, and the prediction of its effects on physiology and behaviour.

Acknowledgements

This thesis was made possible thanks to funding received under the Oxford Medical Research Council Doctoral Training iCASE Programme. The project was supported by the MRC, Hoare Lea, and Pembroke College.

I must first say a massive thank you to my supervisor, Professor Hannah Smithson, for her continued support throughout this project. I am extremely grateful for the way Hannah has believed in me and pushed me to better my work throughout the last four years. I've thoroughly enjoyed developing this project, and I'm proud of the work we have achieved together. Thank you to Professor Manuel Spitschan, my co-supervisor, who has also contributed significantly to this work, providing support and advice throughout. I would also like to thank Ruth Kelly Waskett, Chris Fox and Jonathan Rush for their supervision throughout this project. Our discussions have significantly benefited this work.

I would like to extend my thanks to all of the Oxford Perception Lab members whose paths have crossed mine throughout my DPhil. You have all made this project an enjoyable one, and the experience wouldn't have been the same without any one of you. A huge thank you to Allie and Shiwen, for your help getting my first experiments off the ground and welcoming me to the lab. Thank you to Dana, for your fierce skribbl competition. Thank you to Mengxin for your belief and your kind words. Thank you to Sirui for reminding me to not go too crazy with the caffeine. Thank you to Marianne for your kindness throughout. Thank you to Zahra for the amazing cakes and pastries. Thank you to Jiahe for reminding me to get enough sleep (especially before being imaged). Thank you to Damian for your words of wisdom. I owe a huge thank you to Maria, for helping to keep me sane and always making me laugh. I also owe a huge thank you to Takuma. Thank you for your inspiration, mentorship and kindness, which helped me find this project in the first place. You have all been not only great colleagues, but great friends. Thank you all.

Thank you to my amazing friends to whom I owe a great deal throughout the process of

writing this thesis. Thank you to PCBC for the amazing times and thank you especially to Nik for the most intense games of squash of my life. A massive thank you to Teddy for the support and generosity over the last four years; I appreciate it more than I could ever say. A huge thank you also to Clare, Tats, Emma and Belinda for your support in Regent Street and throughout. I'm still sorry about the fridge. Thank you also to Lucy and Colette for your support. A huge thank you to Ed, Charlie and Andy for being there when I've needed you the most. I appreciate you all so much.

I would also like to thank Leslie Ferguson, whose inspiration and belief helped me with my first application to Oxford, without which I could never have imagined completing this DPhil.

I owe my biggest thanks to mum, dad and Toby. You are my biggest inspirations and have driven me through countless late nights and rough patches, while helping me achieve something of which I'm so proud. There aren't enough words to explain how thankful I am for your unquestioning love and support. This wouldn't have been possible without it. I love you all.

Contents

1 Introduction **19**

1.1 Overview 20

1.2 Physical properties of light 20

1.3 Photoreceptor transduction 22

1.4 Cones 25

1.4.1 Individual differences 26

1.4.2 Cone pathways 26

1.4.3 Colorimetry 30

1.4.4 Daylight modelling 31

1.4.5 Colour appearance 32

1.5 Rods 33

1.6 Intrinsically photosensitive retinal ganglion cells 35

1.7 Non-visual effects of light 37

1.8 The CIE S 026 technical report and standard observer model 40

1.9 Colorimetry for the display of colour 41

1.10 Silent substitution for photoreceptor isolation 42

1.11 Light measurements in natural and built environments 43

1.12 Thesis overview 44

1.12.1 Chapter 2 44

1.12.2 Chapter 3 44

1.12.3 Chapter 4 45

1.12.4 Chapter 5 45

1.12.5 Chapter 6 46

1.13 Acknowledgements 46

2 Prediction of the visual and non-visual effects of short-wavelength blocking optical filters 47

2.1 Introduction . . . 48

2.2 Methods . . . 50

2.2.1 Approach . . . 50

2.2.2 Empirical dose-response curve . . . 50

2.2.3 Naturalistic scene spectra . . . 51

2.2.4 Dataset of commercially available short-wavelength blocking filters . . . 51

2.2.5 Parametric short-wavelength blocking filter . . . 52

2.2.6 Dose response delta . . . 52

2.2.7 Optimisation metrics . . . 52

2.2.8 Effects of ageing . . . 52

2.2.9 Viewing regime changes . . . 53

2.2.10 Contrast sensitivity . . . 54

2.3 Results . . . 54

2.3.1 Dose response delta . . . 54

2.3.2 Parametric filter effects . . . 55

2.3.3 Age effects . . . 57

2.3.4 Viewing regime changes . . . 58

2.3.5 Contrast sensitivity . . . 59

2.3.6 Best performing commercially available filters . . . 59

2.3.7 Optimally performing filters . . . 63

2.4 Discussion . . . 66

2.5 Conclusion . . . 67

2.6 Acknowledgements . . . 68

3 Quantifying CIE alpha-opic signals in the indoor built environment 69

3.1 Introduction . . . 70

3.1.1 Statistical regularities in environmental illumination . . . 70

3.2 Methods . . . 72

3.2.1 Data Specification . . . 72

3.2.2 Camera Location . . . 73

3.2.3 Capture Scheme . . . 76
3.2.4 Data Validation . . . 77
3.2.5 Data Analysis . . . 77
3.3 Results . . . 79
3.3.1 Seasonal variation . . . 79
3.3.2 Weather Variation . . . 85
3.3.3 Directional Variation . . . 88
3.3.4 Additional Modelling . . . 90
3.4 Discussion . . . 92
3.5 Acknowledgements . . . 96

4 Measurement of the directional structure of light in the built environment: single-image circadiometric lightprobes 97

4.1 Introduction . . . 98
4.1.1 The benefits of quantifying environmental illumination . . . 99
4.1.2 Existing methods for the quantification of environmental illumination . . . 100
4.1.3 The single-image circadiometric lightprobe imaging system (SICLIS) . . . 101
4.2 Methods . . . 101
4.2.1 Approach . . . 101
4.2.2 Single-image circadian lightprobe imaging system (SICLIS) . . . 102
4.2.3 Image processing . . . 103
4.2.4 Model-based validation . . . 105
4.2.5 Data analysis . . . 106
4.2.6 Experiment 1 . . . 109
4.2.7 Experiment 2 . . . 111
4.3 Results . . . 112
4.3.1 Experiment 1 . . . 112
4.3.2 Experiment 2 . . . 128
4.4 Discussion . . . 132
4.4.1 The single-image circadian lightprobe imaging system . . . 133
4.4.2 Experiment 1 . . . 134
4.4.3 Experiment 2 . . . 135

4.5 Acknowledgements . . . 136

5 Individualised light dosimetry paired with head direction tracking in the built environment 137

5.1 Introduction . . . 138

5.2 Methods . . . 140

5.2.1 Devices . . . 140

5.2.2 Participants, location and procedure . . . 140

5.2.3 Data analysis . . . 141

5.2.4 Additional data and analysis . . . 142

5.3 Results . . . 144

5.3.1 Dosimetry . . . 144

5.3.2 Heading data . . . 145

5.3.3 Lightprobe models . . . 146

5.3.4 Seat-switch models . . . 148

5.3.5 Egocentric lightprobes . . . 151

5.4 Discussion . . . 153

5.5 Conclusions . . . 157

5.6 Acknowledgements . . . 157

6 Hemispheric differences in human chromatic discrimination thresholds 159

6.1 Introduction . . . 160

6.2 Methods . . . 162

6.2.1 Hyperspectral lightprobes . . . 162

6.2.2 Lightprobe analysis . . . 162

6.2.3 Participants . . . 163

6.2.4 Equipment . . . 163

6.2.5 Heterochromatic flicker photometry . . . 163

6.2.6 Procedure overview . . . 164

6.2.7 Adaptation . . . 165

6.2.8 Stimuli . . . 166

6.2.9 Experiment blocks . . . 166

6.2.10 Staircases . . . 167

6.2.11 Psychophysics analysis . . . 167
6.3 Results . . . 167
6.3.1 Hyperspectral lightprobes . . . 167
6.3.2 Human Chromatic Discrimination Thresholds . . . 170
6.4 Discussion . . . 171
6.4.1 Timescale of adaptation . . . 172
6.4.2 Other variables . . . 173
6.4.3 Biased sampling of the chromatic environment . . . 174
6.5 Acknowledgements . . . 175

7 General discussion **177**
7.1 Chapter Summaries . . . 178
7.1.1 Chapter 2: Prediction of the non-visual effects of short-wavelength blocking optical filters. . . . 178
7.1.2 Chapter 3: Quantifying CIE alpha-opic signals in the indoor built environment. . . . 178
7.1.3 Chapter 4: Measurement of the directional structure of light in the built environment: single-image circadiometric lightprobes. . . . 179
7.1.4 Chapter 5: Individualised light dosimetry paired with head direction tracking in the built environment. . . . 180
7.1.5 Chapter 6: Hemispheric differences in chromatic discrimination thresholds. 180
7.2 Future directions . . . 181
7.2.1 Further use of custom light measurement devices . . . 181
7.2.2 Efficient psychophysical research using machine learning . . . 182
7.2.3 A full model of circadian effects of light . . . 183
7.3 Application of circadian science to lighting design . . . 184
7.4 Final summary . . . 187

List of Figures

1.1 Radiance and irradiance measurements 20
1.2 Spectral interaction between illuminants and surfaces 22
1.3 Photoreceptor spectral sensitivities 24
1.4 Optic media differentiation by age, field size and wavelength 27
1.5 The luminous efficiency function 29
1.6 The daylight locus and black body locus in XYZ colourspace 32
1.7 Simultaneous contrast illusion 33
1.8 Scotopic luminous efficiency function 35

2.1 Empirical dose-response curve 51
2.2 Age effects of lens filtering 53
2.3 The effect of luminance on contrast sensitivity as a function of spatial frequency 54
2.4 Dose response delta comparison 56
2.5 Parametric filter comparison 57
2.6 Age effects on circadian response 57
2.7 Age effect comparison 59
2.8 Viewing regime changes 1 60
2.9 Viewing regime changes 2 61
2.10 Contrast sensitivity changes due to filtering 62
2.11 Best performing commercially available filters 62
2.12 Principal components of filter transmittances 64
2.13 Performance of optimised filters 64

3.1 Circadiometer setup information 74
3.2 Artificial light information 75
3.3 Chromatic distribution under artificial light 76

3.4 Window view cartoons . . . 76
3.5 PCA and K-means results . . . 79
3.6 Mean chromatic and luminance values by region . . . 80
3.7 Luminance-weighted chromatic and luminance values by region . . . 83
3.8 Region difference histogram . . . 84
3.9 Difference polarity timeseries . . . 84
3.10 Rod threshold timeseries . . . 85
3.11 Weather-dependent chromaticity plots . . . 86
3.12 Weather-dependent timeseries plots . . . 87
3.13 Indoor correlation between luminance and chromatic statistics . . . 88
3.14 Phase histogram plots . . . 89
3.15 Lightprobe timeseries plot . . . 89
3.16 Sun elevation plots . . . 90
3.17 Daylight spectra plots . . . 91
3.18 Daylight chromaticity plots . . . 91
3.19 S-cone vs. Melanopic light correlations . . . 92

4.1 Circadiometric lightprobe setup . . . 102
4.2 Lightprobe image processing pipeline . . . 103
4.3 Images of modelled lightprobe setup . . . 106
4.4 Structural similarity image metric statistics for modelled data . . . 107
4.5 Spherical harmonic decomposition example . . . 108
4.6 Experiment 1 layout . . . 110
4.7 Experiment 2 layout . . . 112
4.8 Lightprobe Images . . . 113
4.9 Luminance spherical harmonic amplitudes . . . 114
4.10 Luminance spherical harmonic means . . . 115
4.11 Melanopic spherical harmonic amplitudes . . . 117
4.12 Melanopic spherical harmonic means . . . 118
4.13 Light vector direction plotted on lightprobes . . . 120
4.14 Phi direction of first harmonic degree . . . 121
4.15 Luminance diffuseness heatmap . . . 123

4.16 Melanopic diffuseness heatmap . . . 123
4.17 Spectral analysis: August . . . 124
4.18 Spectral analysis: November . . . 125
4.19 Spectral analysis: March . . . 126
4.20 Spectral analysis: June . . . 127
4.21 Photopic-melanopic dominance spherical harmonic amplitudes . . . 128
4.22 Photopic-melanopic dominance diffuseness . . . 129
4.23 Hoare Lea office lightprobe images . . . 129
4.24 Spherical harmonic amplitudes throughout the day . . . 130
4.25 Directionality vector throughout the day . . . 131
4.26 Spherical harmonic amplitudes for melanopic-photopic ratio . . . 132
4.27 Melanopic-photopic ratio directionality vector throughout the day . . . 132

5.1 Custom dosimetry devices . . . 141
5.2 Participant seat locations . . . 142
5.3 Average daily light exposure . . . 145
5.4 Average daily circadian effect . . . 146
5.5 Heading distribution plots . . . 147
5.6 Heading pitch comparison . . . 148
5.7 Lightprobe model regression . . . 148
5.8 Seat switch models: seat 1 . . . 149
5.9 Seat switch models: seat 2 . . . 150
5.10 Seat switch models: seat 3 . . . 150
5.11 Seat switch models: seats 4 and 5 . . . 151
5.12 Egocentric lightprobe example images . . . 151
5.13 Egocentric lightprobe spherical harmonic analysis . . . 152

6.1 Schematic of stimuli colour sampling . . . 164
6.2 Experiment procedure schematic . . . 165
6.3 Hyperspectral lightprobe upper distributions . . . 168
6.4 Hyperspectral lightprobe lower distributions . . . 169
6.5 Ellipse statistic comparisons . . . 169
6.6 Luminance distribution comparisons . . . 170

6.7 Chromatic discrimination thresholds 171
6.8 Discrimination threshold statistic comparisons 171

List of Tables

3.1 Chromaticity Kruskal-Wallis test statistics for seasonal differences 82

3.2 Luminance-weighted Kruskal-Wallis test statistics for seasonal differences 82

4.1 Data collection dates for experiment 1 110

4.2 Data collection locations for experiment 2 111

4.3 Three-way ANOVA results for luminance light scalar value amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; n = 5 per cell.) 115

4.4 Three-way ANOVA results for luminance light vector amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 116

4.5 Three-way ANOVA results for luminance squash tensor amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 116

4.6 Three-way ANOVA results for luminance texture amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 117

4.7 Three-way ANOVA results for melanopic scalar light value amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 119

4.8 Three-way ANOVA results for melanopic light vector amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 119

4.9 Three-way ANOVA results for melanopic squash tensor amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 119

4.10 Three-way ANOVA results for melanopic texture amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.) 120

Chapter 1

Introduction

Contents

1.1 Overview 20

1.2 Physical properties of light 20

1.3 Photoreceptor transduction 22

1.4 Cones 25

1.4.1 Individual differences 26

1.4.2 Cone pathways 26

1.4.3 Colorimetry 30

1.4.4 Daylight modelling 31

1.4.5 Colour appearance 32

1.5 Rods 33

1.6 Intrinsically photosensitive retinal ganglion cells 35

1.7 Non-visual effects of light 37

1.8 The CIE S 026 technical report and standard observer model . . . 40

1.9 Colorimetry for the display of colour 41

1.10 Silent substitution for photoreceptor isolation 42

1.11 Light measurements in natural and built environments 43

1.12 Thesis overview 44

1.12.1 Chapter 2 44

1.12.2 Chapter 3 44

1.12.3 Chapter 4 45

1.12.4 Chapter 5 45

1.12.5 Chapter 6 46

1.13 Acknowledgements 46

1.1 Overview

The purpose of this chapter is to review concepts, methods and existing research that are important to this thesis. It will begin with a broad introduction to the study of human vision, and the methods that are used to measure and display light which are key prerequisites of the experimental chapters that follow. Then, I will introduce the research specifically surrounding the non-visual effects of light, identifying gaps and limitations in the current literature. Finally, I will provide an overview of the experimental chapters that follow.

1.2 Physical properties of light

To quantify the effects of light exposure on various behaviours, researchers require methods to accurately and reliably measure the light that exists in the environment. Recordings of radiant flux, measured in Watts (W), form the basis of the physical quantification of light. Irradiance, measured in Watts per square area (e.g., W/m^2), is a measurement of the density of radiometric flux, or the amount of radiometric energy per unit of area. This means that the irradiance of a light source is inversely proportional to the square of the viewing distance, as the density of light decreases further from its source. On the other hand, radiance measures radiometric flux density per unit of solid angle (e.g., $W/m^2/sr$). [1]. This fixed viewing angle means that, for large diffuse light sources such as monitors displaying a single colour, radiance measurements are distance-invariant as the sampled surface increases in size at the same rate as the losses in irradiance. Figure 1.1 shows the geometry for both radiance and irradiance measurements.

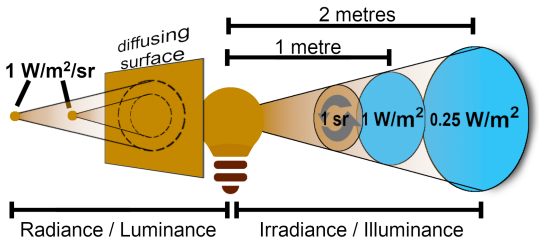

Figure 1.1: Schematic showing radiance and irradiance measurement geometry.

Both irradiance and radiance measurements give rise to intensity values, which correspond

to the quantity of light measured at a particular wavelength (λ). Human vision is broadly characterised as being sensitive to electromagnetic radiation between 380 nm and 780 nm. The range between these bounds is typically called the 'visible light spectrum'. As such, light-sensing devices such as spectroradiometers are designed to record irradiance and radiance across this range, though many are sensitive to infrared and ultraviolet light too. When measuring the spectral properties of a light-reflecting surface or emitter, devices yield a spectral power density function, which describes the contribution of different wavelengths of electromagnetic energy to the light emitted from the target surface. Figure 1.2 shows the spectral power density function of a CIE standard daylight, known as CIE A [2].

An important physical interaction to note is the reflectance of light from surfaces. Depending on their physical composition, different surfaces absorb and reflect different wavelengths of light. As such, a surface can be assigned a reflectance profile as a function of wavelength, $R(\lambda)$, which shows the relative amounts of each wavelength the surface reflects, as shown in figure 1.2. However, the light that a surface is able to reflect is dependent on the light shone on it in the first place. As such, we can consider illumination also as a function of wavelength - $I(\lambda)$. The resultant light reflected by the surface $R(\lambda)$ when illuminated by light $I(\lambda)$ is the product of the two spectra – $I(\lambda)R(\lambda)$, as shown in equation 1.1. As such, the light reflected by a surface is the result of the interaction of the surface properties of an object and the light under which it is illuminated. An example is shown in figure 1.2.

$$L(\lambda) = I(\lambda) \cdot R(\lambda) \tag{1.1}$$

An additional set of physical parameters that affect surface-light interactions is dictated by the geometric structure of the surface. The bi-directional scattering distribution function (BSDF) of a surface defines the angular specificity of light reflected from a surface, balancing directionally-specific light, known as specular reflections, and directionally-broad light, known as diffuse reflections. These features depend on the structure of the surface, with diffuse reflections being the result of multiple intra-surface reflections which result in losses of light in the form of absorption within the surface. This breaks the one-to-one relationship between incident and reflected light of a theoretically-perfect specular surface. These correspond to material-perceptual properties. For example, surfaces with high specular reflections and minimal diffuse reflections are seen as 'mirror-like', while diffuse reflections make surfaces appear matte or dull [3].

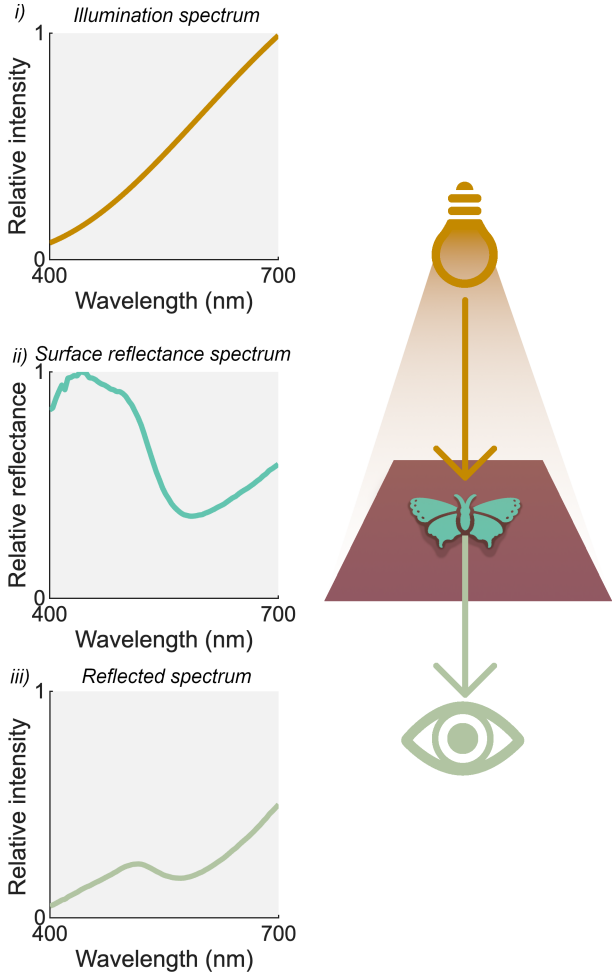

Figure 1.2: Spectral interaction between illuminants and surfaces. Schematic showing that the light that reaches our eyes is the product of illuminant spectra and surface reflectance spectra. A light with significant long-wavelength bias (*i)* shines onto an object with a short-wavelength bias (*ii)*, and the resultant wavelength spectrum is the product of the two input spectra (*iii)*.

1.3 Photoreceptor transduction

Vision is arguably the most important sense that humans possess and is entirely mediated by photoreceptors found within the human eye. The sensing of light relies on the phototransduction of electromagnetic radiation which lies in the range visible to humans, around 380 nm to 780 nm. The primary pathways for vision are initiated by the rods and cones in the photoreceptor layer of the retina [4]. Phototransduction then begins a cascade of effects, as the result of the efferent connections of the photoreceptors. These connect to other layers of neurons in the eye – namely bipolar cells and retinal ganglion cells (RGCs), before being projected to several targets in the brain via the optic nerves. The primary visual pathway takes visual input to the visual cortex via the lateral geniculate nucleus (LGN) through the retinothalamic tract. In the cortex, information is retinotopically mapped, and different areas in the occipital lobe [5]

have been associated with different visual functions, such as V1 showing edge selectivity and V4 showing chromatic selectivity [6–9]. RGC outputs also project to non-cortical areas such as the superior colliculus (SC), which is important for controlling eye movements, as well as to the olivary pretectal nucleus (OPN), which is important for regulating pupillary light reflexes. A small subset of RGCs are themselves photosensitive and have other cortical projections to the suprachiasmatic nucleus (SCN).

There are five spectrally-sensitive classes of photoreceptor in the human eye, which can be split into three functional groups – the cones, rods and intrinsically photosensitive retinal ganglion cells (ipRGCs). These groups largely overlap in their function and physiology but are broadly associated with different components of light-sensing behaviour. Each of the five photoreceptor classes is sensitive to different but largely overlapping parts of the visible light spectrum.

Figure 1.3ii shows the spectral sensitivities for each of the five photoreceptor classes [10–13]. These depict the sensitivity of each photoreceptor class to a given availability of photons at the cornea. As light enters the eye, it travels through various optic media that filter the light spectrum. The photoreceptor spectral sensitivities represent the *effective sensitivity* - the probability of absorbing a photon of a given wavelength - of the photoreceptors given the pre-receptoral filtering that light undergoes in a standard 32-year-old observer. However, these sensitivities are subject to change depending on this pre-receptoral filtering, which changes with age and viewing eccentricity. This concept is explored further later in this chapter. The photoreceptor spectral sensitivities are all characterised by a similar unimodal shape, with differing peak sensitivities. These differing sensitivity peaks are determined by the photopigments by which each class is defined.

The outputs of rod and cone photoreceptors are univariate – they change only in the intensity of graded potentials sent to bipolar cells – which means they are only able to encode the overall intensity of light in the sensitivity spectrum range [14]. ipRGCs also abide by the principle of univariance, though their outputs are action potentials, meaning that only the frequency of depolarisations encodes the overall intensity of light to which they are exposed. Thus, all photoreceptors are unable to encode information about the wavelength distribution of light in their respective sensitivity band. The photoreceptor signal in response to a particular light spectrum is given by the integral of the weighted sum of the photoreceptor's sensitivity spectrum and the light spectrum, as shown in equations 1.2 to 1.6. These represent the photoreceptor response calculations for a given light, $I(\lambda)$, given that $\bar{L}(\lambda)$, $\bar{M}(\lambda)$, $\bar{S}(\lambda)$, $\bar{R}(\lambda)$ and $\bar{U}(\lambda)$ are the L-cone, M-cone, S-cone, rod and ipRGC spectral sensitivities, respectively, as a function of

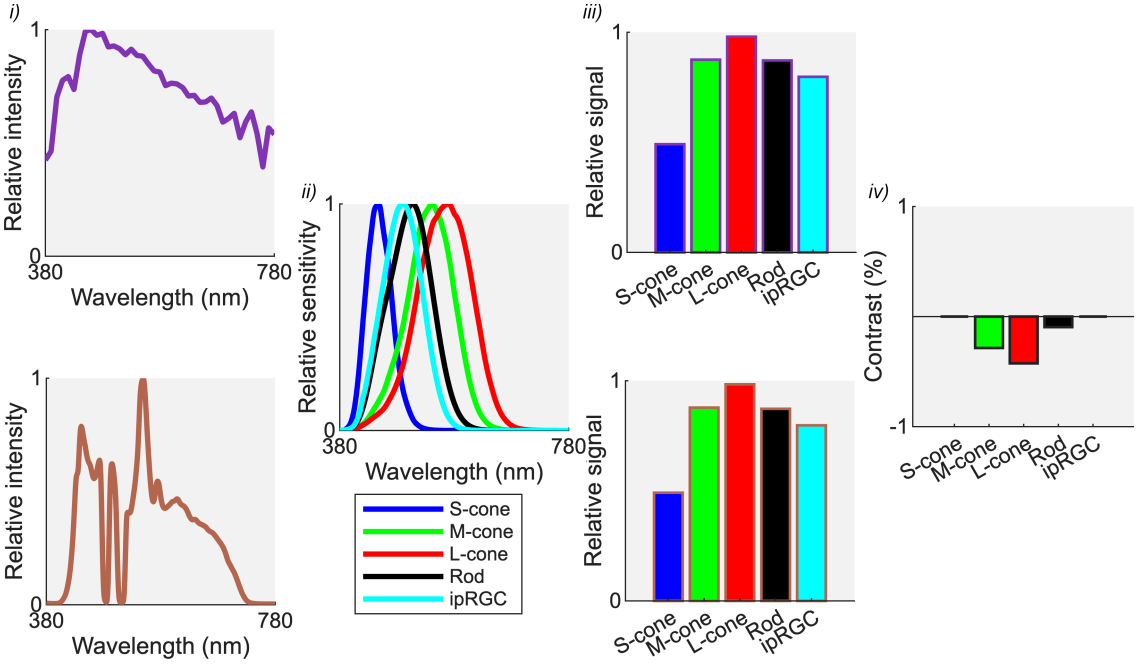

Figure 1.3: Photoreceptor spectral sensitivities. *i)* Two example light spectra. Top shows the D65 light spectrum, and bottom shows an alternative light spectrum that, despite a different spectral distribution, is metameric across the five photoreceptor classes with D65. *ii)* Spectral sensitivities for the S-cones, ipRGCs, Rods, M-cones and L-cones of a standard observer. *iii)* Estimated photoreceptor outputs for the spectra in *i)*. *iv)* Percent contrast between the two sets of cone excitations.

wavelength. L, M, S, R and U are as such the resultant photoreceptor values yielded by the light $I(\lambda)$. This can lead to ambiguities as the multivariate input, which varies in wavelength and intensity, leads to a univariate output. This 'principle of univariance' gives rise to the phenomena of metamerism, where different light spectra give rise to identical photoreceptor outputs. An example pair of metamers and their corresponding cone outputs is shown in figure 1.3ii and 1.3iii. Despite their spectral distributions looking different, these two spectra yield almost identical photoreceptor signals for a standard observer, as shown in 1.3iv. In fact, for any light spectrum, there are infinitely many possible metameric matches, which emphasises the fact that our visual system takes only five wavelength samples from the continuous reality of light spectra.

$$L = \int I(\lambda) \cdot \bar{L}(\lambda) d\lambda \tag{1.2}$$

$$M = \int I(\lambda) \cdot \bar{M}(\lambda) d\lambda \tag{1.3}$$

$$S = \int I(\lambda) \cdot \bar{S}(\lambda) d\lambda \tag{1.4}$$

$$R = \int I(\lambda) \cdot \bar{R}(\lambda) d\lambda \tag{1.5}$$

$$U = \int I(\lambda) \cdot \bar{U}(\lambda) d\lambda \tag{1.6}$$

1.4 Cones

The cones are the most important and widely studied functional group of photoreceptors in the human retina [15]. They contribute to photopic vision - that is, vision in most typical daytime light conditions. There are approximately six to seven million cones in the human retina, with most being densely packed into the fovea, where gaze is focused [15]. The remainder span the full extent of the retina at decreasing densities as the distance from the fovea increases, though at the *far* edge of the retina there is a 'cone-rich rim' which may be useful for quickly detecting stimuli in the periphery [16, 17]. In colour-normal humans there are three classes of cones which differ in the photopigment they express, as well as the post-receptoral channels they feed, but overlap in many of their attributes. These are known as the S-, M- and L-cones, corresponding to the short, medium and long wavelength ranges to which they are sensitive [18]. Typical cone sensitivity spectra are shown in figure 1.3. The S-cone population also differs from the other two cone classes in its distribution across the retina, being largely absent from the foveola - the centremost region of the fovea [19, 20]. The high density of L- and M-cones at the fovea, the three distinct wavelength-dependent populations and the post-receptor processing of their signals mean cones allow humans to detect form, colour and motion.

Cones afford humans colour vision [21], which is extremely helpful for a range of tasks that confer an evolutionary benefit, such as identifying ripe, red fruits amongst green leaves [22]. Significant spectral information is lost due to cone signals being determined by the integration of wavelength and intensity information due to the principle of univariance [14]; however, the comparison of cone signals provides some information about the distribution of light across the visible light spectrum and is thus used to yield the percept of colour. Importantly, because of the information loss inherent in the photoreception process, colour is not a physical, observer-agnostic quantity, and relies on the specific sensitivities of the photoreceptors in the visual system.

Thus, for any spectrum of light observed by an individual, their photoreceptor system affords them three spectral samples which can be compared for colour perception - one from each of the three cone classes. Therefore, any spectrum within the photopic system can be thought of as a point in 3D space according to the ratio of activations it generates from each of the cone classes. Because of this, humans are termed to be trichromatic. This additionally means that spectra which vary in their spectral content but yield identical S-, M- and L-cone responses will be perceived as the same colour, despite differing underlying spectral data. These sets of

spectra are cone metamers, though they may yield different excitations from the rod and ipRGC photoreceptor classes.

1.4.1 Individual differences

There are significant individual differences in cone phenotypes across the population. Because of cones' role in colour vision, many of these differences are associated with colour vision deficiencies [23–25]. There is also significant non-clinical variation in cone spectral sensitivity both between individuals and within an individual's lifetime. One determinant is the changes throughout the lifetime in the eye's optic media through which light has to pass to reach the photoreceptors [26–29]. In particular, the lens selectively attenuates short wavelengths and because the optical density changes throughout the lifetime, the extent of short wavelength attenuation changes. The CIE standard observer model allows cone spectral sensitivities to be predicted for an observer of a particular age, taking into account this age lens density effect. Additionally, the model considers the spectrally-selective filtering of the macular pigment, the optical density of which varies spatially across the retina in a radial pattern such that the eccentricity of viewing with respect to the fovea matters, allowing for spectral sensitivity predictions that consider the location in the visual field at which colour judgements are made [30]. Figure 1.4 shows various components of the individual differences in optic media, and differences across viewing conditions, that result in changes in photoreceptor sensitivities. These components are: *i)* the lens optical density, which is high at short wavelengths and changes with age, such that density is higher in older individuals; *ii)* the peak density of the macular pigment, which changes with field size, with it decreasing as more of the retina is included; *iii)* the wavelength-dependent relative density of the macular pigment, which is high in the S-cone and ipRGC range, while it is almost zero in the longer wavelength region of the spectrum (λ>550 nm); *iv)* changes in peak optical density of the cones themselves due to field size such that with higher field sizes, peak optical density reduces for all cone classes, and there is an overall difference between the S cones and the L and M cones.

1.4.2 Cone pathways

The specific way in which cone outputs are combined has a significant effect on colour vision. There are traditionally understood to be three groupings of retinal ganglion cells that differ in their cone inputs and lead to three independent retinogeniculate pathways to the LGN [31].

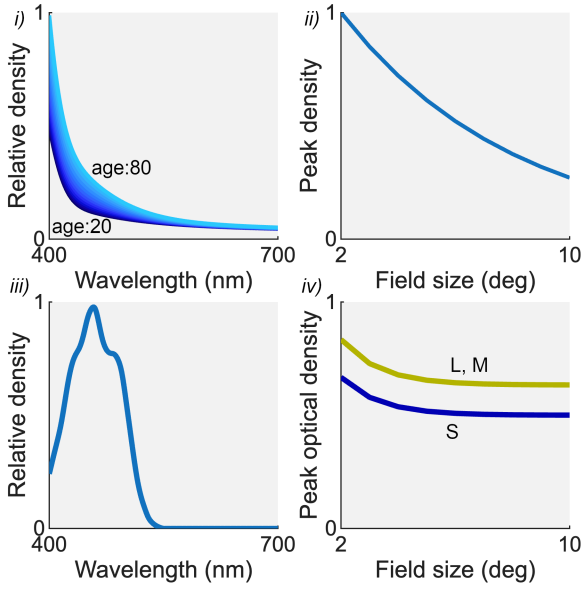

Figure 1.4: Optic media differentiation by age, field size and wavelength. *i)* changes in relative density of the lens due to age. The youngest lenses, from age 20, are shown in dark blue, while the oldest lenses, up to age 80, are shown in light blue. *ii)* changes in peak macular pigment density due to field size. *iii)* wavelength-dependence of macular pigment density. *iv)* changes in peak optical density due to field size. Optical density for L- and M-cones is shown in yellow, and optical density for S-cones is shown in blue.

Cones are traditionally understood to form two main post-receptoral pathways for colour – the L-M opponent pathway [32] and the S-opponent pathway [33] - and one pathway for luminance, the L+M pathway [34–37]. The opponent pathways are so-called opponent as they are excited by some cone signals and inhibited by others. As such, they act as comparators between their input cone classes.

L-M opponent signals are transmitted from midget retinal ganglion cells [38, 39] to layers three to six of the LGN via the parvocellular pathway [40, 41], which is mainly responsible fine detail perception, due to the high density of L- and M-cones in the fovea [42], and the one-to-one synapses between them and bipolar cells [43]. They are also responsible for comparing L- and M-cone signals and thus discrimination between red-ish and green-ish hues. A common architecture for the midget pathway takes input from a local patch of photoreceptors, forming a receptive field which is split into an ‘on centre’ and a concentric ‘off surround’ (known as an ’on midget’), or an ’off centre’ with an ’on surround’ (known as an ’off midget’). This spatial opponency means that activation in the centre increases output, while activation in the surround decreases output. This setup enables structure detection, but these regions are also split by their cone class (e.g., an L-on centre and an M-off surround - though it is likely that there is a mixture of cones in the surrounding region [44]), additionally enabling for chromatic opponency [45]. The receptive fields of these cells vary with eccentricity, with small receptive fields in the fovea and

larger ones in the periphery.

S-opponent signals take input from all three cone classes, opposing excitation from the output of the S-cones to the summed output with inhibition from the L- and M-cones. They are transmitted from small bistratified retinal ganglion cells to the small koniocellular layers between the main LGN layers via the koniocellular pathway [46]. Small bistratified RGCs do not have centre-surround receptive fields [33]. Because of the lower density of S-cones in the fovea [19, 20], the S-opponent pathway is less optimised for fine spatial detail but does enable the discrimination between blue-ish and yellow-ish hues. It is important to note that the traditional picture of post-receptoral pathways for vision is complicated by evidence of S-off inputs to the midget RGC pathway with a range of cone opponency dynamics [47, 48].

The third grouping of retinal ganglion cells, known as the parasol ganglion cells, produce an achromatic pathway that projects to magnocellular layers one and two of the LGN. This pathway additively combines L-cone and M-cone signals and is thus not a chromatically opponent pathway, providing no colour information. However parasol ganglion cells have spatially-opponent centre-surround receptive fields that can take different polarities, such that light in the centre can either increase or decrease firing rate. This channel is also known as the luminance channel. It is additionally useful for the perception of motion as it supports a high temporal resolution compared to the other two pathways [34, 42]. By combining the L-cone and M-cone sensitivity spectra, the luminous efficiency function ($V(\lambda)$) is found and is shown in figure 1.5. To obtain this function for a standard observer, the L- and M-cone sensitivity spectra are typically weighted by 0.6899 and 0.3483, respectively [49]. Since this luminous efficiency function is wavelength-selective, it means that luminous intensity is subject to not only the overall intensity of a light, but also its spectral composition.

Photometric quantities are those that take the luminous efficiency function into account, whereas radiometric quantities do not. Luminance is the amount of light emitted from or reflected off a surface and is measured in candelas per square metre (cd/m^2). Illuminance is the light incident on a surface. The unit for illuminance is Lux, which corresponds to 1 cd/m^2. For a given light spectrum, the corresponding cd/m^2 is calculated by integrating the product of the luminous efficiency function with the light spectrum and multiplying by 683.002. This value is the conversion factor between Watts, a radiometric measurement, and lumens, a photometric measurement. At 555 nm, which is the peak of the luminous efficiency function, $1W$ of radiant intensity is equal to 683.002 lumens. The weighting of the L- and M-cone contributions to the

luminous efficiency function corresponds to the ratio of L-cones to M-cones [50, 51]. However, there is significant variation in this ratio across the population. This is important as widely used units such as cd/m^2 represent a standard observer, and individuals' personal luminous efficiency functions differ from this depending on their L:M ratio.

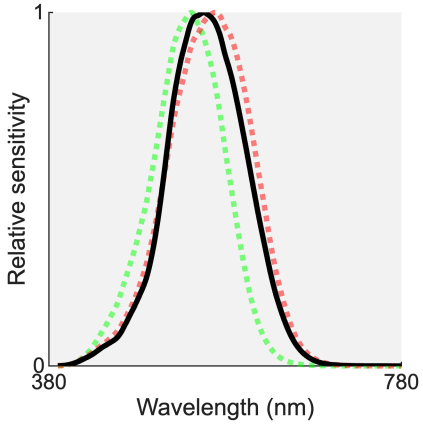

Figure 1.5: The luminous efficiency function, $V(\lambda)$. $V(\lambda)$ is shown in black. L- and M-cone spectral sensitivities are shown in the red and green dashed lines, respectively. The functions are plotted on a linear scale.

To estimate the appropriate L- and M-cone weighting for a given observer of experimental condition, heterochromatic flicker photometry (HFP) techniques have been used. These techniques that leverage the preferential response of the magnocellular pathway to high frequency flicker stimuli compared to the parvocellular pathway. HFP uses stimuli that rapidly oscillate between L-dominant and M-dominant chromaticities at a frequency faster than can be chromatically detected by the parvocellular pathway. Depending on the weighting of L- and M-cones into the magnocellular pathway, the apparent luminance of the two chromaticities may be differing, yielding a flicker percept. Adjustments can be made to the relative intensities of L- and M-cone activating lights in order to achieve equiluminance. This is a commonly used procedure in psychophysics testing when examining responses to chromaticity changes, and is important because of the subtle individual differences of participants with respect to standard observer models. HFP is utilised and explored further in Chapter 6.

While these three pathways consitute the predominant model of colour vision, there is a complex heterogeneity in the morphology of retinal ganglion cells. In addition to the midget, parasol and small bistratified cells mentioned already, there is evidence for up to 18 distinct kinds of RGC [52]. For example, smooth monostratified RGCs observed in primates feature complex receptive fields that may aid spatial vision [53, 54], and recursive RGCs feature complex dendritic webs and are directionally-sensitive [55]. Perhaps the most impactful addition to

our understanding of post-receptoral pathways is the set of intrinsically photosensitive retinal ganglion cells that are associated with wide-ranging impacts on both visual and non-visual function. These are discussed in more depth later in this chapter.

1.4.3 Colorimetry

The architecture of post-receptoral processing described above generates what are termed cardinal axes [56] that, when used as an orthogonal coordinate system, can capture the gamut of human colour signals. An early example of this is the Derrington-Krauskopf-Lennie (DKL) colour space [40], which features three axes that correspond to the three post-receptoral pathways. This allows the physiological response to light to be positioned in the space in terms of both its chromaticity and its luminance. An alternate formulation of these cardinal axes is the MacLeod-Boynton (MB) chromaticity diagram [57]. The MB chromaticity diagram represents an equiluminant plane for a two-axis representation of chromaticity. In this case, the axes are formulated as ratios of cone excitations divided by their luminance in order to achieve equiluminance. By convention, the x axis is generally taken to be $L/(L+M)$, and the y axis is taken to be $S/(L+M)$. As such, the axes capture the red-ish-to-green-ish and violet-ish-to-lime-ish percept of colour variation in the opponent pathways. Units in this space are formally arbitrary, however scaling factors are used to give meaning to particular values in the colour space. For example, standard whites are often scaled to lie at y values of 1, and the L:M ratio of a standard observer is used to weight L and M. Equations of the standard MacLeod-Boynton chromaticity diagram are shown in equations 1.7 and 1.8. For these equations, the y axis is scaled such that equal energy white ($I(\lambda) = 1$)lies at 1.

$$X_{MB} = \frac{0.6899L}{0.6899L + 0.3483M} \tag{1.7}$$

$$Y_{MB} = \frac{1.9354S}{0.6899L + 0.3483M} \tag{1.8}$$

There are several other colour spaces that are frequently used in colour vision research, each with their own historical and functional significance. An early prominent example is the CIE 1931 XYZ space [58], which uses data from colour matching experiments to estimate chromaticity properties for an average observer across three dimensions [59–61]. The Y component of XYZ space is luminance, with X and Z corresponding to chromatic axes. This, too, has an equiluminant

representation, known as the CIE 1931 xy chromaticity diagram.

The above chromaticity spaces use linear combinations of cone inputs. While this makes them easy to formulate and manipulate, their geometry does not map linearly onto perceptual space. This means that a given distance does not correspond to an even perceptual change. The CIE LUV and LAB colour spaces have been devised as perceptually uniform spaces to represent colour, meaning that even steps throughout the space correspond with approximately even steps in perceptual space [62, 63]. Both feature a luminance axis ($L*$ - which can be omitted for isoluminant representations of uniform perceptual space), and two colour axes, which are created from non-linear combinations of cone signals in order to achieve perceptual uniformity. The $a*$ and $b*$ axes of the LAB colourspace and the $u\prime$ and $v\prime$ axes of the LUV colourspace correspond to perceptual changes from red to green from yellow to blue, respectively. In these two colourspaces, a white point is defined and assumed to represent the adaptation state of an observer and since colour percepts change systematically with adaptation states, colours are represented with respect to this point.

1.4.4 Daylight modelling

Daylight constitutes the most ubiquitous illumination and is thus highly important to model and quantify. The spectrum and therefore colour of daylight varies dramatically based on the time of day, season and weather, but is bound by its largest contributors – the long-wavelength rich yellow light of the sun, and the short-wavelength rich blue light of the sky. Of course, the sky itself is illuminated by the sun, but appears blue throughout most of the day due to Rayleigh scattering. The declination in sun angle makes the sky appear yellow or red in the direction of the sun at sunrise and sunset [64, 65], however increased scattering across the full dome of the sky, along with the ozone absorption of longer wavelength light (known as the Chappuis effect), leads to an overall relative increase in the amount of short-wavelength light in the sky at twilight [66]. This chromatic axis – from "cold" blues to "warm" yellows and reds is known as the daylight locus and is shown in the XYZ colourspace in figure 1.6. The daylight locus is very similar the Planckian or black body locus, which shows the chromaticity of a black body when heated [67]. Figure 1.6 shows both the daylight and black body locuses in the CIE 1931 colourspace. Points on this axis are labelled with temperatures in Kelvin (K) to show the temperature to which a chromaticity corresponds. This is known as the correlated colour temperature (CCT) [68]. Low black body temperatures appear as 'warm' colours. High temperatures perceptually correspond correspond

to 'cool' colours. Since the black body and daylight locus are similar, these temperatures also correspond to the range of daylight chromaticities, with cool colours appearing similar to skylight on a clear day, and warm colours appearing similar to sunlight at dawn or dusk. Daylight spectra, modelled as a function of sun elevation, are discussed further in chapter 3. Since this daylight model codifies a standardised range of daylight colours, it is often used to define key illuminants. The most notable of these is the CIE D65 [69]. D65 is frequently used as a standard white point in many psychophysical experiments and calculations of physiological metrics, however the CCT convention is most predominantly used in lighting manufacture and design. In industry and commercial practice, there is not a perfect one-to-one correspondence between CCT and chromaticity, as approximations are used [70]. Various models have enabled researchers to correspond any colour with a colour temperature [71]; however, the values are generally used to quantify chromaticities close to the black body locus [72].

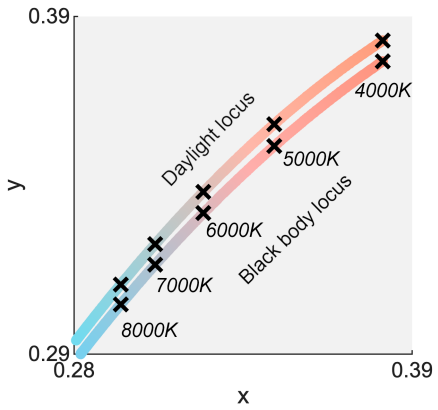

Figure 1.6: The Daylight locus visualised in the XYZ colourspace. Points marked X correspond to correlated colour temperatures of $4000K$, $5000K$, $6000K$, $7000K$ and $8000K$

1.4.5 Colour appearance

Though our colour perception system only has access to cone outputs, colour appearance is complex and there is no one-to-one correspondence between appearance the spectrum profile. Thus, there are several non-linearities that change the appearance of colours based on contextual factors. This can be shown simply in colour patches rendered within surrounds of different colours, as shown in figure 1.7. In the left and right parts of the figure, the central patches contain identical colour information and thus generate the same cone signals. However, due to the different background contexts, they appear different. This is known as simultaneous contrast and serves to show the decoupling of colour appearance from cone signals as a result of context [73]. This is not the only example of the complex modulation of colour percepts. There

are many known mechanisms that modulate colour perception [74–77].

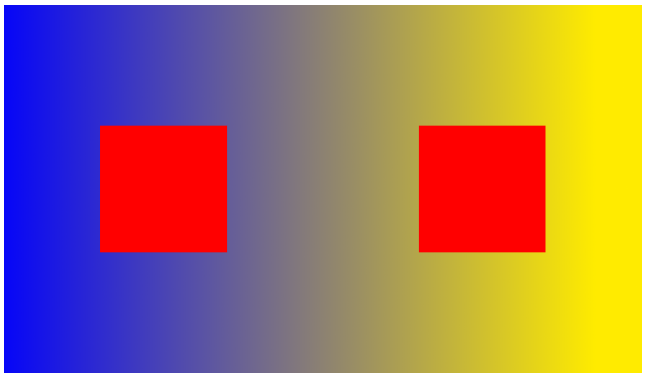

Figure 1.7: Simultaneous contrast illusion. The two square patches have the same RGB values, though they appear different colours due to the different backgrounds.

This concept can be broadened to a wide range of so-called adaptation effects which shift our colour perception away from the raw spectrum input we receive from the world [78–82]. Adaptation can occur on a range of temporal and spatial scales, but in general serves to normalise cone signals over time [83]. For example, S-cone signals will decline over time in response to stable blue light as their responses adapt. Perceptually, this is accompanied by a desaturation of the blue hue, as the adaptation process levels the cone signals towards an achromatic grey [84]. Adaption processes lead to a range of complex colour phenomena. They are thought to be important for colour constancy [85], where estimations of global illumination are discounted from surface percepts in order to better estimate the colour of objects in environments around us. This is a complex process that likely relies on transformations at a number of different stages in the visual system and relies on contextual information [86–91]. Adaptation is an extremely important physiological process, as it means that our perception of light is modulated by our broader illumination environments. Adaptation and the stability of colour appearance across the visual field are explored further in Chapter 6.

1.5 Rods

Rod cells lie in the photoreceptor layer of the retina, distributed in the same retinal mosaic as the cones. Rods are significantly more populous than the cones, with around 120 million rods compared to only 6-7 million cones. The fovea entirely lacks rods, which are distributed across the periphery [92, 93]. The most notable functional difference between rods and cones is that rods are significantly more sensitive to light than cones are – they have been shown to be able to produce signals in response to single photons [94, 95]. Unlike cones, there is only one class of

rod cell, with a peak sensitivity at approximately 500 nm [96]. These features make rods ideal for low-light or *scotopic* vision, which is roughly below -3 log units of luminance (cd/m^2). This means that under low light conditions, the scotopic luminous efficiency function - known as $V'(\lambda)$ and defined by the rod fundamental - shows the range over which the human visual system is sensitive to light. $V'(\lambda)$ is shown in figure 1.8.

Phenomenologically, scotopic light levels correspond to late evening or nighttime environments, perhaps only dimly lit by distant artificial lights such as street lights, or natural sources such as the moon and stars. Compared to cones, rod outputs are widely pooled into bipolar cells and then retinal ganglion cells. This convergence amplifies rod signals but reduces the resolution. Scotopic vision also operates at a slower temporal resolution [97, 98]. While rod signals and thus scotopic vision are theoretically achromatic, there is some evidence that colour percepts do occur at light levels in the scotopic range [99, 100]. However, it is difficult to induce rod signals while fully isolating cone signals [101]. There is much less known about genetic variation leading to individual differences in rod spectral sensitivities [102], though they are subject to the same pre-receptoral filtering as the cones described above.

There is evidence for two main pathways that mediate rod-based vision. In scotopic conditions, rods primarily contribute to vision via rod-on bipolar cells and a subset of amacrine cells known as AII amacrine cells. This system is slow, with a maximum flicker frequency of 15 Hz, but highly sensitive [94]. However, there is also evidence of a second pathway for rod signals that pass through gap-junctions that directly connect the rods to cone cells, bypassing much of the traditionally-understood rod circuitry and instead using cone pathways [96]. This second pathway is active in the upper range of rod function, where both rods and cones are active - between roughly -3 and 0.47 log units of luminance. This overlapping viewing regime is known as mesopic vision [103, 104]. Operating at higher light levels, this system is much less sensitive to light, but has a faster flicker visibility threshold of around 28 Hz.

Above this light level, at around 0.47 log units of luminance, rods are traditionally understood to be saturated and not contribute to photopic vision [105–107]. However, recent evidence has suggested that many rods are still active well into the photopic light range [108, 109]. The way in which rod and cone signals interact is not fully understood - particularly those through gap-junctions and other other potential rod pathways [110–112]. The role of rods in many typical lighting environments is likely to be underestimated.

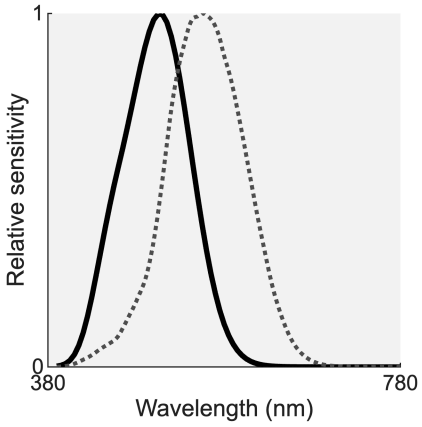

Figure 1.8: Scotopic luminous efficiency function plotted on a linear scale. For comparison, the photopic luminance efficiency function is plotted as a dashed line.

1.6 Intrinsically photosensitive retinal ganglion cells

The retinal ganglion cell (RGC) layer lies in front of the photoreceptor layer of the retina. It serves to combine signals from several rods and cones, via bipolar cells, in order to create receptive fields with particular chromatic and spatial features such as chromatic opponent centre-surround comparators. However, a small population of ganglion cells observed in mouse, primate and human research are themselves photosensitive, and have been shown to sense light in the absence of rod and cone input in mice lacking the photoreceptors typically associated with image-forming photoreceptors [113]. These intrinsically photosensitive retinal ganglion cells (ipRGCs) express the protein opsin melanopsin, giving them a unique sensitivity spectrum compared to rods and cones, peaking at around 480 nm [114]. In humans, ipRGCs make up an estimated 1% of the 1.2 million retinal ganglion cell population making them by far the least populous photoreceptor [115, 116], and like S-cones and rods, ipRGCs are largely absent from the fovea and are distributed across the peri-fovea and periphery [117]. As with the other photoreceptors and the overall retinal ganglion cell population, their density decreases with increasing eccentricity with respect to the fovea [118]. Compared to the rods and cones, ipRGCs are much less sensitive and have a slower temporal response. ipRGCs adapt much more slowly than rods and cones, which makes them sensitive to longer light exposures [119, 120]. This makes them able to better represent the absolute intensity of light, while rods and particularly cones quickly adapt to the ambient levels of stimulation. Importantly, since these cells are part of the retinal ganglion cell layer, they receive cone and rod signals in addition to being photosensitive themselves [121].

In addition to being a physiological departure from rod and cone cells, ipRGCs are functionally distinct from the other two photoreceptor classes. There is growing evidence that ipRGCs

contribute to vision - particularly in brightness perception [122–127] - but the main thrust of ipRGC research has focused on non-visual physiological changes they mediate. In brief, this includes the incorporation of light as an external timekeeper into the body clock [128–130]. ipRGCs allow for the signalling of information about environmental light levels which have a wide array of psychological and hormonal effects which better prepare humans for the current time of day. For example, reducing sleepiness and increasing alertness in response to daytime-like light, and prompting sleepiness in the absence of light. These processes also have significant effects on mood [131].

ipRGCs were originally discovered in the mouse retina. In mice, six subtypes labelled M1-M6 are known to exist. M1 and M2 are the most widely studied subgroups, and are understood to account for around 74-90% of ipRGCs in mice [52]. Each subtype has differing anatomy, melanopsin expression, downstream pathways and apparent functions [132–135]. For example, M1 ipRGCs express high levels of melanopsin and have high intrinsic photosensitivity and seem to play a key role in photoentrainment. On the other hand, M4 ipRGCs express less melanopsin and seem to play a role in mood [136]. Brn3b is a transcription factor that, in ipRGCs, is associated with different levels of melanopsin expression as well as different projection targets. M4, M5 and M6 cells generally have high brn3b expression but low melanopsin expression, while M1 ipRGCs are split into two subpopulations depending on whether or not the cell expresses the transcription factor brn3b. Brn3b(+) M1 ipRGCs project to a range of structures, but primarily the olivary pretectal nucleus (OPN) and dorsal LGN. Brn3b(-) M1 ipRGCs, on the other hand, have almost no projections to the OPN and dorsal LGN, but instead project to the suprachiasmatic nucleus (SCN). M2 and M4 ipRGCs project similarly to the brn3b(+) M1 cells - to the OPN and LGN - but also the superior colliculus (SC) [137, 138]. Projections of the M3 ipRGCs have been less widely studied, though there is some evidence that they also project to the superior colliculus.

In humans, most ipRGCs feature morphology similar to that of the mice M1 and M2 cells [139, 140], though there is some evidence for the existence of M3 and M4 ipRGCs [141–143]. There are some notable anatomical differences between human and mice subpopulations. Namely, there is evidence for four subclasses of M1 ipRGC in primates and humans. These subtypes are typically referred to as M1, displaced M1 (dM1), gigantic M1 (GM1) and displaced gigantic M1 (dGM1). These delineations refer to the size and position of the cell soma. Non-displaced M1 cell somas lie in the ganglion cell layer (GCL), alongside almost all other retinal ganglion cells. Displaced M1 cells have their somas in the inner nuclear cell layer (INL). This is where many of

the off bipolar cells innervate, and there is evidence in macaques to suggest that these displaced cells respond to decrements in light [144, 145]. Physiological analysis of the giant M1 cells has shown large responses to cone and rod inputs, as well as S-off, (L+M)-on opponency [121, 146]. They also are shown to project to the LGN and ultimately V1, suggesting that image-forming and non-image forming signals are merged. In macaques, LGN projections are seen in a wider range of ipRGC subgroups than in mice, perhaps suggesting an increased role of ipRGCs for vision in primates. M5 and M6 ipRGCs have not been found in humans, though known shortcomings in the experimental detection of ipRGCs mean that current experiments are likely biased away from detecting M3-M6 cells [147].

Owing to the many brain regions targeted by ipRGCs, they have wide-ranging physiological effects. Research into the function of ipRGCs, since their discovery in mice and later in humans, has mainly focused on their non-visual effects. However, there is growing knowledge of their image-forming effects, evidenced by their input to both the LGN and SC. ipRGCs have been widely shown to play a role in perceived brightness, but have also been related directly to yellow-ish colour percepts when other photoreceptors are silenced [124]. By virtue of their slow resolution and adaptation speed, ipRGCs have also been shown to play a role in the slowing of Troxler fading, suggesting a role in the perception of persistent scenes, colour constancy and long-term adaptation [119, 120]. All ipRGCs receive input from some combination of rods and cones and often do so in an opponent formulation (e.g., ipRGC-on, S-off). While it is not known precisely how this contributes to human vision, opponency in M5 ipRGCs appears to contribute to mouse colour vision [135]. In humans, S-ipRGC opponency is speculated to be useful for spectral changes to the environment that occur at sunrise and sunset [148–150]. In general, it is difficult to assess ipRGC contributions to colour because of the presence of rods and cones. Since ipRGCs themselves receive input from the other photoreceptor classes, it is difficult to assess the contributions of each component of signals sent to visual areas. Section 2.10 below discusses one of the main methods for examining ipRGC contributions while silencing rods and cones.

1.7 Non-visual effects of light

The phenomenology of the non-visual effects of light is well known to most people - winter lethargy induced by short days and dark commutes or the invigorating feeling of a summer morning walk under a clear blue sky are all familiar and emotionally poignant experiences that

deeply connect us to the kinds of light environments we live in, whether coming from natural or artificial sources. These experiences are largely driven by outputs from ipRGCs and are in particular believed to be related to their connections to the SCN [151]. The suprachiasmatic nucleus lies above the optic chiasm between the eyes and the LGN and is widely understood to play the role of internal timekeeping [152]. Utilising positive feedback loops at the cellular level, the SCN generates oscillating signals that last roughly a day - hence the term “circadian rhythms” – from "circa" and "diem", meaning “approximately one day”. These circadian signals trigger psychological changes throughout the body via projections to several brain regions, including the pineal gland which triggers a wide array of hormonal changes. The purpose of such a neural circuit is to entrain the body to the time of day, enacting physiological changes appropriate for different goals throughout the diurnal cycle [153]. Input from photoreceptors is therefore essential to mediate information about the time of day – based on the available light – in order to adjust the internal timekeeper to ensure that physiological states match the time of day. This process is known as photoentrainment and combines the endogenous rhythms in the SCN with the exogenous timekeeping cues of light to appropriately adjust physiological processes [154–158]. This process has been successfully modelled by researchers as a van der Pol oscillator with a light input that phase-advances the timekeeper when the function is increasing, and phase-delays it when the timekeeping function is decreasing [159, 160]. Phenomenologically, this corresponds to scenarios such as the awakening effect of a morning walk (phase-advance) or the inability to sleep after looking at an LED screen late in the day (phase-delay). Notably, there is some evidence that cross-talk between rods and cones depends on the time of day, which suggests a feedback connection from the SCN back to the retinal level [161].

Perhaps the most widespread mechanism through which the SCN changes circadian rhythms is through melatonin suppression [162]. The pineal gland secretes melatonin, which has been identified as a physiological signal of darkness and increases feelings of sleepiness as well as lowering body temperature in preparation for sleep [163]. The SCN’s input to the pineal gland can alter the secretion of melatonin based on the internal clock state. As such, in natural lighting conditions, melatonin suppression is caused by the illumination of the environment in the morning and throughout the day, while at night the lack of light allows secretion of melatonin to continue, which facilitates sleepiness.

Disruption to this rhythm has been shown to have a wide range of negative effects such as slowed reaction times, depleted cognitive capacity and reduced mood [164–168]. The 21st century

has bought about unprecedented changes in the typical light exposure of individuals throughout the day and year, with much of our daytime spent indoors away from natural light, rich in melanopic wavelengths, and much of our evenings occupied by artificial light that can cause phase delays that can interfere with a regular sleep schedule. This has generated significant public health interest in the way that the built environment and broader "spectral diet" is affecting our physiology and mental health [169]. Many studies seeking to evaluate the non-visual effects of melanopic light have quantified changes in melatonin suppression, circadian phase shifting, as well as self-report measures of wellbeing, alertness, cognitive capacity and subjective sleepiness as a result of the availability of light in chronobiological sleep studies [170]. In typical studies, indoor spaces such as offices have been altered in their light distribution to assess the effects of light intervention in the built environment [171–173]. As previously mentioned with visual effects, it is important to note that both rods and cones contribute to ipRGC signals and thus are likely to contribute to the non-visual effects of light [174, 175]. However, circadian changes have been shown to broadly match the spectral sensitivity function of melanopsin and are thus understood to be predominantly mediated by light in this range, and therefore ipRGCs [176, 177]. The precise effect of rods and cones on non-visual effects of light has not been quantified. While this is a fast-growing field, at the time of writing there are currently no models that can fully predict the non-visual effect of light as the effects are multi-dimensional, depending on the wavelength and intensity of light exposure, the timing and duration of exposure, as well as previous light exposure and adjustments to internal clocks, though several models exist that quantify the effects of some of these variables [159, 160, 170, 178, 179].

As previously detailed, another significant target of ipRGCs is the olivary pretectal nucleus (OPN). This brain region has a key role in the pupillary light reflex – that is: the constriction and dilation of the eye's pupil [180]. Pupil size has an important effect on the amount of light let onto the retina, similar to changing the aperture size of a camera. Pupil size is controlled by many pathways, and can change as the result of alerting effects, but photoreceptors have perhaps the largest effect on pupil size [181]. However, the precise relationship is complex, depending on photoreceptor type and ambient light level, as well as having opponent relationships [129]. ipRGCs play a particular role in long-term pupil size changes which have been extensively investigated in post-illumination pupillary response experiments that show that melanopic light reliably causes constriction, even in the absence of rod and cone signals, while steady-state effects have also been demonstrated [181–189].

1.8 The CIE S 026 technical report and standard observer model

Several standard observer models of human colour vision have been created in order to estimate physiological responses to light, and the CIE S 026 model is the first of its kind to include ipRGC sensitivity spectra and thus allow for the estimation of non-visual responses [190]. The model defines photoreceptor absorptance spectra, age-dependent optic media densities, and field-size-dependent macular pigment density to allow researchers to calculate photoreceptor signals for all five receptor classes, while considering observer age and task field size.

In addition, the CIE technical document proposes a novel set of metrics for measuring the light available to each receptor type in illuminance units based on their contributions to a standard D65 – meaning that comparisons between quantities of lux can be compared. This is useful, as it broadens the concept of a luminance or illuminance measure to be taken for all five photoreceptor classes, using a standard illuminant as the benchmark [69]. The measures, collectively known as α-opic equivalent daylight illuminance (EDI), are calculated by finding the ratio of contributions between the chosen photoreceptor class and luminance. This is known as the efficiency of luminous radiation (ELR), which is constant for each $\alpha - opic$ signal given a particular reference light. These measures allow us to calculate metrics for different parts of the visible light spectrum using lux-equivalent units. Equations 1.9 and 1.10 show that Φ_{α}^{D65} corresponds to the radiance or irradiance of D65 for the photoreceptor α, and Φ_{v}^{D65} corresponds to the radiance or irradiance of D65 weighted by $V(\lambda)$, the luminous efficiency function. The efficiency of luminous radiation is given by the ratio of these two values, as shown in equation 1.11. Then, the radiance or irradiance of the target light is weighted by the chosen photoreceptor class, as shown in equation 1.12, and is divided by the ELR, as in equation 1.13, to give the α-opic EDI. In this equation, Φ_{α} corresponds to the measured radiance or irradiance for photoreceptor α. While D65 is the standard illuminant used for this equation, any standard can be used if contextually appropriate. α-opic EDI has widely been used in research into ipRGCs, where it has become the standard method for quantifying amounts of melanopic light, and so many of the dose-response curves generated from studies of melatonin suppression and circadian phase shifts use lux melanopic EDI as the light quantifying unit.

$$\Phi_{\alpha}^{D65} = \int I_{D65}(\lambda) \cdot \alpha(\lambda)\, d\lambda \tag{1.9}$$

$$\Phi_v^{D65} = \int I_{D65}(\lambda) \cdot V(\lambda)\, d\lambda \tag{1.10}$$

$$_{\alpha-opic}ELR = \frac{\Phi_\alpha^{D65}}{\Phi_v^{D65}} \tag{1.11}$$

$$\Phi_\alpha = \int I(\lambda) \cdot \alpha(\lambda)\, d\lambda \tag{1.12}$$

$$_{\alpha-opic}EDI = \frac{\Phi_\alpha}{_{\alpha-opic}ELR} \tag{1.13}$$

1.9 Colorimetry for the display of colour

Knowledge of colorimetry and metamers is exploited for the display of colours on RGB displays, which is broadly useful for any colour display such as TVs, computers and phones, but is particularly pertinent in psychophysical study where the display of precisely calibrated colours is crucial. This method echos some of the earliest colour vision experiments which were used to propose that humans are trichromatic based on colour match data [191].

RGB displays function by additively mixing three channels – red, green and blue – at a small spatial scale in order to produce a broad range of colours. Three primary lights are used as it is the minimum number required to reproduce any colour using metamerism within a particular gamut, given a trichromatic observer [192]. Display calibration aims to enable the user to target particular points in cone colour space and reproduce them using the three channels of the display. In order to calibrate an RGB display, the primary lights must be measured using a spectroradiometer. The key step in calibration is then to shift from cone space – L, M and S – to primary space – R, G and B. In a linear system, a 3 x 3 matrix is sufficient to transform between these two spaces, and is defined by equation 1.14.

$$T = \begin{bmatrix} Lr & Lg & Lb \\ Mr & Mg & Mb \\ Sr & Sg & Sb \end{bmatrix} \tag{1.14}$$

$$\begin{bmatrix} L & M & S \end{bmatrix} = T \cdot \begin{bmatrix} r \\ g \\ b \end{bmatrix} \tag{1.15}$$

$$\begin{bmatrix} r & g & b \end{bmatrix} = T^{-1} \cdot \begin{bmatrix} L \\ M \\ S \end{bmatrix} \tag{1.16}$$

This conversion matrix turns an RGB triplet – that is; the colour designation of a single pixel on the display – into its corresponding L, M and S values, as shown in equation 1.15. These can then be transformed into any of the colour spaces described above. In order to find the RGB value that yields a given colourspace coordinate, the inverse of matrix T - T^{-1} - is used, as shown in 1.16.

It is important to emphasise that this conversion assumes both spaces are linear, and this is often not the case for displays, which feature gamma curves that represent the mapping from driving voltage to achieved light level. RGB values which are to be converted must be linearised by gamma-correcting the display's output before conversion.

1.10 Silent substitution for photoreceptor isolation

It is not presently possible to study photoreceptor signals *in vivo* in a human observer, so indirect methods must be used to isolate the contribution of each photoreceptor in response to a particular stimulus. One of the most widely utilised methods for achieving this is silent substitution [193, 194]. Silent substitution aims to create pairs or groups of stimuli that differ in their photoreceptor-exciting content for some set of photoreceptors while having no difference for others, thus, leaving them silent. Recently, this has particularly been a hallmark of the study of ipRGCs, where sets of stimuli are silent to the three cone classes and rods, while exciting different outputs from ipRGCs. Thus, the stimuli are chromatically metameric while providing different signals to the ipRGCs. This allows researchers to quantify perceptual or physiological measures caused by a single photoreceptor class.

There are a number of methodological considerations that are important for researchers to use silent substitution. In order to silence N photoreceptors, N+1 primary lights are required for stimulus presentation. This is a product of the linear algebra that underlies stimulus presentation,

as described earlier in this chapter. Thus, typical RGB displays are unable to produce cone metamers, as three primaries are only able to create one triplet that corresponds to a particular set of cone outputs. This is in contrast to light exposure in the real world, where there is not a one-to-one correspondence between colour and spectrum, and many spectra take the same chromatic properties [195]. Recent developments in multi-primary displays have enabled researchers to use five or six primary lights that allow for the silencing of cones and rods while displaying different light levels in the melanopic range [196, 197]. When designing a display for silent substitution, it is important to consider the primary lights used, as the significant overlap between the spectral sensitivities of the photoreceptors reduces the maximally achievable contrast when silencing multiple photoreceptors. As more photoreceptor classes are silenced, this also reduces the effective gamut available to produce metameric pairs. Another consideration is the accuracy of the spectral sensitivity estimates themselves. As explored, the sensitivity of individual's photoreceptors is highly heterogenous, even within the colour normal population, and most observers deviate from standard observer models in some way, even if age and field size are properly considered. 'Splatter' [198] refers to the cross-talk between receptor types that is introduced by these model-observer discrepancies, and has been used to calculate the appropriate level of noise in such photoreceptor-isolating experiments.

1.11 Light measurements in natural and built environments

Quantifying the light that is available in natural scenes is an important step in analysing the human visual system, as it constrains the kinds of stimuli the system is generally used to processing. Many research projects have sought to quantify the distributions of light in natural environments and analyse how they change across daily and seasonal changes, as well as considering the spatial variation within these scenes [83, 150, 199–202]. A widely-used approach for generating these measurements is the use of hyperspectral cameras. These are devices that feature a high spectral resolution and are thus able to capture physiologically-relevant changes in the spectral composition of light. Combining this with a high-spatial-resolution imaging sensor means that researchers have been able to begin to quantify the distribution of light in natural scenes. The Specim IQ [203], for example, uses 204 spectral channels and a 512x512 imaging sensor to produce high-spatial- and spectral-resolution images. Other cameras feature fewer channels that are specifically designed to target human light exposure metrics, such as the

Westboro Photonics Circadiometer, which features five spectral channels that are calibrated to correspond to the five photoreceptor pathways of the human visual system [204]. Recent increased interest in the distribution of light in indoor environments has led to the development of research projects that quantify light in built environments such as office and domestic spaces [205–207].

1.12 Thesis overview

In the following sections I detail the structure and content of the five experimental chapters in this thesis, as well as contextualising their findings amongst the literature summarised above.

1.12.1 Chapter 2

In Chapter 2, I establish a method to predict the non-visual effects of light using published dose-response curves. In particular, I apply this to the problem of evaluating the physiological effects of short-wavelength-blocking filters, which are popularly worn to improve circadian health. However, while these products often report their absorption spectra, it is not possible to directly interpret the non-visual effects of these filters based on this information. In this chapter, I analyse a dataset of commercially available filters to assess their effect on predictions of non-visual behaviour across the range of naturally-occurring light levels, as well as modelling a set of parametric filters. Additionally, I analyse the way different filters compare in their colorimetric and photopic properties, devising metrics to find the filters that achieve the best performance in terms of, for example, reducing circadian effect while minimally affecting the visible colour gamut. Finally, I generate hypothetical filter profiles that optimise these metrics.

1.12.2 Chapter 3

Significant previous literature has explored the chromatic distribution of light in natural environments, and the way these distributions are affected by seasonality and diurnal timing. However, less research has explored the way mixed-illuminant environments – those combining artificial and natural illumination – vary in terms of their chromatic properties throughout the day and year, and yet these environments make up more and more of humans' time budget as we spend more time indoors. In Chapter 3, I utilise a physiologically-calibrated camera - the Westboro Photonics Circadiometer - to measure the distribution of light within a typical office environment at different times of year. I focus on comparing key patterns observed in natural

scene literature to the patterns observed in the office environment, analysing the temporal and spatial characteristics of daylight-artificial light interactions, as well as considering the non-visual and mesopic characteristics of the indoor light.

Associated publication and presentations

Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson H E. (2025). "Quantifying CIE alpha-opic signals in the indoor built environment." *J Opt Soc Am A Opt*; 42(5):B379-B390.

Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson, H E. (2024). "Quantifying CIE alpha-opic signals in the indoor built environment." *International Colour Vision Society Meeting*, Ljubljana, Slovenia (July, 2024). <Oral Presentation> *** W. D. Wright Award**

Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson, H E. (2023). "Quantification of Melanopic Light in an Office Environment." *Applied Vision Association Christmas Meeting*, London, UK (December, 2023). <Poster>

1.12.3 Chapter 4

The use of planar cameras that create a 2D perspective image of one section of the light field are useful as they are similar the directional view of humans. However, they are unable to capture the full nature of the lighting environment of a particular point in space. Being able to characterise the full surround gives a greater insight into the way in which light is distributed throughout the environment and can thus be used to better quantify the spatial and directional structure of light within a scene. In Chapter 4, I detail a method for capturing the full light field using a single image capture. This method builds on previous lightprobe methods, but utilises extra mirrors to fully automate image capture, allowing for the highly mobile and fast temporal capture of high spatial resolution lightprobes that feature physiologically-relevant information. Using a digitally rendered model, I demonstrate that this single image methodology faithfully reproduces the spatial structure of environmental illumination. I present SICLIS, the Single-Image Circadian Lightprobe Imaging System, and two experiments that emphasise the novel benefits of SICLIS – *i*) the mobile flexibility of the setup and *ii*) the high temporal resolution of capture.

1.12.4 Chapter 5

Humans do not necessarily sample their lighting environments evenly. Task-oriented behaviour often significantly limits the spaces from which we sample light in our environment. Additionally,

light in the spaces we occupy can have a direct effect on the way in which we sample light, either through changes in viewing angle or light-altering behaviours such as closing blinds. In Chapter 5, I utilise wearable dosimetry devices to measure the light exposure at the corneal plane of participants freely sampling their light environment while carrying out task-oriented behaviours. I present bespoke dosimetry devices that capture spectral information using a Condor Instruments ActLumus dosimeter, as well as a custom-programmed Arduino device that records heading direction and tilt during free-viewing. I use these data to compare light exposure and viewing behaviours in different seasonal periods and desk locations. I also incorporate lightprobe data collected in Chapter 4 to produce 'egocentric' lightprobes that characterise the structure of an individual's light exposure.

1.12.5 Chapter 6

In Chapter 6 I present a psychophysical experiment investigating chromatic discrimination thresholds in different localised regions of the visual field. Previous chapters in this thesis have established the myriad of ways that light varies in regular patterns throughout our typical lighting environments. This experiment sought to test whether these regularities can have an effect on perceptual mechanisms through adaptation. I show that the chromatic distribution of light significantly differs in the upper and lower hemifields of example illumination maps. By replicating these lighting conditions in a psychophysical experiment, I adapt participants' retinae to the chromatic statistics of outdoor environments, and evaluate chromatic sensitivity under these conditions by measuring chromatic discrimination ellipses using a 2-AFC paradigm.

Associated presentation

Ponting, S; Morimoto, T; Smithson H E. (2022). "Object colour discrimination in directionally asymmetric lighting environments". *International Colour Vision Society Meeting*, Crete, Greece (July, 2022). <Oral presentation>

1.13 Acknowledgements

The work presented in this chapter is entirely that of the author, with supervision by Hannah E. Smithson and Manuel Spitschan.

Chapter 2

Prediction of the visual and non-visual effects of short-wavelength blocking optical filters

Contents

2.1 Introduction 48

2.2 Methods 50

2.2.1 Approach 50

2.2.2 Empirical dose-response curve 50

2.2.3 Naturalistic scene spectra 51

2.2.4 Dataset of commercially available short-wavelength blocking filters . . . 51

2.2.5 Parametric short-wavelength blocking filter 52

2.2.6 Dose response delta 52

2.2.7 Optimisation metrics 52

2.2.8 Effects of ageing 52

2.2.9 Viewing regime changes 53

2.2.10 Contrast sensitivity 54

2.3 Results 54

2.3.1 Dose response delta 54

2.3.2 Parametric filter effects 55

2.3.3 Age effects 57

2.3.4 Viewing regime changes 58
2.3.5 Contrast sensitivity 59
2.3.6 Best performing commercially available filters 59
2.3.7 Optimally performing filters 63
2.4 Discussion 66
2.5 Conclusion 67
2.6 Acknowledgements 68

2.1 Introduction

Over the past 25 years, there has been significant research interest in the non-visual effects of light [129]. The primary physiological pathway mediating these effects originates with the intrinsically photoreceptive retinal ganglion cells (ipRGCs). Comprising less than 1% of the retinal ganglion cell population, these cells express the photopigment melanopsin, which for a standard observer has an in-situ peak wavelength sensitivity of around 480 nm [114, 143, 208, 209]. Activation of these cells through exposure to short-wavelength light initiates a broad array of physiological effects. ipRGCs send signals to the suprachiasmatic nucleus (SCN), which acts as an internal timekeeper by entraining hormonal changes to an internal circadian rhythm, which is modulated by light exposure [162, 210–212]. In particular, the hormone melatonin is key in the regulation of sleep and thus has been found to be related to measures of sleepiness and cognitive capacity [213, 214], as well as having a direct effect on the timing and duration of sleep [154, 215–217].

Manipulation of light exposure can thus modulate these circadian effects [185]. For example, increasing short-wavelength light exposure early in the morning has been related to increased mood and better sleep [218]. However, not all short-wavelength light exposure is desirable.

Short-wavelength blocking filters are frequently used in medical fields where blue light's phototoxicity may become dangerous in high doses [219], but are also becoming increasingly popular tools to reduce the amount of melanopic light exposure in everyday life. This is particularly the case in being able to reduce the amount of short-wavelength light exposure when using devices with LED screens, which can expose users to artificially high levels of melanopic stimulation if used later in the day, causing a circadian phase delay and disrupting sleep.

Many previous studies have demonstrated the benefits to sleep that short-wavelength blocking

filters or short-wavelength reducing color correction on digital displays can provide [220–224]. The action spectra of phototoxicity and circadian effects are different but overlapping, and so commercial products and some research literature conflate the two. This chapter focuses on the non-visual circadian effects of light and will thus consider the effect of filters on circadian responses.

In addition to non-visual effects, the manufacturers of short-wavelength blocking optical filters must also consider the effects of these filters on visual function. While restricting short-wavelength light may lead to increased protection from unwanted circadian effects, it also reduces the range of colours visible to the wearer, which may impede the wearer's ability to carry out tasks that rely on colour discrimination or detection. By attenuating parts of the wavelength spectrum, short-wavelength blocking filters reduce the available colour gamut that wearers are able to see, leading to reduced colour discrimination ability, as shown in several psychophysical tests. Similar effects may also linger after the user has stopped wearing the filter. Spitschan et al. [225] call on authors to report both the visual and non-visual effects of filters in future publications in order for readers to be able to interpret the filters' effects. Specifically, they name luminous transmittance, melanopic transmittance, white shift and gamut reduction as key metrics to be reported.

The functional changes caused by manipulation of light exposure are complex and often non-linear, so simple metrics are incomplete in their description of the visual and non-visual effects of short-wavelength blocking filters. For example, the reduction in luminance caused by light filtering has complex effects on visual acuity and contrast sensitivity, and has the capacity to modulate the viewing regime that a particular illumination yields. Reduction in melanopic irradiance as a result of filtering has widespread effects on circadian function, and this relationship is known to be non-linear.

The preferred methodology for measuring the non-visual effect of wearing a short-wavelength blocking filter is running a randomised control trial, where participants wear either filter-containing glasses or a control filter. Non-visual response is then measured either directly, in terms of sleep timing and duration, or through physiological measures such as melatonin suppression, or indirectly through alertness questionnaires. However, these experimental designs can be costly, time-consuming and may require dedicated facilities and labs where sleep timing and light exposure can be controlled. Dose-response curves generated from previous experiments can be used to relate the intensity of melanopsin-weighted light (quantified as melanopic equivalent

daylight illuminance – melanopic EDI) to normalised measures of circadian response. As such, it is possible to model the relative effect of different short-wavelength blocking filters on the available melanopsin-weighted light, and use dose-response curves to quantify and compare the circadian effect of different filters, without the need for large scale comparisons of the effects of each filter under controlled conditions. Additionally, using parametric models of individual variation in visual sensitivity [27], it is possible to incorporate individual differences, such as the yellowing of the ocular lens with age.

This chapter details a method for modelling the physiological effects of the use of short-wavelength blocking filters and their functional consequences, considering the ways in which this can be modulated by design and observer age. We additionally present analysis of the filters that perform the best, by maximally reducing circadian effect and minimally affecting colour gamut, as well as optimised filters that would perform better than the commercially available ones presently analysed.

2.2 Methods

2.2.1 Approach

Our model seeks to estimate the effects of short-wavelength blocking filters on visual and non-visual function using biologically plausible means. We calculate the normalised circadian responses to a range of naturalistic scene spectra, using an empirical dose-response curve obtained from Brown et al [170]. We then calculate the circadian response for the same set of spectra when passed through an optical filter. The resulting differences in circadian response denote the circadian effect of the filter across the range of naturally obtained light spectra. We additionally considered the effects on spatial vision with reference to modelled changes in contrast sensitivity produced by the filter.

2.2.2 Empirical dose-response curve

In order to model physiological responses to different levels of melanopic light, we used the normalised global dose-response curve presented in Brown et al [170]. It is shown in figure 2.1. This curve is generated by normalising the physiological responses observed in a range of dose-response studies, using measures of circadian activation such as phase shift [154, 185, 215–217], melatonin suppression [185, 216–218] and subjective alertness [221]. The light in these studies

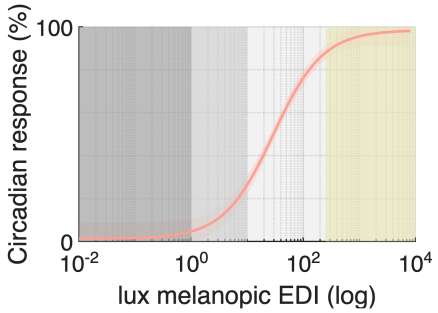

Figure 2.1: The empirical global circadian dose-response curve [170]. Shaded regions show standard error. Dark shaded region corresponds to night time sleep environment conditions. The lighter grey region corresponds to evening domestic conditions. The region shaded yellow corresponds to indoor daytime environments. The unshaded region corresponds to environments that are considered sub-optimal for daytime light exposure [170]

is quantified in melanopic equivalent daylight irradiance (melanopic EDI). Melanopic EDI is calculated by scaling melanopic-weighted irradiance by the melanopic efficacy of luminous radiation (melanopic ELR) – the ratio of melanopic irradiance and illuminance present in a D65 spectrum, as shown in Equations 2.1 and 2.2 and described in the CIE S026: 2018 technical report [190]. The melanopic ELR for D65 is 1.326. A cumulative logistic function is then fitted to the dose-response curve, to predict normalised circadian response from melanopic EDI.

$$melanopic\ ELR = \frac{melanopic\ irradiance}{illuminance} \tag{2.1}$$

$$melanopic\ EDI = \frac{melanopic\ irradiance}{melanopic\ ELR\ for\ D65} \tag{2.2}$$

2.2.3 Naturalistic scene spectra

We estimated the normalised physiological response to a large dataset of 10,000 typical daylight spectra collected using corneal-plane dosimetry [226]. These spectra spanned a wide array of intensities, with melanopic EDI values between 0.0001 and 10,000 lux melanopic EDI and had a range of spectral variation across this intensity range.

2.2.4 Dataset of commercially available short-wavelength blocking filters

We analysed an existing set of commercially available short-wavelength blocking filters, obtained from Spitschan et al.'s analysis [225]. This dataset ($N = 120$) ranges broadly from medical and safety filters to filters designed specifically for use during sport, driving, and work with visual display units.

2.2.5 Parametric short-wavelength blocking filter

To model the effects of a range of possible short-wavelength blocking filters, we defined a set of transmittance functions parametrically, using a logistic curve, as in equation 2.3.

$$T(\lambda) = \frac{\alpha}{1+e^{-\beta\cdot(\lambda-\lambda_0)}} \tag{2.3}$$

Where α represents the asymptote of the filter transmittance, β represents the slope, and λ_0 represents the cutoff wavelength. By examining filters with a range of cutoffs, we can assess the effect of changing filter characteristics on physiological responses.

2.2.6 Dose response delta

We calculated the normalised physiological response for each spectrum by calculating their melanopic EDI values, both before and after filtering and then using the dose-response curve to calculate the reduction in circadian response after filtering. The difference in this normalised response is then referred to as the dose response delta.

2.2.7 Optimisation metrics

A number of metrics were used to measure the performance of the filters by finding a balance between maximally reducing non-visual effects while minimally affecting visual percepts. The maximum dose response delta calculated from the dataset of natural spectra was used to quantify non-visual effects. For visual effects, in addition to luminance, the reduction of 2D chromatic gamut, calculated by the percentage shrinkage of the area of a convex hull encapsulating 99 reference spectra plotted in u'v' space was used. For a third measure, these 99 reference spectra were plotted in a 3D chromaticity and luminance space, using u'v' as well as luminance for the three axes. The percentage shrinkage of the volume of this convex hull was used.

2.2.8 Effects of ageing

Throughout the lifetime, the optical properties of the lens change, allowing less short-wavelength light to reach the retina later in life. A schematic of the eye in figure 2.2. 2.2i shows the eye looking through an external filter, demonstrating that light goes through both the external filter and the eye's lens before reaching the retina. 2.2ii-iii show the effect of age on lens density, and therefore ipRGC sensitivity spectra in modelled observers aged 20 (dark blue) to 80 (light

blue). Reduction in circadian signalling in elderly individuals has previously been measured [227] The changes in ipRGC absorption caused by age-related lens changes were modelled according to the procedure outlined in Asano et al (2016) [27]. Additionally, we used the lens filtering procedure outlined in the CIE S026 document in order to make age-based modifications to the melanopsin spectral sensitivity [190, 228]. Using these modelled spectral sensitivities, we calculated the circadian and chromatic properties of short-wavelength blocking filters for a range of ages between 20 and 80.

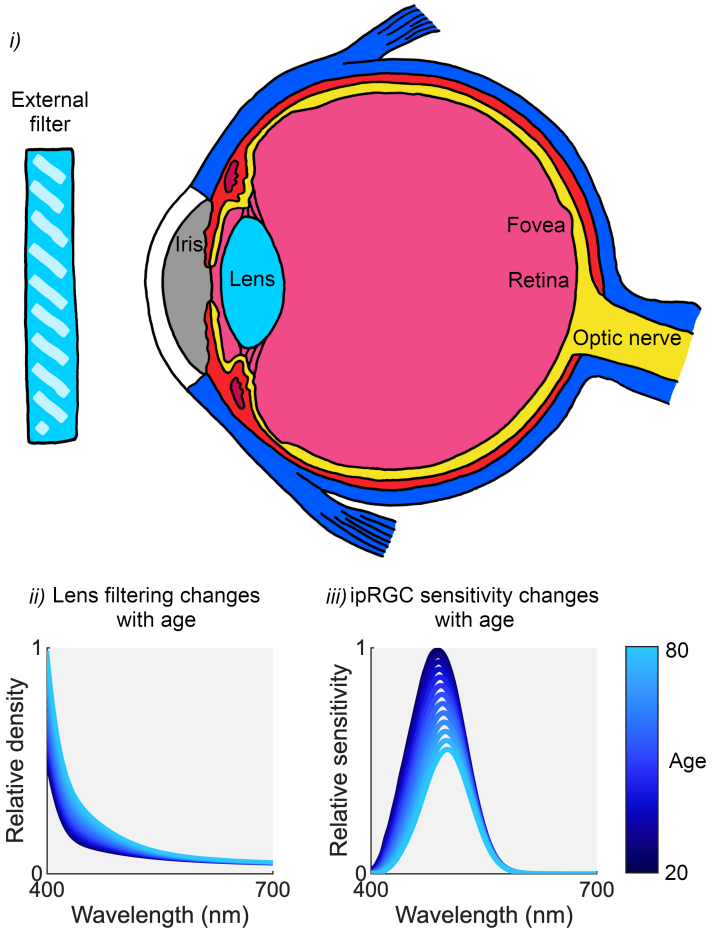

Figure 2.2: Schematic of lens filtering effects due to age. *i)* a schematic of the eye, showing the lens in front of the retina. *ii)* the relative density of the lens across different wavelengths. Dark blue curves are for 20-year-old observers and light blue curves are for 80-year-old observers. *iii)* ipRGC sensitivity spectra accounting for lens filtering.

2.2.9 Viewing regime changes

We additionally analysed the changes in viewing regime as a result of filter use. To do this, we analysed the luminous intensity of D65 spectra of a range of intensities after being filtered by each filter and categorised the intensities based on the CIE-defined intensity boundaries [107, 229]. This means that scotopic vision is defined as occurring below -3 log units of luminance, mesopic

vision occurs between -3 and 0.47 log units of luminance, and photopic vision occurs above this. We calculated the range of intensities that fall into each category after filtering, thus finding the changes in boundaries that each filter causes.

2.2.10 Contrast sensitivity

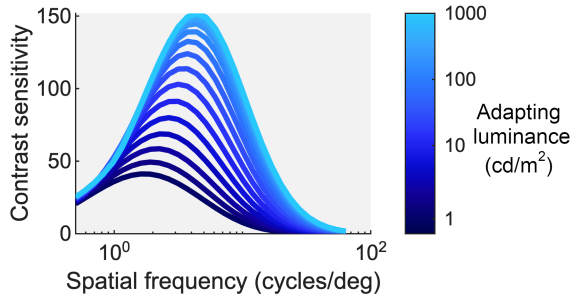

Figure 2.3: Contrast sensitivity as a function of spatial frequency for different luminance backgrounds according to the castleCSF model [230]. Dark backgrounds are shown in dark blue, and light backgrounds are shown in light blue.

Finally, we analysed changes in contrast sensitivity as a result of the different amounts of light allowed through the different filters. To do this, we utilised the castleCSF unified model [230], which combines the results from many psychophysical experiments to create a parametric model of the contrast sensitivity function. Figure 2.3 shows contrast sensitivity as a function of spatial frequency for stimuli at a range of luminances (darkest stimuli, 1 cd/m^2, shown in dark blue to lightest stimuli, 1000 cd/m^2, shown in light blue). To score each filter based on its overall contrast sensitivity, we calculated the area under the contrast sensitivity function curve for a 100 cd/m^2 light. We then calculated the proportion of this global contrast sensitivity score retained after viewing the light through each filter. We additionally calculated the proportional decrement in high spatial frequency cutoff - known as acuity - for the same light viewed through each filter.

2.3 Results

2.3.1 Dose response delta

Figure 2.4 shows a comparative analysis of two commercially available short-wavelength blocking filters, the *Brain Power Inc. Tennis Tint*, shown in brown, and the *DAO-AG L480-90*, shown in teal. Figure 2.4i shows the transmittance spectra for the two filters, with the L480-90 restricting much more light than the Tennis Tint, particularly in the short-wavelength range. Figure 2.4ii shows the circadian dose response delta curves for the two filters, calculated for a

standard observer (age = 32 years). In this figure, each datapoint represents the difference in normalised global circadian response for a given natural daylight spectrum when viewed with and without the filter. The darkened lines overlapping each set of datapoints shows the same data smoothed through a Savitsky-Golay filter. Both filters show significant decrements above 10 lux melanopic EDI, however the difference in magnitude of these effects is large. The Tennis Tint filter maximally reduces the circadian effect of light by 23% at 70 lux melanopic EDI. This corresponds to the Tennis Tint filter being maximally effective at reducing circadian effects in typically domestic evening light levels. At light levels higher than this, the effect of the filter is reduced, and provides no reduction in circadian effect for light levels above 10,000 lux melanopic EDI. By comparison, the L480-90 produces its maximal effect at around 1,500 lux melanopic EDI, reducing the normalised circadian effect by 94%. Compared to the Tennis Tint filter, it also has a much greater impact at higher light intensities, still reducing the normalised circadian response by around 87% at 10,000 lux melanopic EDI. Sufficiently intense light spectra would yield a reduction in effect, similar to that seen in the Tennis Tint curve, however the typical range of natural daylight spectra do not span to this intensity. Figure 2.4iii shows the luminant and melanopic transmittance, and chromatic properties of both filters. They show that the Tennis Tint (left column), compared to the L480-90 (right column), permits significantly more luminous and melanopic transmittance (85.5% vs. 6.7% and 63.7% vs. 3%, respectively). The L480-90 has a much larger chromatic effect also, shifting D65 by 0.1 Euclidean units in u'v' space and reducing the gamut size, calculated for 99 representative spectra in u'v' space [225] with and without the filter, by 74.4%. The chromatic effects of the Tennis Tint are smaller, with D65 shifting only 0.06 u'v' units, and only a 49.1% reduction in colour gamut. Figure 2.4iv shows the same circadian effect plots normalised by the Savitzky-Golay curve ($Order = 3$, $Frame\,Length = 21$), thus showing the spread around the smoothed data. For the Tennis Tint filter, it shows a spread of around 2% at the peak of the function. This variation is the result of the spectral differences in the tested spectra, with those being more affected than baseline having more short-wavelength light, and those that are less affected than the baseline having relatively more long-wavelength light. These spectral effects are greatly diminished for the L480-90 filter.

2.3.2 Parametric filter effects

Figure 2.5 shows the non-visual effects of a range of parametrically-defined filters. Figure 2.5i shows a range of filters, whose cutoff, λ_0, ranges from 380 nm to 550 nm. These filters have

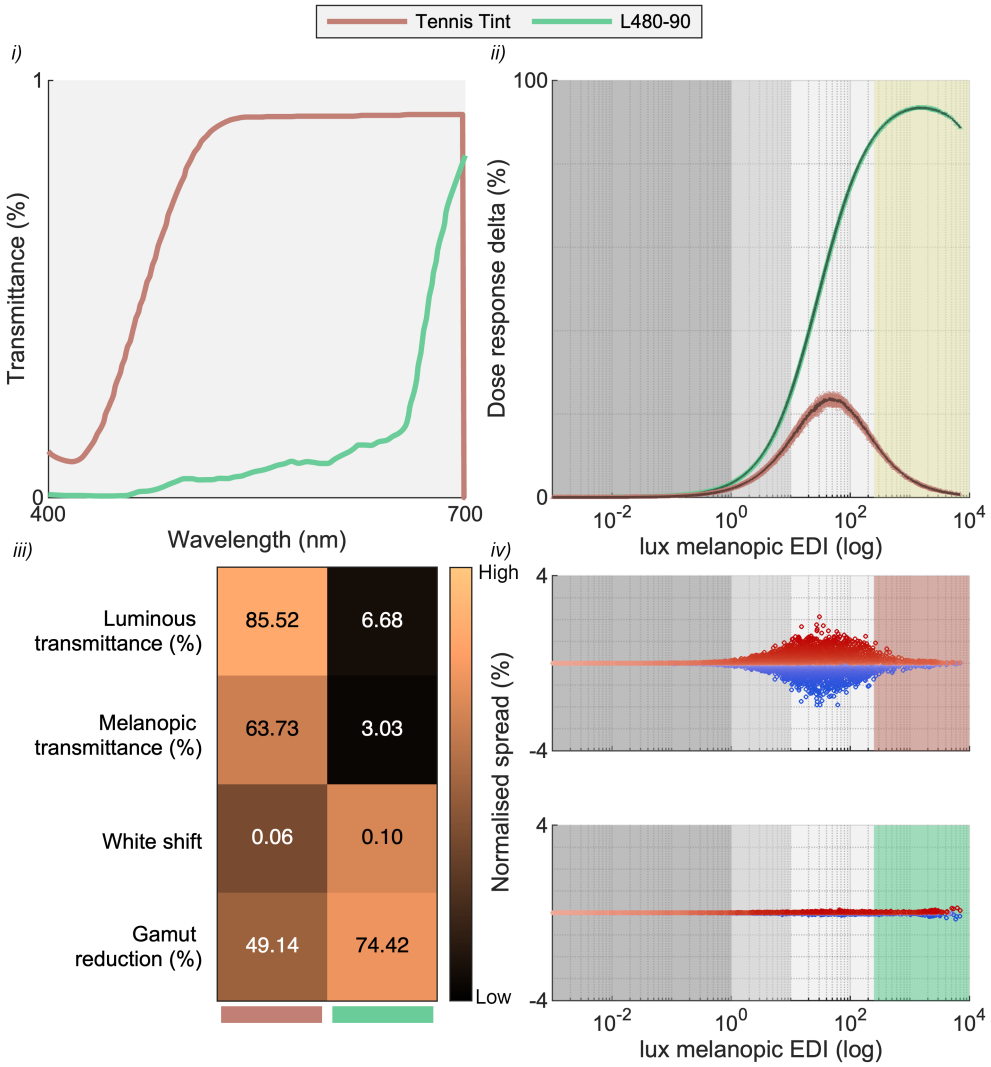

Figure 2.4: Comparison of the dose-response delta of two example filters, with the Tennis Tint filter marked by teal data points and the L480-90 filter marked by brown data points. *i)* transmittance spectra of two short-wavelength blocking filters. *ii)* dose response delta for the two filters. Dark lines show the Savitzky-Golay smoothed data. Dark shaded region corresponds to night time sleep environment conditions. The lighter grey region corresponds to evening domestic conditions. The region shaded yellow corresponds to indoor daytime environments. The unshaded region corresponds to environments that are considered sub-optimal for daytime light exposure [170] *iii)* table of visual metrics for both filters. Cell colour maps are normalised by metric, with percentage metrics normalised between 0 and 100, and white shift values normalised between 0 and 0.15 *iv* spread around the Savitzky-Golay estimate for both filters. Upper plot shows the Tennis Tint filter, and the lower plot shows the L480-90 filter. Positive spread is shown in red and negative spread is shown in blue.

parameters $\alpha = 0.95$, and $\beta = 0.5$. Figure 2.5ii shows the corresponding circadian deltas caused by each of these filters. It shows that as the filter cutoff increases, as does the maximum circadian effect of the filter. Filters with the lowest cutoff have maximum effects of around 4%, and filters with cutoffs around 550 nm have a maximum effect of around 92%. Filters with cutoffs higher than this show no additional reduction. In addition to this increase in maximum effect, filters with higher cutoffs have effects over larger melanopic EDI ranges than ones with lower cutoffs. Figures 2.5iii and 2.5iv show the effects of slope modulations. These filters have slopes ranging

from $\beta = 0$ to $\beta = 0.5$, $\alpha = 0.95$ and cutoffs $\lambda_0 = 500$ nm. The effect of this parameter is much smaller than cutoff, but a steeper slope still increases the maximum circadian effect from 39% to 50%. Figures 2.5v and 2.5vi show the corresponding results of asymptote modulations. These filters have asymptotes between $\alpha = 0$ and $\alpha = 1$, slopes of $\beta = 0.5$ and cutoffs $\lambda_0 = 450$ nm. These similarly show increased maximal circadian effect, up to 96%, when the asymptote is lowest and the most light is blocked by the filter.

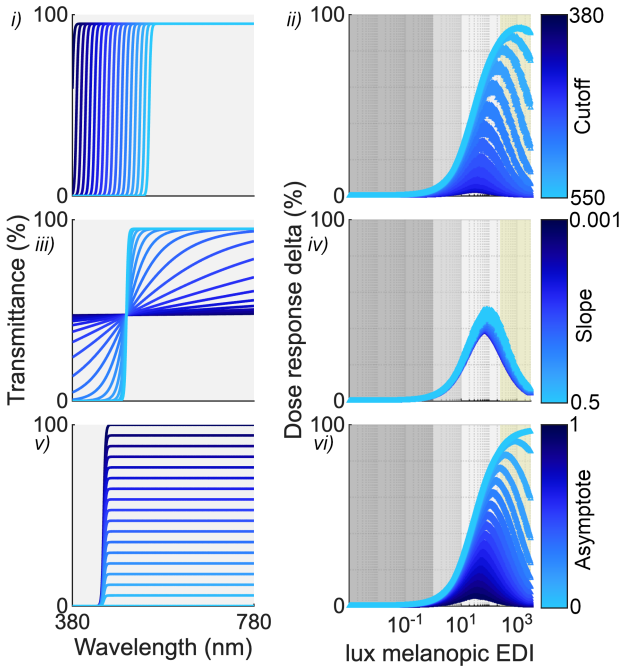

Figure 2.5: Comparison of parametric filters. subfigures *i)*, *iii)* and *v)* show parametric filters resulting from varying one variable - cutoff, slope and asymptote respectively. Subfigures *ii)*, *iv)* and *vi)* show the resultant dose-response deltas for each modelled filter.

2.3.3 Age effects

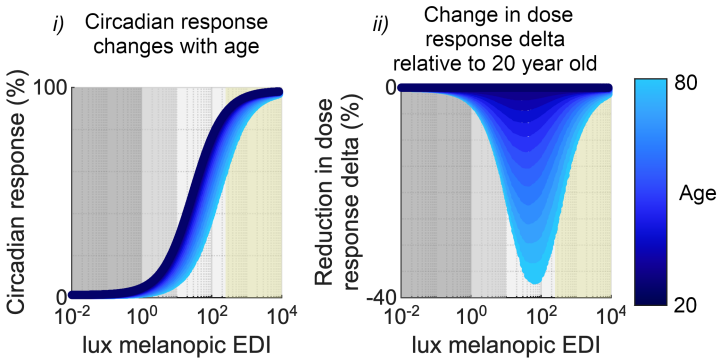

Figure 2.6: Age effects on circadian response. *i)* Change in circadian response across melanopic EDI for observers of different ages. *ii)* Dose response delta for different observers relative to a 20-year-old observer.

Figure 2.2 shows the changes in ipRGC sensitivity that occur as an observer ages. The

circadian effects of this are shown in figure 2.6. Subfigure 2.6i shows changes in circadian responses to the dataset of natural light spectra. Subfigure 2.6ii shows the reduction in dose-response delta for all modelled observers relative to a 20-year-old. This shows the significant decrement in circadian response that occurs for observers as they age.

To investigate the way observer age modulates the effect of short-wavelength blocking glasses, we modelled observers of different ages to observe changes to the dose-response delta curve as a result of prereceptoral filtering. Increased optical filtering leads to a reduction in the sensitivity of the resultant melanopic action spectrum. Thus, a reduction in the maximal effect of short-wavelength blocking filters is observed. This is shown in Figure 2.7i, which shows the relative reduction in maximal circadian effect for the 120 commercially available filters. There is a large variation in the extent to which the effect of the filters is modulated by observer age, with some yielding a consistent circadian effect throughout the lifetime, and others reducing their effectiveness by up to 12%, when comparing 20- and 80-year-old observers. Figure 2.7ii shows the transmission spectra for filters that are maximally affected by age variation. They all show a signature high-pass filter shape, with a cutoff of around 480 nm. As short-wavelength light at the retina is attenuated in old age, the effective melanopsin spectral sensitivity is maximally weighted around longer wavelengths. This means that, for high-pass filters with cutoffs in this range, the effect of the filter is reduced as wavelengths above the cutoff become the dominant contributors to melanopic light. Figure 2.7iii shows the transmission spectra for filters that are minimally affected by age variation. In comparison to the filters shown in 2.7ii, the filters in 2.7iii are significantly flatter around 480 nm, which makes them resistant to the effects of filtering due to age. It is important to note that these effects show a decrease in the difference between filtered and non-filtered viewing throughout the lifetime. As shown in 2.2, there are additional effects throughout the lifetime that reduce circadian responses. The results in figure 2.7 demonstrate that the expected effect of filters reduces throughout the lifetime in addition to these changes to non-filtered viewing.

2.3.4 Viewing regime changes

Figures 2.8 and 2.9 show the viewing regime changes for each of the 120 filters. Each row corresponds to a filter. The x axis shows log luminance in candelas per square metre. On the x axis, the non-filtered viewing regimes are shown, with mesopic vision spanning from -3 to 0.47 log units, and photopic vision being above 0.47 log units. For each filter, the blue bar shows

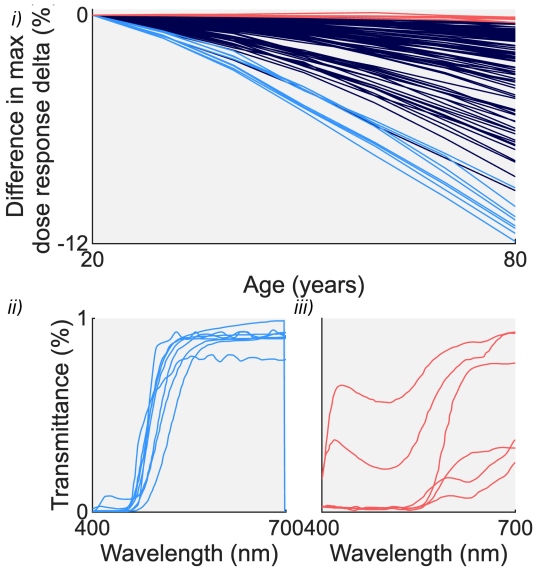

Figure 2.7: Comparison of the effects of filters when worn by observers of different ages. *i* relative change in maximum dose-response delta for each filter as a result of age. Maximally changing filters are highlighted in pale blue and shown in *ii)*, and minimally changing filters are highlighted in red and shown in *iii)*

the intensity range for D65 spectra that becomes scotopic after filtering, the red bar shows the light range that becomes mesopic after filtering, and the yellow region shows the intensity range that remains photopic after filtering. While there is some variance depending on the individual filter characteristics, these figures broadly show that much brighter environmental lights are brought into the scotopic viewing regime with any of the filters. Likewise, photopic environmental conditions are significantly shifted into the mesopic viewing range.

2.3.5 Contrast sensitivity

Figure 2.10i shows the proportion of global contrast sensitivity (area under the contrast sensitivity function) each filter has with respect to a 100 cd/m^2 light seen without a filter. There is a wide range of reductions in contrast sensitivity as a result of the reduction of luminance. Most filters retain around 80% of the contrast sensitivity, while the most strict filters reduce sensitivity to around 30%. Figure 2.10ii shows the reduction in spatial acuity - the high spatial frequency cutoff - as a result of filtering. Most filters reduce acuity by around 80%, with the strictest filters reducing acuity by around 50%.

2.3.6 Best performing commercially available filters

We sought to identify which of the commercially available filters provided the best overall performance, which we defined as maximising the circadian reduction while minimising the visual

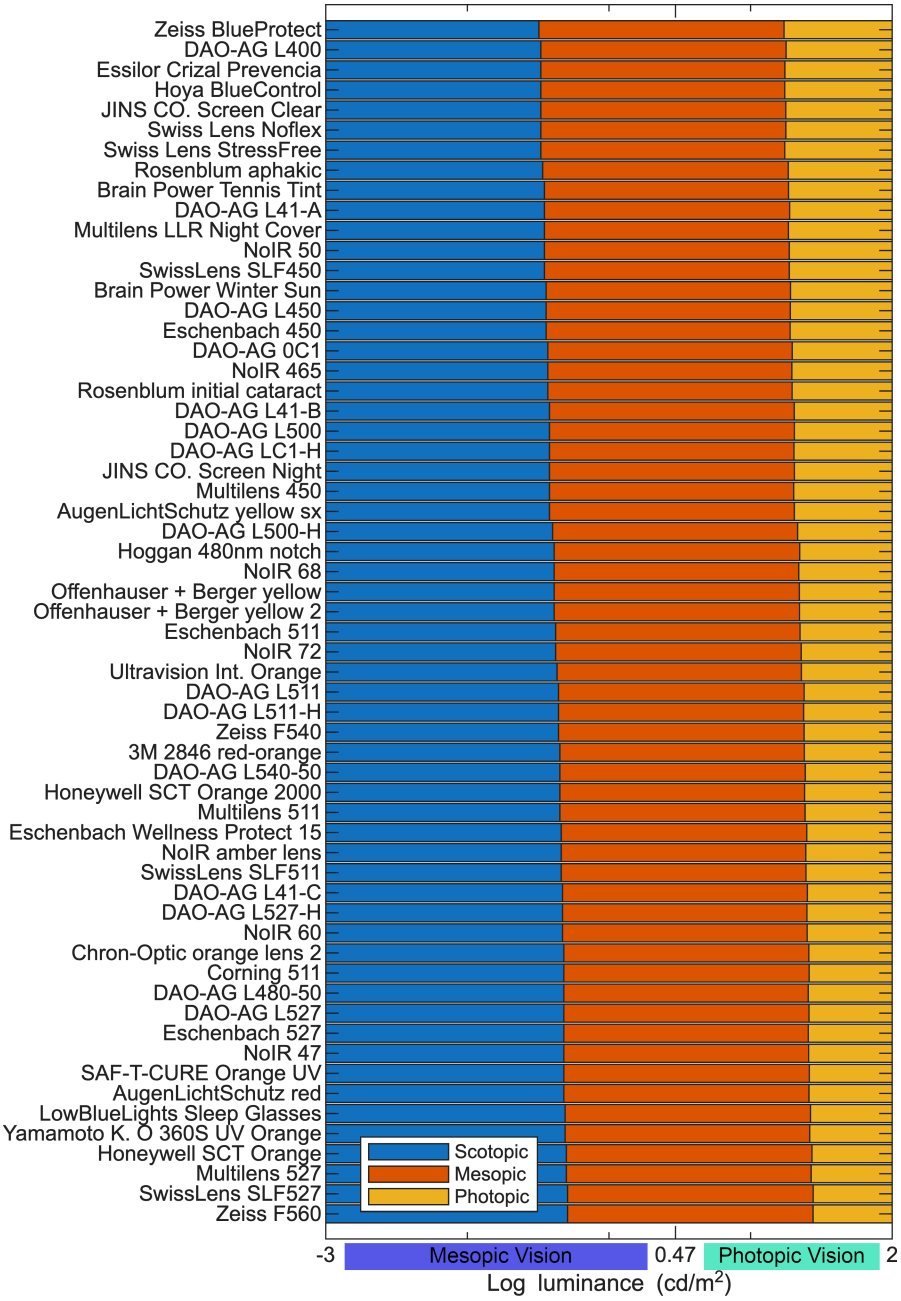

Figure 2.8: Changes in viewing regime due to filter use. Each row shows the viewing regime range for each filter. The x axis shows log luminance with the non-filtered mesopic range shown in purple and the non-filtered photopic range shown in teal. Blue ranges show scotopic vision, red ranges show mesopic vision and yellow ranges show photopic vision after filtering.

distortions introduced by the glasses. We weighted these two criteria equally.

Figure 2.11 compares commercially available filters (N = 120) in their ability to balance maximising circadian effect reduction (x-axes) and minimising chromatic and luminance distortions (y-axes). Each row of Figure 2.11 corresponds to a different metric of visual distortion that is compared against the circadian effect for each filter.

Row one shows the 2D u'v' colour gamut against circadian reduction. Figure 2.11i shows

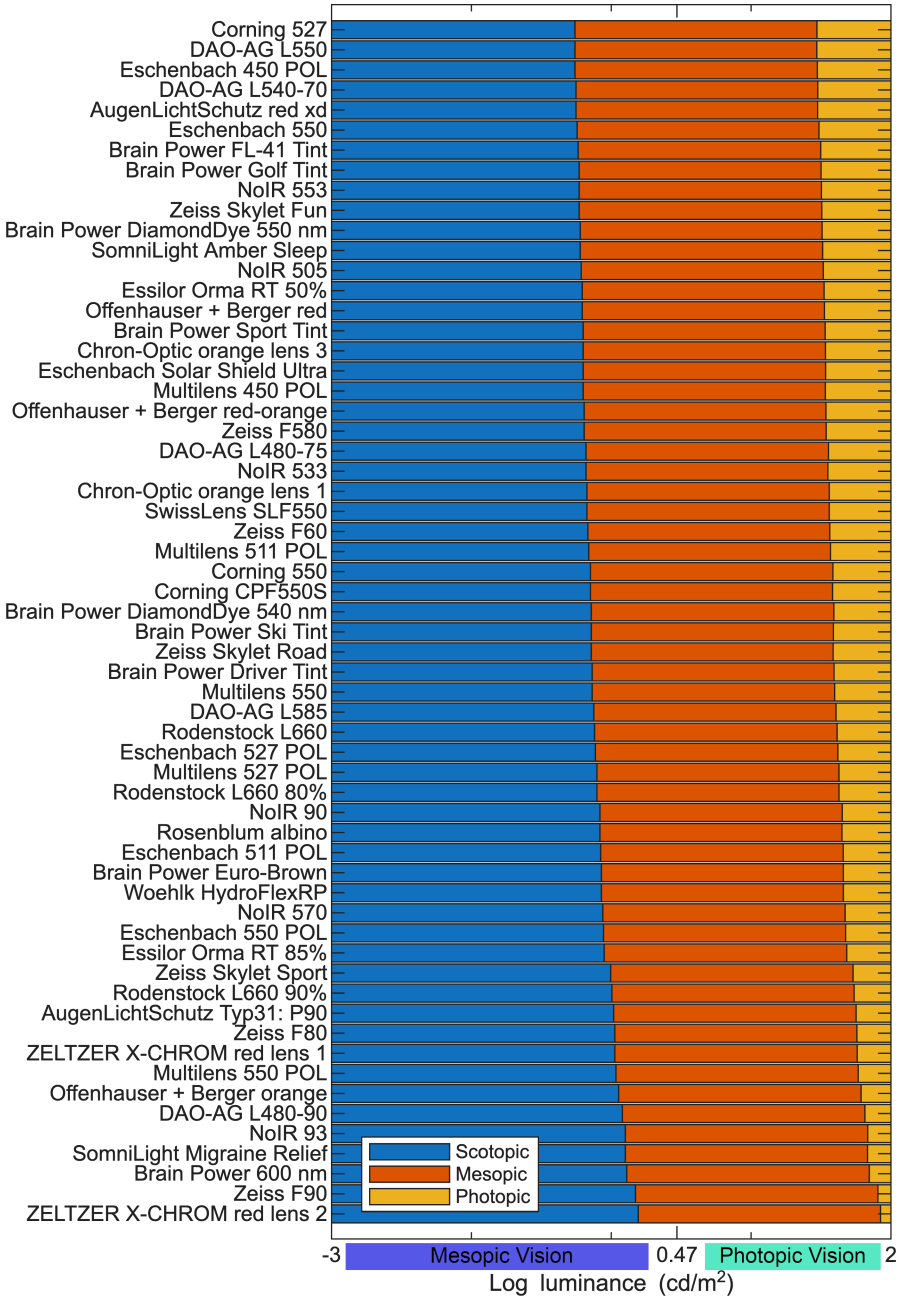

Figure 2.9: Changes in viewing regime due to filter use. Each row shows the viewing regime range for each filter. The x axis shows log luminance with the non-filtered mesopic range shown in purple and the non-filtered photopic range shown in teal. Blue ranges show scotopic vision, red ranges show mesopic vision and yellow ranges show photopic vision after filtering.

there is a significant negative correlation between the remaining gamut size and the circadian effect of the filter – those that block the most short-wavelength light also shrink the gamut the most. The best performing filters – those that minimise visual distortions while having significant non-visual effects – appear towards the top right corner of the scatter plots, while the worst performing ones appear towards the bottom left. Figure 2.11ii shows the transmission spectra for the two points highlighted in blue in Figure 2.11i. These spectra are characterised by largely

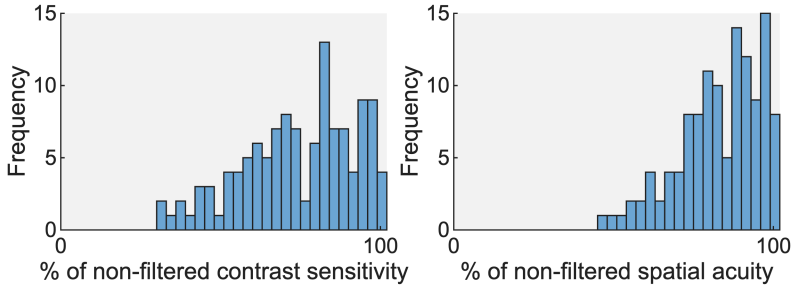

Figure 2.10: Histogram showing contrast sensitivity changes due to filtering. x axis shows the proportion of contrast sensitivity retained after viewing a 100 cd/m^2 through each of the filters

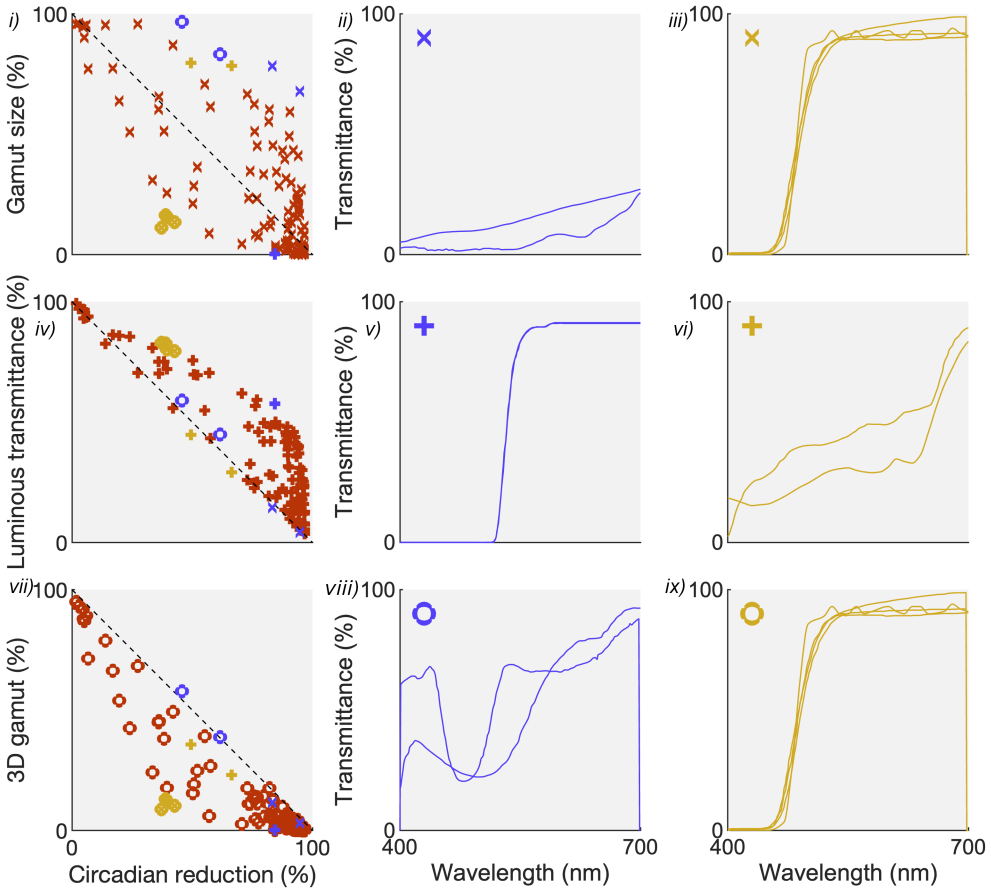

Figure 2.11: Best performing filters from the commercial dataset. Row 1 shows the 2D gamut size model. *i* each filter plotted in terms of its circadian reduction and gamut size reduction. Best performing filters are highlighted in blue and shown in *ii)*. Worst performing filters are highlighted in yellow and shown in *iii)*. Similarly, row 2 shows the luminous transmittance model, and row 3 shows the 3D colour model.

flat responses across the visible spectrum, which preserves a large 2D colour gamut. Importantly, this model does not consider the luminance of the resultant spectra, and consequently both of these identified 'best performing' filters significantly cut out light of all wavelengths. Figure 2.11iii shows the four filters that perform worst on this metric, highlighted in yellow in Figure 4a. These filters all have a characteristic low pass filter shape, with a cut off at around 480 nm. The sharp attenuation of short wavelengths in this range significantly reduces the blue-yellow

colour axis, while not fully extinguishing melanopic activation, leading to poor performance in the combined metric.

Row two shows the luminous transmittance of each filter plotted against their circadian effect. 2.11iv shows there is also a significant negative correlation between circadian reduction and luminous transmittance – filters that have a greater circadian effect also reduce luminance. Figure 2.11v shows the spectra of the two best performing filters, highlighted in blue in 2.11v. Both have very similar high-pass filter characteristics, with cutoffs at around 500 nm. This maximally blocks melanopic transmittance while maximising photopic luminous transmittance. Figure 2.11vi shows the transmittance spectra of the two worst performing filters for this metric. Both feature steadily sloping spectra, which allow moderate melanopic and photopic transmittance.

Finally, row three shows the change in 3D colour and luminance gamut, calculated as the volume of the convex hull around the u'v'L coordinates (two chromatic dimensions and one luminance dimension) of 99 representative spectra before and after filtering for each filter against its circadian effect. Figure 2.11vii shows that generally, filters perform worst at this 3D gamut metric as it combines both 2D colour gamut and luminance restrictions. There is also a significant negative correlation between circadian effect and 3D colour gamut. Figure 2.11viii shows the transmittance spectra of the two best performing filters for this metric. Both feature spectra that are flatter than the high-pass filters, both with a marked dip in the 480 nm range, and more transmittance in the extremely short wavelengths. This optimises the colour gamut by preserving some of the short-wavelength light that is captured by the S cones, while also dipping in the range dominated by the melanopsin spectral sensitivity. Figure 2.11ix shows the transmittance spectra of the four worst performing filters for this metric, all of which feature high-pass filter shapes with cutoffs at around 480 nm – notably these are the same filters that performed poorly in the 2D gamut model, chiefly because of the collapsing of the blue-yellow colour axis.

2.3.7 Optimally performing filters

In order to find the best performing short-wavelength blocking filters, optimisation functions were used to find the model parameters that maximise circadian dose-response effects while minimising chromatic and luminance effects. This optimisation process was completed for the logistic model, and also on a set of basis functions derived from the 120 filters from Spitschan et al.'s [225] dataset of commercially available filters.

To derive these basis functions, a principal component analysis (PCA) was performed on the

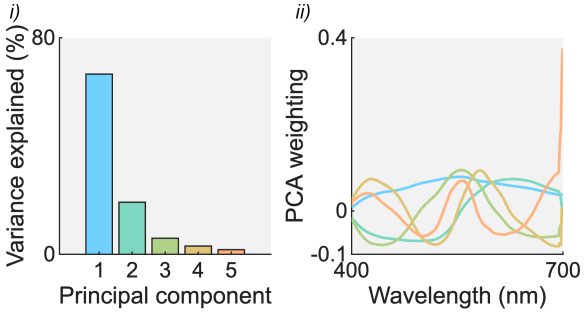

Figure 2.12: Principal components of filter transmittances. *i)* shows the amount of variance explained by the top five principal components. *ii)* shows the spectral shape of these principal components which are subsequently used as the basis functions for optimisation.

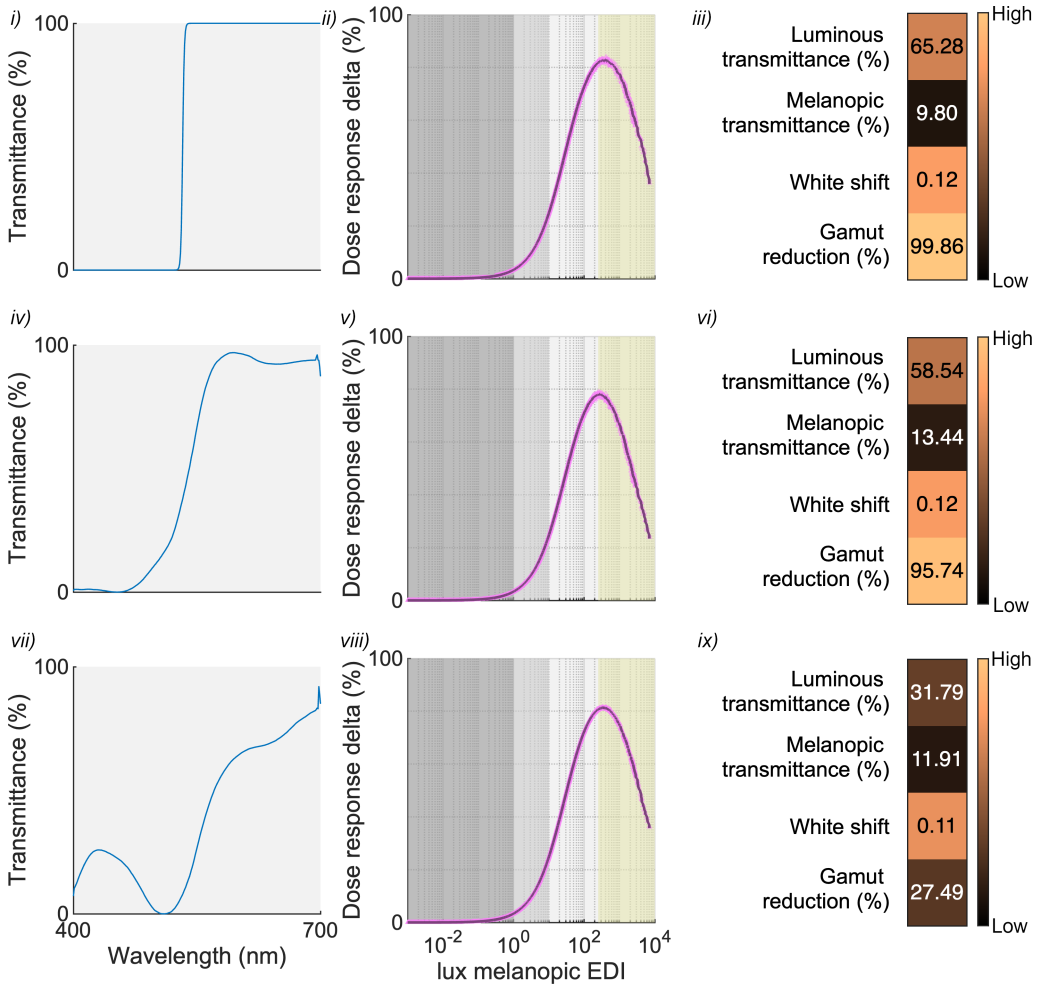

Figure 2.13: Performance of optimised filters. Row 1 shows the parametric model optimising luminous transmission and circadian effect. *i)* shows the optimal transmittance profile. *ii)* shows the dose-response delta for this model. *iii)* shows the visual characteristics of this model. Similarly, row 2 shows the PCA-derived model optimising luminous transmission and circadian effect. Row 3 shows the PCA-derived model optimising the 3D colour and luminance gamut and circadian effect.

dataset, which treats each transmittance wavelength (400 nm-700 nm in 1 nm spacings) as an independent variable, and each filter as an instance of observations for each variable. The PCA calculates the underlying principle components that best describe the original transmittance

data. We took the five principal components that explain the most variance and found optimal weights for these based on the maximisation of non-visual and minimisation of visual effects of the resultant filter. Figure 2.12i shows the amount of variance explained by each of the top five principal components. Figure 2.12ii shows the basis functions of these five components.

In total, six optimisations were performed: three using the parameters of a logistic model, and three using the weightings of five principal components. Each model sought to maximise the circadian effect of the filter while minimising a different visual component, in accord with the metric combinations shown in figure 2.11. One function minimised two-dimensional (u'v') gamut reduction. A second function sought to maximise luminous transmission. A third maximised the volume of a three-dimensional (u'v'L) gamut, taking into consideration both the chromatic and luminance properties of the filter. Some of these optimisations led to trivial results. Namely, parametric high-pass filters optimising 2D or 3D gamut along with circadian effect always converge to flat transmittances, either with very low or very high overall transmittance for 2D and 3D metrics, respectively. Additionally, the PCA-derived optimisation using only the 2D colour gamut converged on a flat spectrum, as luminance was not considered. The results of the remaining three non-trivial optimisation procedures are shown in Figure 2.13.

Figure 2.13 (row 1) shows the result of the high-pass parametric optimisation maximising luminous transmission and circadian effect. It is characterised by a cutoff at $\lambda_0 = 530$ nm, an asymptote of 1 and a slope of 1. Figure 2.13i shows its transmittance spectrum. Figure 2.13ii shows the circadian response using this filter. The maximum effect of this filter occurs at around 360 lux melanopic EDI, reducing circadian effects by 83%. The filter does this while achieving 65.3% luminous transmittance, as shown in Figure 2.13iii. The gamut is reduced to just 0.14% of its original size.

Figure 2.13 (row 2) shows the result of the PCA-derived optimisation maximising luminous transmission and circadian effect. Figure 2.13iv shows that the resultant spectrum resembles a low pass filter with a cutoff in the region of 530 nm, though it appears more sloped than the parametric optimisation. As a result, Figure 2.13v shows a similar but slightly reduced curve when compared with Figure 2.13ii. This filter has its largest circadian effect at around 260 lux melanopic EDI, reducing circadian response by 78%. Figure 2.13vi shows that this filter retains marginally lower luminous transmittance than the parametric filter, at 58.5%, but retains 4.26% of the original gamut.

Figure 2.13 (row 3) shows the result of the PCA-derived optimisation maximising the 3D

gamut and circadian effect. Figure 2.13vii shows that the transmission spectrum of this filter replicates the dip around 480 nm observed for the best performing commercial filters seen in Figure 2.11viii, with some short-wavelength transmittance retained. Figure 2.13viii shows that this filter also performed similarly to the other optimised filters, obtaining a maximum of 81% circadian reduction at 350 lux melanopic EDI. Figure 2.13ix shows that this filter performs worse in luminous transmittance than the other two optimised filters, allowing only 31.8% of luminous intensity through, however this is balanced by the large remaining gamut size, which is 72.51% the size of the unfiltered gamut.

2.4 Discussion

This analysis has provided a framework to model and estimate the physiological effect of short-wavelength blocking filters. By calculating the change in normalised circadian response using the dose-response curve reported by Brown et al [170], it is possible to estimate the intensity range over which a filter reduces circadian effects and to what extent it does this, thus accounting for a non-linearity previously not incorporated into models of the effect of short-wavelength blocking filters. We encourage authors to report these physiological metrics in future work. Shortly before the submission of this thesis, a paper describing a similar procedure was published by Glickman et al [231]. While there is some overlap in the content of the two papers, our method considers the performance of filters over the range of naturally occurring light spectra, rather than singly the maximal effect of the filter.

Our results also highlight the heterogeneity in the effects of commercially available short-wavelength blocking filters. All of the filters considered here have some effect on circadian, luminousity and chromaticity metrics, and while these are correlated there are large differences in the benefits and trade-offs of different transmittance profiles. High-pass-like profiles can significantly reduce circadian exposure while retaining luminous intensity, but they significantly distort the available colour gamut. They are also more susceptible to age-based variation, as shown in figure 2.7. By comparison, flatter filters that feature a dip around 480 nm, where the melanopsin spectral sensitivity at its highest, are much better at retaining the colour gamut and are less susceptible to changes due to age. However, these filters are often worse at luminous transmission. We also specifically note the importance of considering age when selecting short-wavelength blocking filters. Some filters vary by as much as 12% in terms of

their maximal circadian effectiveness, depending on the age of the user. It is thus important to consider the effect of short-wavelength blocking filters for an individual wearer, in addition to asking whether the effect of the glasses will change over time. We additionally show that the viewing regime a particular light intensity yields varies greatly when filters are worn. This indicates another consideration for users of short-wavelength blocking filters, as the introduction of rods into mesopic vision at much higher light levels may lead to differences in perceptual experience [99, 100]. By utilising physiologically-relevant models of functional changes to viewing regime, contrast sensitivity, acuity and circadian responses, the modelling methods presented here allow us to predict the many ways in which the use of short-wavelength blocking filters affect the perceptual experiences of the users who wear them.

Mathematical models have been devised to predict the effect of different light levels on circadian responses, such as the work by Skeldon et al [159] which uses van der Pol oscillators to predict sleep-wake state changes controlled by an internal circadian clock, which is itself nudged by light input. This model is able to account for the temporal dependence of circadian effects; however, the model quantifies light input in photopic, rather than melanopic units. Further sleep lab experiments are required to help generate models that incorporate this temporal dependency, as well as clarify the precise photoreceptor inputs to the suprachiasmatic nucleus.

In addition to short-wavelength blocking filters, "night mode" is a popular method for many who wish to reduce their circadian light exposure late in the day while still being able to use devices with LED displays. This works by rescaling the red, green and blue channels of the screen in order to reduce the user's short-wavelength exposure. This can also be used in conjunction with films placed in front of displays to alter their appearance. Analysis of the primaries used in displays along with our methodology would allow for night modes scaling, film filters and even primaries to be selected in a way that optimises visual metrics like luminance and colour gamut while reducing adverse effects on circadian health.

2.5 Conclusion

Commercially available and parametrically modelled short-wavelength blocking filters can have a range of physiological effects in both the visual and non-visual domain, and there are a number of ways in which these optics can be successful - by maximising a particular metric, or by balancing between desired non-visual effects and unwanted visual distortions. We encourage authors to

report on both the visual and non-visual properties of filters and consider the ways in which their effects might be modulated by individual- and task-based characteristics.

2.6 Acknowledgements

The contents of this chapter are entirely that of the author, with supervision by Hannah E. Smithson and Manuel Spitschan. Thank you to Rafael Lazar for providing the dataset of spectral recordings used in this paper.

Chapter 3

Quantifying CIE alpha-opic signals in the indoor built environment

Contents

3.1 Introduction 70

3.1.1 Statistical regularities in environmental illumination 70

3.2 Methods 72

3.2.1 Data Specification 72

3.2.2 Camera Location 73

3.2.3 Capture Scheme 76

3.2.4 Data Validation 77

3.2.5 Data Analysis 77

3.3 Results 79

3.3.1 Seasonal variation 79

3.3.2 Weather Variation 85

3.3.3 Directional Variation 88

3.3.4 Additional Modelling 90

3.4 Discussion 92

3.5 Acknowledgements 96

3.1 Introduction

As humans spend more time in mixed-illuminant 'built' environments, it is important to quantify how light in indoor spaces differs from naturalistic scenes. Previous studies have quantified light across many natural environments and shown regularities in the chromatic variation across different seasons, times of day and weather patterns. This study measures light in a typical mixed-illuminant office space in the Northern Hemisphere (51.76° N, -1.27° W) and finds that it shares some regularities of chromatic variation with naturalistic scenes. In this dataset, such regularities are primarily conveyed through outdoor light entering through East- and North-facing windows, and reflected by surfaces inside the office, rather than by light directly imaged through the North-facing window that was visible in the camera field-of-view. Built environments that combine natural daylight and artificial light to create mixed-illuminant spaces can share many of the statistical regularities that have been found in natural environments.

3.1.1 Statistical regularities in environmental illumination

The chromatic properties of the lighting environments in which we find ourselves have a significant effect on our colour vision, and this is attributed to processes of adaptation that optimise our visual systems to encode useful colour information from our surroundings. As such, these optimisations are malleable and change over time as our environment changes, which in turn affects colour percepts [232–235]. The effects of adaptation to statistical regularities have been measured in infants as young as 4 months old [236], but long-lasting adaptation to different chromatic environments can yield significant changes to colour vision throughout the lifetime [232]. A number of datasets recorded in natural scenes have sought to quantify typical lighting environments in different parts of the world such as Japan, South Africa, Canada, among others [200, 237–240]. These datasets have demonstrated a number of reliably measured statistical regularities, both in the composition of individual scenes, and how these scenes change throughout the year due to seasonal variations in foliage, cloud coverage and sun paths. This study seeks to characterise the available light in an example mixed-illuminant office environment and determine whether some of the statistical features quantified in natural environment recordings are reproduced in indoor environments where the natural world is reduced in its influence over the scene, and artificial lights and surfaces additionally play a role in determining our light exposure.

One example of a significant statistical regularity observed in natural scenes was shown

by Webster and colleagues [240], who found significant chromatic differences when comparing the balance of light activating the long-wavelength cone population and light activating the medium-wavelength cone population ('L' and 'M' hereafter) in 'arid' and 'lush' environments from images taken across different seasons in the Sayadhri Mountains (∼10° N, ∼77° E) in India and the Sierra Nevada Mountains (∼37° N, ∼3° W) in the USA. They demonstrated significantly lower ratios of L-activating light to M-activating light in lush environments, owing to the 'greening' of foliage during spring or wet seasons. Webster et al. also found a number of other chromatic changes that occurred throughout the year, such as a shift in the distribution of colours in the MacLeod-Boynton chromaticity space – scenes examined in dry seasons showed a negative correlation between the S/(L+M) and L/(L+M) axes (where 'S' refers to the activation of the short-wavelength cone population), which was not observed in wet seasons. While the aforementioned 'greening' finding may be expected, it has significant implications for colour appearance. Welbourne and colleagues [241] examined unique yellow settings in a group of participants at multiple times throughout the year, finding that unique yellow matches made in the winter (in York, UK 54° N, -1° W) had a higher long-wavelength intensity than matches made in summer.

Other statistical regularities observed in natural environments have been linked to psychophysical results, such as those found by Vincent and colleagues [242], who showed L:M excitation ratio differences in the lightest and darkest regions of natural scenes depended on the seasonal characteristics of the scenes. Buck and DeLawyer showed that participants, when adapted to achromatic backgrounds of different luminances, made different unique yellow settings [243]. This suggests that the visual system may have different L:M adaptation states in different luminance levels, which may be the result of the kinds of variation in chromaticity across different luminance levels observed in natural scenes.

While the above examples provide the basis for the development of long-term adaptation states driven by changes in natural statistics, it may be important to consider whether these statistics are relevant in chromatic environments typical of the modern industrialised world. Wide ranging surveys of American people's time budgets suggest many spend around 90% of their time indoors [244] and thus quantifying the chromatic distributions of typical indoor environments becomes an important research goal for evaluating present-day adaptation stimuli, as well as the effects these might have on colour appearance. Indoor mixed illuminant scenes present an interesting chromatic environment, as they feature complex interplay between natural outdoor

lighting visible through windows, indoor artificial lighting and the many surfaces inside the indoor environment that are lit by a combination of the two. Some research, particularly that motivated by lighting design, has sought to quantify light in so-called 'built environments', though this often focuses on the non-visual effects of light in office spaces [206, 207, 245], rather than chromatic adaptation to indoor environments. Developments in our understanding of the non-visual effects of light, through the discovery of intrinsically photosensitive retinal ganglion cells (ipRGCs), have created significant research and commercial interest around 'Human Centric Lighting' [246], and while this has led to the use of some dynamic lighting systems that reduce blue light intensity in the afternoon [247], longer-term seasonal effects are not widely incorporated. Combining visual and non-visual approaches that consider these factors across different time scales to humans' 'spectral diet' [169] may be a useful way to answer questions about the effect of the built environment on those who use it.

We sought to collect a physiologically relevant dataset of indoor images that would allow us to gain some insight as to whether indoor environments behave similarly to naturalistic scenes previously reported. Specifically, we sought to evaluate whether the aforementioned differences in chromatic distributions were observed in a mixed-illuminant indoor environment that combined static indoor lighting with dynamic natural light from multiple windows. We considered two main predictions that would be affirmed if the influence of daylight entering the scene through windows was sufficient to impart natural statistical regularities into the indoor environment. First, we predicted that the ratio of L-cone activating light to M-cone activating light will be lower in spring and summer months compared to autumn and winter, coinciding with the seasonal 'greening' effect [240]. Secondly, we predicted that the regions of the scene with the highest luminance will feature lower ratios of L- vs. M-cone activating light during autumn and winter, though we might expect this effect to be in the opposite direction in spring and summer months [242].

3.2 Methods

3.2.1 Data Specification

Data were collected using the Westboro Photonics WP690E-Circ Imaging Circadiometer, controlled by the Westboro Photonics Photometrica software. The Circadiometer's sensor generates 3388x2712 pixel maps using 5 filters which are calibrated to conform to the CIE S 026/E:2018

CIE System for Metrology of Optical Radiation for ipRGC-Influenced Responses to Light [190]. This CIE report defines a standard observer model for the five photoreceptor classes of the human visual system, including sensitivity curves for the S-, M- and L-cone, as well as rods and ipRGCs. The term "alpha-opic" is used to identify the outputs of each of these receptors. Thus, each sample is comprised of five 3388x2712 photoreceptor maps of the alpha-opic radiances across space. Additionally, each photoreceptor map is the result of a High Dynamic Range (HDR) stack of up to five images, determined by the Photometrica software, which selects from a range of exposure times in order to prevent over-/under-exposure in each sample. The data values are then scaled to represent radiance in W/sr/cm2.

3.2.2 Camera Location

The office space chosen for data collection was an open-plan office area on the second floor of New Radcliffe House in Oxford, UK (51.76° N, -1.27° W). The camera was positioned in front of a desktop computer workstation at seated head height to mimic the point-of-view of an individual sat at the computer workstation. The office area was freely used by members of the Oxford Perception Lab throughout the duration of the image capture schemes. Ethical approval for this was granted by Oxford's Medical Sciences Interdivisional Research Ethics Committee with approval reference R81269/RE001. From the point of view of the camera, the desk, keyboard and workstation monitor are visible, along with other furniture inside the office space such as other desks, monitors, chairs and cupboards. A prominent North-facing window is visible in the scene on the far side of the office. Additionally, three East-facing windows were located to the right of the camera. While these were not captured in the camera's images, significant light entered the workspace from these windows. Notably, the artificial lights used throughout the office space were motion-activated, switching on whenever motion was detected in the area around the office, regardless of the ambient light level. The availability of artificial light is therefore linked to core working hours of the users of the office space (08:30-18:00) and decoupled from the availability of natural light. Throughout the data collection periods, the screen remained on, in order to mimic the additional light provided by computer use in the office. The light emitted from the screen remained constant throughout the data collection period. Figures 3.1a, b and c contain images depicting the location and setup of the camera, as well as details about the surrounding office.

Figure 3.2 contains information about the artificial lights in the scene. These lights are motion sensitive, turning on when movement in the office is sensed, regardless of ambient light level.

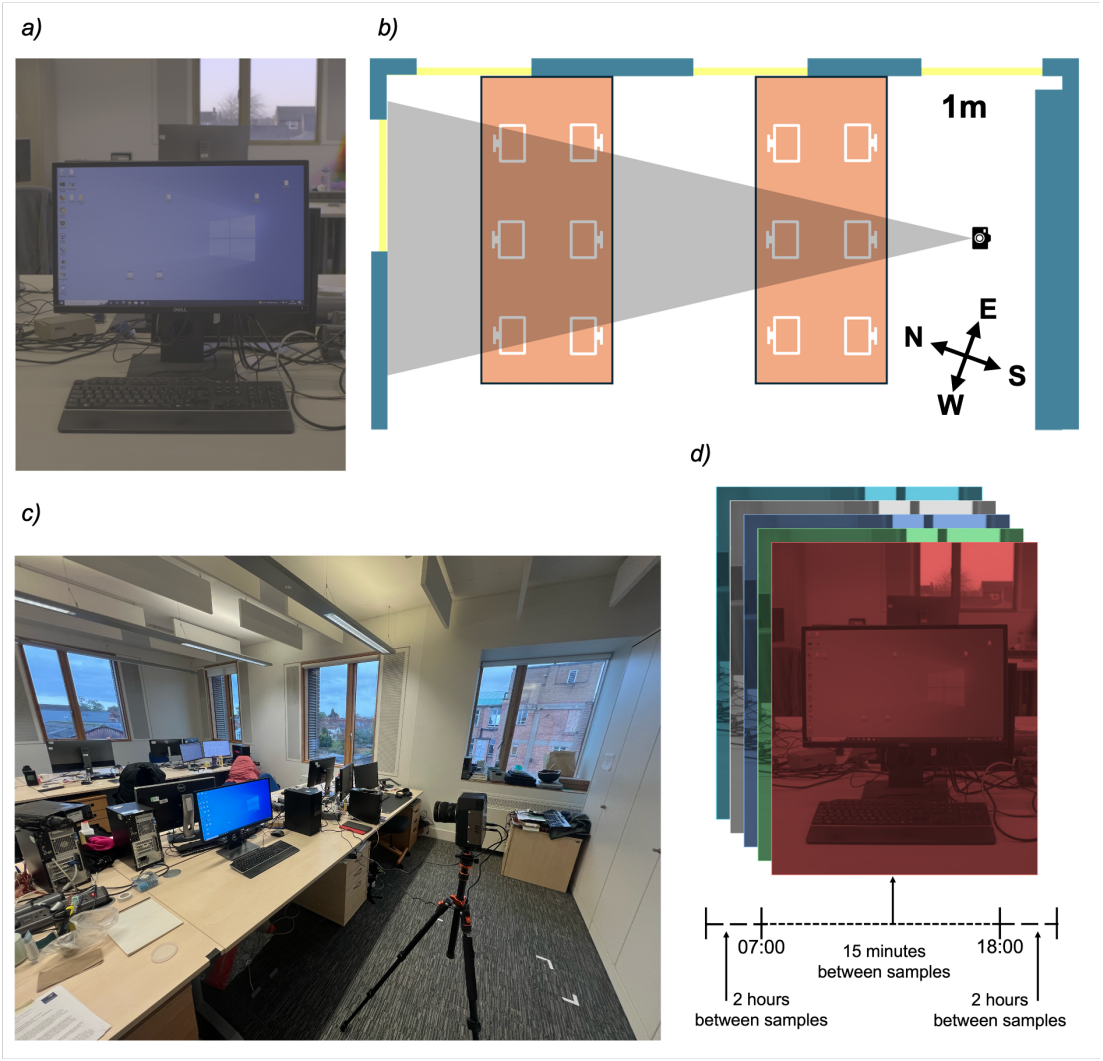

Figure 3.1: Circadiometer setup information. a) Re-colorised image of the point of view of the Circadiometer. b) Bird's-eye schematic of the office space. Blue bars denote walls, and the yellow gaps in these walls represent windows. The orange blocks show the positions of the desks, with the computer workstations atop them. The camera icon shows the location of the camera. The shaded triangle approximates the area that is visible to the camera. The compass icon shows directions relative to the office space. c) Side view of the location of the Circadiometer. d) Schematic of the data that are collected throughout the capture scheme, with 15-minute sampling between 07:00 and 18:00, and 2-hour sampling overnight. Each recording represents a five-layer map of each of the CIE alpha-opic radiances, each with a resolution of 3388x2712 pixels.

Plots a) and b) show the spectrum and relative photoreceptor responses to the overhead lights. Plots c) and d) show the spectrum and relative photoreceptor responses to the light emitted by the screen throughout the duration of the experiment. Schematic e) shows a schematic of the office space, with the locations of the overhead lights shown in purple. Table f) shows chromatic information for each of the light types.

Figure 3.3 shows a scatter plot of the chromatic distribution of the indoor scene with artificial illumination. This data is taken from a sample from the January dataset after dusk, such that there was no light entering the scene via the windows, but the indoor lights and computer monitor

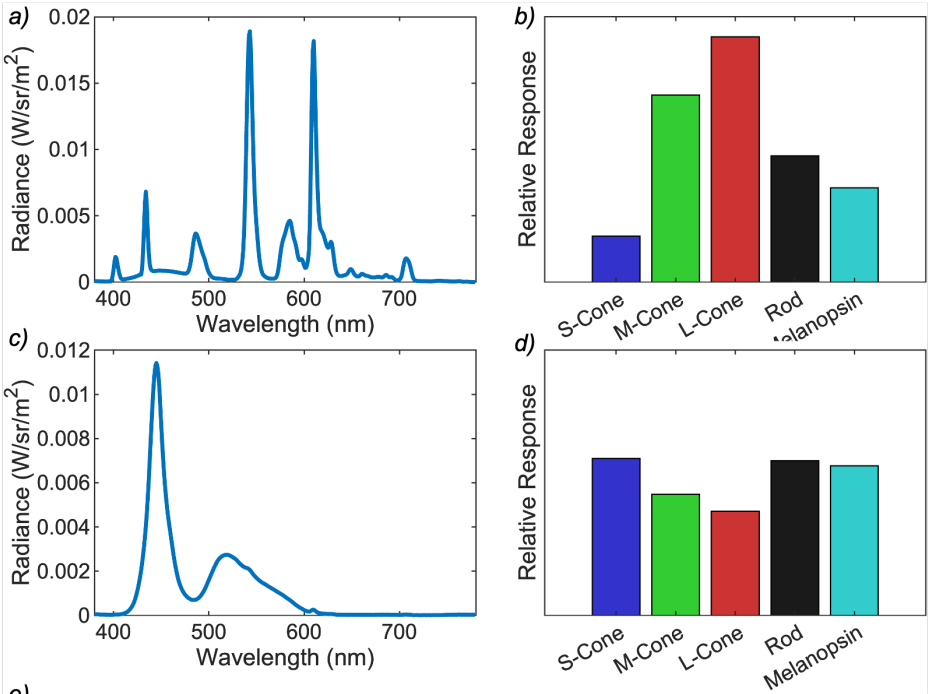

f)

	Room Lights	Monitor
Light Type	Fluorescent	LED
CCT	3337K	N/A
CIE 1931 xy	x = 0.4212	x = 0.1892
	y = 0.4110	y = 0.1960
Intensity	345 lux	97 cd/m^2

Figure 3.2: Schematic providing information about the artificial lights in the office space. a) shows the light spectrum emitted by the overhead lights, and b) shows the relative photoreceptor responses to this spectrum. c) shows the light spectrum emitted by the monitor, and d) shows the relative photoreceptor responses to this spectrum. e) shows a schematic of the office space, with the three purple bars denoting the locations of the lighting strips. f) is a table featuring additional colorimetric information about the lights. The intensity of the room lights is quantified in horizontal illuminance, measured from the desk in front of the camera. The intensity of the monitor is quantified in luminance.

were still on.

Figure 3.4 shows cartoon images of the view out of the North (plot a) and East (plot b)

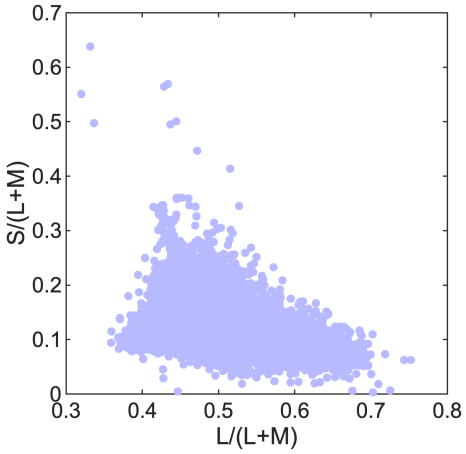

Figure 3.3: Scatter plot showing the chromatic distribution of pixels in the indoor region under only artificial light, plotted in MacLeod-Boynton chromaticity space, where the x axis corresponds to L/(L+M), and the y axis corresponds to S/(L+M).

windows. The images are split into sky (blue), foliage (green), building (orange) and tarmac (black) sections. Where applicable, a sun path is shown. Specifically, the sun path is visible from the East window at around 07:00 each day in the May dataset.

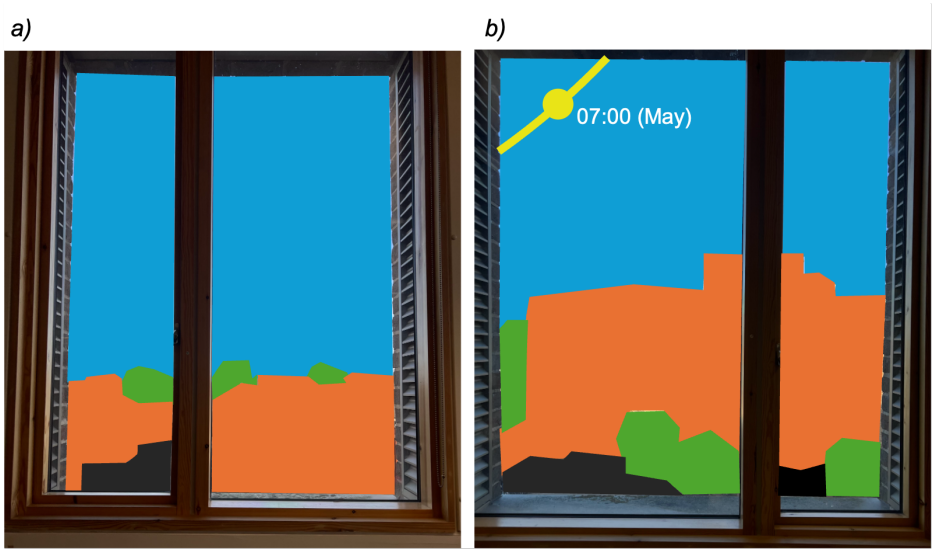

Figure 3.4: Cartoon images of the terrain visible out of the North (a) and East (b) windows. The regions are categorised into sky (blue), buildings (orange), foliage (green) and tarmac (black).The East window shows a partial sun path in yellow, with the sun being directly visible at 07:00 each day of the May dataset. It was not directly visible in January.

3.2.3 Capture Scheme

Figure 3.1d shows the capture scheme for a 24-hour period within the 7-day data collection period. Data captures were triggered automatically throughout the data collection periods every 15 minutes between 07:00 and 18:00 and every two hours overnight. These timings were chosen

to capture data at a high frequency during the early morning and 'working day' hours, and at a lower frequency outside of this time. Each HDR-stacked photoreceptor map was captured sequentially (in the order: S-cone, M-cone, L-cone, rhodopsin, melanopsin). Each image capture took between 30 seconds and two minutes to complete, depending on the light intensity at the time of capture. This 24-hour capture scheme was repeated over seven days, giving 329 time points, once in January 2023 and once in May 2024. These timings were selected in order to match the Autumn/Winter and Spring/Summer differences identified in previous literature demonstrating seasonal differences.

3.2.4 Data Validation

The Circadiometer's output values were validated by comparing the values recorded in an image taken of an X-rite ColorChecker in a Judge QC lightbox to measurements made by a PR-670 spectroradiometer of the same test patches in the same light setup. For each Circadiometer filter, this validation check confirmed the linearity and appropriate spectral filtering of Circadiometer outputs.

3.2.5 Data Analysis

To split the images into data-derived regions, a spatial principal component analysis (PCA) was performed on each dataset individually. Each timepoint in the dataset was spatially downsampled by a factor of four before the PCA was conducted to reduce the number of variables in the analysis. Each downsampled pixel location was then treated as an independent variable, and each timepoint within the season as a separate observation of each of the variables. The PCA was conducted on log transformed radiance data, as this provided a more reliable and interpretable result. K-means clustering was applied to the resultant PCA scores, with a prior K value of three. Three was selected for the prior K value as it corresponds to three kinds of regions within the scene – the window, which provides a direct view to natural light, the indoor region, which reflects light emitted by the indoor artificial lights as well as the natural light let in by the window, and the screen, which itself emitted light throughout the duration of the data capture periods.

In order to evaluate the chromatic characteristics of the images, spectral statistics were calculated using the L-cone, M-cone, S-cone, and melanopic activation maps. Scaling coefficients (0.689903 and 0.348322) were applied to L- and M-cone activations, respectively. This rescales

the summation of L- and M-cone activations to the luminous efficiency function, V(λ). No scale factor was used for S-cone and melanopsin activation maps. In addition to calculating luminance, defined as (L+M) * 683.002, we calculated chromaticity statistics for each pixel - L/(L+M), S/(L+M) and I/(L+M), where L and M refer to the scaled L- and M-cone activation maps, and S and I refer to the unscaled S-cone and melanopsin activation maps.

When comparing these statistics, only data collected between sunrise and sunset for each dataset (so-called 'daytime hours') was considered as the nighttime samples had low light levels and were therefore noisy. Within the January dataset, sunrise shifted from 08:01 to 07:54 and sunset shifted from 16:30 to 16:40. Within the May dataset, sunrise shifted from 05:00 to 04:54 and sunset shifted from 21:03 to 21:11. Core office hours remained between 08:30 and 18:00 throughout both datasets. Luminance-weighted chromatic statistics were also calculated to weight chromaticity information by a measure of intensity to better capture the light energy contributing to adaptation.

Additionally, to analyse the rod data we calculated the proportion of the scene intense enough to fully saturate the rods. To do this, we modelled pupil size using the average luminance across the scene and a replication of Watson and Yellott's pupil response model [181, 248]. We then down-sampled the image dataset to match the acuity limit of the rods at 6.5 cycles per degree [249] and then scaled scotopic luminance values by pupil area, and then by the peak scotopic efficiency of 1700 lm/W to find scotopic Trolands. We used a threshold of 3000 scotopic Trolands to determine the saturation cutoff [105, 250]. We then calculated the proportion of saturated pixel locations over the time series of data collection.

Finally, we analysed the indoor region more closely to identify the patterns that arise from the complex interplay of light across different surfaces. To do this, we calculated correlations between luminance and the stated chromatic statistics across pixels within the indoor region over time. Additionally, the time series of luminance variation in each pixel location was fitted with a sinusoidal function in order to compare phase deviations between the indoor region and the window. The cosine fitting model was given a fixed periodicity corresponding to the 24-hour diurnal cycle, fitting amplitude and phase in order to find the strength and delay on the 24-hour variation in each pixel location.

3.3 Results

3.3.1 Seasonal variation

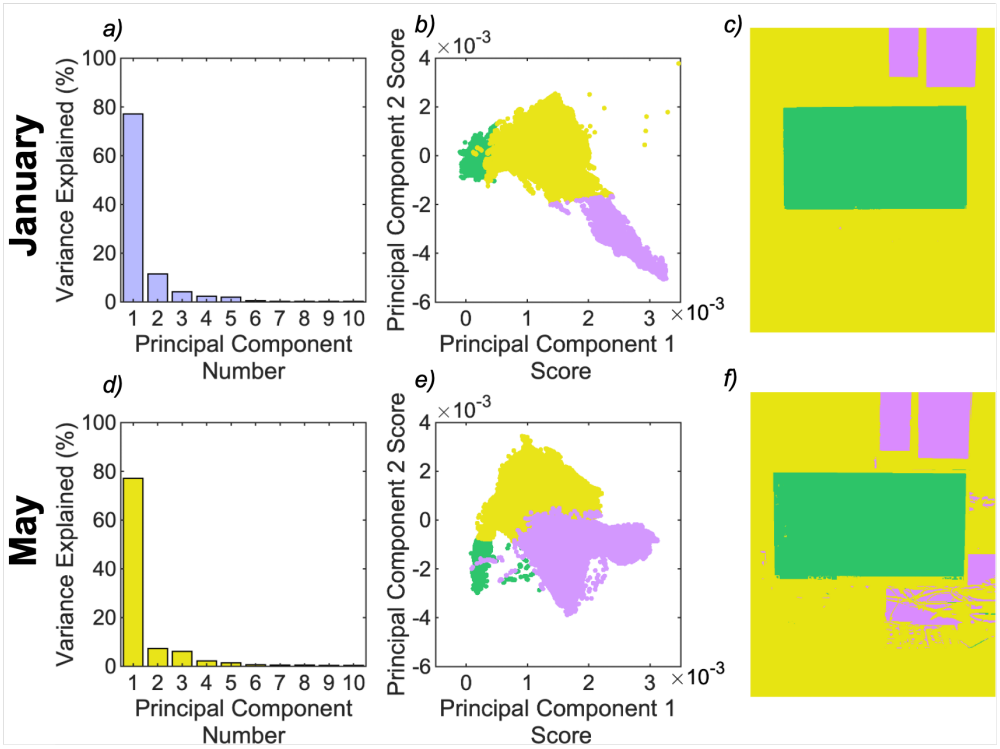

Figure 3.5: PCA and K-means results. a) and d) the amount of variance explained by each principal component for the January and May datasets, respectively. b) and e) the distribution of pixel weightings for principal components 1 and 2, with colours reflecting the K-means clusters. c) and f) clusters plotted back onto their position in the original scene.

Figure 3.5 shows the results of the PCA and K-means clustering. 3.5a and 3.5d show the amount of variance explained by each principal component for the January and May datasets respectively. In both PCAs, the first principal component explains around 80% of the variance. In the January PCA, one additional component contributes more than 10% variance and, in the May PCA, there are two components both contributing more than 5% variance. 3.5b and 3.5e show the distribution of pixels in the scene according to their scores on principal components 1 and 2. 3.5c and 3.5f show the K-means clusters replotted onto the original scene. Both K-means classifications segment the scene into three regions broadly characterised by the 'screen' (green), 'indoor' (yellow) and 'window' (purple) regions. Notably, the two classifications differ in their categorisation of some pixels on the desk – being clustered with other 'indoor' surfaces in January, but with the 'window' pixels in May. Subsequent analysis described below revealed that this was the result of the changes in the sun path, where highly directional sunlight causes bright reflections from the desk.

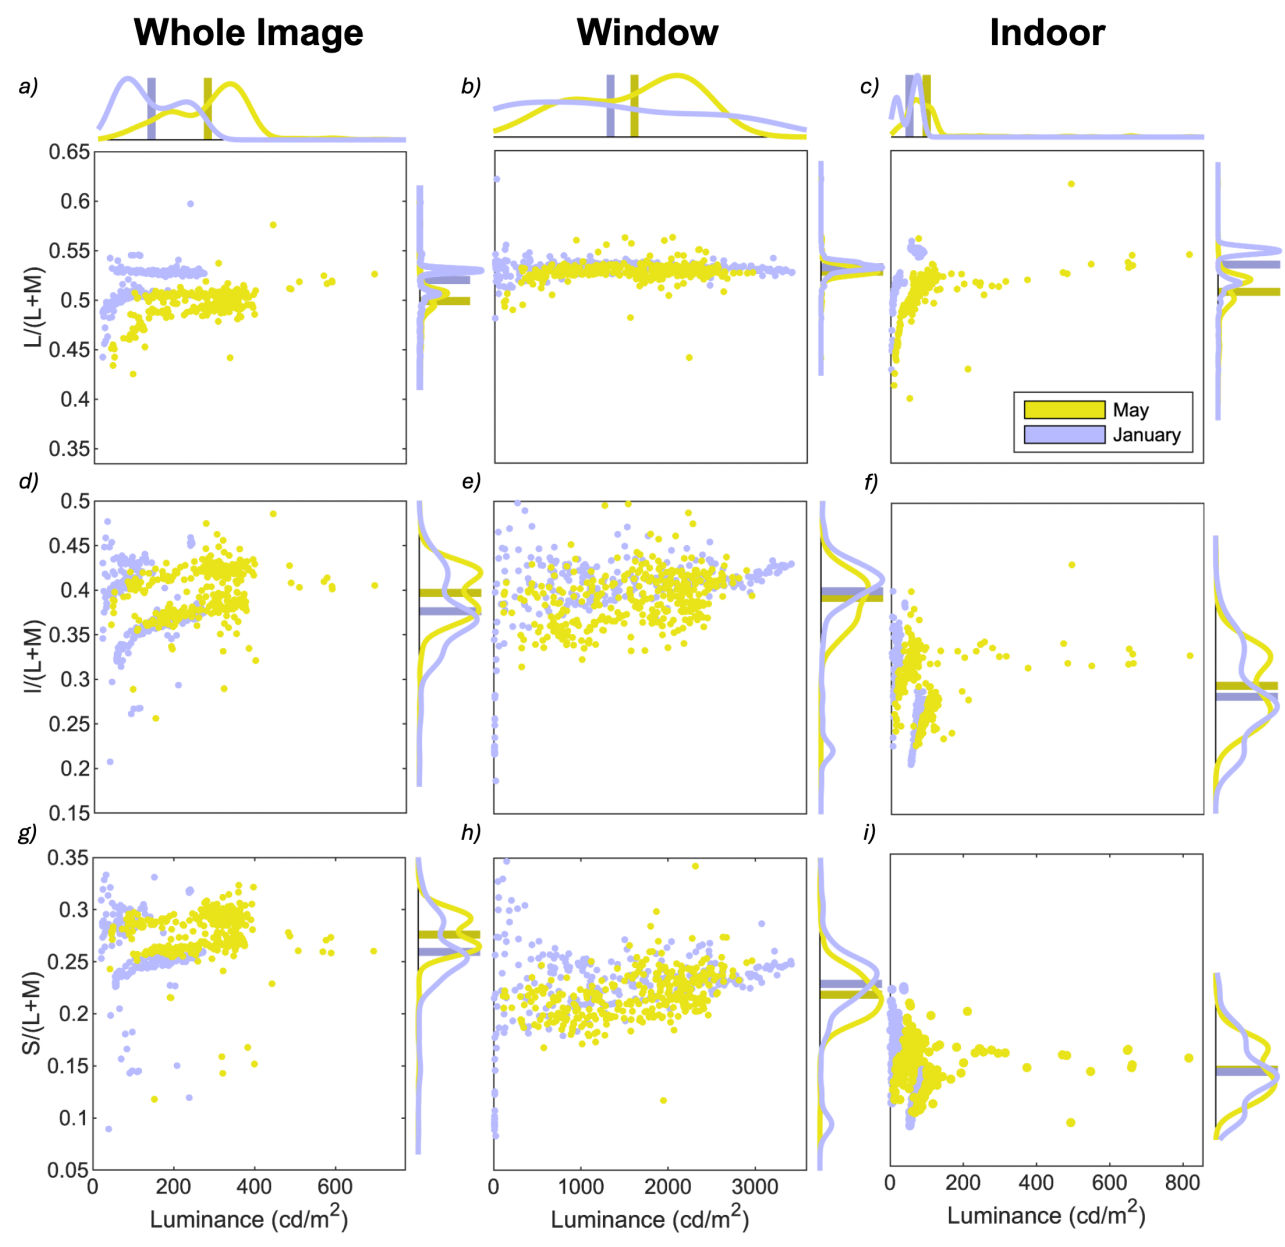

Figure 3.6: Scatter plots and histograms for mean chromatic and luminance values across different regions comparing the January and May datasets. January data are shown in blue, and May data in yellow. Each subplot shows a scatter plot and two kernel-smoothed histograms showing the distribution of each variable. Darkened bars on the histograms denote the mean for each variable. The left column shows whole image data, the middle column shows the window region data, and the right column shows the indoor region data. The top row shows L/(L+M) statistics, the middle row shows I/(L+M) statistics and the bottom row shows S/(L+M) statistics.

Figure 3.6 shows chromatic statistics across daytime images in both seasonal datasets. Row 1 (3.6a-3.6c) denotes L/(L+M) statistics, row 2 (3.6d-3.6f) denotes I/(L+M) statistics and row 3 (3.6g-3.6i) denotes S/(L+M) statistics, each plotted as a function of luminance. Within each figure, the January dataset is plotted in blue, and the May dataset in yellow. Each column relates to a spatial averaging of the data – data in column 1 is averaged across the whole image, meaning

that it includes data from the indoor, window and screen regions. Data in column 2 is averaged only over the window region and data in column 3 is averaged only across the indoor region. The curves to the right of each scatter plot and at the top of each column show the kernel-smoothed distributions of the data for each variable in each plot. The darkened bars denote the means for each distribution. For example, datapoints in the middle column are calculated by averaging across the window, with each datapoint corresponding to a different point in time. Equally, the whole image column features statistics averaged across the indoor, window and screen regions. Thus, the lowest luminance observed across the whole image is higher than that observed only in the window, as a result of the constant light emitted by the screen, incorporated into the whole image mean but not the window mean. This also means that the whole image statistics show a positive shift in S/(L+M) values compared to either of the other columns, as the blue light of the screen is included in the whole image column but not the other two. The whole-image analysis (column 1) shows that datapoints from the May dataset have lower L/(L+M) values, and higher I/(L+M) and S/(L+M) values. Conversely, the window analysis (column 2) shows minimal seasonal differences in the window region. Notably, the whole image and indoor regions feature a small number of May datapoints with significantly higher luminances than the rest of the data. These datapoints – that appear above 450 cd/m2 in the whole image data and above 200 cd/m2 in the indoor data – all occur shortly after sunrise, before 09:00.

The parametric Kruskal-Wallis test was used to compare each of these variables across the two seasonal datasets, as the presented data are not normally distributed. Results are shown in Table 3.1. Signs in Table 3.1 denote the direction of significant differences where they exist, such that "+" signs mean that a given region-statistic pair was higher in May than January, and "-" signs mean that the region-statistic pair was lower in May. Significant differences were found across all statistics when the images were averaged as a whole. Where window region differences were found between the two seasons, they were significantly smaller, and no significant difference was found in the L/(L+M) statistic. The indoor region also showed significant differences between L/(L+M), I/(L+M) and Luminance statistics, though not in S/(L+M).

To ensure that the differences in pixel location categorisation did not significantly influence our comparisons of the January and May datasets, we conducted the presented statistical tests on data grouped into these clusters, as well as manually defined 'window', 'indoor' and 'screen' regions, finding no difference in the season- and region-based effects. We additionally calculated luminance-weighted statistics for each region and season and conducted similar statistical testing.

Table 3.1: Chromaticity Kruskal-Wallis test statistics for seasonal differences

	Luminance	L/(L+M)	I/(L+M)	S/(L+M)
Whole Image				
χ^2	257.76^{+}	24.55^{-}	44.11^{+}	59.89^{+}
p	**<.001**	**<.001**	**<.001**	**<.001**
Window				
χ^2	30.21^{+}	1.77	11.14^{-}	6.25^{-}
p	**<.001**	.184	**<.001**	**.012**
Indoor				
χ^2	153.14^{+}	15.57^{-}	4.56^{+}	0.24
p	**<.001**	**<.001**	**.033**	.621

Figure 3.7 has the same data structure as figure 3.6, but for luminance-weighted chromaticity statistics. Here, each datapoint relates to a statistic averaged across the region defined at the top of each column, weighting pixels within the region based on their L+M value. Similar results are observed from these data as were observed in the non-luminance weighted data – the May dataset shows a 'greening' compared to the January dataset, such that L/(L+M) statistics decrease, while I/(L+M) and S/(L+M) statistics increase. However, the spread of differences is different for these data. Notably, differences in the 'whole image' datapoints are greatly reduced, while differences in the 'window' datapoints are increased. Kruskal-Wallis tests were also conducted for these data, the results of which are shown in Table 3.2.

Table 3.2: Luminance-weighted Kruskal-Wallis test statistics for seasonal differences

	Luminance	L/(L+M)	I/(L+M)	S/(L+M)
Whole Image				
χ^2	257.76^{+}	1.97	42.68^{+}	26.4^{-}
p	**<.001**	.160	**<.001**	**<.001**
Window				
χ^2	30.21^{+}	29.29	40.32^{+}	12.84^{+}
p	**<.001**	**<.001**	**<.001**	**<.001**
Indoor				
χ^2	153.14^{+}	25.26	18.55^{+}	4.12^{+}
p	**<.001**	**<.001**	**<.001**	**.0423**

Figure 3.8 shows L/(L+M) differences between indoor and window regions at the same time point. A positive difference means that the window region has a higher L/(L+M) value than the indoor region. January data are shown in blue and May data are shown in yellow. Figure 3.8 shows that for most of the May dataset, L/(L+M) is increased in the window region compared to the indoor region, however the January dataset is bimodal, with some of the differences being positive and some negative, though the group with negative differences is larger than the group with positive differences. This bimodality is associated with weather effects (see further

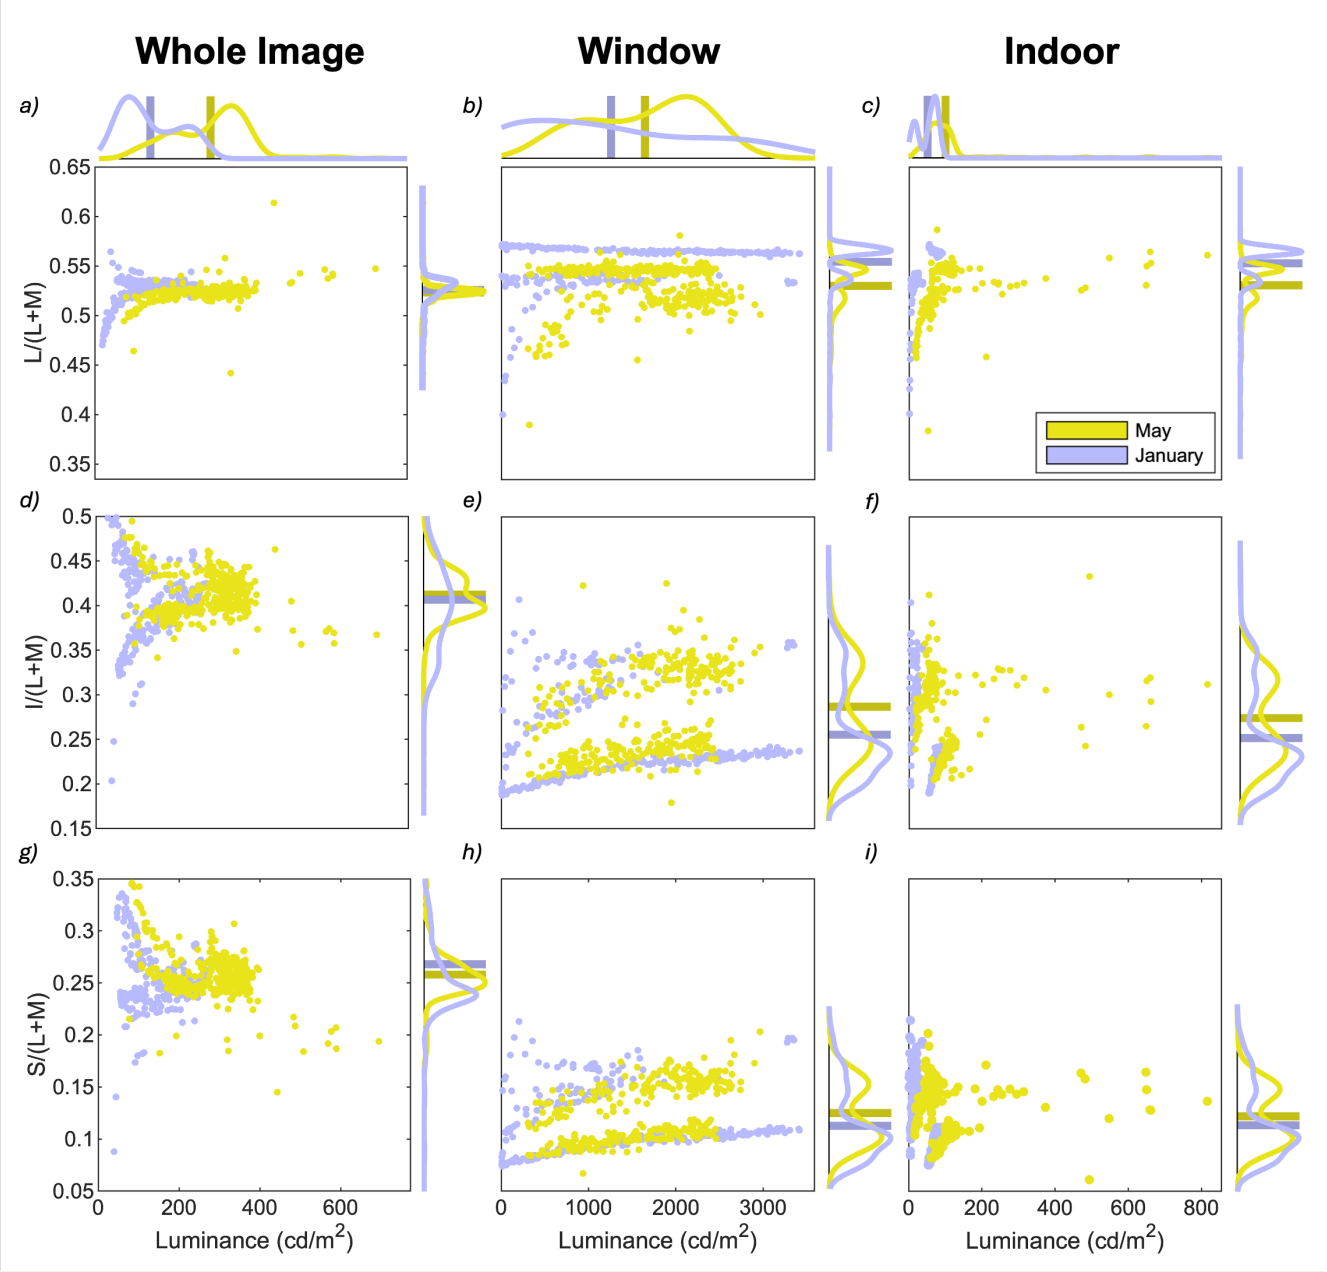

Figure 3.7: Scatter plots and histograms for luminance-weighted mean chromatic and luminance values across different regions comparing the January and May datasets. January data are shown in blue, and May data in yellow. Each subplot shows a scatter plot and two kernel-smoothed histograms showing the distribution of each variable. Darkened bars on the histograms denote the mean for each variable. The left column shows whole image data, the middle column shows the window region data, and the right column shows the indoor region data. The top row shows L/(L+M) statistics, the middle row shows I/(L+M) statistics and the bottom row shows S/(L+M) statistics.

discussion below).

Figure 3.9 shows a timeseries plot for the January dataset, where timepoint data are grouped based on whether the difference in L/(L+M) between the indoor and window regions was positive or negative (see Fig. 3.8). Orange timepoints indicate that the window L/(L+M) is lower than the indoor L/(L+M), while blue timepoints indicate that the window L/(L+M) is higher.

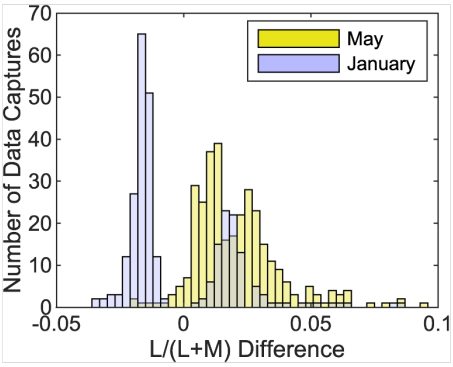

Figure 3.8: Difference in L/(L+M) statistic between window and indoor regions within time points. Positive differences denote that L/(L+M) was higher in the window region. January data are shown in blue and May data are shown in yellow.

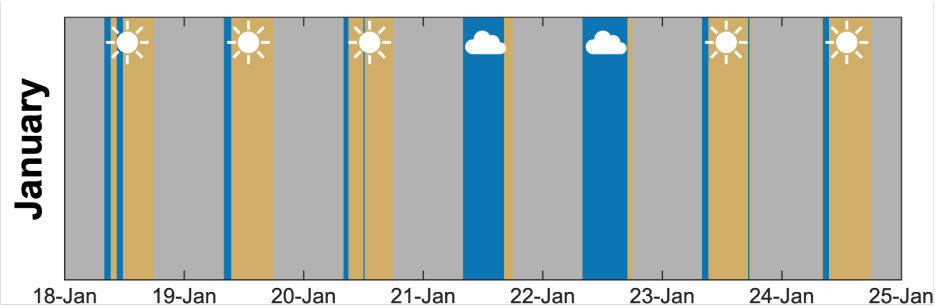

Figure 3.9: Timeseries plots where timepoints are colour-coded by the polarity of the difference in L/(L+M) between the window and indoor regions. Timepoints before sunrise and after sunset are coloured grey. Orange timepoints indicate a time where L/(L+M) was lower in the window region than the indoor region, and blue timepoints indate that L(L+M) was higher in the window region compared to the indoor region. White weather icons denote the weather grouping for each day. See section 3.3.2 for further explanation.

Figure 3.10 shows the percentage of the scene in which rods are saturated over time. The data show consistent peak at around 6% of the scene throughout the January dataset. An example image from January demonstrates that this reflects the outdoor region of the scene being almost entirely rod-saturated, with no rod saturation occurring in the indoor portion of the scene. In May, there are persistent fluctuations above this level, an example image from May 28th shows that around 10% of the scene is rod-saturated. In addition to the window region, parts of the desk are also rod-saturated. This is the result of increased directional sunlight in the May dataset, which leads to bright reflections as the angle of direct sunlight approaches the surface normals of objects within the scene. Further examples show that on three occasions in the May dataset, the rod-saturation percentage spikes at around 26%, each time at 07:00. In these cases, the desk and some parts of the south-facing wall become rod-saturated, further demonstrating the effects of directional sunlight entering the scene.

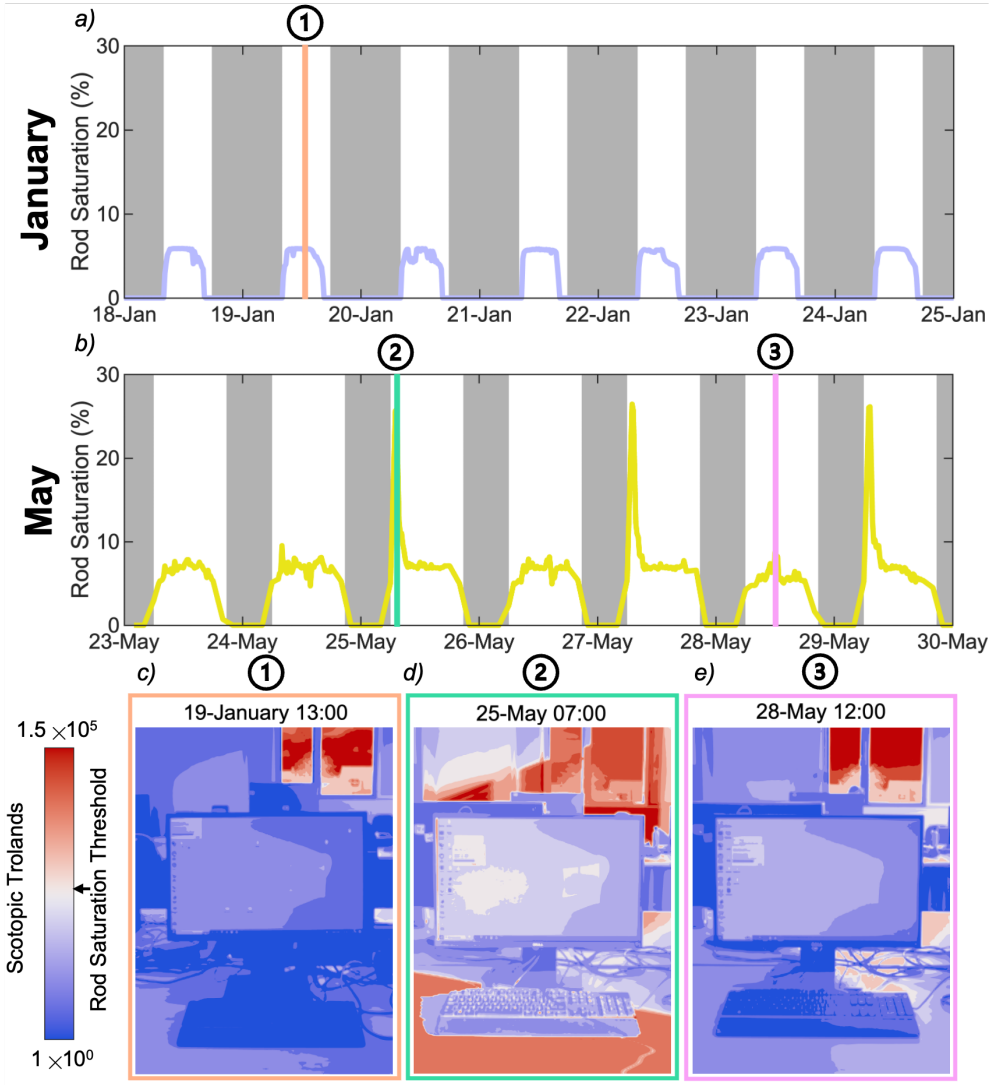

Figure 3.10: Timeseries plots showing the percentage of the scene which saturates the rods over time in January (top row) and May (second row). Shaded time periods show the time between dusk and dawn. The bottom row shows contour plots of example time points from each of the dataset. In these schematics, red regions show where rod saturation has occurred and blue regions show where the scotopic troland level is below the saturation threshold.

3.3.2 Weather Variation

Figure 3.11 shows chromaticity plots for daytime samples across each season, averaged across the whole image, and the indoor and window regions. In each case, the data appear split into two groupings. A K-means (K = 2) clustering was used to classify these groupings, and the data are colorised to reflect them. The differences in these groups are characterised by a shift in both L/(L+M) and S/(L+M), such that the 'orange' grouping has a lower S/(L+M) ratio and a higher L/(L+M) ratio, meaning they have less short wavelength light and more long wavelength light. Notably, data in the orange grouping, which make up the larger proportion of data across all conditions, have different relationships in January and May. In January, the data in this

group have a negative slope, while the data in May show no obvious slope. Whole Image data are significantly higher in S/(L+M) than the two region-based calculations as the screen, which is omitted in the region-based calculations, emitted predominantly blue light.

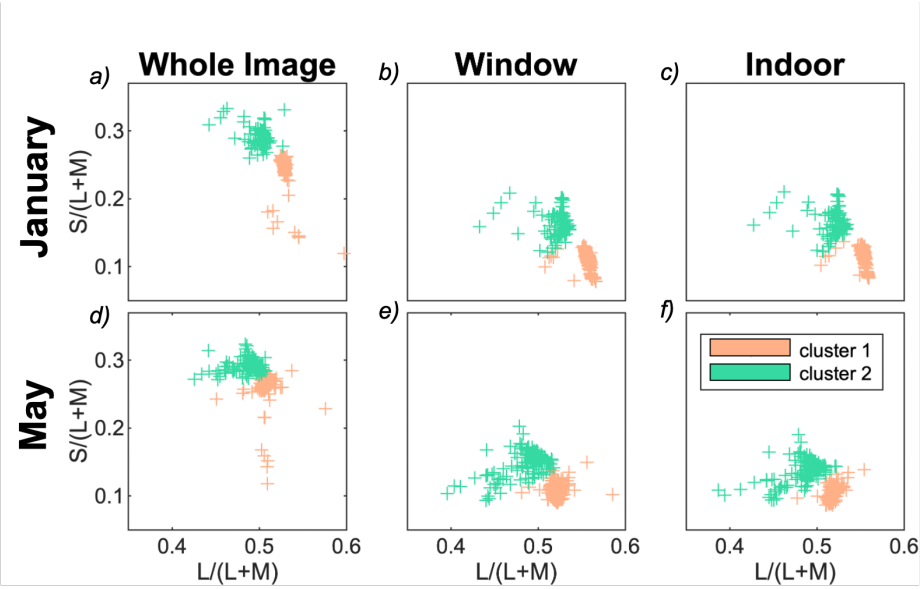

Figure 3.11: Chromaticity plots for daytime samples across January (top row) and May (bottom row), averaged across the whole image, the indoor region and the window region (columns). Data are coloured by a K-means clustering.

Figure 3.12 is a timeseries plot showing the classification of samples over time, according to the K-means clustering of L/(L+M) vs. S/(L+M) data presented in figure 3.11. Thus, daytime timepoints appear either orange or teal and nighttime timepoints appear grey. Each row corresponds to the clustering within a particular season-region combination. While almost all days in this time series contain both orange and teal times, the days can be categorised by their predominant grouping – January features five predominantly 'orange' days and two 'teal' days, while May features four 'orange' days and three 'teal' days. Notably, the teal grouping appears consistently at dawn and dusk on 'orange' days. The classifications between groups are largely consistent across regions.

To investigate what might underlie the chromatic change between the different days, we investigated historical weather data for Oxford during the capture periods. Weather data was accessed on timeanddate.com [251], where weather data is provided by CustomWeather, who report data from weather stations run by the World Meteorological Organisation (WMO). Nearby weather reports suggest that throughout both of our dataset capture periods, days classified into the orange group featured some direct sunlight throughout the day. However, the days largely identified as teal had no recorded period of direct sunlight, often being characterised by 'fog',

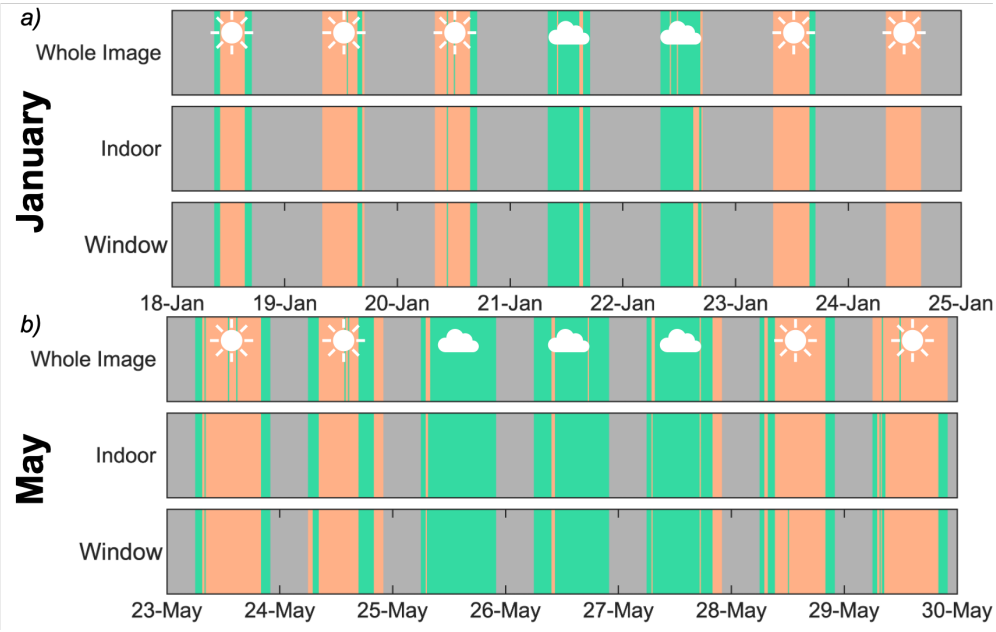

Figure 3.12: Timeseries plots where timepoints are colour-coded by the K-means cluster classification, based on the subplots in figure 3.11. Each row corresponds to a different season-region combination, with the top three rows corresponding to the Whole Image, Indoor and Window region classifications for the January dataset, while the bottom three rows correspond to the May dataset. The grey shading appears before sunrise and after sunset. White weather icons denote the weather grouping for each day.

or 'overcast'. This suggests that the earlier classification of days into 'orange' and 'teal' groups presented in figure 3.12 corresponds to 'sunny' and 'overcast' days, respectively.

Figure 3.13 shows Spearman's Rho statistics for correlations between luminance and different chromatic statistics within the indoor region. Here the data are split into 'sunny' and 'overcast' days based on the data in figure 3.12. 'Sunny' day data are shown in light colours, and 'overcast' day data are shown in dark colours. Throughout both datasets, and in both weather conditions, L/(L+M) has a positive relationship with luminance, such that pixels with higher intensity contain more long-wavelength light. However, both season and weather have different effects on the shorter-wavelength dependent statistics. In the January dataset, 'sunny' days show a persistent negative relationship between luminance and these statistics, while 'overcast' days show a relationship close to 0 in the mornings, which decreases throughout the day. In the May dataset, 'sunny' days also show less correlation than cloudy days, with S/(L+M) being more strongly negative than I/(L+M). 'Overcast' days in May show less variation than 'overcast' days in January, having a consistent positive relationship for both statistics.

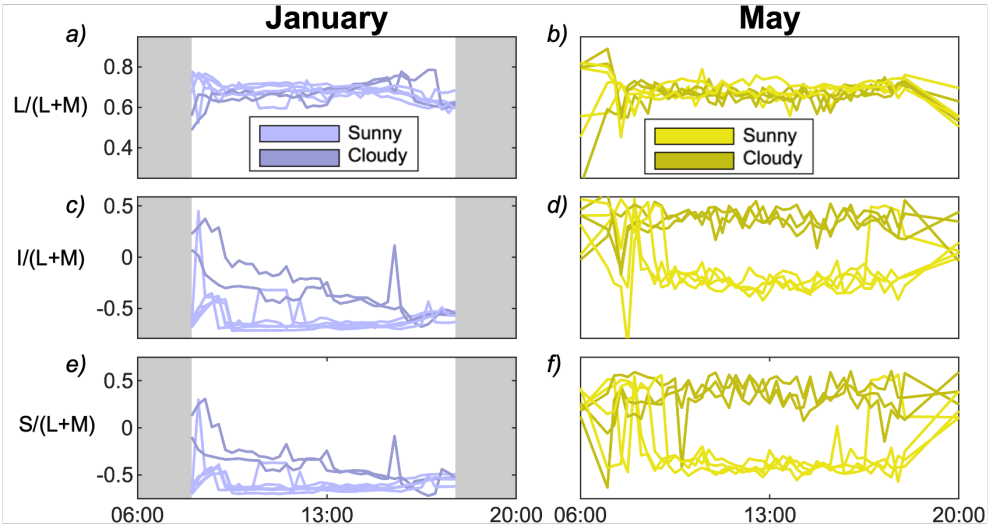

Figure 3.13: Spearman's Rho statistics for correlations between chromatic statistics and luminance for pixels within the indoor region, averaged across specific days within the 7-day capture period. Each subplot shows a combination of season and chromatic statistic, where the left column shows correlations in January and the right column shows correlations in May. The top row shows correlations between luminance and L/(L+M), the middle row shows correlations between luminance and I/(L+M) and the bottom row shows correlations between luminance and S/(L+M). Within each plot, the light-coloured lines show data for the 'sunny' days, and the dark-coloured lines show data for the 'overcast' days. The grey shading appears before sunrise and after sunset.

3.3.3 Directional Variation

The direct sunlight that enters the office appears to be a major source of chromatic changes but is highly time dependent as the position of the sun changes throughout the capture periods. To capture this variation, we have investigated luminance as a function of time for each pixel, fitting a cosine model for each luminance profile. This allows us to determine the size of changes in the magnitude and phase of diurnal luminance variation across parts of the scene. Figure 3.14 shows relative phase delay histograms for pixels in the window and indoor regions, corresponding to the distribution of phases from the cosine fits to each pixel. Data from the January dataset are shown in blue and data from the May dataset are shown in yellow, with means shown by the darkened bars. The data in this figure have been aligned such that the peak of the January window data is at 0. There are many more datapoints in the 'Indoor' subplot as a result of this region being markedly larger than the window section and thus featuring many more pixels. The data show different effects for the two regions. In the window region, there is a two-hour delay between the means of the January and May distributions, with the May distribution being later. On the other hand, for the indoor region, there is a 1-hour-and-40-minute delay between the January and May distributions, though here the May distribution is earlier. The data also show significantly more phase variation in the indoor region, and especially during May.

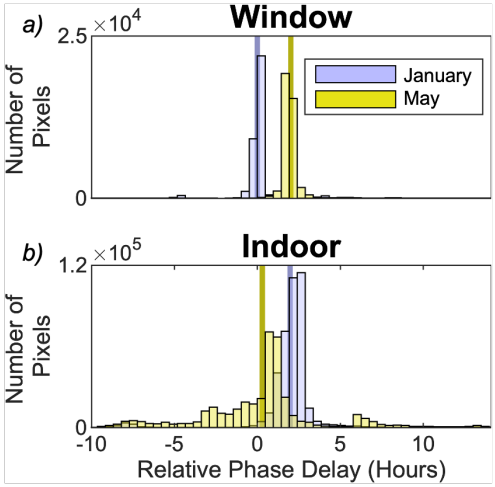

Figure 3.14: Histograms showing the relative phase delay in hours of cosine-fitted pixels in the window and indoor regions. January data are shown in blue and May data are shown in yellow. Data are aligned such that the peak in the January window data appears at 0. A negative phase delay with respect to this denotes a 'phase advance'. Means for each distribution are shown in the darkened bars.

In order to test if the 'greening' observed in the indoor region may be attributable to the East-facing windows (see floorplan in figure 3.1), we used a lightprobe setup, similar to that used in previous experiments [201,252]. By combining this setup with the Circadiometer, we can create a 'Circadiometric lightprobe'. This allowed us to capture light from a full 360-degree surround, including light entering the workspace from both the North-facing window that appeared in the January and May datasets, as well as the most-visible set of East-facing windows to the right of the camera. Using the same 15-minute interval capture scheme between 7:00 and 18:00 and 2-hour interval overnight we captured one day of lightprobe data on 28th June 2024.

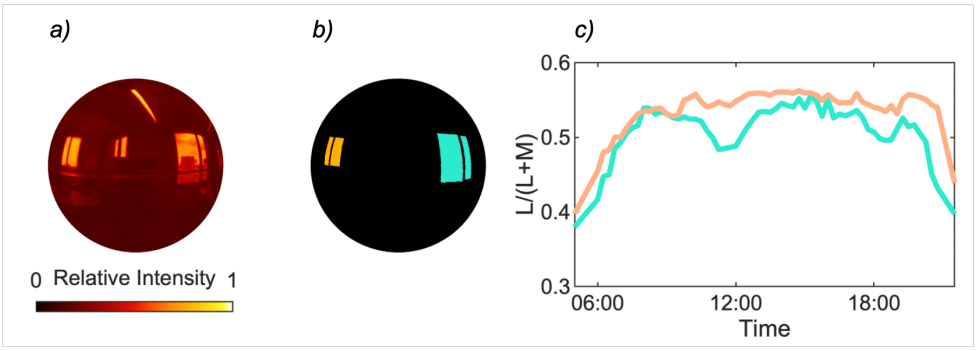

Figure 3.15: Lightprobe data. a) An example lightprobe, mapped onto a sphere, showing the North and East windows. b) North (orange) and East (teal) target windows are highlighted. c) Timeseries plot for the day of lightprobe data, showing the mean L/(L+M) statistics over time for the two target windows. The North window is shown in orange, and the East window is shown in teal.

Figure 3.15 shows an example lightprobe projected onto a sphere, and the North and large East windows highlighted in orange and teal, respectively. We compared the L/(L+M) statistics of each of these window regions over time. 3.15c shows changes in L/(L+M) over the course of

the day recorded using the lightprobe setup, with the North target window shown in orange and the East target window shown in teal. The graph shows that the large East window has a consistently lower L/(L+M) statistic throughout the day (thus being "greener"), and a matched-pairs Friedman test showed that this difference was significant ($\chi^2(63) = 104.02$, p=.001).

3.3.4 Additional Modelling

Figure 3.16 shows the sun elevation angle throughout the daytime period in each data capture sequence. There is a marked difference in the range of daytime sun elevations across different seasons, with the highest elevation angles in January being greatly reduced compared to those observed in May.

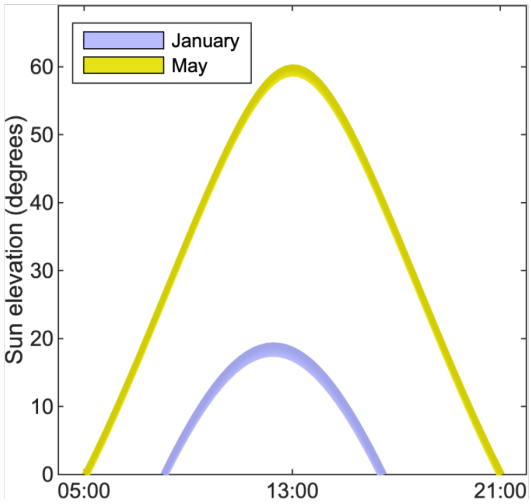

Figure 3.16: Sun elevation angles for January and May datasets at the geographic location of the test site (51.76° N, -1.27° W) across the daytime samples. January data are shown in blue and May data are shown in yellow.

Figure 3.17 shows the range of modelled sun spectra for the January and May datasets respectively, modelled in Smithson et al. (2014) [253]. The model considers the effect of Rayleigh scattering, aerosol scattering, Ozone Chappuis band absorption and sun elevation on the observed sun spectrum. The model does not account for molecular oxygen and water vapor absorption. Spectra are colored according to their respective sun elevation. The May dataset shows increased intensity, particularly in shorter wavelengths.

Figure 3.18 shows the chromaticity of the spectra plotted in 3.17 in a MacLeod-Boynton chromaticity diagram. Datapoints are colored according to their respective sun elevation. Black lines join datapoints from the same time of day (e.g., a datapoint from 9:00am in January

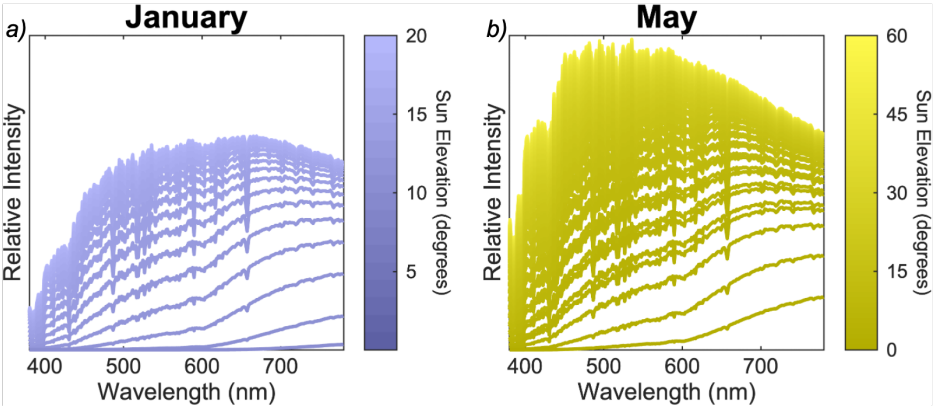

Figure 3.17: Spectral plots of modelled daylight spectra [33]. Left figure shows January data, and right figure shows May data. Spectra lines are colored according to the sun elevation, shown by the color bars to the right of each plot.

and 9:00am in May). Different day lengths mean these timepoints have different sunlight chromaticities. Chromaticities corresponding to the sun elevation during May dataset are greener and bluer than the chromaticities from January. 3.18b shows the size and direction of these differences, with each line coloured according to its time of day. The largest chromatic differences occur early and late in the day, with the smallest differences occurring at 12:15.

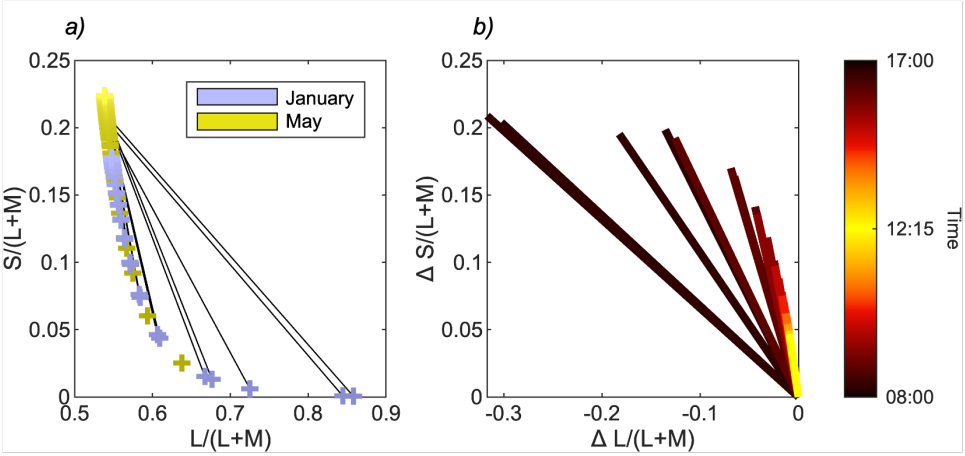

Figure 3.18: a) Chromaticity plot of modelled sunlight spectra across January and May datasets. January data are plotted in blue and May data are plotted in yellow. Datapoints are darkened based on their corresponding sun elevation. Black lines join datapoints taken from the same hour of day. b) Difference vectors connecting chromaticities from the same time of day in January and May. Lines are coloured according to their time of day.

We have additionally investigated environmental signals for the decorrelation between S-cone and melanopic light intensity. These two photoreceptor populations provide antagonistic signals to pupil control [186], and so their decorrelation in natural scenes is of physiological interest. Figure 3.19 shows the Spearman's Rho statistics for correlations between S-cone intensity and melanopic light intensity across regions and seasons, with darker and lighter lines representing

overcast and sunny day data, respectively. In all cases, S and I are highly correlated, and differences between regions, seasons and weather conditions are minimal. However, correlation statistics are notably reduced at sunrise in the window and indoor regions. This reflects times in which there are greater directional light signals due to the lowered sun elevation.

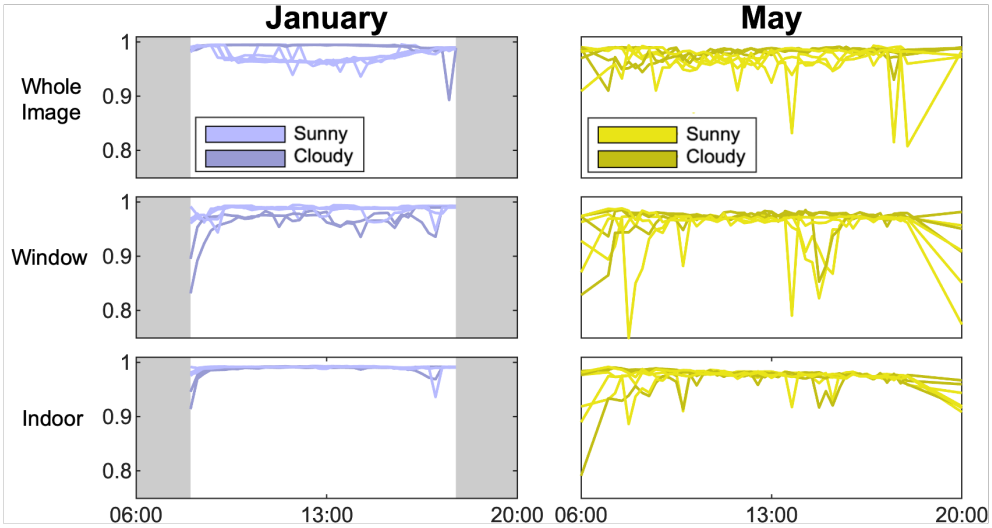

Figure 3.19: Spearman's Rho statistics for correlations between S-cone intensity and melanopic intensity across different regions and seasons, and in different weather conditions. Overcast days are shown in darker colours, and sunny days in lighter colours. The grey shading appears before sunrise and after sunset.

3.4 Discussion

The aim of this study was to collect a dataset of calibrated images taken from a typical indoor office environment across different seasons in order to assess whether some chromatic characteristics of natural environments carried over into an example mixed-illuminant indoor environment. Our data affirmed the influence of external natural scene statistics on the indoor space.

First, we showed lower L/(L+M) statistics in our May dataset compared to our January dataset (see figure 3.6), which corroborates Webster et al.'s finding [240] in a mixed-illuminant environment. Changes to this statistic in natural environments are associated with the 'greening' of foliage in wet seasons, however our PCA-derived segmentation that split the scene into the key 'indoor' and 'window' sections, allowed us to demonstrate that the lowering of L/(L+M) was due to changes in light observed in the indoor region rather than directly from the window (where foliage was directly visible), despite the fact that there were no changes in the chromaticity of the artificial illumination between the two datasets. That these differences are largely driven by the indoor region suggests that chromatic changes to the light entering the scene via the

windows can markedly alter the chromatic profile of light reflected from indoor surfaces to the extent that they can reproduce the seasonal changes that occur outside, even if the windows themselves are not directly visible. We observed seasonal changes in the opposite direction for the shorter-wavelength statistics I/(L+M) and S/(L+M). Previous analysis of natural scene statistics using Principal Component Analysis to derive opponency channels has identified S and I opponency to L and M to explain a significant amount of variance across natural scenes [254], and our analysis demonstrates that this variance can also be observed indoors through seasonal variation.

Secondly, we showed that different segments of the scene, associated with higher and lower luminances, had different chromatic properties, but that this effect depended on the season (see figure 3.8). For most of the May dataset, there was an increase in the L/(L+M) statistic in the higher luminance region. For the January dataset, the effect was much more varied and was most typically in the opposite direction. This suggests the regularities observed by Vincent et al. [242] in natural scenes can be carried over into mixed-illuminant indoor environments. Analysis provided in figure 3.9 also reveals that the bimodality observed in the January dataset was attributable to weather effects, with negative differences being observed on sunny days, and cloudy days reversing the direction of the difference.

Other than the greening of foliage, another significant way in which the outdoor environment changes across different seasons is the path the sun takes throughout the day. At the latitude of data collection, there are large differences in sun elevation in January and May, which in turn leads to significantly different chromaticities. Modelling of sun spectra at different sun elevation levels reveals that these changes in sun path could provide a major source for the decrease in L/(L+M) in Summer [253] (See figures 3.16-3.18). The change in sun elevation had further effects on the light in the office too, such as in bright regions in the indoor region as a result of intense directional sunlight. This is exemplified by the desk-containing pixel locations that behave more similarly to the window region in the May dataset PCA, but also by the group of high-luminance datapoints that remain in the indoor classification, shown in the right column of figure 3.6, all of which occur between 06:00 and 09:00, suggesting that the early morning sun from the East-facing windows produces these bright patches. Previous experiments exploring indoor spaces using only light from windows have noted the dependence of indoor spaces on sun elevation for access to this light [255], and our findings extend these results to our mixed illuminant environments.

Weather also had a significant effect on the chromaticity not only of the window region, but the indoor space too (See figures 3.11, 3.12 and 3.13). The whole image, window and indoor region classifications all demonstrated a significant shift in chromaticity when the weather was overcast compared to when it was sunny. In all region and season combinations, overcast days were bluer and less red. For data collected on the sunny days, a negative slope was observed in L/(L+M) vs S/(L+M) plots in January but not in May, similar to the distributions observed by Webster and colleagues, and other previous research showing blue-yellow bias in natural scenes [199]. Previous work has demonstrated how cloud coverage can affect the direction and strength of diffuse components of light [256], which in turn affects the light available to the indoor region.

We also investigated correlations between S-cone and melanopic light intensity across each region and season (See figure 3.19), as well as across weather conditions. While there were no clear weather or region effects, the strength of the correlations decreased at sunrise and sunset – particularly in May. Previous research has demonstrated increased variance of cone excitation ratios in natural scenes in the evening [257]. The increased variance in short-wavelength light that we observe is likely partially due to changes in the directionality of illumination in the scene. Specifically, when the sun is lower in the sky, illumination is highly directional. This increases local discrepancies in the illumination of surfaces, thus decorrelating photoreceptor signals. Notably, the correlations are otherwise largely invariant. The motion-activated artificial lights are activated at around 08:30 across both datasets, though there is no clear step change observed, and the decorrelations align more with the rising and setting of the sun. This dissociation between melanopic light and S-cone signals at specific times in the diurnal cycle may provide a useful physiological signal for dawn and dusk. Recent animal studies using bespoke cages with programmable lighting profiles have shown the ability of spectral changes to induce locomotion in mice [258].

We sought to specifically investigate light in the indoor region, as this represents the complex interplay of natural and artificial light, and surfaces inside the office (See figure figure 3.13). Correlations between luminance and different chromatic statistics within different seasons and weather conditions revealed further regularities that appear in the indoor region. Interestingly, correlations between luminance and L/(L+M) were relatively stable across all conditions, but there were large differences for the short-wavelength photoreceptor correlations. On sunny days in January, luminance correlated negatively with these statistics, suggesting that the brightest

pixels in the indoor region were more yellow. The strength of overcast days in January ended similarly negative to sunny days, but began with a strength closer to 0, decreasing throughout the day. Meanwhile, weather differences in May were more pronounced, with overcast days showing a persistent positive correlation, while sunny days continued to have a negative correlation. The differing effects of weather may be due to the differences in sun elevation (figure 3.16). Correlations in January may be more sensitive to the directional effect of the sun throughout the day as it stays lower in the sky, while high path of the sun is blocked by the ceiling for most of the daytime in May. Additionally, the increased variance in the mornings of the May dataset might reflect the increased influence of natural light on the indoor region before the artificial lights turn on.

The range of phases in which light increases and decreases across different surfaces in the indoor region, as a result of the differing surface normals and locations with respect to the North and East windows, demonstrates the complex chromatic interactions that occur between the natural outdoor lighting and the large, reflective indoor surfaces, specifically as a result of directional light like the sun (See figure 3.14). By capturing a 360-degree lightprobe of the office environment, we were able to record differences in the chromaticities of the illumination entering the office from the North windows and East windows, and found that the main East windows – those not originally visible to the Circadiometer in our January and May datasets – featured significantly lower L/(L+M) values than the North windows, which may provide an explanation as to how the L/(L+M) statistics of the indoor region decrease in the Summer (See figure 3.15). Chromatic changes in the natural environment are very directionally sensitive and are prominently influenced by the path the sun takes and the geometric properties of objects and surfaces from which the sun reflects. In addition to the sunlight directly entering the office, the spectra reflected by natural and built terrain around an office building change throughout the day and the year as the location of the sun and visibility of the sun due to clouds and foliage changes [259]. Because of this directional sensitivity, it is plausible that the lower L/(L+M) ratio light entering the indoor office from the East was a significant determinant in changing the chromatic properties of the indoor region, while the light entering the scene from the North window was less affected by the seasonal changes typically observed between January and May. This result emphasises the need for more mixed-illuminant datasets that quantify physiologically relevant light across the full light field using multi-directional setups.

A limitation of the data presented here is the extent to which it is generalisable to other

indoor spaces due to the specificity of indoor furniture, geometry, and window locations and directions. While we note the specificity of indoor spaces, and the myriad factors that affect their chromatic environments, it is significant that changes to outdoor seasonal characteristics can have such an effect on an indoor environment, and the factors that allow this to occur may be shared by many typical indoor spaces. This prompts further inquiry into different mixed-illuminant spaces, such as those with more limited access to natural light, single-direction window visibility, or different kinds of artificial lighting. This research also prompts behavioural studies, which might investigate the chromatic characteristics that lead to light-manipulating behaviours such as closing curtains or blinds, and the effects this has on adaptation states.

To summarise, indoor, mixed-illuminant scenes represent complex interactions between outdoor (natural) and indoor (artificial) chromatic environments. Despite this, mixed-illuminant environments may undergo many of the same seasonal changes that naturalistic scenes experience, and this may facilitate seasonal adaptational changes that are understood to be the result of regularities in the seasonal natural environment.

3.5 Acknowledgements

The work presented in this chapter on the temporal, spatial, and spectral characteristics of indoor lighting is part of a collaborative project led by Samuel Ponting, with Professor Hannah Smithson, Professor Manuel Spitschan and Ruth Kelly Waskett. I would additionally like to thank Robert Fosbury for providing code and data for the calculation of elevation changes in sun spectra.

The work contained in this chapter in different form has been published as:

Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson H E. (2025). "Quantifying CIE alpha-opic signals in the indoor built environment." J Opt Soc Am A Opt; 42(5):B379-B390. Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson, H E. (2024). "Quantifying CIE alpha-opic signals in the indoor built environment." International Colour Vision Society Meeting, Ljubljana, Slovenia (July, 2024). <Oral Presentation> * W. D. Wright Award Ponting, S; Kelly Waskett, R; Spitschan, M; Smithson, H E. (2023). "Quantification of Melanopic Light in an Office Environment." Applied Vision Association Christmas Meeting, London, UK (December, 2023). <Poster>

Chapter 4

Measurement of the directional structure of light in the built environment: single-image circadiometric lightprobes

Contents

4.1 Introduction 98

4.1.1 The benefits of quantifying environmental illumination 99

4.1.2 Existing methods for the quantification of environmental illumination . 100

4.1.3 The single-image circadiometric lightprobe imaging system (SICLIS) . . 101

4.2 Methods 101

4.2.1 Approach 101

4.2.2 Single-image circadian lightprobe imaging system (SICLIS) 102

4.2.3 Image processing 103

4.2.4 Model-based validation 105

4.2.5 Data analysis 106

4.2.6 Experiment 1 109

4.2.7 Experiment 2 111

4.3 Results 112

4.3.1 Experiment 1 112

4.3.2 Experiment 2 128

4.4 Discussion 132

4.4.1 The single-image circadian lightprobe imaging system 133

4.4.2 Experiment 1 134

4.4.3 Experiment 2 135

4.5 Acknowledgements 136

4.1 Introduction

The lighting environments in which we find ourselves throughout our lifetimes have profound visual and non-visual effects on us [129,169,232], and increasingly over the course of human history these lighting environments are indoor ones. Recent surveys estimate that the average American spends around 90% of their time indoors [244], and so it is crucial to consider indoor lighting environments when attempting to quantify the physiological effects that our light exposure has on us. Indoor lighting environments are often complex and dynamic, especially when they feature windows that let in natural daylight from outside. Sun position is a strong driver of changes in daylight illumination and undergoes significant changes throughout the day – known as diurnal motion, owing to the earth's rotation – but also throughout the year, known as annual solar motion, owing to the elliptical orbit of the earth. These complex motion patterns affect the intensity, directionality and spectra of daylight across time [253]. For example, the sun's low elevation at sunrise and sunset leads to increased Rayleigh scattering, which increases the longer-wavelength light available in daylight spectra at these times. Spectral changes due to time of day and time of year are modelled in the CIE daylight model [67]. Sun position, along with the orientation of a building's windows, dictate the natural light entering an indoor scene and can affect its spatial structure by determining, for example, whether sunlight shines directly into a room or merely provides diffuse skylight illumination [260]. Furthermore, the structure and chromatic properties of the space inside a building has a significant impact on the visual scene. The light sensed by a user of the space is the result of the complex interplay between natural and artificial illuminations, as well as the spectral reflectance of surfaces in their local environment [261].

4.1.1 The benefits of quantifying environmental illumination

Characterising the structure of these complex light field interactions is useful as they can have significant implications for the physiological effects of light exposure throughout the day. Light exposure has wide-reaching consequences for visual perception and can thus modulate performance in visual tasks that rely on colour discrimination [262]. In natural environments, spectral content can often vary differently across directional structures, for example, at dawn and dusk, when increased directional intensity of long-wavelength light (red sunrise or sunset), is paired with an increase in short wavelength light (blue-enriched) from the sky dome, as discussed in Chapter 1. Directional spectral effects may be exacerbated in environments with chromatically biased surfaces and objects that introduce wavelength dependent biases in light reflected from them. The complex picture of mixed-illuminant environments painted above, as well as the fast-adapting nature of cone cells demonstrates how these heuristics can change throughout the day and year according to the oscillating changes to the environment [199, 206, 207, 256, 259], as well as across location within an office space [245]. The spatial structure of light can also have implications for visual comfort [263]. Highly directional illumination, such as that caused by direct sunlight, can modulate visual contrast throughout a scene, interfere with illuminant estimation and lead to discomfort. In cases where the source of highly directional illumination is directly viewed, it can cause glare and drastically reduce the visual comfort of the scene [264]. Lighting designers in industry are aware of these considerations, and industry standards such as the WELL Building Standard provide guidelines for minimum illuminance and contrast levels, colorimetric specifications such as colour reproduction index (CRI), and structural considerations such as illuminant diffuseness and solar glare control [265].

Perhaps equally as significant are the non-visual effects of light exposure, mediated by the photopigment melanopsin found in intrinsically photosensitive retinal ganglion cells (ipRGCs). Light in the spectral range to which these cells are sensitive has wide-reaching effects on circadian phase setting, alertness and subjective wellbeing [170]. ipRGCs adapt much more slowly than rod and cone cells and are thus suited to measuring ambient absolute light levels throughout the day [266]. The WELL building standard encourages lighting designers to reach minimum levels of melanopic illumination; however, these standards are not yet widely met in industry.

4.1.2 Existing methods for the quantification of environmental illumination

There is a large literature of previous work that has sought to quantify the spatial and temporal characteristics of typical light environments. These range from natural environments in which the human eye is thought to have developed [200, 240, 242], to recent studies that quantify natural lighting [201, 202, 237, 252, 257], to mixed illuminant environments where artificial and natural illuminants interact [206, 207, 255]. Interest in the directional structure of light fields has led to a number of imaging solutions that allow for the omnidirectional characterisation of light from every angle at a particular location in space. Two prominent examples include environmental illumination maps [201, 252] and the plenopter [259, 267]. Environmental illumination maps, also known as lightprobes, use optical elements such as a camera and mirror sphere or fish-eye lens to create a panoramic image of an environment, such that light incident from every angle is recorded. These methods require an operator to rotate the camera within the scene to capture 4-6 images to derive a full 4π steradian panoramic image. Once processed, these lightprobes can provide a high-resolution image that captures the spatial structure of a scene in high detail. A range of imaging devices can be used to create an environmental illumination map, with previous examples using RGB, hyperspectral, and physiologically-relevant cameras, often allowing for high dynamic range (HDR) capabilities. The plenopter offers another method for capturing the directional structure of a scene. In place of a camera, the plenopter uses a number of wide-aperture photocells in a regular formation to create a low-resolution radiometric measurement of a whole scene in one capture period. When more sensors are used, the device can collect higher angular frequency information. These two solutions offer contrasting benefits and caveats: environmental illumination maps offer high-spatial-resolution detail and can thus be used to target specific locations within a scene and examine high spatial frequency information. However, their capture requires multiple capture periods to complete. A full hyperspectral lightprobe with HDR using six recording angles may take up to half an hour to record and requires an operator to move the camera into each position. On the other hand, the plenopter requires only a single capture period, and so recordings can be automated to create a timelapse of directional information to capture the temporal dynamics of a scene. However, its low spatial frequency of imaging means that it is not capable of capturing high-angular frequency information. Plenopter methods also often feature much lower spectral resolution than hyperspectral methods.

4.1.3 The single-image circadiometric lightprobe imaging system (SICLIS)

Here, we present the single-image circadian lightprobe imaging system (SICLIS), which seeks to mitigate the limitations of environmental illumination maps and the plenopter. It is a capture setup like that of environmental illumination maps, utilising a mirror sphere, but additional mirrors are introduced to allow the capturing of a full 4π-steradian lightprobe in a single acquisition. Thus, it combines the spatial resolution of environmental illumination maps with the unsupervised capture and temporal resolution capabilities of the plenopter. Here, we use the single-image circadian lightprobe imaging system to characterise the directional structure of two example mixed-illuminant office spaces in physiological terms using the Westboro Photonics Circadiometer – a five-channel imager which is calibrated to align with the spectral channels of the human visual system.

4.2 Methods

4.2.1 Approach

To characterise directional illumination in built environments we devised a portable imaging system capable of capturing the 4π steradian light field. The unique multi-mirror geometry of the setup allows for lightprobe generation with a single initial image capture, which increases the speed of lightprobe measurement and allows for unsupervised lightprobe timelapses, which can feature high spatial and temporal resolutions. A sensor with physiologically-relevant filters was also used, meaning each lightprobe gives a spatial map of responses for the five photoreceptors found in the human eye. Reconstruction of lightprobes from the viewing geometry was validated using a model generated using the *Mitsuba 3* physically-based ray-tracing engine, which showed that lightprobes generated using this setup closely match the spatial structure of the lighting environments they capture. We placed the imaging system on a custom-built platform with four castor wheels to allow for the easy manoeuvring of the system.

Two experiments were completed using the SICLIS. The first experiment sought to demonstrate the portability of the imaging system by measuring lightprobes across 11 desk positions in an open-plan office area. Lightprobes were taken across a 12-month period creating a year-long dataset containing spatial and seasonal variation. The second experiment emphasised the unsupervised capabilities of the system. Recordings were taken from a second open-plan office area for a total of 16 days. The system alternated between two locations while taking lightprobes

at a high temporal frequency.

4.2.2 Single-image circadian lightprobe imaging system (SICLIS)

An image of the single-image circadian lightprobe imaging system is shown in figure 4.1. It consists of one tripod supporting a Westboro Photonics WP690E-Circ Imaging Circadiometer, and a second tripod supporting a 76.2 mm diameter AISI 52100 chrome steel ball with an etched equator. The tripod supporting the mirror sphere also supports two 100 x 100 mm Enh/Ali planar mirrors from Knight Optical, each held by a Manfrotto 143N Magic Arm. The two tripods are positioned such that the mirror sphere is ∼650 mm away from the Circadiometer lens, and the planar mirrors are an additional ∼100 mm behind the focal plane of the sphere, angled such that the full sphere is visible to the Circadiometer, thus the sphere was visible to the imager from three angles – once directly and once from each of the planar mirrors. The imaging apparatus is set at seated head height for a hypothetical observer seated on an office chair at a desk workstation.

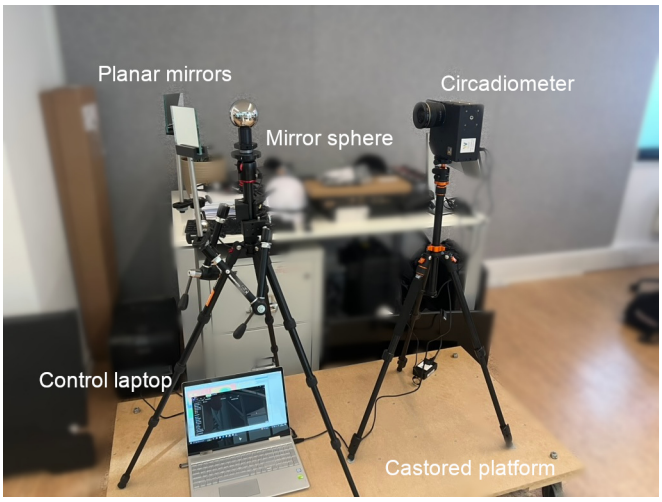

Figure 4.1: An image of the circadiometric lightprobe setup. The image shows the circadiometer on one tripod pointed towards a second tripod holding a mirror sphere and two planar mirrors. Both tripods are set up atop a castored platform. A control laptop is also positioned on the platform

During recording, the Circadiometer is controlled by a laptop using the Westboro Photonics Photometrica software, which automates image capture timing and camera control. For each lightprobe image, five 3388 x 2712-pixel maps are generated, corresponding to five spectral filters inside the camera. These filters conform to the CIE S026/E:2018 technical report for the definition of α-opic spectral sensitivities of the five human photoreceptor types (S-, M- and L-cones, rhodopsin and melanopsin) [190]. Each pixel map is the result of a stack of exposures to create a high-dynamic range (HDR) image. Captures from a range of exposure times are

obtained until either there are no over- or under-exposed pixels, or the minimal exposure time is reached. The values in each HDR image are scaled to represent α-opic radiance in W/sr/cm^2.

4.2.3 Image processing

Figure 4.2 shows the image processing pipeline. The pipeline is designed to turn a single channel image into a 4π steradian lightprobe.

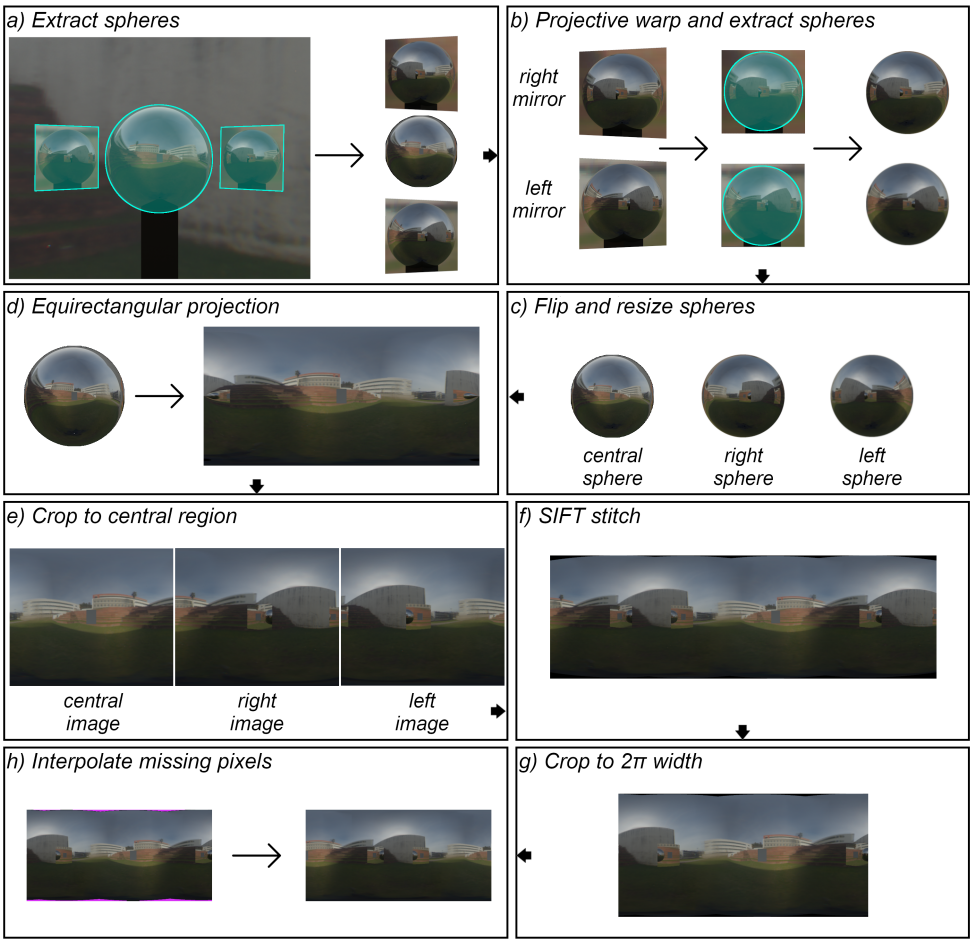

Figure 4.2: The steps of the lightprobe image processing pipeline. a) Images of the mirrors are manually cropped using a binary mask. b) Projective warp is applied to the images of the planar mirrors, aligning them with the plane of the camera. Sphere images viewed through the planar images are then manually cropped using a binary mask. c) Planar-imaged spheres are horizontally flipped and the three spheres are resized to match their dimensions. d) An equirectangular projection is applied to the sphere images. e) The edges of the projected images are cropped to reduce warping in the images. f) SIFT stitching is applied to join the three images. g) The image is cropped to 2π, removing repeated content. h) Missing points from the top and bottom of the image are filled in using horizontal interpolation

First, (fig. 4.2a) the directly viewed sphere region $I_{center}(x, y, \alpha)$, along with the left and right mirror images, $I_{left\ mirror}(x, y, \alpha)$ and $I_{right\ mirror}(x, y, \alpha)$, are manually extracted from the planar image $I_{whole}(x, y, \alpha)$ using binary masks $M_{\{center,\ left\ mirror,\ right\ mirror\}}(x, y)$. Equations are shown in 4.1, 4.2 and 4.3.

$$I_{center}(x, y, \alpha) = I_{whole}(x, y, \alpha) \cdot M_{center}(x, y) \tag{4.1}$$

$$I_{left\,mirror}(x, y, \alpha) = I_{whole}(x, y, \alpha) \cdot M_{left\ mirror}(x, y) \tag{4.2}$$

$$I_{right\,mirror}(x, y, \alpha) = I_{whole}(x, y, \alpha) \cdot M_{right\ mirror}(x, y) \tag{4.3}$$

Second, (fig. 4.2b) the images of the left and right mirrors are warped using projective transformations, $T_{correct\ left}$ and $T_{correct\ right}$, that align the faces of the planar mirrors with the optical plane of the camera. The images of the sphere viewed via the planar mirrors, $I_{left}(x, y, \alpha)$ and $I_{right}(x, y, \alpha)$, are extracted using binary masks $M_{left}(x, y)$ and $M_{right}(x, y)$. Equations are shown in 4.4 - 4.7.

$$I_{corrected\ left}(x, y, \alpha) = I_{left\ mirror}(x, y, \alpha) \cdot T_{correct\ left} \tag{4.4}$$

$$I_{corrected\ right}(x, y, \alpha) = I_{right\ mirror}(x, y, \alpha) \cdot T_{correct\ right} \tag{4.5}$$

$$I_{left}(x, y, \alpha) = I_{corrected\ left}(x, y, \alpha) \cdot M_{left}(x, y) \tag{4.6}$$

$$I_{right}(x, y, \alpha) = I_{corrected\ right}(x, y, \alpha) \cdot M_{right}(x, y) \tag{4.7}$$

Third, (fig. 4.2c) the extracted images of the sphere viewed through the aligned mirrors are horizontally flipped to account for the mirror-flipped view of the sphere. The three spheres are then resized to align their spatial scales. Equations are shown in 4.8 and 4.9.

$$I_{left}(x, y, \alpha) = I_{left}(-x, y, \alpha) \tag{4.8}$$

$$I_{right}(x, y, \alpha) = I_{right}(-x, y, \alpha) \tag{4.9}$$

Fourth, (fig. 4.2d) an equirectangular projection was used to map the pixels from the sphere images onto a rectangular image. This is a standard transformation for globe-to-map projections. Sphere yaw (azimuth) and pitch (elevation) angles are mapped onto horizontal and vertical pixel locations.

$$I_{\{center,\ left,\ right\}}(\lambda, \varphi, \alpha) = Equirectangular\ projection(I_{\{center,\ left,\ right\}}(x, y, \alpha)) \tag{4.10}$$

Where *Equirectangular projection* is defined as:

$$\lambda = \frac{x}{R} + \lambda_0 \tag{4.11}$$

$$\varphi = \frac{y}{R} + \varphi_0 \tag{4.12}$$

Where R is the radius of the mirror sphere, and λ_0 and φ_0 are the central parallel and meridian of the map. λ is the yaw (azimuth) and φ is the pitch (elevation) of the resulting image.

Fifth, (fig. 4.2e) the rectangular images are horizontally cropped. The left and right sides of each equirectangular projection are distorted the most by the transformation and are projected with less distortion in one of the other two sphere images. To minimise distortions, the most distorted parts of each image are removed.

Sixth, (fig. 4.2f) the three rectangular images are stitched together using the scale-invariant feature transform (SIFT) [268] automatic feature detection algorithm to produce a single image. Stitching is based on a pseudo-luminance channel based on a weighted combination of the M and L channel, which has a gamma of 1/3.5 applied in order to amplify dimmer features in the images. This provided more accurate and consistent matches. This transformation is then applied to each α-opic map.

Seventh, (fig. 4.2g) the image is horizontally cropped so that the vertical edges of the image represent a single 2π rotation.

Eighth, (fig. 4.2h) missing pixels in the top and bottom rows of pixels, which can appear during stitching, are filled using horizontal interpolation.

4.2.4 Model-based validation

To validate that a multi-mirror, single-image lightprobe setup can reliably recover the directional structure of an environment, a scene model was created using the Mitsuba 3 physically-based renderer and controlled using Blender. This scene consisted of a model of the single-image lightprobe setup, including a camera, modelled using the parameters of the Westboro Photonics Circadiometer, a mirror sphere and two additional planar mirrors. All mirrors were given a glossy bi-directional scattering distribution function (BSDF) with roughness and anisotropy values of 0. The scene was illuminated by a world environment lighting background. We used 10 lightprobes collected in a previous study [201] to act as these lighting textures. Blender applies images to the environment light map by applying an equirectangular projection, reversing the process that

was used to generate the lightprobe. A rendering of the setup viewed from above is shown in figure 4.3a, and the image taken by the modelled camera is shown in figure 4.3b.

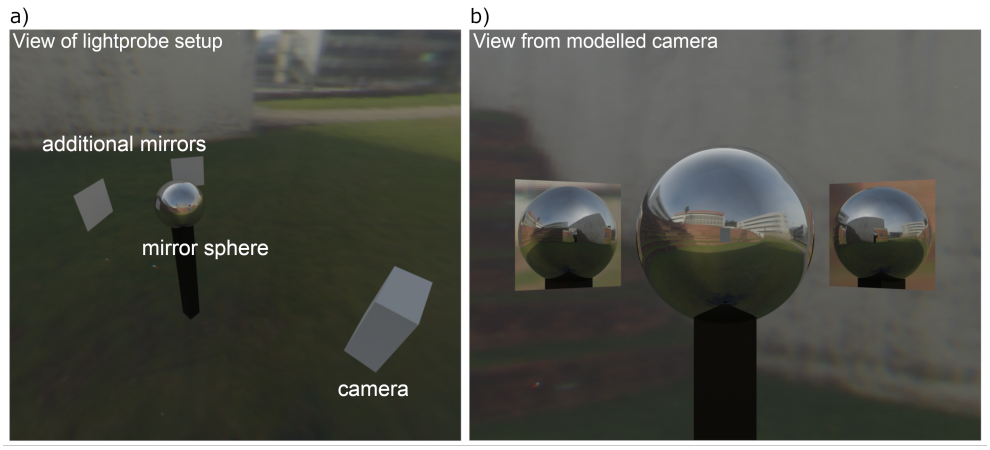

Figure 4.3: Blender images of the modelled lightprobe setup. a) A view of the whole scene model. Image shows the model camera pointed towards a spherical mirror and two planar mirrors. The background is an environmental illumination map. b) A point-of-view image from the modelled camera, showing the three mirrors.

In order to evaluate the performance of the modelled setup in reproducing the 10 original lighting textures, we rendered images of the modelled setup using *Mitsuba 3* and generated lightprobes using the method described above. We then compared the spatial structure of the resultant lightprobes with the objective environment lighting maps using the structural similarity index measure (SSIM). SSIM is an image quality metric that compares the local luminance, contrast and structure of two images [269, 270].

Figure 4.4 shows the results of the modelling analysis. Figure 4.4a shows the mean SSIM values for each of the 10 sample lightprobes. All structural similarity metrics were high, with the average SSIM of 0.91 and a standard deviation of 0.026, showing a consistent fit between the reference lightprobes and the reproductions. Figure 4.4b shows an example reference lightprobe, and figure 4.4c shows the reproduction using our method. One difficulty with a single capture lightprobe method is that since it lacks the camera movement of other previous methods the space taken up by the camera cannot be filled by pixels from alternate angle images. The introduction of two additional mirrors further adds debris into the image, which accounts for SSIM values < 1.

4.2.5 Data analysis

In developing this single-image lightprobe approach, we hope to characterise the directional and spatial structure of lighting environments. For planar images, a Fourier transform can be used to calculate the intensity and phase of separated angular frequency harmonics. For spherical

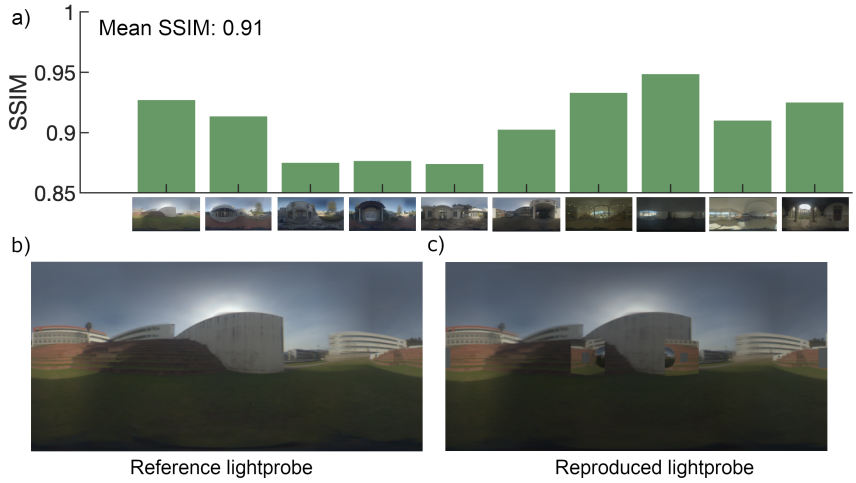

Figure 4.4: Structural similarity image metric statistics for Blender modelled data. a) SSIM values for each of the 10 analysed lightprobes. b) An example reference lightprobe. c) Single-image reproduction of the example reference lightprobe

images such as our lightprobes, spherical harmonic decomposition is used. Figure 4.5 shows an example schematic. Spherical harmonic decomposition quantifies the intensity of information in a spherical image at different angular frequencies, creating a Fourier series of circular functions with increasing angular frequency. For each frequency, known as the degree, a set of basis functions are generated as the group of orthogonal rotations in three dimensions of positive and negative poles in that frequency. Example basis functions for degrees $l = 0,\ 1$ and 2 are shown in 4.5b. Harmonics of higher degrees are described using more basis functions, as there are more possible orthogonal configurations of higher angular frequency patterns. In each case, a particular basis function, known as the order, is used to determine the intensity of angular frequency information along a particular orientation of the basis function. By summing the intensity across order, the intensity at each angular frequency information is calculable. This, we refer to as the *amplitude*. Additionally, the relative intensity across each orthogonal axis allows for the calculation of the primary vector for each angular frequency.

Spherical harmonic decomposition has been used previously to quantify lighting environments [201,271,272], with advanced metrics being developed to interpret spherical harmonic information in industrially-relevant ways. The zeroth degree harmonic, which features only one parameter, describes the intensity of light across the whole light field. It is widely termed as the 'light scalar' or 'density of light' and corresponds to ambient light throughout the scene. The first-degree harmonic features three parameters and identifies the strongest directional component in the scene and is thus called the 'light vector'. As such, the basis functions for this degree have one positive and one negative pole. The ratio of amplitude of these first two harmonics is often used

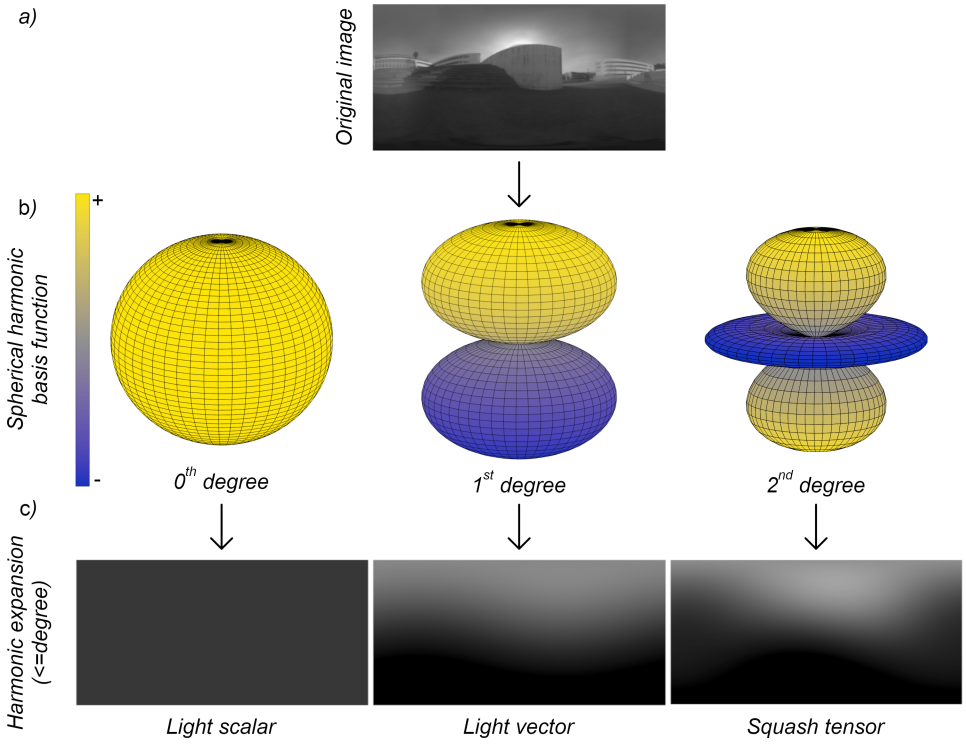

Figure 4.5: Spherical harmonic decomposition example. a) An example lightprobe image. b) Spherical harmonic basis functions for degrees $0, 1$ and 2. Yellow regions have positive values, and blue regions have negative values. c) Decomposed image for degrees less than or equal to degrees $0, 1$ and 2.

to quantify the diffuseness of a scene. The second-degree harmonic is often termed the 'squash tensor' and is again more complex. It features five parameters and defines a light field where poles share the same polarity. When considering a non-negative value such as light, this can take two formations. First, the 'light clamp' formation is observed when a scene has two light sources at opposing poles with respect to the measurement location. Second, the 'light ring' sees the light distributed evenly in a torus formation that extends across two of the three axes. These two formations are shown to be 'opposites' as the negative space around the torus is filled by the two opposing lights in the light clamp, and vice versa. Higher degree angular frequencies are not often analysed as their individual contribution to lighting scenes is diminished and their physical meanings are less immediate. However, their contributions can be summed together to quantify the strength of high frequency information in the scene, which may correspond to dappled or textured light [273]. Here, we define texture as the sum of the amplitude of spherical harmonics 3-30. Using the high-resolution Westboro Photonics Circadiometer and our single-image lightprobe system, we can measure the intensity and alignment of a large number of spherical harmonic components up to and beyond the squash tensor, additionally considering the way different α-opic maps are distributed across time. We also calculate the diffuseness metric

D_{Xia} [272]. D_{Xia} is calculated as the ratio of the amplitude of the first degree harmonic divided by the amplitude of the zeroth degree harmonic. This formulation is shown in equation 4.13, where L_0 corresponds to the amplitude of the zeroth degree harmonic, and L_1 corresponds to the amplitude of the first degree harmonic. This is then scaled and normalised between 0 and 1.

$$D_{Xia(Normalised)} = 1 - \frac{L_0}{L_1}/\sqrt{3} \tag{4.13}$$

As an additional analytic technique, we considered the dominance of photopic and melanopic light in the scene. A number of experiments quantifying physiologically-relevant light in built environments define a *photopic-melanopic ratio* to assess the distribution of this dominance throughout scenes [206, 274–276]. Here, we calculate photopic-melanopic dominance by finding the difference between the two receptor maps, and mapping this difference between 0 and 1, with the "neutral point" at 0.5, meaning that the amounts of photopic and melanopic light in a given region of the lightprobe is equal. We then performed the spherical harmonic decomposition described above on the photopic-melanopic dominance data. This requires a slight re-framing of the interpretation of the metrics outlined above as the decomposition is performed on a "difference" metric rather than scalar values. The zeroth degree harmonic of the photopic-melanopic data corresponds to the overall dominance in the scene. So, higher values correspond to luminance-dominated scenes and low values correspond to melanopic-light-dominated scenes. The $1^{st} - 4^{th}$ harmonic degrees correspond to the amplitude of variance in photopic-melanopic dominance at difference angular frequencies. The direction of the 1st degree harmonic corresponds to the direction of positive variance such that it is biased towards luminance-dominant areas of the scene, and biased away from melanopic-light-dominated areas. Diffuseness, being a scaled ratio of the zeroth and first degree harmonics, shows the comparison of the dominance of a spectral band across the overall scene and a hemifield.

4.2.6 Experiment 1

The first experiment sought to use the SICLIS to characterise the directional distribution of light in an example office environment, and map how position in the office, time of day and time of year change this distribution. The imaging system was set up in an open-plan office area on the second floor of New Radcliffe House in Oxford, UK (51.76° N, -1.27° W). A birds-eye schematic of the office space is shown in figure 4.6. The office features prominent North- and

East-facing windows and is otherwise artificially lit by motion-activated fluorescent pendant lights which remained on for the duration of data collection. The office space features two large desk groupings with several computer workstations on each. The imaging system was used to take measurements in 11 locations around the two desk groupings, corresponding to the location of workstations. Data collection took place at four equally spaced periods throughout the year in order to capture seasonal effects. In each instance, five consecutive days (Monday-Friday) of data were recorded. Table 4.1 shows the date ranges that were used for data collection. On each of the five days, recordings were taken in all 11 locations at 09:00, 12:00 and 17:00. The recording of 11 lightprobes took around 35 minutes in each instance and the order in which the locations were measured from was randomised. During data collection, the office was used freely by lab members. Ethical approval for this was granted by the Oxford Medical Science Interdisciplinary Research Ethics Committee with approval reference R81269/001. A total of 660 images were captured and processed into lightprobes.

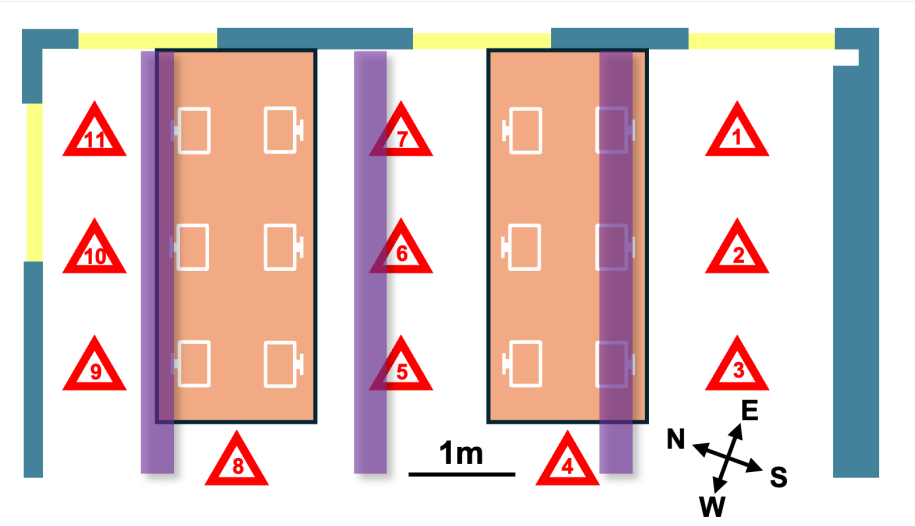

Figure 4.6: Schematic view of the office in which experiment 1 was carried out. Black bar shows 1 metre of physical space. Red triangles show lightprobe recording locations. Orange boxes show desks, with computer icons showing the location of workstations. Blue bars are walls. Yellow bars are windows. Purple bars are artificial pendent lights.

Table 4.1: Data collection dates for experiment 1

Season label	Data collection dates
Summer	12th–16th August
Autumn	19th–23rd November
Winter	3rd–7th March
Spring	10th–14th June

4.2.7 Experiment 2

The second experiment sought to take advantage of the unsupervised capture capabilities of the single-image circadian lightprobe imaging system to create a high temporal resolution dataset of lightprobe recordings. The imaging system was set up in a second open-plan office area on the ground floor of the Hoare Lea offices in Oxford, UK (51.758° N, -1.267° W). The building features a mezzanine floor, under which most of the users of the office space work, as well as prominent skylights in the double-height roof. A schematic of the office space is shown in figure 4.7. Data were collected from two locations in the office. Location one is under the mezzanine floor, in a lighting environment that matches that experienced by the users of the office. Location two is away from the mezzanine floor, with a direct line of sight to the skylights. Data collection took place over the course of 17 days between March 31st and April 16th. Images were collected every 15 minutes between 05:00 and 22:00 and every two hours between 22:00 and 05:00. During data collection the imaging system was positioned at location one and recorded from 12:15 on one day until 12:00 the next day. Subsequently, between the 12:00 and 12:15 recordings, the imaging system was moved from location one to location two, where measurements were then made until 12:00 on the day after. Data collection continued as such, with the location of the imager switching locations between 12:00 and 12:15. Over the weekends, the building was not accessible and thus the imager was not moved between Fridays and Mondays. During data collection, the office was freely used by members of the office, with ethical approval being granted by Oxford's Medical Sciences Interdivisional Research Ethics Committee under reference R81269/002. Table 4.2 shows the location the imaging system was positioned in from 12:15 each day over the course of data collection. A total of 1022 images were captured and processed to create lightprobes.

Table 4.2: Data collection locations for experiment 2

Day (from 12:15)	Week 1	Week 2	Week 3
Monday	Mezzanine	Skylight	Mezzanine
Tuesday	Skylight	Mezzanine	Skylight
Wednesday	Mezzanine	Skylight	Mezzanine
Thursday	Skylight	Mezzanine	–
Friday	Mezzanine	Skylight	–
Saturday	Mezzanine	Skylight	–
Sunday	Mezzanine	Skylight	–

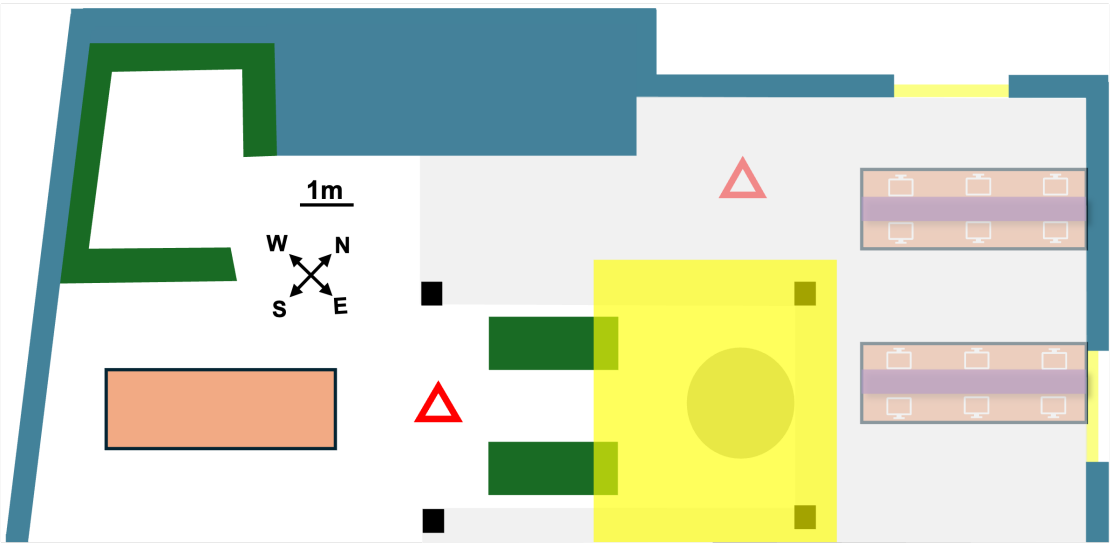

Figure 4.7: Schematic view of the office in which experiment 2 was carried out. Black bar shows 1 metre of physical space. Red triangles correspond to lightprobe recording locations. Grey region shows the mezzanine floor. Yellow regions show windows and skylights. Blue regions show walls. Green regions show sofas. Orange regions show desks. Purple regions show artificial pendant lights. Black boxes show pillars.

4.3 Results

4.3.1 Experiment 1

Lightprobes

Figure 4.8 shows melanopic light maps from the lightprobes taken from each of the 11 locations used in experiment 1. The lightprobes are arranged per the bird's-eye view shown in figure 4.6, with position 1 in the top right and position 11 in the top left. For each location, the unwrapped lightprobe is shown, centred on north, as shown by the teal vertical line. The two spherical images above each lightprobe show the wrapped view for each location while facing north (left) and south (right). The grey vertical line shows south.

Spherical harmonics - luminance

Figure 4.9 shows heatmaps for the luminance light scalar, vector, squash tensor and texture amplitudes across the dataset. Each cell represents the average amplitude across the five days during each seasonal recording period. The data show notable seasonal, location and time of day effects for each statistic.

Light scalar value heatmaps are shown across figure 4.9i. When comparing the seasons broadly, November and March have comparatively low scalar values, while June and August have much higher values. Season additionally has an effect on the way light scalar modulates with

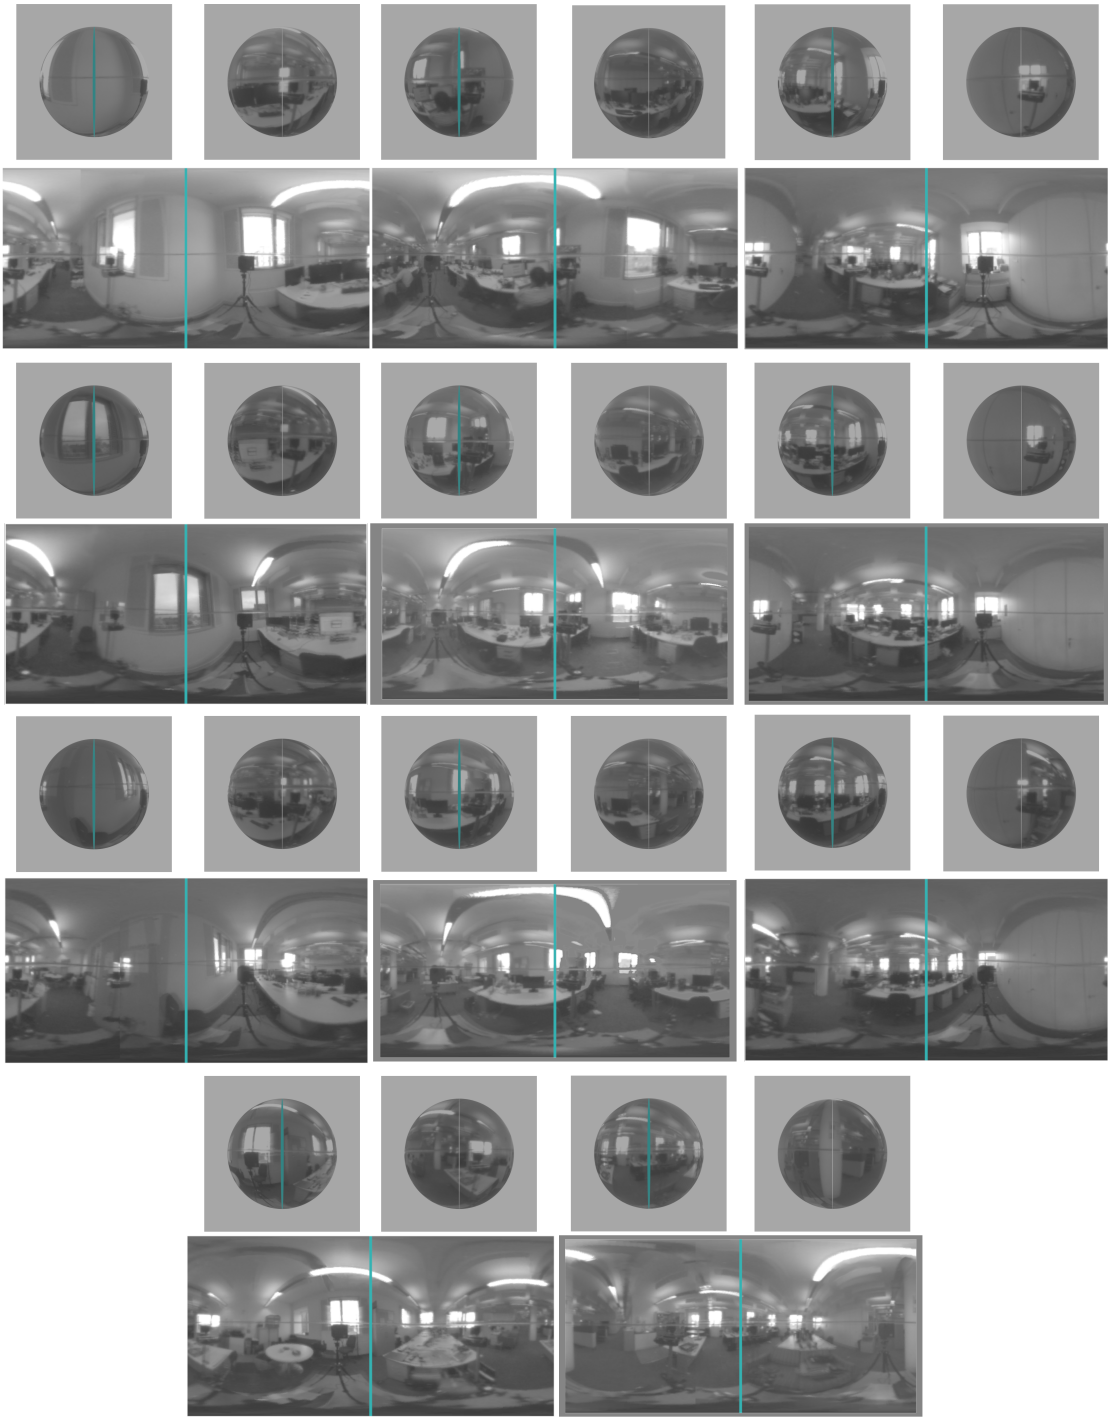

Figure 4.8: Example images of each lightprobe in experiment 1. Lightprobes are arranged according to the bird's-eye-view perspective shown in figure 4.6. Each lightprobe is shown unwrapped, with north shown by the teal vertical line. Additionally, the lightprobes are shown wrapped onto spheres while facing north (left) and south (right). South is shown by the vertical gray line.

time of day. In months where day lengths are longer (August and June), there are a number of time points with much higher scalar values, particularly in the 09:00 and 17:00 samples in locations 1, 7, 10 and 11. As shown in figure 4.6, these locations are closest to the windows. More broadly, the scalar light value is often diminished at 12:00 compared to the morning and

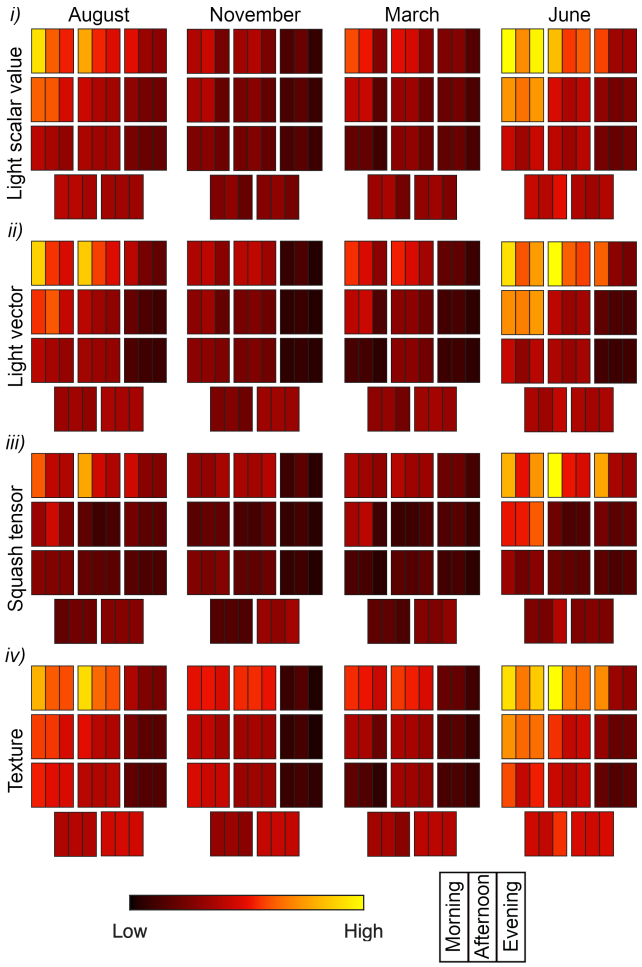

Figure 4.9: Heatmap showing the amplitudes of luminance spherical harmonics. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season, and rows show data for each spherical harmonic statistic. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left of each square sub-component, afternoon recordings in the middle, and evening recordings on the right.

evening samples in these time periods. The 09:00 and 17:00 scalar values are comparatively low in the November and March dataset. This is particularly true for the 17:00 sample, when the sun has often set, leaving little natural light to enter the office. Locations 7, 10 and 11 show increased light scalar values compared to the other locations throughout the dataset, and this difference is modulated by the season, with June and August showing the largest differences between these locations and the others. Results of a three-way ANOVA are shown in table 4.3, with significant effects shown in bold. Figure 4.10i shows the means for each group for the three variables.

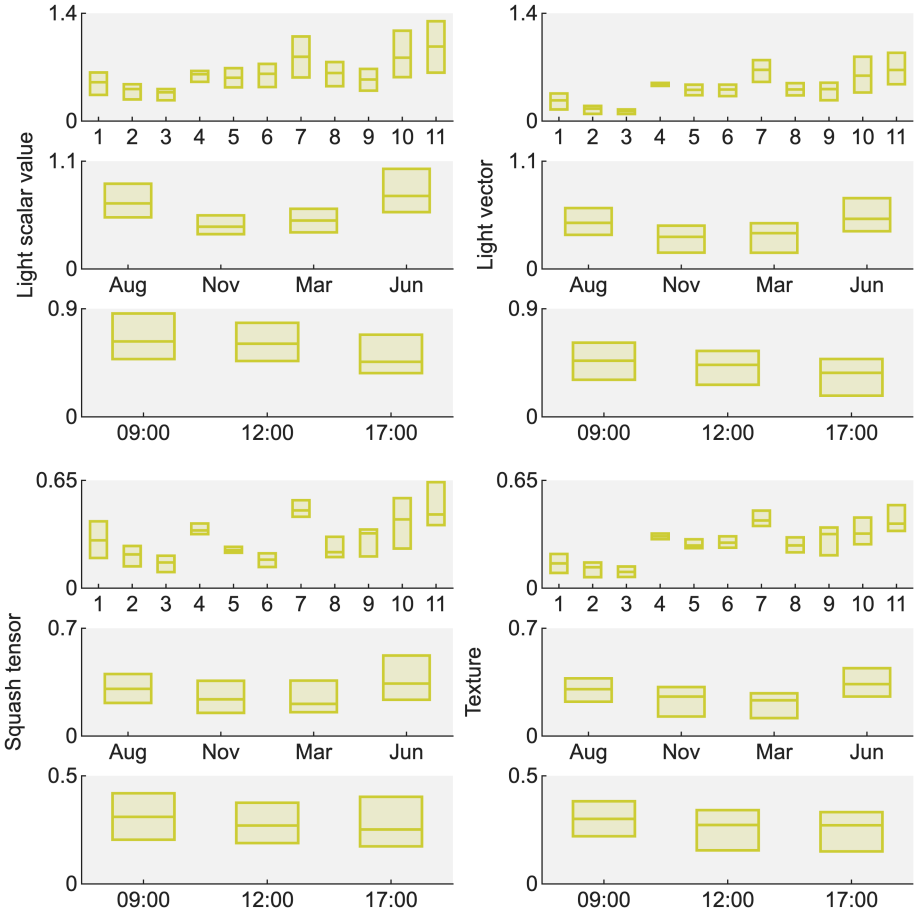

Figure 4.10: Means across each variable for luminance spherical harmonics: *i)* Light scalar value means across location, season and time. *ii)* Light vector means. *iii)* Squash tensor means. *iv)* Texture means. Shaded regions show inter-quartile range.

Effect	df	F	p
Location	**10, 528**	**103.87**	**< .001**
Season	**3, 528**	**209.14**	**< .001**
Time of day	**2, 528**	**53.30**	**< .001**
Location:Season	**30, 528**	**7.82**	**< .001**
Location:Time of day	**20, 528**	**3.20**	**< .001**
Season:Time of day	**6, 528**	**16.30**	**< .001**
Location:Season:Time of day	60, 528	0.96	.5715

Table 4.3: Three-way ANOVA results for luminance light scalar value amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; n = 5 per cell.)

Light vector heatmaps are shown in figure 4.9ii, showing similar patterns to the light scalar value. June and August generally show stronger directionality. Seat locations nearest the windows - namely locations 7, 10 and 11 - have the strongest directionality, particularly in the summer months. Directionality is also shown to be strongest in the morning, decreasing in the afternoon and being at its lowest in the evening. Light vector values are persistantly low in seats 1, 2 and

3 during November and March. Results of a three-way ANOVA are shown in table 4.4. Figure 4.10ii shows the means for each group for the three variables.

Effect	df	F	p
Location	**10, 528**	**158.85**	**< .001**
Season	**3, 528**	**156.71**	**< .001**
Time of day	**2, 528**	**57.25**	**< .001**
Location:Season	**30, 528**	**9.35**	**< .001**
Location:Time of day	**20, 528**	**5.57**	**< .001**
Season:Time of day	**6, 528**	**10.46**	**< .001**
Location:Season:Time of day	**60, 528**	**1.41**	**.029**

Table 4.4: Three-way ANOVA results for luminance light vector amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Figure 4.9iii shows the amplitudes of the second-degree harmonic degree, known as the squash tensor. Again, there are clear seasonal, time-of-day and location effects. Squash tensor amplitude is significantly higher in August and June than November and March, and significantly higher in the morning and afternoon than the evening. Locations closest to windows – 7, 10 and 11 – also have significantly higher amplitudes than locations further from windows. Similar to the light vector, desks 1, 2 and 3 have particularly low values in November and March. Figure 4.10iii shows the means for each group for the three variables. Table 4.5 shows the results of a three-way ANOVA analysing the squash tensor results.

Effect	df	F	p
Location	**10, 528**	**89.50**	**< .001**
Season	**3, 528**	**93.44**	**< .001**
Time of day	**2, 528**	**22.53**	**< .001**
Location:Season	**30, 528**	**8.52**	**< .001**
Location:Time of day	**20, 528**	**5.15**	**< .001**
Season:Time of day	**6, 528**	**8.00**	**< .001**
Location:Season:Time of day	**60, 528**	**2.16**	**< .001**

Table 4.5: Three-way ANOVA results for luminance squash tensor amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Figure 4.9iv shows the amplitudes of texture values, calculated as the mean of the third to thirty-first degree harmonics. In addition to the seasonal differences shown in the lower frequency harmonics, the texture amplitudes divide seating locations into low-texture areas - seats 2 and 3 - and high-texture areas - seats 7 and 11. Seat 1 lies where these two groups meet, and shows high texture in the morning and low texture in the evening. Generally, texture is higher in the morning and decreases throughout the day. Figure 4.10iv shows the means for each group for

the three variables, and table 4.6 shows the results of a three-way ANOVA.

Effect	df	F	p
Location	**10, 528**	**285.48**	**< .001**
Season	**3, 528**	**254.90**	**< .001**
Time of day	**2, 528**	**57.60**	**< .001**
Location:Season	**30, 528**	**11.94**	**< .001**
Location:Time of day	**20, 528**	**4.53**	**< .001**
Season:Time of day	**6, 528**	**14.74**	**< .001**
Location:Season:Time of day	**60, 528**	**1.70**	**.0013**

Table 4.6: Three-way ANOVA results for luminance texture amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Spherical harmonics - melanopsin

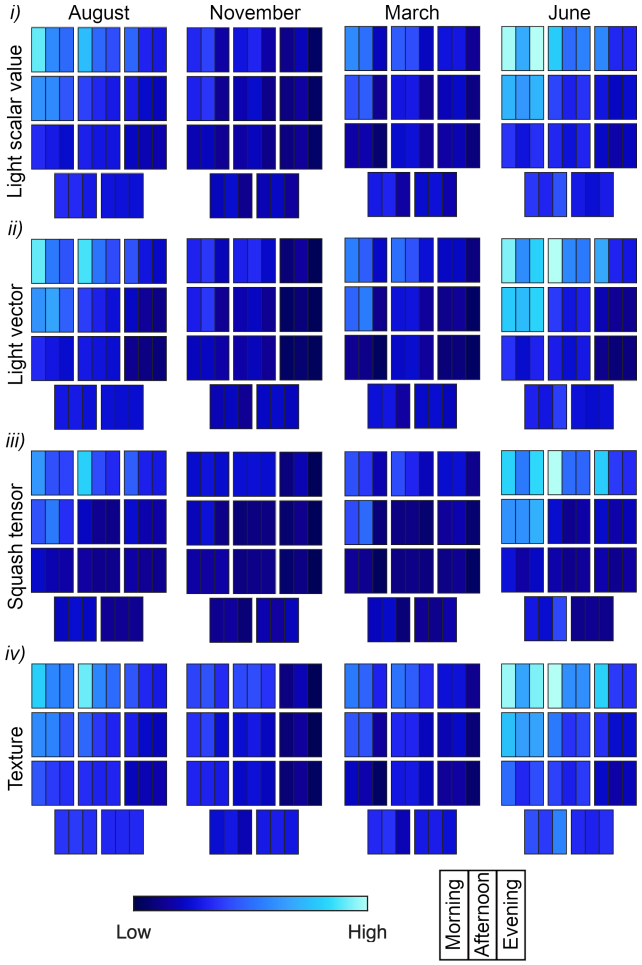

Figure 4.11: Heatmap showing the amplitudes of melanopic spherical harmonics. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season, and rows show data for each spherical harmonic statistic. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left of each square sub-component, afternoon recordings in the middle, and evening recordings on the right.

Figure 4.11 shows heatmaps for the melanopic light scalar, vector, squash tensor and texture

amplitudes across the dataset. Each cell represents the average amplitude across the five days during each seasonal recording period. The data show notable seasonal, location and time of day effects for each statistic. Figure 4.12 shows the averages across each variable type for the melanopic data.

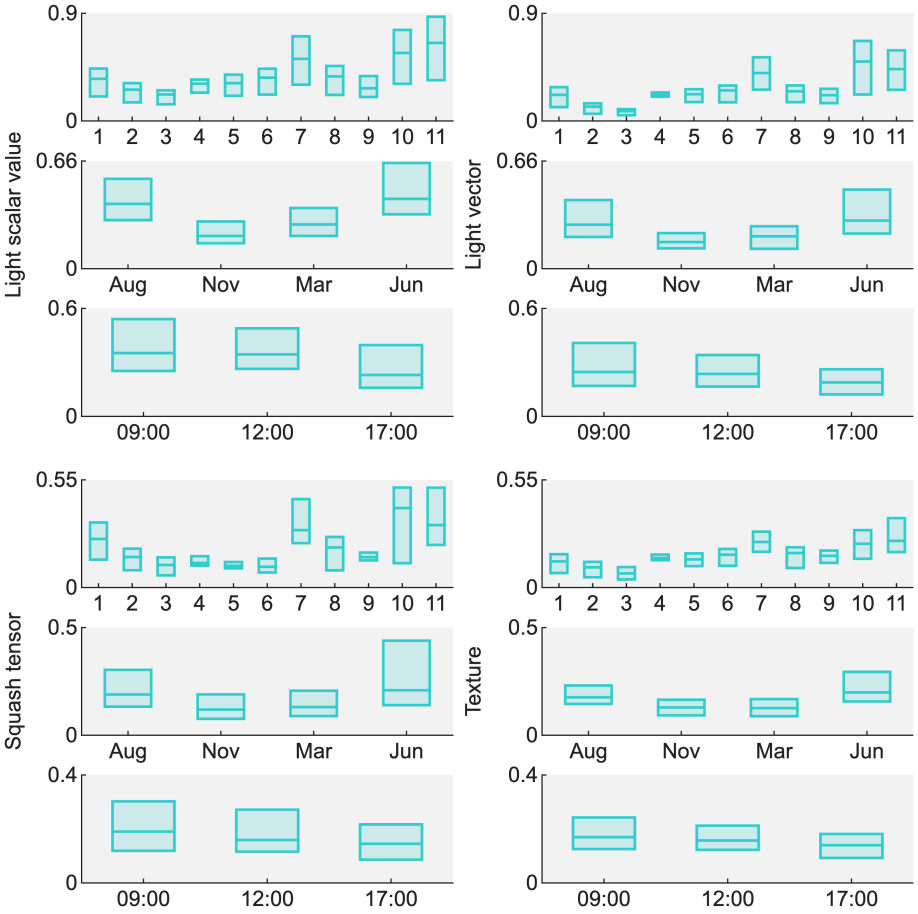

Figure 4.12: Means across each variable for melanopic spherical harmonics: *i)* Light scalar value means across location, season and time. *ii)* Light vector means. *iii)* Squash tensor means. *iv)* Texture means. Shaded regions show inter-quartile range.

Largely, melanopic light varies throughout location, time and season similarly to the luminance data presented above. Light scalar values, shown in figure 4.11i are highest in August and June, and particularly so for seats nearest the windows. Light scalar values decrease throughout the day, with morning recordings having the highest amplitude. Means for the melanopic light scalar values are shown in figure 4.12i, and the results of a three-way ANOVA are shown in table 4.7.

Effect	df	F	p
Location	**10, 528**	**90.49**	**< .001**
Season	**3, 528**	**183.45**	**< .001**
Time of day	**2, 528**	**55.73**	**< .001**
Location:Season	**30, 528**	**7.34**	**< .001**
Location:Time of day	**20, 528**	**3.17**	**< .001**
Season:Time of day	**6, 528**	**17.69**	**< .001**
Location:Season:Time of day	60, 528	1.03	.4134

Table 4.7: Three-way ANOVA results for melanopic scalar light value amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Light vector values also follow suit. Seats near to the windows are also most exposed to directional light, and this is particularly the case in the morning. Directionality is low in November and March. Means for the melanopic light vector amplitudes are shown in figure 4.12ii, and the results of a three-way ANOVA are shown in table 4.8.

Effect	df	F	p
Location	**10, 528**	**120.25**	**< .001**
Season	**3, 528**	**147.18**	**< .001**
Time of day	**2, 528**	**60.63**	**< .001**
Location:Season	**30, 528**	**8.68**	**< .001**
Location:Time of day	**20, 528**	**5.59**	**< .001**
Season:Time of day	**6, 528**	**12.00**	**< .001**
Location:Season:Time of day	**60, 528**	**1.46**	**.0174**

Table 4.8: Three-way ANOVA results for melanopic light vector amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Melanopic squash tensor amplitudes are particularly concentrated in seats 7 and 11, close to the East-facing wall. Again, these are most pronounced in the morning and in the months of June and August. The results of a three-way ANOVA are shown in table 4.9 and means for the melanopic light vector amplitudes are shown in figure 4.12iii.

Effect	df	F	p
Location	**10, 528**	**84.66**	**< .001**
Season	**3, 528**	**115.14**	**< .001**
Time of day	**2, 528**	**31.38**	**< .001**
Location:Season	**30, 528**	**8.95**	**< .001**
Location:Time of day	**20, 528**	**5.12**	**< .001**
Season:Time of day	**6, 528**	**10.34**	**< .001**
Location:Season:Time of day	**60, 528**	**1.89**	**.0001**

Table 4.9: Three-way ANOVA results for melanopic squash tensor amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Finally, melanopic texture showed a similar but less intense pattern compared with the luminance texture. Low-texture seats 2 and 3 are less separated from the other seats, though high-texture seats 7 and 11 are still present. The results of a three-way ANOVA are shown in table 4.10 and means for the melanopic light vector amplitudes are shown in figure 4.12iv.

Effect	df	F	p
Location	**10, 528**	**116.89**	**< .001**
Season	**3, 528**	**210.50**	**< .001**
Time of day	**2, 528**	**64.77**	**< .001**
Location:Season	**30, 528**	**7.02**	**< .001**
Location:Time of day	**20, 528**	**4.13**	**< .001**
Season:Time of day	**6, 528**	**17.33**	**< .001**
Location:Season:Time of day	**60, 528**	**1.42**	**.0265**

Table 4.10: Three-way ANOVA results for melanopic texture amplitudes. (Location = 11 levels; Season = 4 levels; Time of day = 3 levels; $n = 5$ per cell; error df = 528.)

Vector direction

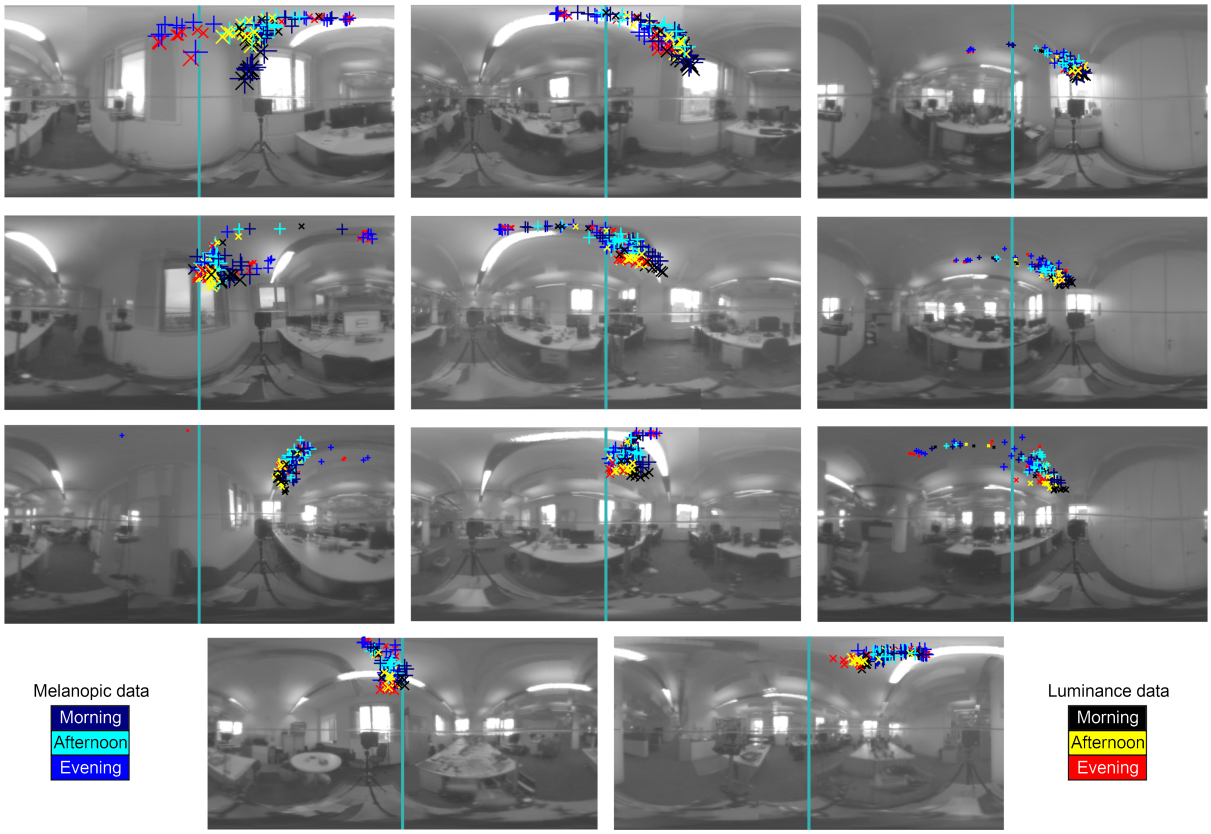

Figure 4.13: Light vector directions plotted on lightprobes. Cross symbols coloured black, yellow and red correspond to luminance light vector locations taken in the morning, afternoon and evening, respectively. Plus symbols coloured dark, light and navy blue correspond to melanopic light vector locations taken in the morning, afternoon and evening respectively. Vertical teal lines show north. Datapoint size corresponds to light vector strength.

Figure 4.13 shows the light vector directions plotted onto the lightprobes. All points are located in the upper hemifield of the lightprobes, most frequently oriented around ceiling lights and windows. Vectors for seat 1, for example, cluster around the large eastern window, while vectors for seat 5 align more closely with the ceiling light.

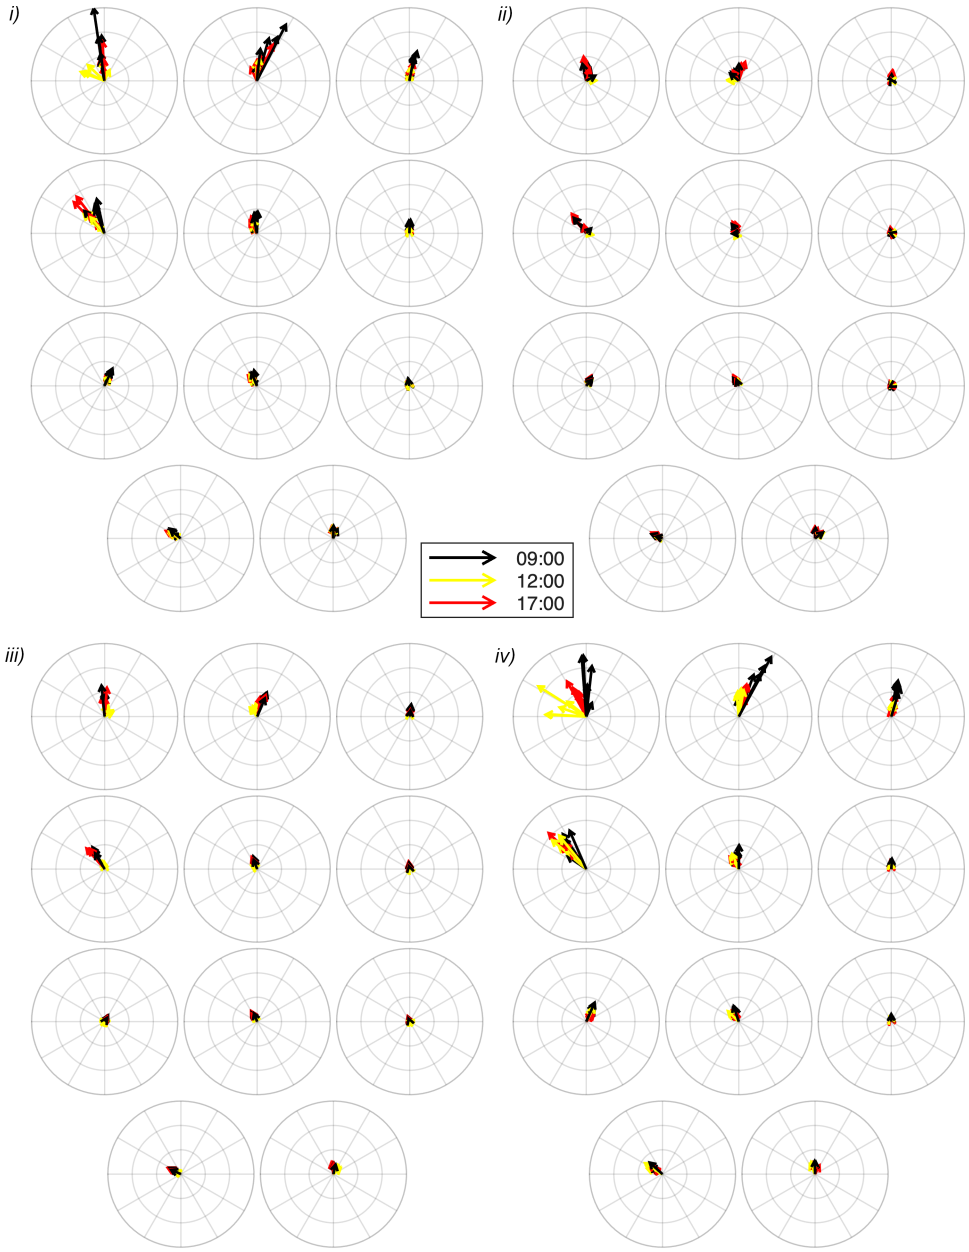

Figure 4.14: Phi plots of first harmonic degree. Figures show phi direction and amplitude for light vectors across the dataset. Morning samples are shown in black, afternoon samples in yellow and evening samples in red.

Figure 4.14 shows the amplitude and azimuth (λ) of first-degree spherical harmonics throughout the dataset. Each subfigure corresponds to a seasonal period, and within each subfigure the harmonic amplitudes are shown as vectors from the position of the recording. Subfigure *i)* shows August recordings, subfigure *ii)* shows November recordings, subfigure *iii)* shows March

recordings and subfigure *iv)* shows June recordings. 09:00 samples are coloured in black, 12:00 samples are coloured in red and 17:00 samples are coloured in yellow. The vectors show a range of behaviour depending on the season, time of day, and location. August and June show the strongest directional amplitudes, and these vectors are nearly always pointed towards the windows. This is the case for November and March too, though the directionality in these seasons is generally diminished. Location seems to be particularly important for the directionality vector. Locations near the windows (1, 7, 10 and 11) feature the strongest directional characteristics in all four seasonal datasets, with the vectors generally angled toward the nearest windows. Within this, there is also a strong effect of time of day on the amplitude and direction of the vectors. In many cases, vectors moved based on the position of the sun with respect to the building. Key examples of this include locations 7 and 11 during June and August. These months feature long, sunny days where the sun provides direct light to the windows on the east side of the building in the morning, moving towards the west throughout the day. At locations 7 and 11, morning vectors are often the strongest, pointing east towards the morning sun, before being reduced dramatically and pulled anti-clockwise towards the northern-westerly window, which lets some of the evening sun into the office. Here, the sun is a key driver for the direction of the first-degree harmonics to move throughout the day in certain seasons. Other locations are largely unaffected by the directionality of the sun. For example, locations 2 and 3 show minimal directional amplitude throughout the duration of data collection. Locations 4, 5, 6 and 8 show relatively small directional amplitudes, though they are clearly affected by both season and time of day. The combination of artificial light and daylight seems to play a large role in modulating the direction of the light vector in these cases.

Diffuseness

Figures 4.15 and 4.16 show D_{Xia} (see equation 4.13) diffuseness statistics for each scene for luminance and melanopic data, respectively. The metric D_{Xia} is calculated as a scaled ratio between the zeroth and first spherical harmonics, normalised such that 0 corresponds to fully collimated light and 1 corresponds to fully diffuse light [272]. For both spectral bands, locations 2 and 3 are the most diffuse throughout all seasons, and location 7 is the least diffuse. Light is most directional in the 09:00 samples, and most diffuse in the 12:00 samples. A three-way ANOVA showed significant main effects of season ($F(3, 528)=29.7$, $p<.001$), time of day ($F(2, 528)=10.88$, $p<.001$) and location ($F(10, 528)=483.99$, $p<.001$). There were also significant

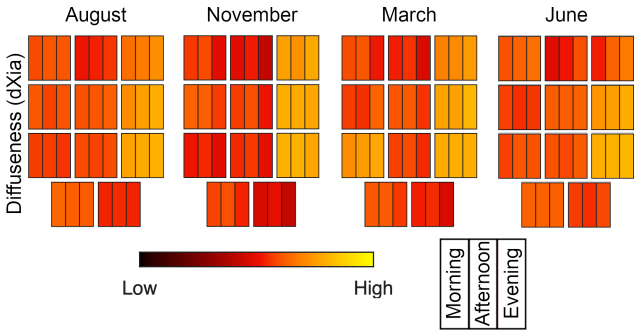

Figure 4.15: Luminance diffuseness heatmap. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left, afternoon recordings in the middle, and evening recordings on the right.

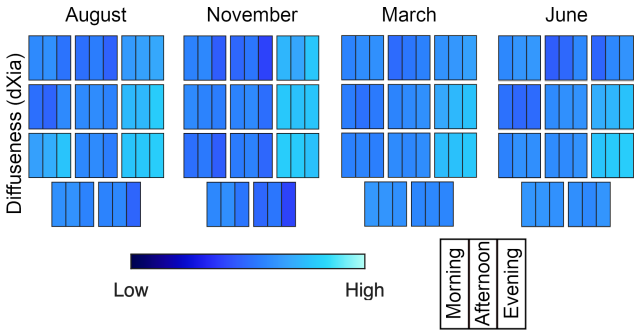

Figure 4.16: Melanopic diffuseness heatmap. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left, afternoon recordings in the middle, and evening recordings on the right.

interactions between season and time of day ($F(6, 528)=32.64$, $p<.001$), season and location ($F(30, 528)=31.3$, $p<.001$) and time of day and location ($F(20, 528)=9.74$, $p<.001$). There was also a significant interaction of all three variables ($F(60, 528)=3.99$, $p=<.001$).

Spectral analysis

Figures 4.17, 4.18, 4.19 and 4.20 show the power spectra of spherical harmonic degrees up to 30 for each location, time, season and photoreceptor type. In all cases, amplitude is highest for low-degree harmonics and decreases as the degree increases. Here, amplitude is plotted on a log scale. Notably, scenes previously identified as having high directionality – those in June and August, in locations near to the windows - most notably 7, 10 and 11 – tend to have broadly complex spatial structures, with much more high angular frequency information. Locations 2 and 3 by comparison show the most texture fall off. In many cases, there is a significant spectral wavelength dependency, which is emphasised in the locations with high directionality. These same scenes with high directionality, likely due to the sun, feature greater amplitudes across the

degree spectrum for receptor maps with longer wavelengths. Seats furthest from the windows - 3, 4, 5 and 8 - show very little time-based variation, which likely shows the dominance of static artificial lighting dominating those scenes. Seats 4 and 8 also show a bumpy structure, which likely reflects the appearance of the multiple far away windows as small light sources that "phase" across the harmonic degrees.

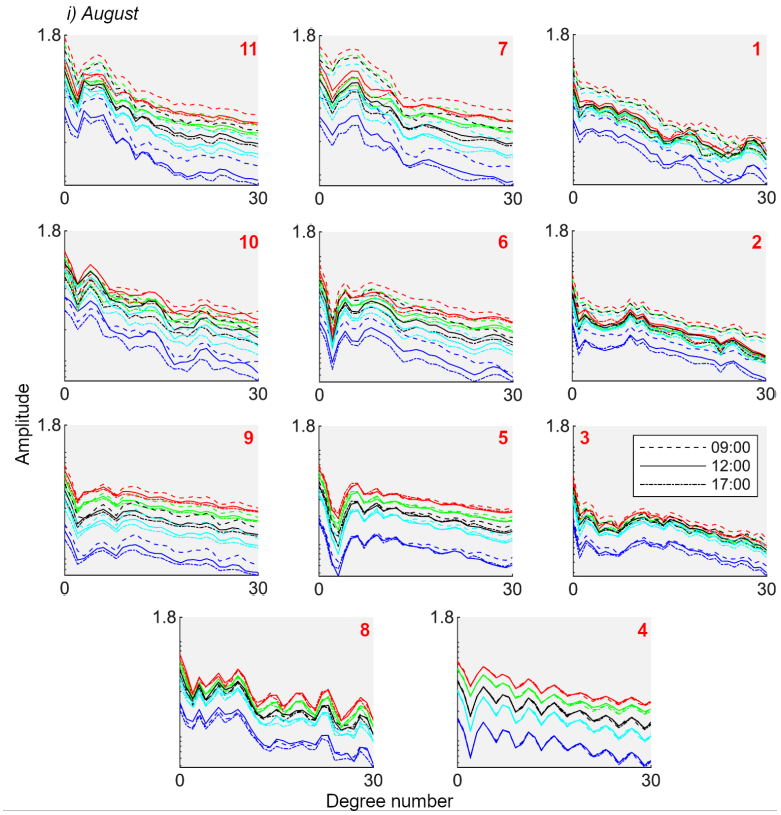

Figure 4.17: Amplitude spectra for spherical harmonic degrees 0-30 for each alpha-opic channel across location and day time in August. Colours correspond to alpha-opic channel - with S-cone values shown in blue, M-cone values shown in green, L-cone values shown in red, rod values shown in black and melanopic values shown in teal. Morning samples are shown with a short dashed line, afternoon samples with a solid line and evening samples with an unevenly-dashed line. The Y axis shows a log scale.

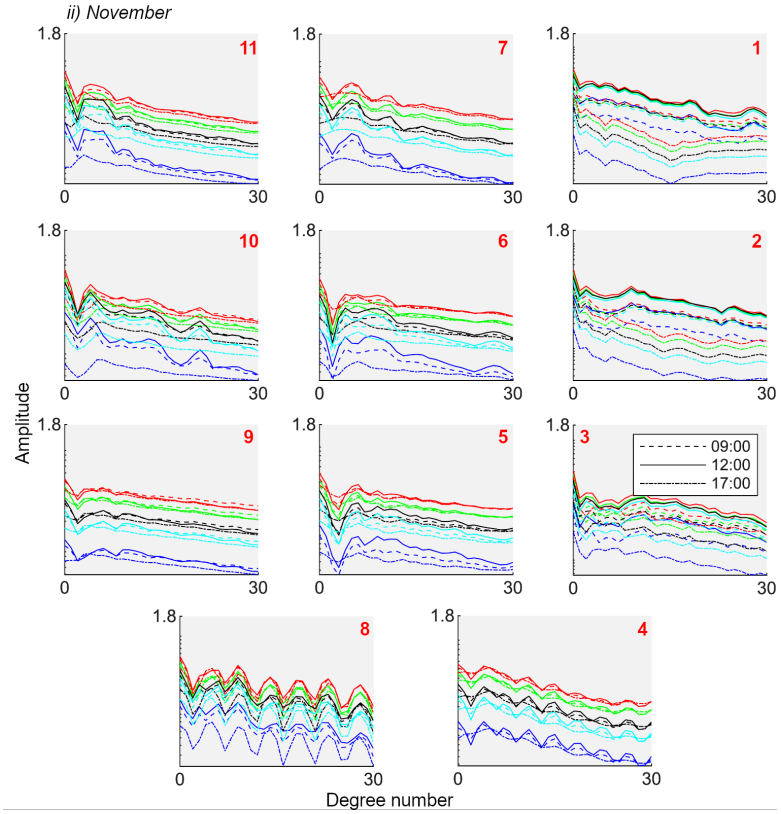

Figure 4.18: Amplitude spectra for spherical harmonic degrees 0-30 for each alpha-opic channel across location and day time in November. Colours correspond to alpha-opic channel - with S-cone values shown in blue, M-cone values shown in green, L-cone values shown in red, rod values shown in black and melanopic values shown in teal. Morning samples are shown with a short dashed line, afternoon samples with a solid line and evening samples with an unevenly-dashed line. The Y axis shows a log scale.

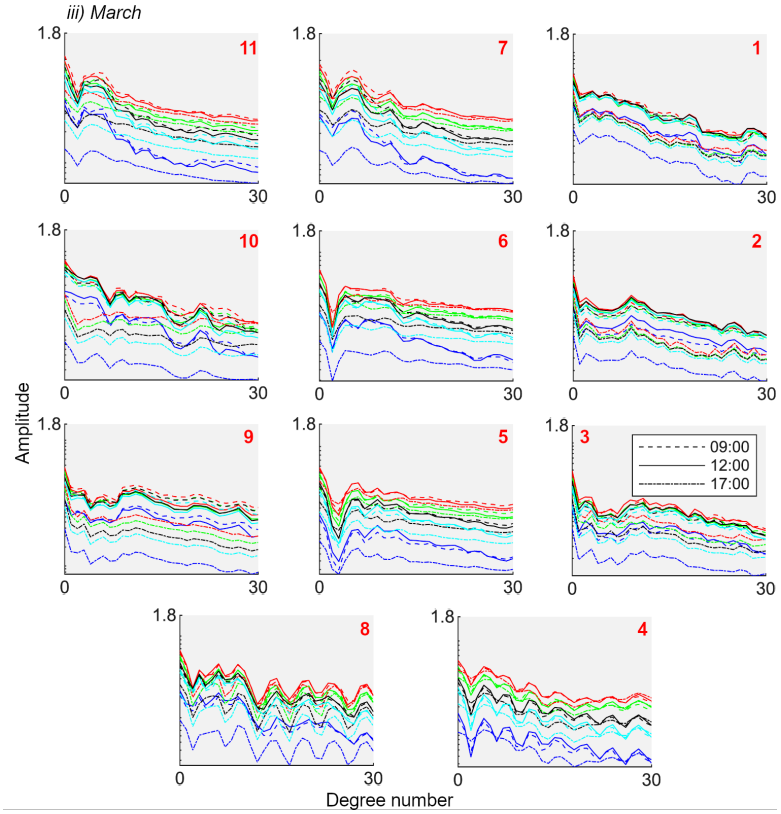

Figure 4.19: Amplitude spectra for spherical harmonic degrees 0-30 for each alpha-opic channel across location and day time in March. Colours correspond to alpha-opic channel - with S-cone values shown in blue, M-cone values shown in green, L-cone values shown in red, rod values shown in black and melanopic values shown in teal. Morning samples are shown with a short dashed line, afternoon samples with a solid line and evening samples with an unevenly-dashed line. The Y axis shows a log scale.

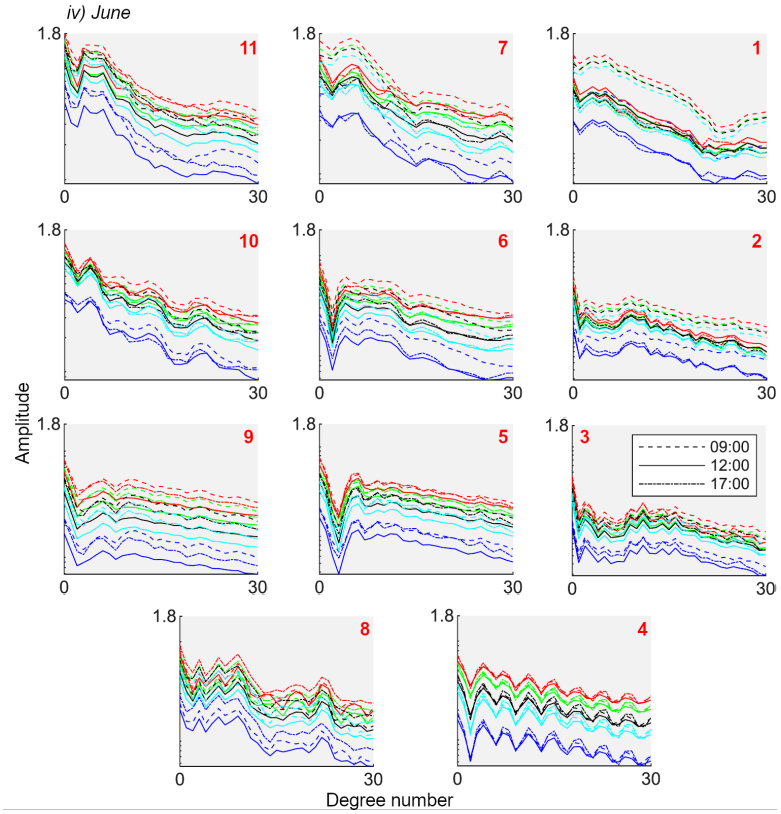

Figure 4.20: Amplitude spectra for spherical harmonic degrees 0-30 for each alpha-opic channel across location and day time in June. Colours correspond to alpha-opic channel - with S-cone values shown in blue, M-cone values shown in green, L-cone values shown in red, rod values shown in black and melanopic values shown in teal. Morning samples are shown with a short dashed line, afternoon samples with a solid line and evening samples with an unevenly-dashed line. The Y axis shows a log scale.

Photopic-melanopic dominance

Figure 4.21 shows photopic-melanopic dominance heatmaps across the dataset. Figure 4.21i shows the light scalar value of the photopic-melanopic dominance map, and shows that all of the samples across the dataset were dominated by luminous irradiance. This is particularly true of seat locations 1 and 2 throughout the year. Across the set of locations, luminance is least dominant in November. Figures 4.21ii-iv show the higher degree spherical harmonic distributions for photopic-melanopic dominance, which correspond to the strength of variance in spectral dominance at higher angular frequencies. This means a value of zero corresponds to a "perfectly spherical" dominance map - meaning that the level of dominance is equal from all angles. Higher values mean that the there is large variation in the spectral dominance at that particuar angular frequency. They show that seat location 2 has the lowest spatial variance in photopic-melanopic dominance across many of the seasonal conditions, however seats 1, 7 and 11 also have low variation at the highest angular frequency. The evening recordings from seat 10 in march show

the highest variation in spectral dominance across all of the angular frequencies.

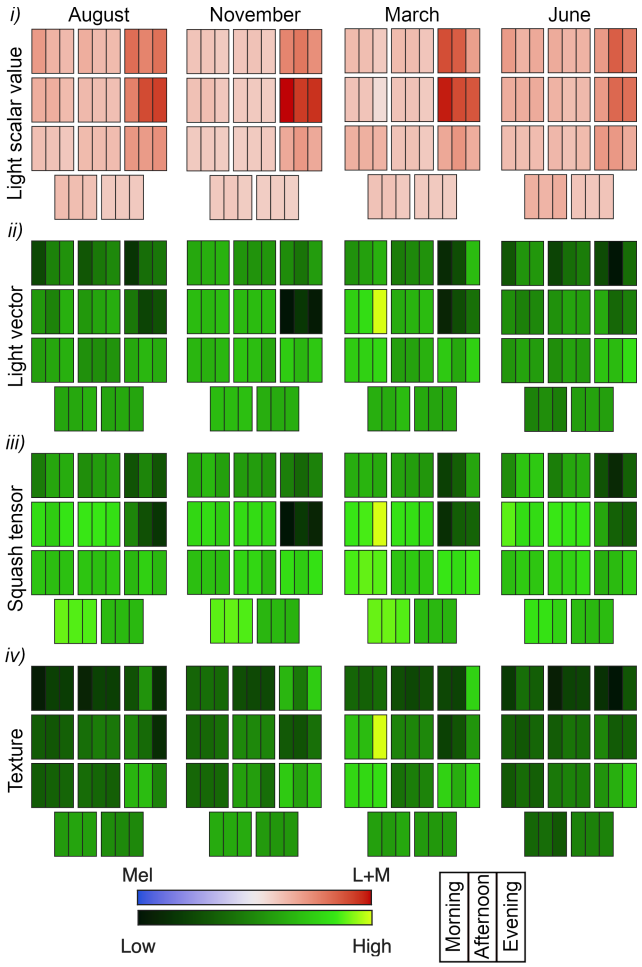

Figure 4.21: Heatmap showing the spherical harmonics of photopic-melanopic dominance maps. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season, and rows show data for each spherical harmonic statistic. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left of each square sub-component, afternoon recordings in the middle, and evening recordings on the right. The heatmap showing light scalar values is blue for melanopsin-dominated samples, red for luminance-dominated samples, and white for neutral samples.

Figure 4.22 shows the diffuseness of the photopic-melanopic dominance map, otherwise interpreted as the ratio of the overall and hemifield-localised spectral dominance. Seats 1, 2, 7 and 11 show the highest diffuseness throughout the year. Seat 3 has the lowest signal to noise ratio across the whole of the dataset.

4.3.2 Experiment 2

Lightprobes

Figure 4.23 shows two example lightprobes from the mezzanine (4.23i) and skylight (4.23ii) locations. Teal vertical lines show north, and grey vertical lines show south. The lightprobes are

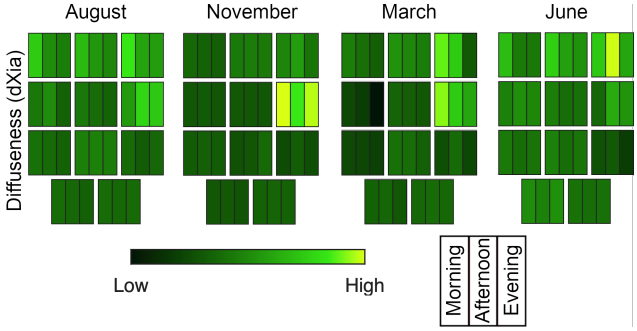

Figure 4.22: Photopic-melanopic dominance diffuseness heatmap. Each cell corresponds to the average of five recordings across each time-location-season combination. Columns show data for each different season. Within each heatmap array, cell groups are organised according to the layout shown in 4.6. Morning recordings are shown on the left, afternoon recordings in the middle, and evening recordings on the right.

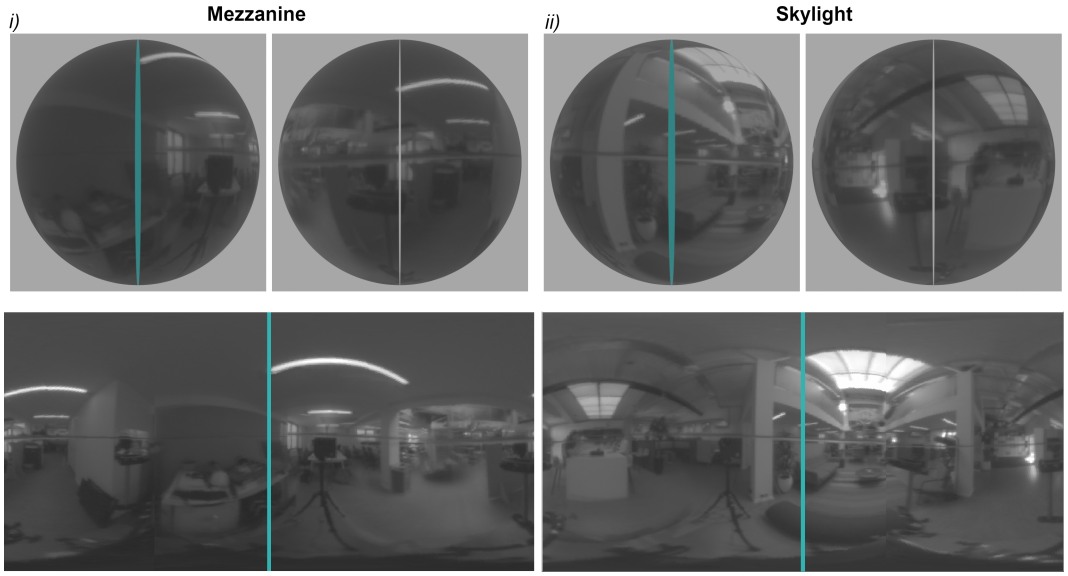

Figure 4.23: Example images of the two lightprobe locations used in experiment 2. Each lightprobe is shown unwrapped, with north shown by the teal vertical line. Additionally, the lightprobes are shown wrapped onto spheres while facing north (left) and south (right). South is shown by the vertical gray line.

also wrapped onto a sphere and viewed from north and south for visibility.

Spherical harmonics - luminance and melanopsin

Figure 4.24 shows the spherical harmonic amplitudes for light samples recording at the Hoare Lea office for experiment 2. Data are averaged across the length of the day for each location - with eight samples for the skylight location and nine samples for the mezzanine location. They show that the skylight position has persistently more intense ambient light throughout most of the day, and that this light is more directional at low and high angular resolutions too. Notably, there is a two-hour period starting at around 12:00 where the intensity and directionality of the light visible to the mezzanine location increase significantly before reducing at around 14:00. Luminance signals are persistently stronger than melanopic signals throughout the data collection

period. Diffuseness is much lower in the skylight area than the mezzanine area, demonstrating the relative intensity of light under the available natural illumination.

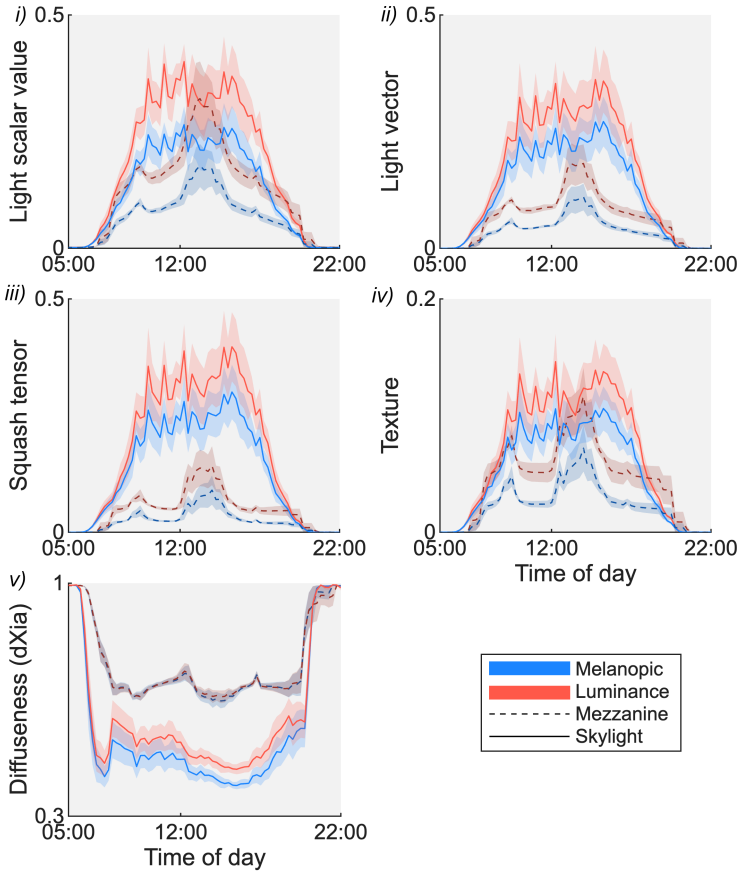

Figure 4.24: Spherical harmonic amplitudes throughout the day. *i)* Light scalar value *ii)* Light vector *iii)* Squash tensor *iv)* Texture *v)* Diffuseness. Luminance data are plotted in red, and melanopic data in blue. Data from the skylight recording area are plotted in straight lines while mezzanine area data are shown in dashed lines and darkened. Shaded area shows standard error across the nine days under the mezzanine, and eight days under the skylight.

Luminance and melanopic light vectors

The angular distribution of luminance and melanopic light vectors is shown in figure 4.25. They show notable groupings and intensity differences based on the the time of day. Both spectral bands and locations receive directional light from the east during the later morning and afternoon. In the skylight location this is highlighted by the vector pointing at the skylight itself. For the mezzanine location, the vector is also pulled towards the position of the skylight and artificial lights. Early morning and later evening samples group more towards surfaces more likely to be lit by artificial lights for both locations.

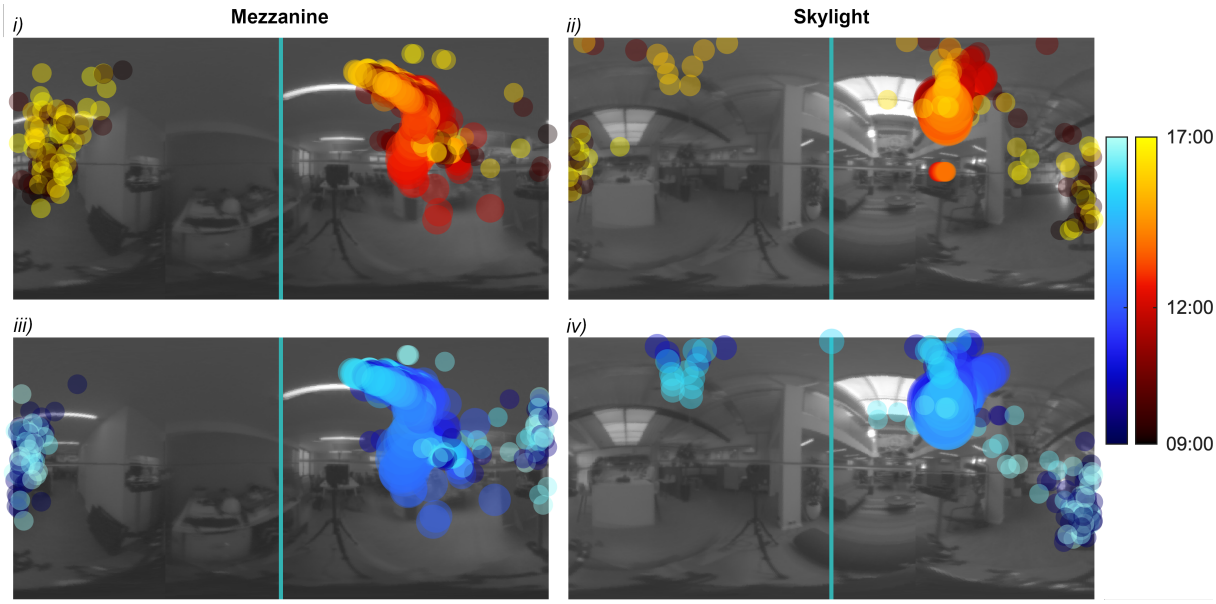

Figure 4.25: Light vector directions plotted on mezzanine and skylight lightprobes. Points coloured from black to yellow to red shown in the top row (*i) and ii)*) show luminance light vector directions from morning to evening. Points coloured from black to light blue shown in the bottom row (*iii) and iv)*) show melanopic light vector directions from morning to evening. Vertical teal lines show north. Datapoint size corresponds to light vector strength.

Photopic-melanopic dominance

Figure 4.26 shows the spherical harmonic amplitudes throughout the day for the photopic-melanopic dominance map. Figure 4.26i shows that for both locations, the ambient light is marginally dominated by the luminance band throughout the day. The mezzanine region becomes markedly more luminance dominated in the morning and evening. The other metrics demonstrate that the angular variation in this dominance increases throughout the day, peaking at around midday. This variation occurs across all of the tested angular frequencies.

Photopic-melanopic dominance light vector

Figure 4.27 shows the distribution of photopic-melanopic dominance light vector angles in the two recorded locations. They show a strong clustering pointing downwards towards the west, notably opposing the scalar light vector positions in the upper east. Because of the luminance-dominated profile of light in this scene, this demonstrates the balancing of light from the opposing direction of the light vector location. Namely, the strong short-wavelength influence of natural light from the skylight. These vectors are strongest in the late morning and early afternoon, with early morning and later evening vectors distributed more widely across the scene.

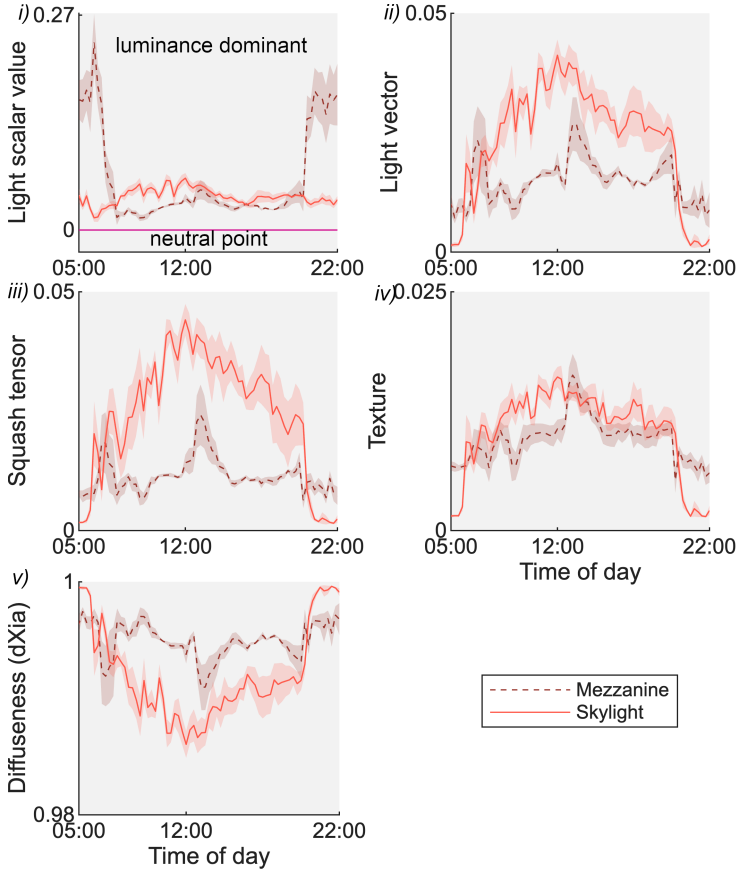

Figure 4.26: Spherical harmonic amplitudes throughout the day for the photopic-melanopic dominance map. *i)* Light scalar value *ii)* Light vector *iii)* Squash tensor *iv)* Texture *v)* Diffuseness. Data from the skylight recording area are plotted in straight lines while mezzanine area data are shown in dashed lines and darkened. Shaded area shows standard error.

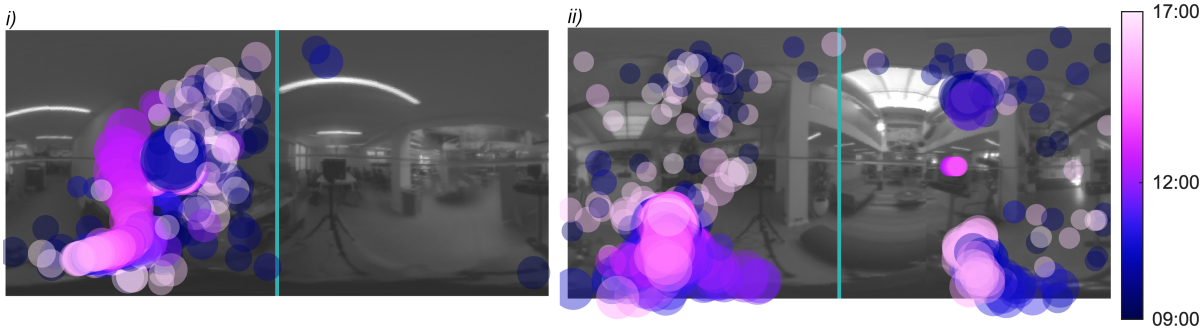

Figure 4.27: Light vector directions for photopic-melanopic dominance plotted on mezzanine and skylight lightprobes. Points coloured from black to purple to white shown in the top row show photopic-melanopic dominance light vector directions from morning to evening. Vertical teal lines show north. Datapoint size corresponds to light vector strength.

4.4 Discussion

The present study sought to develop and validate a single-image lightprobe imaging system and use it to quantify changes to the angular structure of light in example office spaces, leveraging the high spatial, temporal and spectral resolution of the setup to gain insights into the distribution

of physiologically-relevant light in the built environment.

4.4.1 The single-image circadian lightprobe imaging system

We have developed and validated a lightprobe capture system that allows for the generation of high spatial resolution, high dynamic range lightprobes using a single image. Modelling using the Mitsuba physically-based renderer demonstrated that that this single-image approach preserves the spatial and contrast structure of the surrounding scene well. This allows for the flexible characterisation of lighting environments as the setup is mobile and can capture a series of lightprobes with a high temporal resolution without requiring supervision during recording and opens the door for future experiments that seek to quantify the light field in different environments. The flexibility of the imaging system means that it can be used in conjunction with a range of cameras, such as hyperspectral, RGB or ultraviolet imagers. The Westboro Photonics Circadiometer is powered using mains electricity, but battery-powered cameras would allow the setup to be taken outside and to a range of locations.

A limitation of the single-image methodology presented here is that the lack of movement of the camera system means that the planar mirrors and camera appear in the resultant lightprobes as there are not alternate camera angles to fill in the space. In the experiments presented here, the camera did not obscure direct light sources, overlapping with the indoor environment, and thus had no measurable effect on the distribution of lower-angular frequency spherical harmonics. However, this may present an issue in lighting environments with more variegated illumination maps. Future experiments will have to weight these considerations based on the characteristics of the environment, and the kind of analysis targeted by the experiment. Future developments for this kind of imaging system include using computer vision techniques to automate data preprocessing. While many stages of the preprocessing pipeline are automated, the binary masks used to warp and crop the mirror images are manually-defined. This approach was taken as it allowed us to compensate for small changes in the location of the mirror system or camera during recording. However, future work may utilise additional shape and edge detection algorithms to automate this process. For single-exposure systems, this could enable near-real-time analysis of the spatial structure of an environment, which will unlock further capabilities of the system. For example, the direct effect of behavioural changes to the environment such as adding plants, drawing blinds or adjusting a monitor could be assessed.

4.4.2 Experiment 1

In experiment one, we measured spherical harmonic amplitudes in a typical office environment, measuring the effect of season, time of day and location on the way that light is distributed in a scene, with all three significantly modulating the spatial structure of the scene. Most notably, the light field is quite different in positions within one metre of the windows compared to those further away, with significantly more ambient light and directional light. These positions are also more susceptible to changes due to the time of day and time of year, likely due to the increased visibility of direct sunlight when close to the windows. These locations saw the position of the directional component of light shift significantly throughout the day, particularly in June and August, in a way that mimicked the east-to-west diurnal path taken by the sun. By comparison, seats far from the windows changed very little due to time of day and time of year. This demonstrates that being as little as three metres away from windows can be enough to make indoor artificial lights the dominant illuminant in one's field of view and reduce the directional effects of sunlight entering an office space. We also found a significant effect of season. Predictably, ambient and directional light were most intense during the months June and August, and this was also true for the squash tensor metric, due to the multiple windows providing a strong light stimulus into the office. Interestingly, the diffuseness metric showed much smaller differences between the months, in fact finding that November was the least diffuse month. The reduced effect size for diffuseness suggests that across the months, ambient and directional light change at similar rates. There were also significant diurnal changes to the structure of the light field. The 09:00 samples featured the most directional light, due to the highly directional influence of the sun visible through the east windows in the morning. Light vector values were low in the evening, likely due to the lack of west-facing windows and the early sunset times in the November and March datasets. Directionality was strong, however, for some locations close to the windows in the evening, during June and August.

We additionally analysed the α-opic dependence of the spatial structure of the scene, consistently finding that longer wavelengths have finer spatial structures. This concurs with previous research finding a similar long-wavelength fine spatial structure in natural outdoor scenes [201]. Chapter 2 [277] and other previous work [207] has demonstrated the ability for mixed illuminant scenes such as this office to replicate the natural scene statistics such as seasonal changes when sufficiently supplied with natural light. Our results fall in line with these findings, suggesting that the fine spatial structure of long wavelength light, which is likely promoted by dappled

sunlight that is reflected by the many surfaces inside the office space while retaining its directional structure and the sunlight spectrum profile, is replicated indoors.

4.4.3 Experiment 2

Experiment 2 demonstrated the ability of the SICLIS setup to record light values at a higher temporal resolution within a single seasonal timeframe, comparing the light profile of two nearby locations that featured different primary illuminants. Indeed, the two locations differed vastly in the makeup of their light fields, with the region exposed to direct natural light having much more sustained light intensity throughout the day, and also featuring a much more directional light field. The mezzanine location, however also demonstrated a significant uptick in intensity and directionality in the middle of the day, which reflects the capacity that environmental inter-reflections have to influence the light exposure of an individual sat with no direct line of sight to natural light. This emphasises the importance of comprehensive measurement of the built environment, as for some parts of the middle of the day, the differences between the mezzanine and skylight regions are significantly reduced and in some cases are non-existent, while for much of the morning and later afternoon, the mezzanine had much less overall and directional light. This can also raise important considerations for lighting and building designers, as it demonstrates that small changes in location can yield very different light exposures. Being able to quantify these for office spaces is useful, and could enable dynamic behavioural nudges that improve circadian health. Users may, for example, opt to sit under the skylight in the mornings before moving to the mezzanine area in the afternoon, to encourage healthy photo-entrainment.

The skylight proved to be the strongest single directional influence on the scene, as well as being detectable as a clear driver of short-wavelength light in a scene that is broadly dominated by long-wavelength light. Its modulation of the spectral properties of the scene changed throughout the day, and critically drive a dissociation between the photopic and melanopic light in the scene. The fact that this dissociation was apparent in an office with a large number of artificial lights demonstrates that photopic measurements may not be sufficient to quantify the circadian effects of light exposure in built environments. This experiment provides a clear impetus for future experiments to consider how light exposure in different locations of typical office environments can change based on simple characteristics such as direct line-of-sight to windows.

4.5 Acknowledgements

The contents of this chapter are entirely that of the author, with supervision by Hannah E. Smithson, Manuel Spitschan and Ruth Kelly Waskett. Thank you to members of the Oxford Perception Lab for agreeing to allow data collection in their workspace. Thank you to Ben Acton for allowing data collection in the Hoare Lea Oxford office, and to the members of the Hoare Lea Oxford office for agreeing to allow data collection in their workplace.

Chapter 5

Individualised light dosimetry paired with head direction tracking in the built environment

Contents

5.1 Introduction . . . 138

5.2 Methods . . . 140

5.2.1 Devices . . . 140

5.2.2 Participants, location and procedure . . . 140

5.2.3 Data analysis . . . 141

5.2.4 Additional data and analysis . . . 142

5.3 Results . . . 144

5.3.1 Dosimetry . . . 144

5.3.2 Heading data . . . 145

5.3.3 Lightprobe models . . . 146

5.3.4 Seat-switch models . . . 148

5.3.5 Egocentric lightprobes . . . 151

5.4 Discussion . . . 153

5.5 Conclusions . . . 157

5.6 Acknowledgements . . . 157

5.1 Introduction

As discussed in Chapters 3 and 4, quantifying the light available in a range of illumination environments is essential for understanding the possible physiological responses they yield. Previous experiments have quantified light in both natural and artificial lighting environments [278–280]. While it is important to quantify the physical properties of light, as they set limits on the light that is 'available' in a particular environment, an important consideration is what light is actually sensed by individuals using the space [281]. In previous chapters, I emphasised the specification of alpha-opic measurements that capture the down-sampling inherent in our sensory systems. Here, we additionally consider directional sampling by an active human agent embedded in the environment. Individuals rarely sample the lighting environment evenly, for a multitude of reasons. First, task-dependent behaviour significantly limits gaze direction. Our sampling of light from the environment depends on the task we are doing, which typically involves us fixating on surfaces as we manipulate them as part of a task, rarely needing to look directly at lights. There is also typically little need to gaze directly upwards or downwards when carrying out task behaviour. Furthermore, particularly in the context of built environments, task-oriented behaviour may further limit gaze directions to a very small part of the global light environment, as individuals sitting at desks are limited to a range of possible head positions, and often look down to printed materials on their desk or computer displays in front of them. Second, light sampling behaviour is likely directly altered by the light available in the environment. Directionally-biased light or specular highlights causing glare, can have significant implications on visual comfort, and may lead to behaviours that bias environmental light sampling. This may simply involve turning away from a bright specular reflection but also extends to environmentally-altering behaviour such as closing blinds or curtains to avoid a strong directional bias in illumination [282, 283]. Biller et al. (2024) [284] describe light behaviour in a hierarchical structure wherein individual preferences and choices interact with the built environment – factors such as local architecture and lighting design – which in turn are part of wider cultural and location contexts that determine norms around available light. Thus, to quantify individualised light exposure, it is not appropriate to only quantify the *available* light, but rather the *sampled* light. These measures can then give a greater insight into the non-visual effects of light incurred by a particular illumination environment. This is particularly relevant for quantifying the effectiveness of lighting design or lighting-based clinical interventions, as traditional approaches assume that design decisions

result in the desired outcomes.

Many devices for recording individualised light dosimetry have become popularised in the advent of research into the non-visual effects of light, with a particular research and commercial focus on promoting wellbeing through appropriate exposure to light and these devices can help analyse the specific questions about the light that is sampled from the environment [285]. These devices can take many different form factors; some being wrist-worn, others worn on pendants or brooches, and others still being attached to glasses [286, 287]. Dosimetry devices are typically used to measure illuminance and melanopic EDI and have been used to validate that light interventions have resulted in a change in light exposure [288]. Previous experiments have explored the effect of window location, artificial lighting fittings and many other factors [289–292]. Several review papers have devised and suggested best practices for the design and analysis of dosimetry experiments [279, 280, 286].

The different form factors of dosimetry devices confer different advantages and disadvantages – wrist-worn dosimeters have a high participant acceptability, while devices worn on the head can lead to non-compliance. However, as determined above and from the significant spherical harmonic amplitudes showing directional variation in lightprobes, gaze direction is an important determinant of light exposure, and thus head-worn dosimeters likely better reflect the actual light exposure of users by aligning the light sensor with their corneal plane. Some current head-worn dosimeters log information about head tilt to enable researchers to analyse light exposure according to different head pitches [287], however full heading direction – taking into account both head tilt but also compass heading information – is not currently logged in popular light dosimetry devices. In this experiment, we detail the design and use of a novel corneal-plane light dosimetry device that, in addition to recording 8 narrowband channels of spectral irradiance, logs heading direction using a 9-axis inertial measurement unit (IMU), allowing for the characterisation of heading direction in terms of both pitch and yaw. We utilise these devices to compare light exposure across seating positions and seasons, as well as to derive further insights about light exposure using simultaneously-collected lightprobes, as described in Chapter 4.

5.2 Methods

5.2.1 Devices

Light exposure and head direction data were recorded using custom dosimetry devices, shown in figure 5.1. Our dosimetry device was mounted onto a plastic glasses frame, with the lenses removed, such that the sensors moved with the participants' head movements. Light exposure was recorded using a Condor Instruments ActLumus. These are small actigraphy devices that are commercially pitched as off-wrist sensors that measure light using ten spectrally-tuned sensors. The sensor picks up light from a 2π steradian hemisphere which is cosine weighted, giving irradiance measurements. They additionally feature accelerometers and ambient temperature sensors. The ActLumus devices are powered by a rechargeable internal battery. In order to measure corneal-plane light exposure, the ActLumus sensors were mounted to the centre of the glasses frames such that the sensors aligned with the direction of the wearer's head direction. A 3D-printed mount was screwed into the glasses frames to attach the ActLumus device. In order to determine heading direction, an Arduino Nano 33 BLE Sense was used. This small form factor microcontroller features a 9-axis inertial measurement unit (IMU). The IMU features an accelerometer, gyroscope and magnetometer. Each sensor type has three axes – pitch, yaw and roll. The Arduino microprocessors were attached to the arm of the glasses frame, such that their long axes were aligned with the wearer's eyeline. The Arduino Nano 33 BLE was powered by a 10,000 mAh 5 V portable charger. However, Arduino microprocessors in the Nano form factor draw a very low amperage, such that standard portable chargers do not continually power them without being 'woken'. To combat this, we added a USB 'KeepAlive' module to the Arduino's power supply line. This small USB board periodically incurs a 100 mA draw from the power bank, which allows the portable power supply to continuously power the Arduino. Data is logged from the IMU onto a microSD card via a microSD card adapter soldered to the microprocessor. Both the ActLumus and the Arduino devices log data every 15 seconds throughout the duration of experimentation. At this rate, the power bank is able to keep the Arduino charged for around six days, and the ActLumus device can be powered for over a month.

5.2.2 Participants, location and procedure

The experiment took place over two five-day periods: $3^{rd} - 7^{th}$ March and $10^{th} - 14^{th}$ June 2025. Participants were recruited from the test office area in the New Radcliffe House building in



Figure 5.1: Custom dosimetry device components. Top image shows all components of the custom dosimetry devices. It shows a glasses frame with the 3D-printed dosimeter holder screwed into the front. The ActLumus dosimeter fits into this during use. On the left arm of the glasses frame, the Arduino circuit is attached. The Arduino circuit is shown in the bottom image. It shows an Arduino Nano 33 BLE Sense soldered to a microSD card reader. A short micro USB to USB cable connects the Arduino to a USB Keep-alive. The USB Keep-alive connects the Arduino circuit to the power bank via the USB cable shown in the top image.

Oxford, UK (51.76° N, −1.27° W). Five participants were recruited in total, with four completing the experiment concurrently during each period and three participants taking part in both periods. Ethical approval for this was granted by the Oxford Medical Science Interdisciplinary Research Ethics Committee under approval reference R81269/004.

During this period of time, participants were instructed to wear the devices when seated at their desks and remove them when getting up from their desk and moving around the office, placing the device onto the desk. While wearing the devices, the participants continued their work as normal which comprised of working at computer workstations and laptops, and they freely viewed the environment. Figure 5.2 shows the location of each participant's desk. Desk locations 1-4 were used in the March dataset, and locations 1, 2, 3 and 5 were used in June. This coincided with the time periods of capture of the lightprobes in the previous chapter.

5.2.3 Data analysis

Light exposure data were analysed over the time course of the experiment. A linear weighted summation of the sensor outputs was used to estimate physiologically relevant metrics, such as α-opic lux and α-opic EDI. Additionally, a linear weighted summation is used to estimate irradiance values at 415, 445, 480, 515, 555, 590, 630 and 680 nm. The weightings are shown in

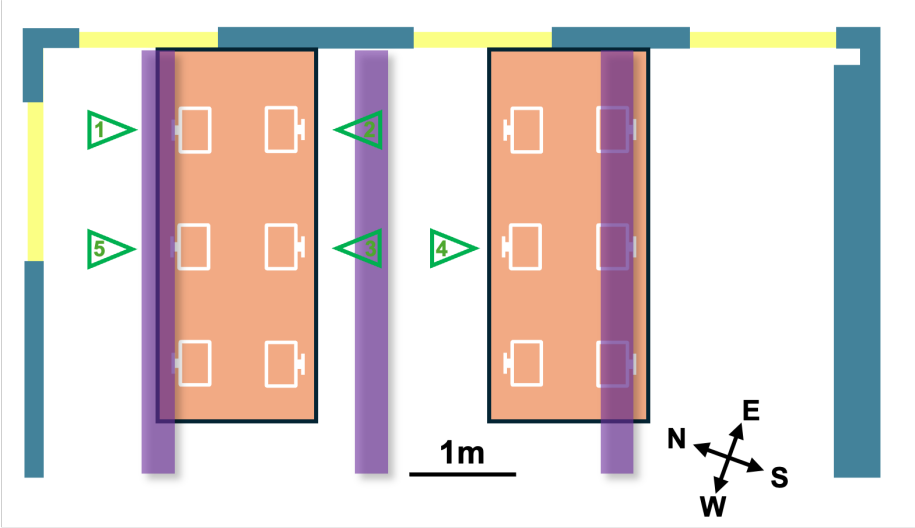

Figure 5.2: Birds-eye view of participant seat locations for dosimetry experiment. Black bar shows 1 metre of physical space. Green triangles show participant seat locations. Orange boxes show desks with computer workstations. Blue boxes show walls, and yellow gaps are windows. Purple boxes show artificial pendant lights.

equation 5.1. Equation 5.3 shows that melanopic EDI is calculated as the vector multiplication of the sensor values (shown in equation 5.2) by the weightings. Spline interpolation is then used to estimate a radiance information between 415 and 680 nm. From this spectral estimation, CCT is calculated. Melanopic EDI was also used to calculate relative circadian effect, using the 'global' circadian dose-response curve, as per the procedure explained in Chapter 1.

$$W_{melanopicEDI} = [\,1454.6\ 3000.4\ 4976.9\ 3324.2\ 131.0\ 196.9\ 7.6\ 132.3\,] \tag{5.1}$$

$$Dosimeter\ Sensors = [\,S_1\ S_2\ S_3\ S_4\ S_5\ S_6\ S_7\ S_8\,] \tag{5.2}$$

$$MelanopicEDI = Dosimeter\ Sensors \cdot {W_{melanopicEDI}}^T \tag{5.3}$$

In addition to light data, viewing direction was calculated using the magnetometer and accelerometer sensors in the Arduino microcontroller [293].

5.2.4 Additional data and analysis

During both data collection periods, circadiometric lightprobes were collected in several locations throughout the office space, as described in Chapter 4. In order to compare the lightprobe data with the dosimetry data, the lightprobes were transformed into irradiance data. The lightprobes

represent all of the light incident on one point in space, creating a 4π steradian field. To match this with the irradiance measurements made by the dosimeter devices, a cosine-weighted hemifield was applied to the lightprobe aligned with a target direction, which reduced the field of view to a 2π steradian hemisphere, and also weighted the light based on its angular deviation from the target direction with respect to the cosine filter. The data in this field were then summed to create an irradiance measurement for a given target angular direction in the lightprobe field. Lightprobes were collected at 09:00, 12:00 and 17:00 for all five days of the week during which dosimetry data were collected. Irradiance data were extracted for each lightprobe image using the sampled heading directions from the custom dosimetry devices, such that light in the lightprobes scene was sampled as if by an individual free viewing in that space.

Lightprobe irradiances were calculated for each daily lightprobe (09:00, 12:00 and 17:00). Since the dosimetry data have a higher temporal resolution than this, 10 minutes worth of dosimetry data - the nearest 10 minutes of the data collection period to each lightprobe sample time - was used to resample the lightprobe and calculate an irradiance sample.

To evaluate the effect of heading direction behaviours on light exposure, we used the lightprobe irradiance models to predict light exposure changes if the lightprobe irradiance for a given seating location was sampled using the heading directions from a different seat in what we call "seat-switch" models. This enabled us to calculate the reduction in correlation between dosimeter and lightprobe data - and thus the change in light exposure - as a result of different a change in heading behaviours.

Finally, we present egocentric lightprobes that combine a short time period of heading data with the lightprobes taken throughout the day to capture the lighting environment experienced by an individual based on their heading behaviours. We take five minutes of dosimetry data starting as close as possible to 09:00, 12:00, and 17:00 respectively, given that all the participants were free-viewing. All sets of samples were initiated within 195 seconds of the target times. We then realigned the corresponding lightprobes such that the heading directions match, and create an averaged image of the resulting illumination maps. As in Chapter 4, we analyse the directional qualities of these lightprobes using spherical harmonic decomposition.

5.3 Results

5.3.1 Dosimetry

Figure 5.3 shows the mean lux melanopic EDI values for each participant throughout the day in the March and June time periods averaged over each day, with standard error shown in the shaded bars. These data include variation over the day in availability of light, as well as light sampling. The error bars show variation across the 5 days and differences between the plots that exceed this variation are likley to reflect time of day and seasonal effects. They show notable variation between seasons, time of day and seating locations. Comparing the March dataset to the June dataset, lux Melanopic EDI is broadly higher in June than in March across all locations and for most of the day, concordant with previous findings that there is generally higher light intensity in June than in March (see Chapter 4). These seasonal differences are most pronounced the mornings and evenings, particularly between 09:00 and 10:00, where the data recorded in June are markedly higher than the rest of the day. As such, the range of values is higher in June, too, with March recordings remaining stable broadly across the day and also showing less variation between seating positions.

There are several ways to compare seating location. One dimension that is pertinent to light exposure is orientation of the desk, and the heading directions this yields. Per figure 5.2, our seating locations can be divided into north-north-west facing seats (2 and 3) and south-south-east facing seats (1, 4 and 5). This classification of seating arrangements demonstrates a difference in light exposure in June, dependent on prevailing heading direction. North-north-westerly seats (2 and 3) show a 'dip' throughout the day, with the highest values occurring in the morning and the evening, and a noticeable reduction throughout the day. By comparison, south-south-easterly seats (1, 4 and 5) feature a negative slope throughout the morning, before an increase in light exposure between 13:00 and 15:00. There is also larger standard error in these time periods. These locations don't feature an increase in the evening.

Figure 5.4 shows the data from figure 5.3 transformed into circadian response data using the procedure outlined in Chapter 2. It shows that for almost the entirety of the data collection period, participants were predicted to have a high level of circadian response throughout the day with minimal variation between the days and across seasonal datasets.

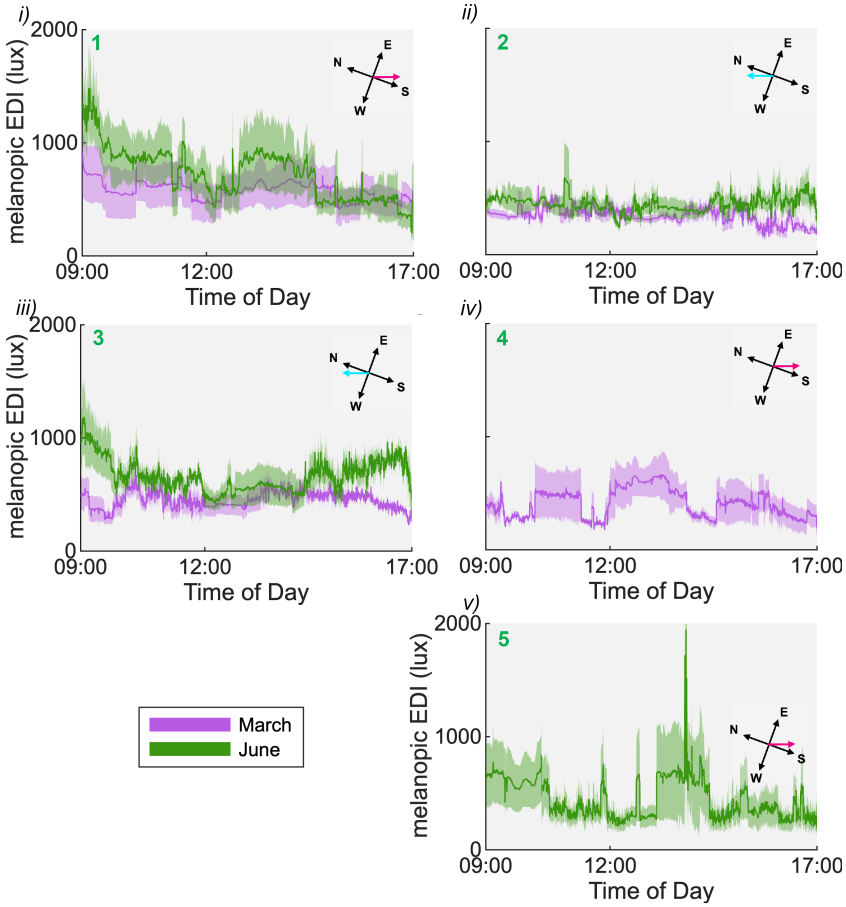

Figure 5.3: Average melanopic equivalent daylight irradiance across the day for each participant in March and June. Each subplot shows exposure for a different participant, with March traces shown in purple and June traces shown in green. Shaded regions show standard error across the five days over which data collection was carried out. Compasses show the the orientation of the desk for each participant.

5.3.2 Heading data

Figure 5.5 shows the viewing directions recorded by the Arduino devices while being worn by the participants. Each subplot shows, from top to bottom, the related lightprobe image from Chapter 4, then the viewing directions from March plotted onto the lightprobe image, then the viewing directions from June plotted onto the lightprobe image. Figure 5.5i corresponds to seat 1, 5.5ii to seat 2, 5.5iii to seat 3, and 5.5iv shows lightprobes for seats 4 and 5, each of which only appeared in one season. The lightprobes have been horizontally aligned such that the normal in the centre of the equirectangular projection is perpendicular to their desk. Overall, participants kept their head direction in a very limited range of angles, owing to the task-oriented behaviour the participants were carrying out during data collection. For each repeated participant, heading distributions are largely consistent between the March and June datasets. For example, the linear grouping shown by the participant in seat 1 and the triangular pattern shown by the participant in seat 3 remain consistent and idiosyncratic.

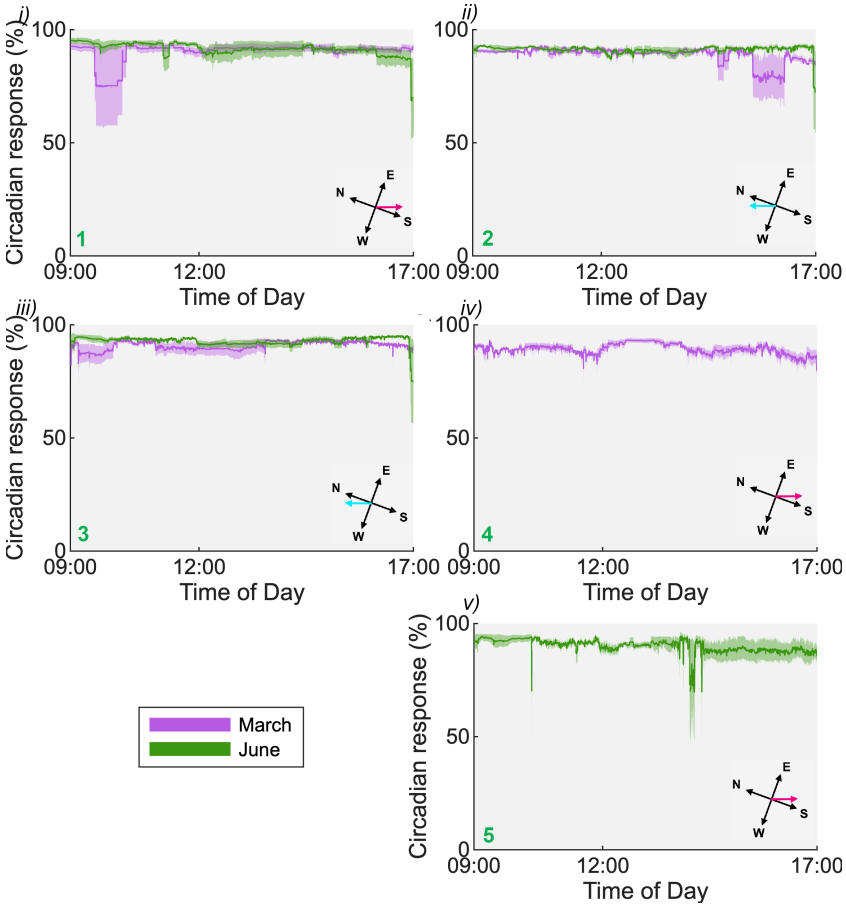

Figure 5.4: Average global circadian response across the day for each participant in March and June. Each subplot shows exposure for a different participant, with March traces shown in purple and June traces shown in green. Shaded regions show standard error across the five days over which data collection was carried out. Compasses show the the orientation of the desk for each participant.

Figure 5.6 shows the mean pitch of the heading directions across seat locations 1, 2 and 3 for March and June. The error bars show standard error. This plot shows significant differences of both season and seating location. Notably, average gaze locations are lower in June, though the range over which this occurs depended on the seat location.

5.3.3 Lightprobe models

We used lightprobes collected simultaneously with dosimetry data collection to model light exposure in each seat location. To do this, we used the heading direction of dosimetry samples and a cosine filter to turn the lightprobes into irradiance measures sampled at the dosimeters' sampling rate, which were then compared with the dosimetry data. On each day, three lightprobes were collected, at 09:00, 12:00 and 17:00. We used the heading data from the dosimetry recordings generate lightprobe estimates for the 10 minute period closes to each lightprobe sampling time. The fit of each lightprobe model was evaluated by finding the r^2 score when using the dosimetry

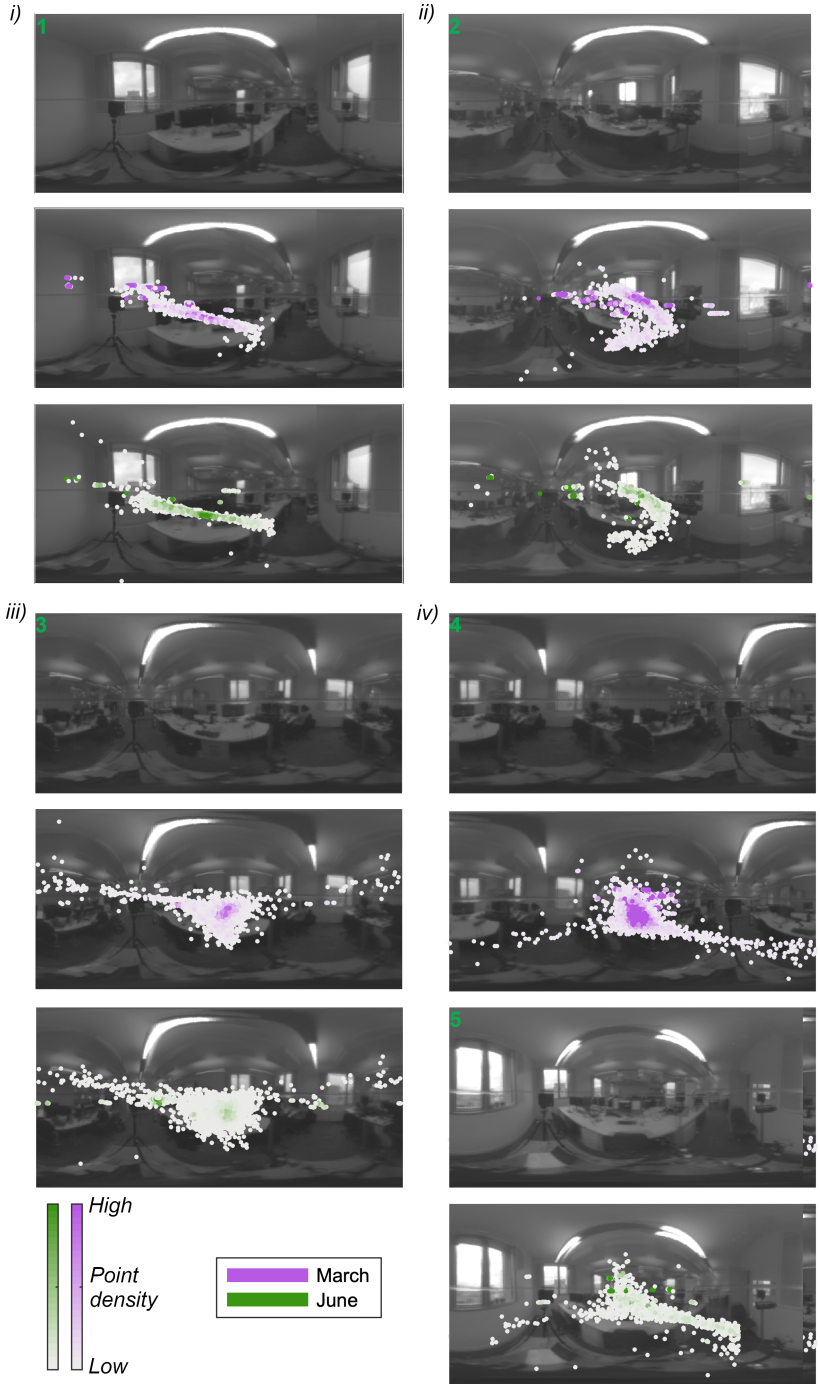

Figure 5.5: Heading distribution plots. Each set of subplots shows first, the lightprobe associated with each seating location; second, the distribution of heading directions in March; third, the distribution of heading directions in June. Points are white but appear more saturated purple and green for March and June, respectively, for the highest density regions of gaze samples. Lightprobes measured at 12:00 are used as the backgrounds for visibility, but all datapoints from active viewing are shown.

data to predict the output of the lightprobe model. Figure 5.7 shows the performances of these models against the original lightprobe data. Generally, the dosimetry data was able to explain a modest amount of the variation in the lightprobe model outputs. Models for seat 1 were best predicted by dosimeter data, and March models performed better than June models. Models for

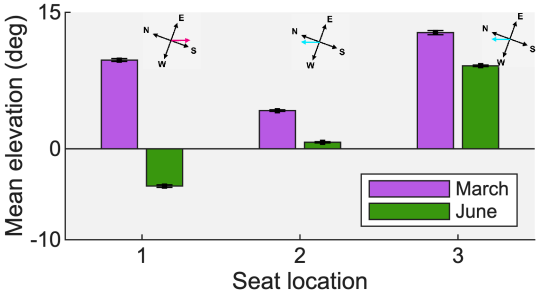

Figure 5.6: Comparison of heading pitch across seats with repeated recording. Purple bars correspond to heading in March and Green bars correspond to heading in June. Error bars show standard error.

seats 4 and 5, and the June model for seat 3 performed poorly, with an r^2 score very close to 0.

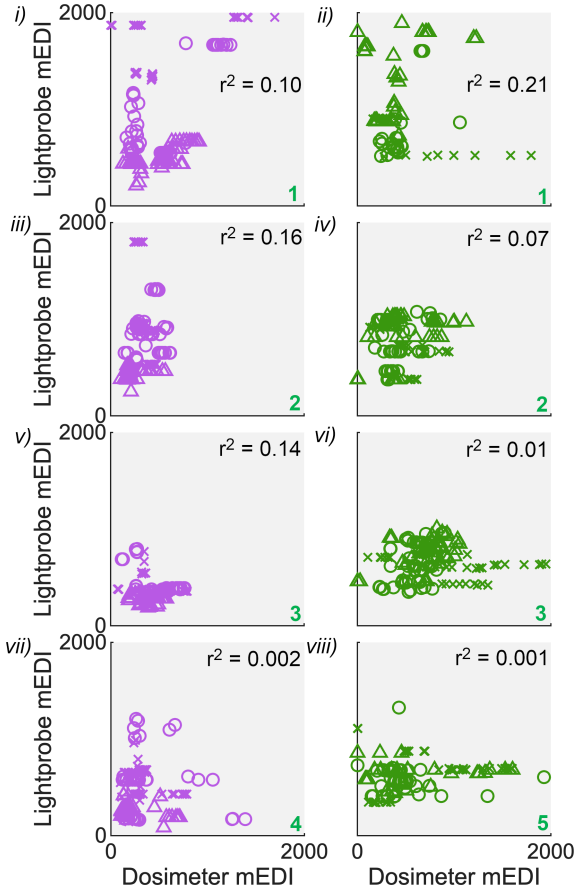

Figure 5.7: Lightprobe model irradiance predicted by dosimeter irradiance data. Each subplot corresponds to a month-seat pair. Left column corresponds to March data, and right column corresponds to June data. Each row shows a different seat location, with the top row showing seat 1, the second row showing seat 2, third row showing seat 3 and the bottom row showing seat 4 for March and seat 5 for June. Morning datapoints are plotted as crosses, afternoon datapoints are plotted as rings and evening datapoints are plotted as triangles.

5.3.4 Seat-switch models

To investigate whether different heading behaviours are sufficiently large to make a meaningful difference on the light exposure for a human agent at a given location, we generated a further

set of models by switching the heading data associated with each seating location. Thus, we analysed, for example, the light exposure of seat 2 using heading data from the participant in seat 3. This allows us to evaluate the extent to which particular gaze behaviours have an effect on light exposure. Figures 5.8, 5.9, 5.10 and 5.11 show this set of seat-mismatched models for seat locations 1, 2, 3, and 4 and 5, respectively. Each individual figure shows the regressions for a particular reference seat, with the top row showing models of March data and the bottom row showing models of June data. Within rows, each subfigure corresponds to a different mis-matched viewing direction. This mis-matched viewing direction is shown in teal in the bottom right corner of each subplot. For example, in figure 5.8, the top row shows the lightprobe models attempting to predict light exposure in seat 1 using the heading directions from seat 2, 3 and 4, respectively. In all cases, the seat-switch models reduce the r^2 score of the models presented in figure 5.7. Models predicting light exposure in seat 2 using the heading directions from other seats performed the best of the seat-switch models, and models for seats 4 and 5 performed the worst. In many of the models, particularly in June, the data that seem to disrupt the correlation the most are from the morning samples (shown in crosses). This is likely due to the lower sun angle leading to higher directionality, and often specular highlights at these times. These specular reflections are intense and highly directional and may only be captured by one of the lightprobe models.

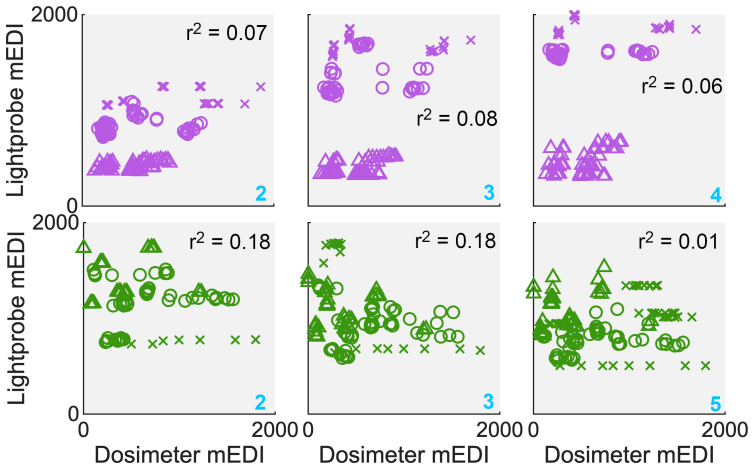

Figure 5.8: Seat switch models: seat 1. Within each figure, each row corresponds to the month, with top rows showing March data and bottom rows showing June data. The seat location from which the heading direction is derived for each model is shown in teal in the bottom right corner of each subplot. Annotations show the r^2 score for each lightprobe model. Morning datapoints are plotted as crosses, afternoon datapoints are plotted as rings and evening datapoints are plotted as triangles.

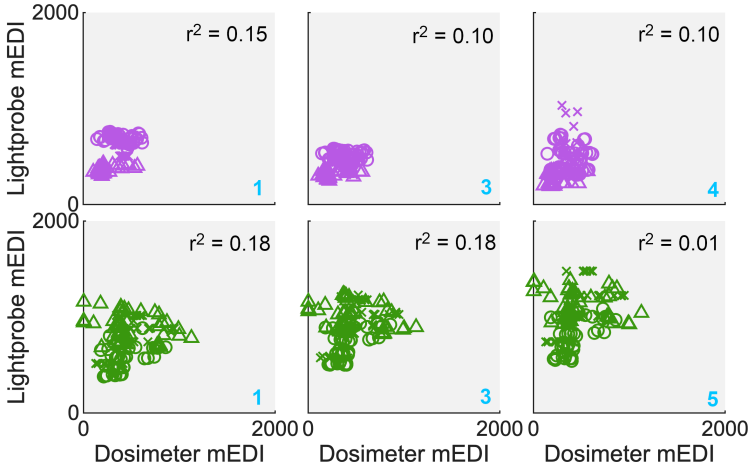

Figure 5.9: Seat switch models: seat 2. Within each figure, each row corresponds to the month, with top rows showing March data and bottom rows showing June data. The seat location from which the heading direction is derived for each model is shown in teal in the bottom right corner of each subplot. Annotations show the r^2 score for each lightprobe model. Morning datapoints are plotted as crosses, afternoon datapoints are plotted as rings and evening datapoints are plotted as triangles.

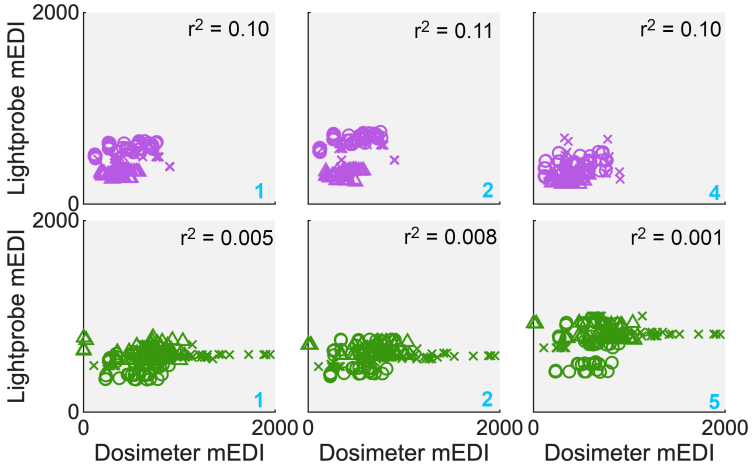

Figure 5.10: Seat switch models: seat 3. Within each figure, each row corresponds to the month, with top rows showing March data and bottom rows showing June data. The seat location from which the heading direction is derived for each model is shown in teal in the bottom right corner of each subplot. Annotations show the r^2 score for each lightprobe model. Morning datapoints are plotted as crosses, afternoon datapoints are plotted as rings and evening datapoints are plotted as triangles.

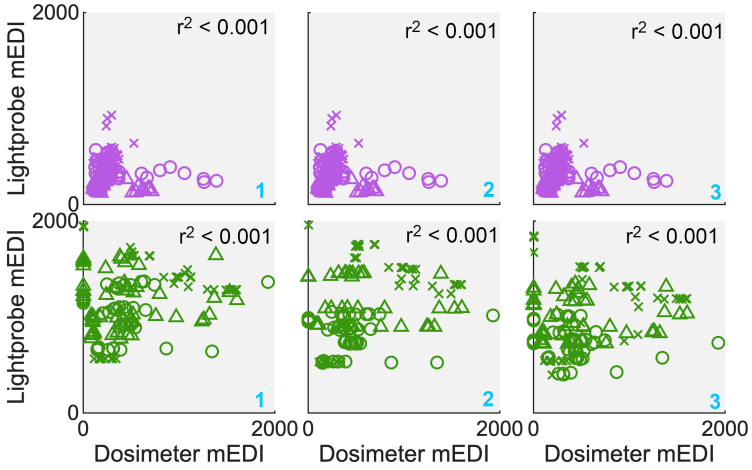

Figure 5.11: Seat switch models: seats 4 and 5. Within each figure, each row corresponds to the month, with top rows showing March data and bottom rows showing June data. The seat location from which the heading direction is derived for each model is shown in teal in the bottom right corner of each subplot. Annotations show the r^2 score for each lightprobe model. Morning datapoints are plotted as crosses, afternoon datapoints are plotted as rings and evening datapoints are plotted as triangles.

5.3.5 Egocentric lightprobes

Finally, we present *egocentric* lightprobes that represent the mean light exposure for different times of day using the dosimetry data. Figure 5.12 shows four examples of these egocentric lightprobes. The top left example shows the egocentric lightprobe from seat 4 in the afternoon, collated over the five minute period with valid heading data closest to the lightprobe capture time. It shows fairly minimal distortions, corresponding to limited changes in gaze during this time period. By comparison, the top right example shows an egocentric lightprobe from seat 1 in the morning in June. This lightprobe is significantly distorted, corresponding to the lateral sampling shown in figure 5.5.

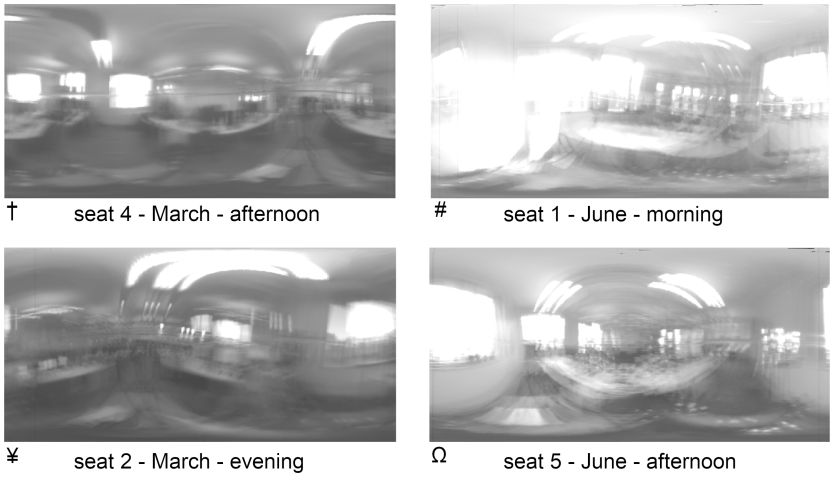

Figure 5.12: Example egocentric lightprobe images. Captions show the seat location, season and time of day the lightprobes are taken from. Symbols are used to relate these examples with the spherical harmonic data below.

Figure 5.13 shows the differences in the spherical harmonic characteristics of egocentric lightprobes compared to the original allocentric lightprobes across each seat, time of day and season. Each subplot shows a different spherical harmonic metric, as explored in Chapter 4. Within each subplot, each row corresponds to a different seat and each column refers to the different seasons. Black bars correspond to morning lightprobes, red bars correspond to afternoon lightprobes and yellow bars correspond to evening lightprobes. Symbols show datapoints that correspond to the egocentric lightprobes shown in figure 5.12. This shows that for most cases, head movements increased the monopole and dipole directionality (light vector and squash tensor) values, particularly in June. This coincides with a reduction in diffuseness compared to the allocentric means. One counter example for this is the recording for seat 4, where values stay largely similar. This is due to the participant making minimal gaze movements in this time, as shown in figure 5.12i.

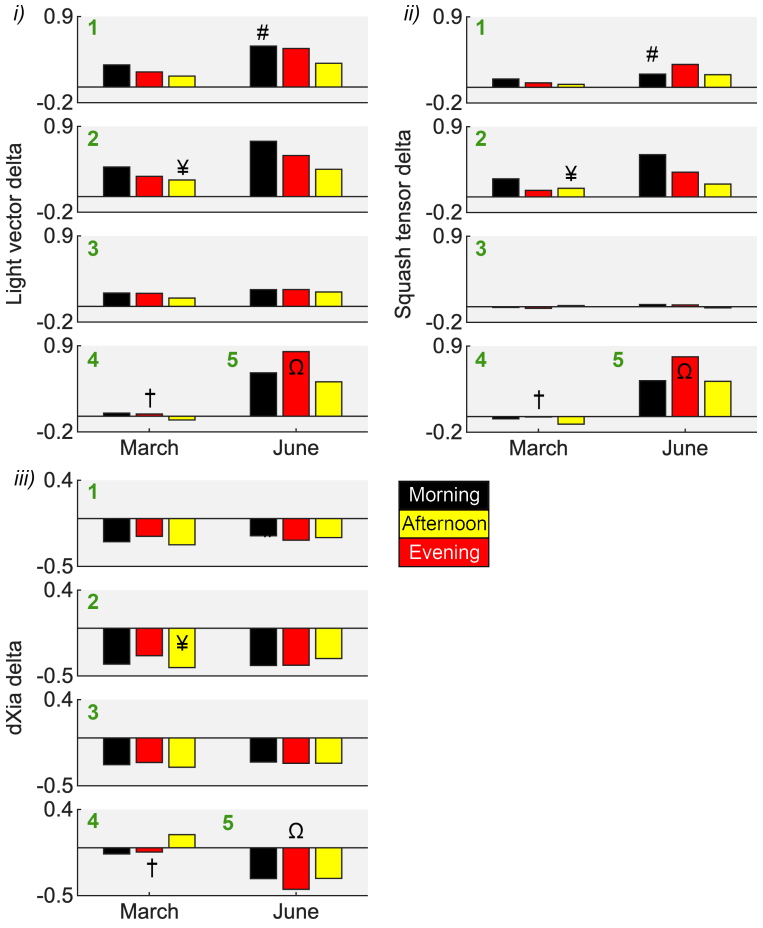

Figure 5.13: Spherical harmonic differences between egocentric and allocentric lightprobes. Each section shows a different spherical harmonic statistic. Within each section, each row shows a different seat location, apart from the bottom row which shows both seats 4 and 5. Left columns show March data and right columns show June data. Morning data are shown in black, afternoon data in red and evening data in yellow. Symbols show the bars that correspond to the example lightprobes in figure 5.12. Positive values mean that the egocentric value was higher than the allocentric mean.

5.4 Discussion

This chapter sought to combine wearable light dosimetry device measurement with head position tracking and observer-agnostic measurements to characterise the way in which light behaviours affect light exposure in the built environment. We present custom light dosimetry devices that allow for the simultaneous measurement of both light exposure and compass heading direction. As far as we are aware at the time of writing, these devices are the first of their kind to combine a high spectral resolution light dosimeter with a full 9-axis IMU. We deployed these devices in two week-long field studies where four participants wore the devices while at work.

First, figure 5.3 shows season and heading direction differences in light exposure in our example office environment, which support findings of previous literature [245,294]. Namely, that light exposure was generally highest in the morning in June, while exposure profiles in March were generally flatter throughout the day. Participants sat at north-north-westerly-facing seats tended to have flatter light exposure profiles too, with south-south-easterly-facing seats being exposed to more intense light in the morning. As explored in Chapter 4, this is likely the result of direct sunlight from the east-facing windows which are visible to this subset of participants. Utilising the method for predicting non-visual effects of light established in Chapter 2, figure 5.4 shows that circadian effects were high for all locations throughout the working day. Analysis of the heading directions across the data collection periods, shown in figure 5.5, shows that individual's gaze behaviours are remarkably similar across time periods, with similar characteristic shapes being shown by participants 1, 2 and 3 in both March and June. This likely reflects individual differences in seating posture, chair height and display setup that remain constant for participants. However, further analysis reveals that participants spend more time in different parts of these distributions, showing differences in average elevation across the two seasons. Overall, these results reinforce location, time and season differences identified in Chapters 3 and 4.

To determine whether the temporal sampling mis-match has a significant effect on light exposure data we compared the light exposure predicted by lightprobes recording at three times each day from behind the participants' seating positions to their actual exposure by sampling irradiance values from the lightprobes using the heading data. We found mixed evidence for the efficacy of these models. r^2 scores showing the predictive value of dosimetry data on data generated using these allocentric methods were generally modest or low, suggesting that the models are not able to fully capture light exposure when accounting for heading behaviours.

There are a number of potential reasons for this. First, the discrepancy in temporal sampling between the dosimetry and the lightprobes was vast, with the dosimetry sample recording every 15 seconds, and the lightprobes only being recorded three times per day. Dosimetry data shown in figure 5.3 show that changes in light exposure changes fluidly throughout the day, and so even though data for the regression models were limited to only 10 minutes around the lightprobe recording time, it is possible in many cases that this time variation is sufficiently large to change lightprobe estimates. Another possible reason is that the mismatches in camera location and head position were sufficient to occlude key elements of the light scene. Lightprobes were recorded from behind the seating locations of the participants, meaning that there is an inherant mis-match in the position of the cosine filter used to generate the irradiance estimates. This shift in the cosine filter may be sufficient to change the lightprobe irradiance estimates. More significantly, the difference in the origin from which recordings are made may actually mean that light elements that reach the participants do not reach the camera and vice-versa. As discussed, free-viewing human observers avoid uncomfortable light fields, and it is possible that the position of the allocentric model is one that would not naturally be occupied by a human observer, given the option to view a scene from elsewhere.

We additionally quantified the effect of heading directions unique to each seating position during data collection and, using a range of seat-switch models, showed that mis-matching the heading direction and light exposure profile leads to a reduction in the correlation strength between the real data and the model. This demonstrated that heading direction does play a key role in light exposure in the built environment, and a change in viewing behaviour is sufficient to meaningfully alter the light to which an individual is exposed, even if the original lightprobe estimates did not have a strong relationship with the dosimetry data. Often the datapoints that showed the least concordance between models were from the morning time period, suggesting the directionality of light at this time enhances the importance of heading directions. Overall, this collection of models shows that there is a complex interaction between the complexity of a particular lighting environment and the extent to which static, observer-agnostic measurements can predict individualised light exposure. While observer-agnostic measurements may be appropriate for estimating light exposure in more controlled environments, they may not be suitable for predicting light exposure in complex, mixed-illuminant scenes that are subject to dynamic daylight illumination.

Additionally, we present egocentric lightprobes, which provide an analytic framework for

considering the observer-specific exposure at different times of day by combining heading direction for a short period of time with lightprobe data. These egocentric lightprobes corroborate the importance of sampling behaviours' ability to alter the experience of a lighting environment by showing different spherical harmonic characteristics to the original observer-agnostic lightprobes collected in Chapter 3. They also showed a seasonal dependants, with the effects of heading movements being larger in June than in March, likely due to the overall intensity increase.

In its present form, this chapter offers a first look into the analytic potential of combining egocentric and allocentric measurements, as well as the capabilities of light dosimetry devices with full 9-axis IMU logging. However there is significant work to be done to fully realise the potential demonstrated here. As with any light exposure experiment, these measurements should be expanded to a range of environments, both in similar office environments but also a wider range of indoor and outdoor environments. It is important to acknowledge that the results presented here may not generalise to all mixed illuminant environments as their geometry, compass orientation and global location all have a significant bearing of the available light. This chapter provides a framework on which future experiments can expand. An additional limitation on the present work is the applicability of the static measurements to the dosimetry data. Lightprobes were collected at roughly seated head-height ($\sim$1.20 m from the floor) in order to roughly align the equator of the lightprobe with a '0 pitch' recording on the IMUs. However, as previously mentioned, there are likely differences in posture and chair height that differ both between participants and across time. The lightprobes were collected from behind the seated location for each desk, so angular recordings do not perfectly align with positions within the lightprobe positions, as they would if the lightprobe was taken from exactly where the participants' heads were. This is a compromise central to any attempt to combine allocentric and egocentric light recordings, as any static recording relies on some set of heuristics about the way in which light is sampled, often ignoring the complexities of sampling behaviour that we have attempted to capture here. These differences, though small, could be the difference between seeing or missing a specular reflection from a glossy surface, and thus it would be useful for the effect of these differences to be quantified in future studies.

Another important consideration here is the interaction between heading direction and fixation. Fixation location is bounded by head position but may also significantly alter light exposure. In a case study of a patient who lacked the ability to saccade, head movements were shown to behave similarly to saccadic eye movements, which raises the suggestion that both may

be controlled by overlapping neural mechanisms [295, 296]. More broadly, the distribution of saccadic eye movement sizes is predicted by natural scene statistics and receptive field sizes [297], which suggests that saccadic eye movements do play an important role in the way natural scene statistics are sampled in free viewing conditions [298]. However, it is also likely that, particularly when considering the low-angular-frequency components of light, head movements present a reasonable approximation of the central tendency of looking behaviours.

Future work may utilise additional cameras with 3D pose estimation techniques to track participants wearing dosimetry devices. Models such as SMPL [299] can reproduce fine details in pose structure and could therefore be used to calculate the offset between the lightprobe recording location and the participants' heads, as well as directly calculating head orientation without the need for IMU devices. This offset could then be corrected to potentially increase the performance of lightprobe models of individualised dosimetry recordings. It could further extend characterisation of light exposure by enabling the mapping of light exposure while participants are moving throughout a pre-set space, rather than requiring all measurements be made in one static location, as we present here. Some head-worn devices are also able to combine scene cameras with eye-tracking, such as the Pupil Labs Neon [300]. While this device lacks the wavelength specificity of other dosimetry devices, its ability to track eye movements while recording corneal-plane light exposure allows it to incorporate an additional dimension into quantifying individualised light exposure not captured by present dosimetry methods. Finally, another consideration important to individualised light exposure is idiomatic characteristics in the light that hits the retina. There are many factors that affect this, but two that have been quantified but are not often incorporated into dosimetry calculations are lens changes due to age [27], and head shape [301]. Chapter 2 has already quantified changes in light exposure due to lens yellowing throughout the lifetime. To properly account for this, dosimeters with higher wavelength resolutions are required, but this would enable analysis of changes to gaze behaviours and light exposure throughout the lifetime. Variation in face shape has also been shown to reduce the amount of light reaching the retina. Future individualised light exposure models may also take into account whether facial structure has an effect on light-seeking or avoiding behaviours.

5.5 Conclusions

This chapter presented custom dosimetry devices that also log heading direction. In an experiment using them, we showed the effect of season, time of day and seating direction on egocentric light exposure. We showed mixed performance of lightprobe models, demonstrating the effect of proximity to windows on the volatility of lighting environments in typical office spaces. We also show *egocentric* lightprobes which demonstrate the way that gaze behaviours affect time-averaged directional light exposure throughout the day. Finally, we discuss future directions for the individualised measurement of light exposure.

5.6 Acknowledgements

The contents of this chapter are entirely that of the author, with supervision by Hannah E. Smithson and Manuel Spitschan. Thank you to the participants who took part in this experiment.

Chapter 6

Hemispheric differences in human chromatic discrimination thresholds

Contents

6.1 Introduction 160

6.2 Methods 162

6.2.1 Hyperspectral lightprobes 162

6.2.2 Lightprobe analysis 162

6.2.3 Participants 163

6.2.4 Equipment 163

6.2.5 Heterochromatic flicker photometry 163

6.2.6 Procedure overview 164

6.2.7 Adaptation 165

6.2.8 Stimuli 166

6.2.9 Experiment blocks 166

6.2.10 Staircases 167

6.2.11 Psychophysics analysis 167

6.3 Results 167

6.3.1 Hyperspectral lightprobes 167

6.3.2 Human Chromatic Discrimination Thresholds 170

6.4 Discussion 171

6.4.1 Timescale of adaptation 172

6.4.2 Other variables 173

6.4.3 Biased sampling of the chromatic environment 174

6.5 Acknowledgements 175

6.1 Introduction

Humans' colour vision is malleable and adapts to the chromatic environment [80, 83, 199]. Thus far, this thesis has considered the non-visual effects of exposure to light in typical lighting environments, however an important physiological impact of our spectral diet is also the chromatic adaptation that we experience as the light around us changes. This adaptation can affect a range of properties of colour vision, from colour discrimination to colour appearance. The ability to discriminate between similar colours is highly important for adaptive behaviour, and humans have finely-tuned perceptual systems for making colour judgements. For example, being able to detect slight colour changes in food that has spoiled or judging subtle chromatic cues from clouds to anticipate whether one will need a coat or not. Many experiments have previously explored colour discrimination thresholds for coloured patches as well as 3D rendered objects [235, 302]. Discrimination thresholds have long been shown to be different depending on the part of colour space in which they are measured, and therefore how the participant is adapted [56, 303, 304]. The more recent of these experiments have demonstrated that realistically-rendered chromatic environments can also affect chromatic discrimination thresholds. For example, Morimoto and Smithson [235] demonstrate that chromatic discrimination thresholds for sets of similarly coloured objects are tuned to the distribution of the lighting environment those objects are found in, however this tuning was more pronounced in lighting environments that were typical of everyday scenes (falling along the daylight locus). This suggests that there may be some kind of long-term adaptation to the daylight locus that occurs, increasing the sensitivity of our discrimination mechanisms to colours that deviate from this norm.

Discrimination threshold changes that result from adaptation to environmental illumination are important as they have significant implications for colour appearance and colour constancy [85]. This is the name given to the phenomenon wherein surface colour appears to stay constant in different lighting environments, thus, the illuminant is 'discounted' from estimations of surface colour. This is a useful adaptive behaviour, as it allows perceptual representations of objects to more closely match their surface properties, which may confer important information about them. In the proximal retinal image, this is confounded by lighting changes. Being able to discount the

illuminant allows humans to learn fine perceptual associations in cases where small chromatic differences are important, such as correctly identifying the ripeness of fruit at different times of day where the lighting environment might be drastically different. Colour constancy is a complex process which appears to arise from compensation at many levels of the visual system – from adaptation at the ganglion cell level to perceptual adjustments at the cortical level [305]. There also appear to be many heuristics that the visual system may use to 'discount the illuminant'. One candidate cue is the 'gray world hypothesis', which suggests that global deviations from an average gray perceived colour are representative of biased illumination, which can be corrected by assuming that, if the chromatic average was set to gray, object colour can be retrieved. Another is the 'brightest is white hypothesis' which assumes that, given the multiplicative interaction between surfaces and lights, the brightest-appearing part of a scene best represents the illuminant, as white surfaces reflect the broadest range of wavelengths [86, 87, 90, 91]. Clearly, no simplistic heuristic in this vein is sufficient to account for the complexity of object colour in real-world scenes, where we often experience multiple illumination sources and surface types. Researchers have proposed Bayesian models that combine many scene heuristics, which have demonstrated that we likely use many clues, weighted based on the apparent reliability of their informativity [88, 89]. Real-world environments with multiple light sources have demonstrated that there are stable individual differences in colour constancy, even in the colour-normal population, as shown in the notable viral example *#TheDress* [306–309].

Given the effect of chromatic adaptation on discrimination thresholds, and the complexity of environmental illumination as discussed throughout this thesis, it is conceivable that directional regularities within those environments yield different adaptations on different parts of the visual field. Perhaps the strongest example of this would be differences between the upper and lower visual hemispheres. It seems that the upper hemisphere is dominated by daylight-aligned light, while the lower hemisphere is dominated by light being reflected by surfaces whose spectra more frequently deviate from the daylight axis [201]. These previous findings also underline a number of other differences, such as the mean chromaticity of the two hemifields, and an order of magnitude difference in luminance, with the upper hemifield being much more intense. As such, the upper visual hemifield may be adapted to daylight locus-aligned light, leading to more elongated discrimination ellipses that are tuned in the direction of the daylight locus, while the lower visual hemifield is adapted to a wider range of colours and thus appear less directionally tuned. Hemispheric differences in retinal activity have been observed in mouse

retina; such differences may be measurable in human discrimination thresholds too [310, 311]. This chapter details an experiment in which chromatic discrimination thresholds were measured while participants were adapted to chromatic distributions matching the upper and lower regions of typical lighting environments to investigate whether these conditions yield differences in chromatic discrimination thresholds across corresponding retinal areas.

6.2 Methods

6.2.1 Hyperspectral lightprobes

In order to test the intuition that this directional distinction exists, I analysed an existing dataset of hyperspectral lightprobes [201]. These lightprobes have many features typical to an outdoor European environment: generally having visible blue skies with some cloud cover, and featuring a range of terrain such as tarmac, grass and buildings in the lower half of the images. Two lightprobes were collected from indoors locations, so these featured ceilings and windows that provided access to natural light. To analyse the chromatic properties of these images, they were converted to LMS images using the CIE standard observer cone fundamentals [27]. The images were then converted to the MacLeod-Boynton (MB) colourspace [57], which gives each pixel an L/(L+M) value and an S/(L+M) value.

6.2.2 Lightprobe analysis

The hyperspectral lightprobe data were analysed in MacLeod-Boynton space by fitting bivariate Gaussian ellipses to the chromatic distribution of pixels in the 5° slice centring around 10° above and below the horizon. The area, eccentricity, orientation, mean chromaticity, x range and y range were all compared using ANOVAs. The eccentricity of the fitted ellipses was calculated using the standard equation shown in equation 6.1. In this equation, e corresponds to the ellipse's eccentricity, a corresponds to the semi-major axis of the ellipse and b corresponds to the semi-minor axis. Orientation is calculated as the angle of the semi-major axis relative to the positive x axis measured counter-clockwise.

$$e = \frac{\sqrt{a^2 - b^2}}{a} \tag{6.1}$$

6.2.3 Participants

Five individuals completed the experiment, aged between 24 and 40. Before participating in the experiment, participants completed visual acuity and HRR colour vision tests to check for colour deficiencies. Any participants with colour deficiencies were to be excluded from the study. Ethical Approval for this study was granted by the Medical Sciences Interdisciplinary Research Ethics Committee under reference code R77333/RE003.

6.2.4 Equipment

The experiment was run using MATLAB with Psychtoolbox3 [312]. Stimuli were displayed on a Cambridge Research Systems Display++ LCD display using the CRS ViSaGe stimulus generation mode. Participants' gaze was tracked using an Eyelink 1000. Participants were sat with their head resting on a chinrest for the duration of the study, 110 cm from the display. Before data collection , the Display++ was gamma-corrected using a Jadak PR670 spectroradiometer. After linearisation, a colorimetric calibration was carried out to find the transform between the RGB primaries of the display and the predicted LMS for a standard observer. A bespoke conversion matrix was obtained for each participant based on their age at the time of participation, accounting for individual differences in cone sensitivity spectra [27].

6.2.5 Heterochromatic flicker photometry

In the main experiment, discrimination thresholds in eight chromatic directions are measured. Before the main experiment began, participants completed a short heterochromatic flicker photometry (HFP) procedure in order to ensure that the stimuli across these eight directions were functionally isoluminant in the stimulus presentation locations for each observer [313]. For the HFP procedure, participants sat in a darkened room in front of the Display++. The stimuli used were gaussian blobs, subtending 5° of visual angle that were centred 10° above or below fixation in the middle of the display. Stimuli were set to have the same photopic luminance as a 100 cd/m^2 D65 background. For each HFP stimulus, two colours were cycled at 50Hz. The chosen colour pairs lay on the same axis in a scaled Macleod-Boynton space, passing through the D65 used for the background. The four colour axes (L/(L+M), S/(L+M), (L+S)/(L+M) and (L-S)/(L+M)), corresponding to four angle pairs (0° and 180°, 45° and 225°, 90° and 270°, and 135° and 315°), match the stimuli used later in the main experiment, and are schematised in figure 6.1. In this figure, the outer ring of stimuli represent those that were used in the HFP. The

chosen colour pairs were equidistant on opposite sides of D65 in Macleod-Boynton colourspace (e.g., 1 step from D65 at 0° and 1 step from D65 at 180°, where a step is an arbitrary distance in the unitless Macleod-Boynton colourspace). The Macleod-Boynton chromaticity space was scaled such that the the L/(L+M) axis was scaled by 0.0071 and the S/(L+M) axis was scaled by 0.1325. Both axes were scaled with respect to the D65 coordinate, meaning that it stayed in the same location. This equates the discrimination thresholds along the two axes measured in a separate experiment, meaning that the colour gamut was sampled in units of discrimination threshold across the two axes. These scale factors were derived from the results of previous unpublished psychophysical studies [314]. Participants were asked to fixate on the cross while adjusting the luminance of the two stimuli at the same time using the arrow keys until any perception of flickering was minimised (key presses increased the luminance of one stimulus while decreasing the luminance of the other. Once the participant felt they had reached equiluminance, these luminance values were saved. Each judgement was repeated three times. Each stimulus presentation was initialised randomly to encourage participants to explore the extent of the range of stimulus weightings. The resultant luminance values for each of the stimuli were used to calibrate the stimuli for each participant throughout the experiment that followed. An Eyelink 1000 eyetracking system was used to detect the participants' gaze and ensure that they were fixated on the middle of the screen throughout the HFP procedure.

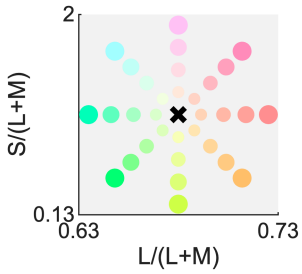

Figure 6.1: Schematic of stimuli colour sampling. Macleod-Boynton chromaticity diagram showing the chromatic coordinates from which stimuli were sampled. Points are pseudo-coloured, and their size corresponds to their saturation. The black cross denotes D65. Points are plotted in scaled Macleod-Boynton space.

6.2.6 Procedure overview

After HFP was completed, the main section of the experiment consisted of four experimental blocks, each of which lasted around an hour. Each block began with the re-calibration of the Eyelink, a 180-second adaptation period and then a series of interleaved 2-alternative forced-choice trials to measure chromatic discrimination thresholds. In total, we measured thresholds for eight

chromatic directions at two locations, giving us 16 conditions split evenly across the four blocks. Figure 6.2 shows a schematic example of two trials one in the upper hemifield, and one in the lower hemifield, further details about each component are given below.

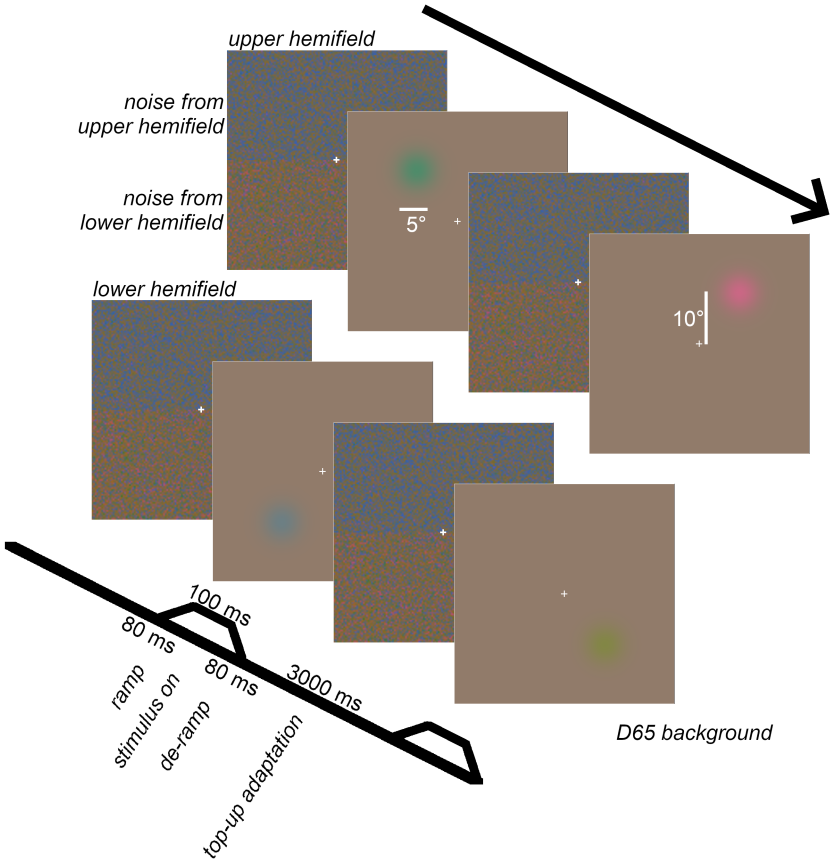

Figure 6.2: Schematic of the chromatic discrimination experiment. Two example trial sequences are shown, one from a block where the stimuli are shown in the upper hemifield and one where stimuli are shown in the lower hemifield. Images are pseudo-coloured. Stimuli, background and adaptation fields appeared isoluminant. Bottom bar shows the timescale for events during each trial.

6.2.7 Adaptation

Each block started with a 180-second adaptation period using chromatic noise. The chromatic noise was defined using the probability density function of the chromaticity distribution of pixels across the $+10° \pm 2.5°$ and $-10° \pm 2.5°$ regions of the hyperspectral lightprobes such that the adaptation stimulus in the top and bottom halves of the screen matched the chromatic distribution of these regions, respectively. Luminance values were randomly sampled from a Gaussian distribution with a mean of 100 cd/m^2 and a standard deviation of $5cd/m^2$, and *not* taken from the distributions from the lightprobes. Panels of noise were cycled at 50 Hz, and noise pixels were downsampled such that blocks of four pixels shared the same value. Stimuli for each trial were shown on a flat D65 background with no noise, so the adapting stimuli ramped

into the D65 noise at the end of each trial, and ramped back in for the next round of adaptation.

6.2.8 Stimuli

The stimuli for each discrimination trial appeared as Gaussian blobs subtending 5° of visual angle, isoluminant with the D65 background, presented at 100 cd/m^2 and scaled per participant according to the HFP setting. In each trial, the stimulus was presented in a 2-alternative forced-choice arrangement, appearing randomly on either the left or the right of fixation, with the opposing side remaining as the D65 background. Participants were asked to identify on which side the stimulus appeared and indicate this using the arrow keys – responding with a left-right judgment.

A total of 16 conditions were tested, with eight chromatic directions appearing in two locations. The colours used during the experiment were geometrically calculated using a scaled Macleod-Boynton colour diagram described in the HFP section above. We investigated chromaticities along eight equally spaced hue angles – 0, 45, 90, 135, 180, 225, 270 and 315 degrees. Stimulus constants along these hue angles were determined by the staircase procedure. For the two location condition sets, stimuli appeared either 10° above or 10° below fixation. Despite this vertical difference across blocks, participants only reported the horizontal position of the stimuli in each trial. Figure 6.1 shows an example schematic for how the chromaticities were sampled.

Stimuli were linearly ramped from the background across 80 ms, reaching the desired saturation for 100 ms before linearly de-ramping into the background for 80 ms. Once the stimuli disappeared, the chromatic noise stimulus described earlier appeared for three seconds, affording the participant 'top-up' adaptation throughout the experiment. The participant responded during this time. Participants then heard either a high-pitched or a low-pitched beep indicating whether they had correctly responded on that trial or not. The next trial then began.

6.2.9 Experiment blocks

The sessions were divided into four blocks, based on on- and off-axis colours, as well as stimulus location. Thus, in a counter-balanced order, the stimuli tested were on-axis in the upper hemifield, off-axis in the lower hemifield, off-axis in the upper hemifield, on-axis in the lower hemifield. In each block, the location of the stimulus was predictable from trial to trial, so to ensure that the participants remained fixated and were not looking directly at the stimulus presentation area, the Eyelink 1000 was used to delay the onset of the next trial until the participant had returned

to fixation for 500 ms. Fixation was recorded for the duration of each trial. At the end of each block, participants were encouraged to get up from the experimental setup and have a more extended break before returning for the next block.

6.2.10 Staircases

For each stimulus hue angle and position, a Palamedes staircase function [315] was used to estimate the discrimination threshold. Correct responses would yield a less saturated stimulus (one closer in hue to D65) on the subsequent trial of that staircase. The block would continue until all of the four interleaved staircases had converged, however the number of trials per colour-location pairing was limited to 50 in each session.

6.2.11 Psychophysics analysis

The psychophysical data were analysed by plotting the threshold estimates for each hue angle on a scaled MB diagram separately for upper and lower visual fields. Then, after finding the averaged thresholds for each participant, ellipses were fitted to the eight hue angles for each location, giving a chromatic discrimination ellipse for each participant's upper and lower hemispheres. Ellipses were also fitted to the data when the thresholds were averaged across participants. Similar to the bivariate Gaussian ellipses fitted to the hyperspectral lightprobe chromaticity distributions, threshold ellipse statistics were compared using ANOVAs. We analysed area, eccentricity, orientation, x range and y range.

6.3 Results

6.3.1 Hyperspectral lightprobes

Figure 6.3 and 6.4 show the chromatic distribution of pixels in the hyperspectral lightprobes. Figure 6.3 shows the distribution in the upper hemifield slice, and figure 6.4 shows the distribution in the lower hemifield slice. The distributions are plotted on Macleod-Boynton chromaticity diagrams [57] scaled according to the scale factors mentioned previously. The density of the distribution is shown by the bivariate Gaussian ellipse shown in each plot and the blue cross shows the mean of the distribution. The blue dashed line shows the black body locus and the red cross shows D65.

Statistical analysis was conducted on the ellipses fitted to each of the distributions, comparing

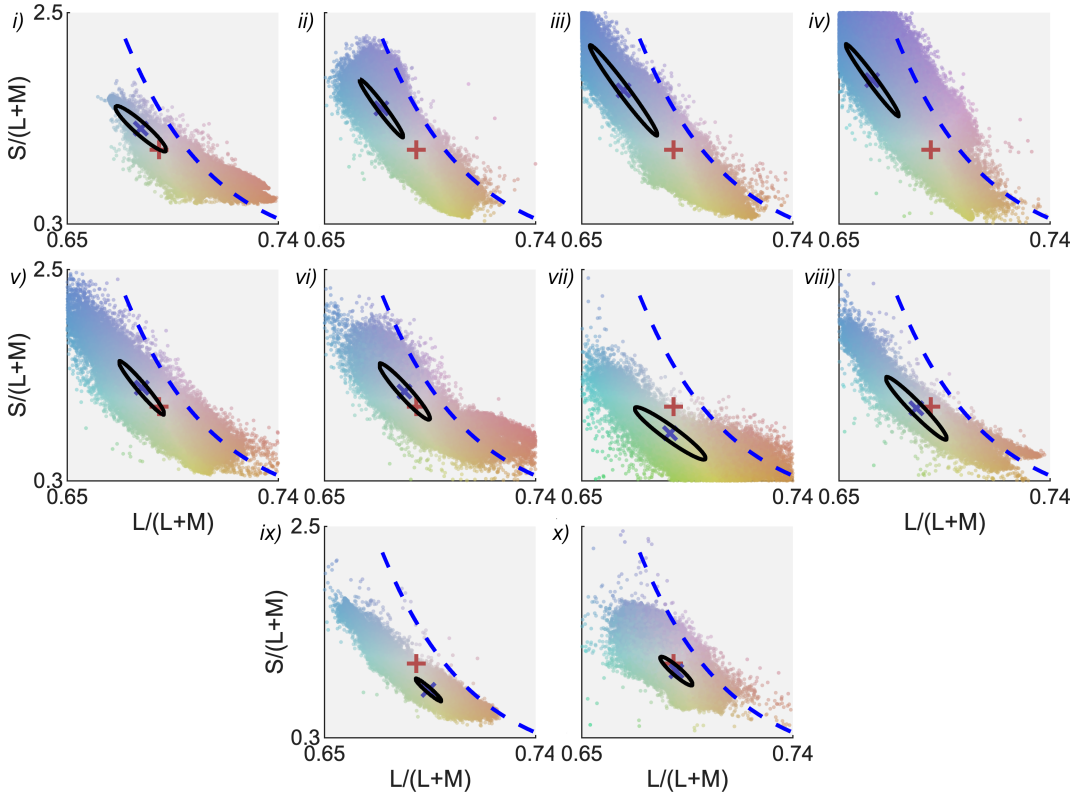

Figure 6.3: Hyperspectral lightprobe upper hemifield slice distributions. Distributions are plotted in unscaled MB space. Black body locus is plotted in blue. D65 is shown by the red cross. Distribution means are shown by the blue cross. Bivariate Gaussian ellipses are plotted in black.

the 5° thick slices that appear 10° eccentric in the upper and lower hemifields. Comparisons are shown in figure 6.5. No significant statistical difference was found between the areas of the ellipses ($F(1, 9) < 0.0001$, $p = .9972$), nor the orientation ($F(1, 9) = 0.2236$, $p = .6420$). A significant difference in eccentricity was found, with upper hemifield slice distributions being significantly more eccentric ($F(1, 9) = 8.204$, $p = .0103$). There was a significant difference in mean chromaticity ($F(1,1) = 7.2506$, $p = .0053$), with the lower hemifield slices featuring higher L/(L+M) and lower S/(L+M) values. There was no significant difference in the x range of the two distributions ($F(1, 9) = 1.8941$, $p = .1856$), but there was a significant difference between the y ranges ($F(1, 9) = 5.7678$, $p = .0273$) with the upper hemifield slices having a larger y range than the lower hemifield slices.

Differences between the slices in the upper and lower hemifields of the lightprobes are characterised by the upper region distributions featuring larger variation along the negative diagonal in MB space - which corresponds to the black body locus - leading to more eccentric ellipses. The centre of the distributions is also different, with upper hemifield distributions being pushed more towards the blue-appearing end of the black body locus, while the lower hemifield

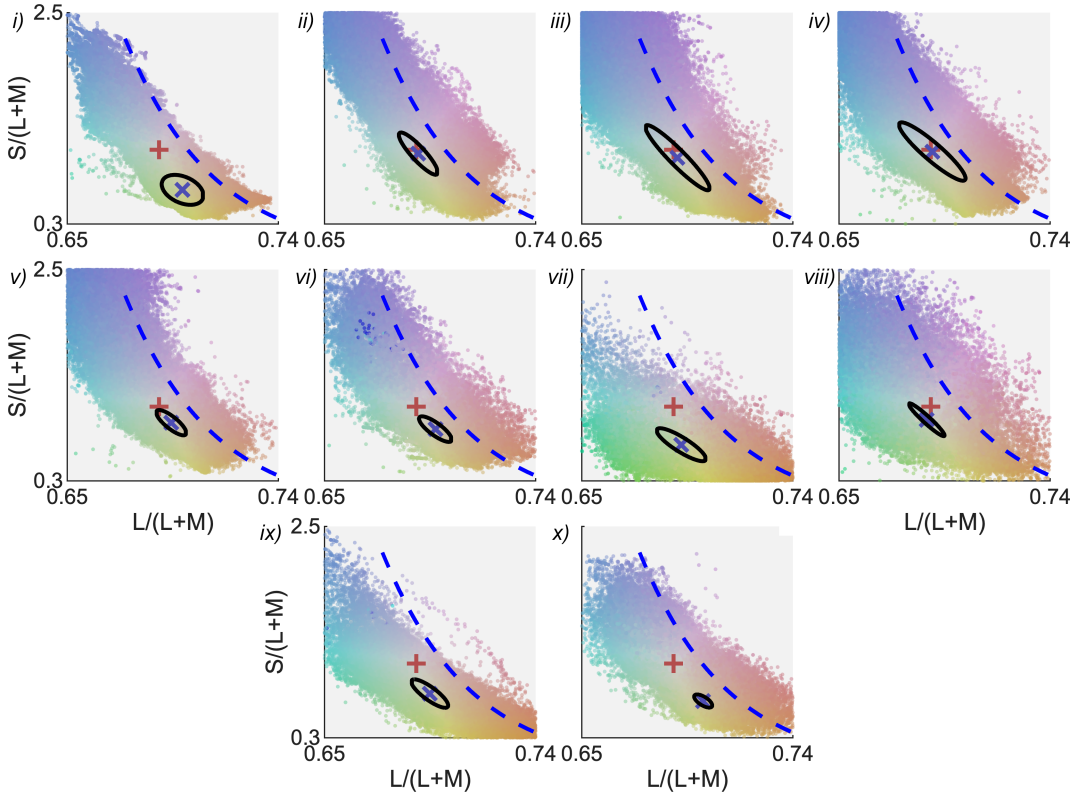

Figure 6.4: Hyperspectral lightprobe lower hemifield slice distributions. Distributions are plotted in unscaled MB space. Black body locus is plotted in blue. D65 is shown by the red cross. Distribution means are shown by the blue cross. Bivariate Gaussian ellipses are plotted in black.

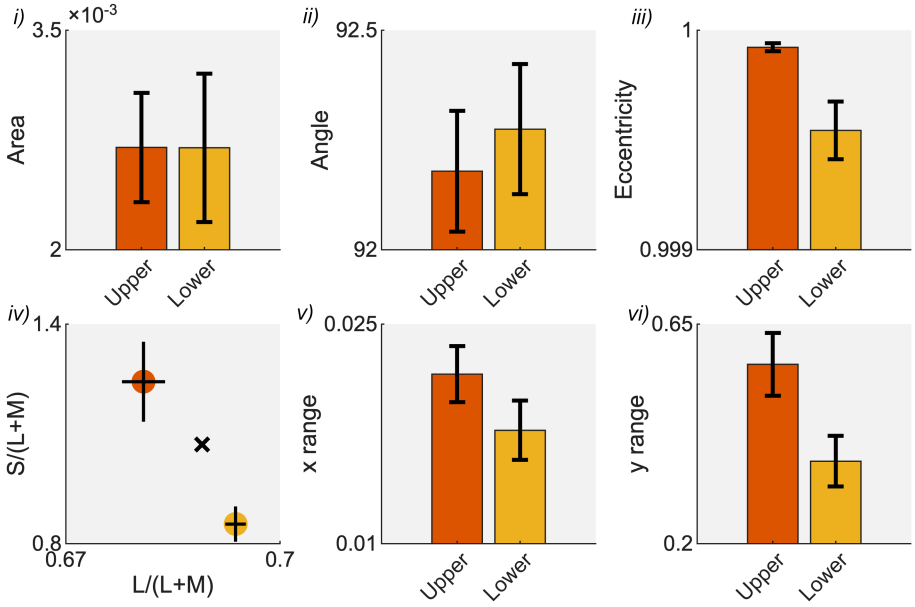

Figure 6.5: Bivariate Gaussian ellipse comparisons. *i)* area comparisons. *ii)* orientation comparisons. *iii)* eccentricity comparisons. *iv)* mean chromaticity comparisons. *v)* L/(L+M) range comparisons. *vi)* S/(L+M) range comparisons.

distributions lie in the warmer end.

Additionally, we analysed the distribution of luminance across the two hemifields. Figure 6.6 shows these distributions. The mean luminance of pixels in the upper hemisphere is 1279 cd/m^2,

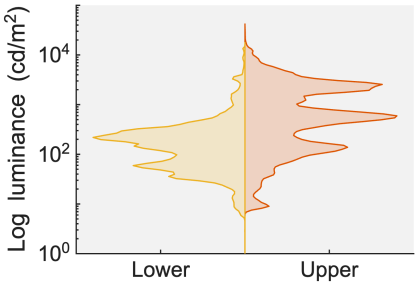

Figure 6.6: Luminance distribution comparisons. Violin plot shows the density of log luminance values for upper and lower hemifields.

while the mean luminance of pixels in the lower hemisphere is 305 cd/m^2. Overall, there is a much wider range of luminances in the upper hemisphere.

6.3.2 Human Chromatic Discrimination Thresholds

Figure 6.7 shows the chromatic discrimination thresholds for each participant, as well as the mean for all participants. These thresholds are plotted in scaled MacLeod-Boynton space. Red and yellow data points and ellipses come from upper and lower hemifield recordings respectively. Each data point corresponds to a threshold for a given hue direction, averaged across the five sessions that each participant completed. As such, better performance means a data point plotted closer to the white point, plotted in black. Figures 6.7i - 6.7v show each participant's discrimination ellipses, and figure 6.7vi shows the overall mean for all the data collected. The plots show individual differences between each of the participants, but an overall pattern of smaller thresholds for the lower hemisphere than the upper hemisphere.

Figure 6.8 shows the comparison of ellipse statistics for the chromatic discrimination threshold ellipses. Lower hemifield threshold ellipses are shown to have significantly smaller areas than the upper hemifield threshold ellipses ($F(1,4) = 5.9545$, $p = .0183$). There is no significant difference between the orientation of the discrimination ellipses in either hemifield ($F(1,4) = 0.0154$, $p = .9018$). There is also no statistically significant difference between the eccentricities of the threshold ellipses ($F(1,4) = 2.6705$, $p = .1085$). Lower hemifield threshold ellipses have a significantly smaller x range than the upper hemifield ellipses ($F(1,4) = 12.021$, $p = .0011$), but there are no significant differences between the y ranges of the threshold ellipses ($F(1,4) = 0.9716$, $p = .3290$).

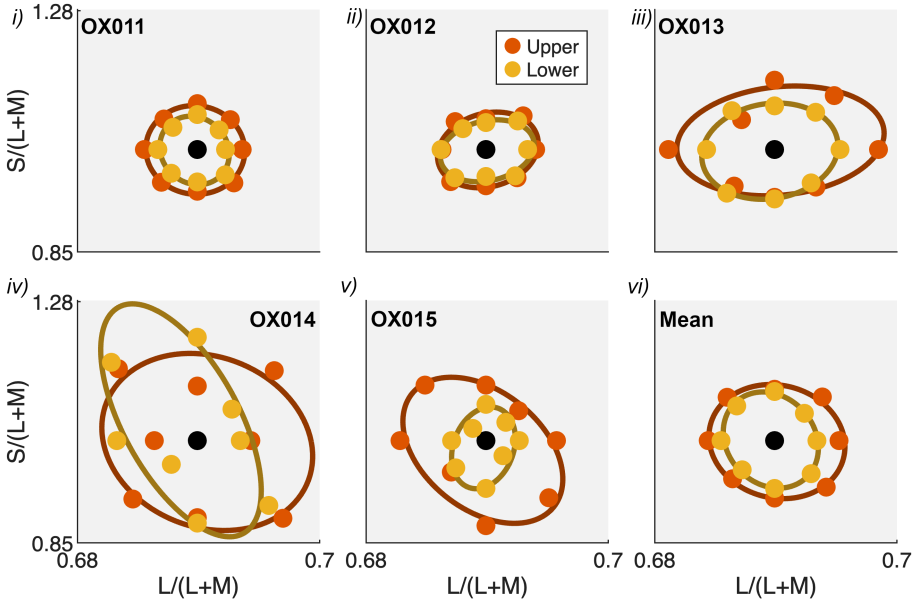

Figure 6.7: Chromatic discrimination thresholds. Red and yellow points show mean thresholds in each chromatic direction, and darkened bands show fitted ellipses. Black dot shows D65. *i)-v)* show individual participant data, and *vi)* shows mean across all participants.

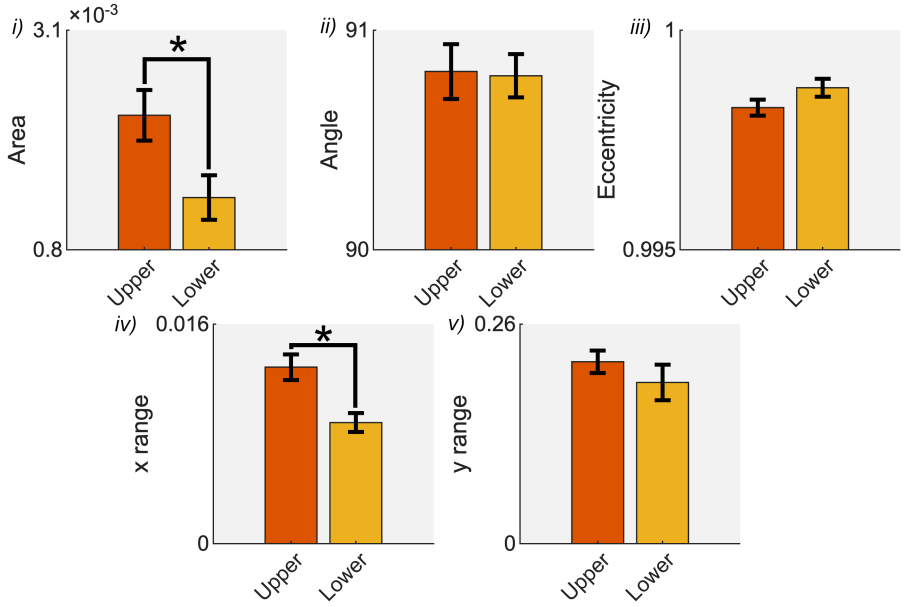

Figure 6.8: Discrimination ellipse comparisons. *i)* area comparisons. *ii)* orientation comparisons. *iii)* eccentricity comparisons. *iv)* L/(L+M) range comparisons. *v)* S/(L+M) range comparisons. Asterisks show significant differences.

6.4 Discussion

The present study aimed to determine whether differences in the chromatic statistics of environmental lighting between upper and lower hemifields coincide with differences in chromatic discrimination thresholds in corresponding localised regions across the retina. Analysis of hyperspectral lightprobes found that ellipses fitted to the distribution of colours in the upper hemifields

were more eccentric and centred more towards the blue (high temperature) end of the black body locus than ellipses derived from the lower hemifield of the scenes. These differences are consistent with the "daylight" content of light from the upper regions compared to reflected light with more diverse spectral content from lower regions which supports the notion of systematic differences between the upper and lower hemifields of our chromatic environments.

We tested the hypothesis that this may lead to differences in chromatic discrimination thresholds. Indeed, there were differences in chromatic discrimination thresholds for Gaussian blob stimuli presented above and below fixation when participants were adapted to chromatic distributions that mimicked hemisphere-specific variations. Thresholds were higher in the upper hemisphere, demonstrating a reduction in chromatic sensitivity for stimuli presented above fixation. It may be that the adaptation state as a result of typical upper-hemifield light exposure, which features significantly higher variation along the daylight locus, and is much bluer on average, gives rise to poorer chromatic discrimination than the more evenly distributed chromaticities typically seen in the lower hemifield. There are two factors to consider: the difference in distributions and the difference in mean adaptation. Since discrimination ellipses showed no systematic difference in eccentricity, it may be that the difference in area of discrimination ellises was driven primarily by differences in mean adaptation. This is supported by previous work that has demonstrated increased discrimination thresholds as a result of higher S-cone adaptation [304]. It is also possible that the neural locus of adaptation was in mechanisms aligned with the cardinal axes, giving a generalised loss of sensitivity to obliquely tuned adaptation [56], rather than narrowly-tuned chromatic mechanisms that would show specificity [316].

6.4.1 Timescale of adaptation

The adaptation of the human visual system to the environment it finds itself in is a widely documented phenomenon, and chromatic discrimination thresholds are no different [83]. Several previous studies have recorded changes to the size and shape of discrimination threshold ellipses in response to different background adaptation stimuli [56, 235, 303]. However, the question of localised differences based on statistics of particular regions of the visual world seems to be more complicated. Some studies have provided supporting evidence for regular hemispheric differences [317], though others have not [318, 319]. This mixed picture is likely due to methodological differences across these kinds of experiments [320]. One particular methodological question that affects the interpretation of these results is the temporal adaptation scale used in experiments

looking to assess underlying differences in chromatic discrimination thresholds across the retina. In the formulation presented here, participants are regularly readapted to the spatially-varying chromatic stimulus. The differences we found, therefore, may reflect differences in chromatic discrimination around D65 while an observer is embedded in an environment that is similar to the chromaticity distributions we analysed. An alternative investigation may instead consider the extent to which there are longer-lasting chromatic differences that persist when making chromatic judgements in chromatically uniform environments. Previous experiments and discussions have identified multiple temporal scales over which adaptation effects seem to occur, likely corresponding to different stages of neural circuitry in the visual system [84, 321, 322]. Future experiments may consider building on the methodology used here while manipulating the period of time for which the re-adaption stimuli are displayed or perhaps removing experiment-based adaptation entirely.

6.4.2 Other variables

In addition to the question of the temporal specificity of chromatic adaptation, there are other considerations for variation in the visual field. In the example lightprobes we analysed – particularly those outside – the sky is the key region that drives chromatic differences in the upper and lower hemifields, but it is important to note that the sky is not visible in the entirety of the upper hemifield, mainly appearing in the top third of the lightprobes. The upper and lower slices tested here is a demonstrative example that these differences do exist; however, the spatial dependence of discrimination differences may be more complex based on the localised differences across the extent of the retina. It would be interesting to measure discrimination thresholds across a range of positions on the retina in order to test the ways in which these vary at a higher spatial resolution. However, extending this further than 10° this may yield greater concerns about lapses in fixation from the participants, as they will likely saccade towards the stimuli more due to the difficulty of the task. Furthermore, adding additional measurement locations will drastically increase the number of trials required to record psychophysical thresholds, which are already long and often difficult for participants to complete without fatigue. However, recent developments in applying machine learning techniques to psychophysical study propose to reduce the number of trials needed to measure discrimination thresholds across the gamut by two orders of magnitude [323]. These approaches, utilising dynamic sampling across a variable range, could enable the efficient measurement of psychophysical thresholds across a range of vertical

eccentricities to elucidate the way in which colour discrimination is distributed across the retina.

An additional difference observed across the two hemifields of hyperspectral lightprobe data not investigated in the psychophysical experiment captured here is the distribution of luminance. Figure 6.7 shows that luminance in the upper hemifield is notably higher than the lower hemifield, and this is supported by previous analysis of hyperspectral lightprobes of typical environmental illumination maps [201]. This feature seems to be much larger than the size of the chromatic differences, and importantly the intensity of backgrounds seems to play a role in colour appearance and discrimination [243, 324]. The current experiment presented stimuli and adaptation fields on an isoluminant plane in order to specifically capture chromaticity differences, however future experiments may consider adapting participants to distributions that consider both the chromatic and luminance properties of upper and lower hemifield distributions to test colour discrimination and appearance differences. In particular, research-grade high dynamic range displays make this possible [196].

6.4.3 Biased sampling of the chromatic environment

Another potential consideration for interpreting the relationship between differences in adaptation to environmental light is that the environment is not evenly sampled, and so the information in the lightprobe does not represent the chromatic environment presented to the visual system, as discussed in Chapter 5 previously. Bright lights that cause glare will cause an observer to avert their eyes and reposition their head in order to be more comfortable. Interfacing with the tactile world also means that observers are frequently looking down at laptops, books or food, for example. As such, the spatio-chromatic distribution of light visible to the visual system may not be well represented by observer-agnostic measurements like hyperspectral lightprobes – there is a difference between 'what light is there' and 'what light is seen'. Data reported in Chapter 5 showed that, when freely viewing from a desk, participants' heading directions were largely limited to a small range of vertical tilts, which would be consistent with maintaining adaptation states in different vertical bands of the retina. However, environments that encourage different kinds of sampling behaviour may lend themselves to reduced differentiation across the retina to those presented here. It is important not to conflate the precise distribution of light in each hemifield from the observer-agnostic lightprobes with actual light exposure data, though it likely concords with general regularities seen by particular parts of the retina. Continued development of wearable – and particularly corneal-plane – dosimetry devices [285, 287] including

those introduced in chapter 5, and those that record eye movements [300] may allow us to identify patterns that may arise as a result of sampling biases.

6.5 Acknowledgements

The contents of this chapter are entirely that of the author, with supervision by Hannah E. Smithson and Manuel Spitschan. Thank you to the participants who took part in this experiment.

Chapter 7

General discussion

Contents

7.1 Chapter Summaries 178

7.1.1 Chapter 2: Prediction of the non-visual effects of short-wavelength blocking optical filters. 178

7.1.2 Chapter 3: Quantifying CIE alpha-opic signals in the indoor built environment. 178

7.1.3 Chapter 4: Measurement of the directional structure of light in the built environment: single-image circadiometric lightprobes. 179

7.1.4 Chapter 5: Individualised light dosimetry paired with head direction tracking in the built environment. 180

7.1.5 Chapter 6: Hemispheric differences in chromatic discrimination thresholds. 180

7.2 Future directions 181

7.2.1 Further use of custom light measurement devices 181

7.2.2 Efficient psychophysical research using machine learning 182

7.2.3 A full model of circadian effects of light 183

7.3 Application of circadian science to lighting design 184

7.4 Final summary 187

This thesis has presented a number of new methodologies, analytical approaches and results concerning the characterisation of light to which humans are exposed in the modern world. This final chapter seeks to briefly summarise the findings presented herein, contextualise these findings in the broader non-academic context of lighting design, and look forward to the future of perceptual and circadian science.

7.1 Chapter Summaries

7.1.1 Chapter 2: Prediction of the non-visual effects of short-wavelength blocking optical filters.

Short-wavelength blocking filters have become a popular wellbeing product that promise to improve sleep by reducing melatonin suppression caused by light which activated intrinsically photosensitive retinal ganglion cells (ipRGCs). These products are significantly heterogenous, yet there are seldom precise estimations of non-visual effects ascribed to them. It is also difficult to assess these effects, as they depend on the environment in which they are used. In Chapter 2, I develop a method to predict the non-visual effects of wearing these filters across the range of naturally-occurring light intensities. I use this method to compare a dataset of 121 filters, as well as a parametric high-pass model. As expected, the filters show a range of behaviours, with some performing better in lower light conditions compared to others, and some almost entirely blocking circadian-affecting light while others provide only a mild attenuation. I also quantify and analyse the effects of natural optical filters as they change with age and assess the interactions this has with external filters. The natural yellowing of the lens changes light exposure in a way that will have significant circadian effects. The additional effect of external filters is reduced, and the optimal filter for an older individual in a given environment is likely to be different from the optimal filter for a younger individual. I compare the non-visual effects of short-wavelength filters with their visual effects – two factors that need to be balanced in practical filter design. I identify the best performing filters based on a number of composite metrics, showing that the ideal filter depends on the task. Finally, I propose hypothetically optimised filter profiles that maximise non-visual effects while minimising colour and luminance distortions. The optimal filters share features with the best-performing commercial filters, but indicate that there are in principle still improvements that could be made to filter design.

7.1.2 Chapter 3: Quantifying CIE alpha-opic signals in the indoor built environment.

While there is a long history of research quantifying light in naturalistic environments, such as those where the human eye is thought to have evolved, there is little research into the light to which humans are exposed in typical office or domestic environments, yet they comprise the light we experience up to 90% of the time. These environments often present unique and complex

illumination patterns, with a mixture of natural and artificial light, as well as a complex array of surfaces with different material properties that dynamically interact across space and time. In Chapter 3, I utilised a specialised camera with physiologically-relevant filters to quantify alpha-opic signals in a typical office environment at two time points throughout the year. This revealed several interesting regularities in the indoor environment. Seasonal regularities observed in naturalistic environments, such as 'greening' observed in lush seasons similar to British summer, were reproduced in the indoor environment. This was the result of the changing outdoor environment due to changes in sun path, the effects of which were mediated by subtle changes in the light reflected by surfaces within the office space, despite stable artificial illumination. Chapter 3 also showed weather effects that significantly increased the amount of short wavelength light available indoors on overcast days. I also demonstrated a decorrelation of S-cone and melanopsin signals at dawn and dusk – specifically in May – which supports suggestions that S-cone-ipRGC opponency can signal sunrise and sunset times, and interestingly that this signature is available indoors in the location tested.

7.1.3 Chapter 4: Measurement of the directional structure of light in the built environment: single-image circadiometric lightprobes.

Chapter 4 built on the investigation of light available in the built environment established in Chapter 3 through the creation of the Single-Image Circadian Lightprobe Imaging System (SICLIS). The SICLIS is an imaging setup that utilises a camera pointed at a mirror sphere and two additional planar mirrors. This setup is capable of reproducing a full 4π steradian environmental illumination map with minimal data loss. Compared to previous methods, the SICLIS enables unsupervised image capture, as well as faster image capture times. In Chapter 4 I demonstrate the validity of this method using a Blender reconstruction of the SICLIS to show the accurate reconstruction of environmental illumination maps. Then, I conduct two experiments that demonstrate the advantages of the SICLIS design compared to previous approaches. In the first experiment, I record lightprobes in 11 desk locations in a typical office environment across three times of day in four seasonal periods throughout the year, demonstrating the effect of seating proximity to windows, window orientation and time of day on the distribution of directional lighting in the office. In the second experiment, I record lightprobes from two locations in a second typical office environment, emphasising the temporal resolution possible with the SICLIS. The two locations show idiosyncratic light profiles throughout the day, and crucially, they

show a dissociation between the luminance and melanopic light fields within this, emphasising that photopic measurements alone are not sufficient to capture light patterns that could have a circadian influence.

7.1.4 Chapter 5: Individualised light dosimetry paired with head direction tracking in the built environment.

Chapter 5 sought to further build on the quantification of light to which individuals are exposed when in the built environment, however in this instance focusing on individualised measures of light exposure quantified using mobile light dosimetry devices. Wearable dosimetry devices vary widely in their form factor, but to best capture the light to which individuals are exposed, they must sit on the corneal plane so that head movements are accounted for. Present wearable dosimeters rarely log full inertial information, so limited information about head position is known during dosimetry studies. In Chapter 5, I present custom light dosimetry devices featuring 9-axis inertial measurement units (IMUs) that combine magnetometer, accelerometer and gyroscope data to calculate heading direction during light measurement. Using these devices, I collect two weeks of light dosimetry data from four individuals while at work in the same office environment from which lightprobes were collected in Chapter 4. The dosimetry data showed significant seat location, season and time of day effects. I additionally analysed the heading direction data, which showed significant individual differences, though remarkably consistent behaviour across the two weeks analysed, suggesting consistency in gaze location, likely determined by individual preferences in desk setup. Utilising the lightprobes collected at the same time as the dosimetry data, I predict dosimetry data by mapping heading location back onto the lightprobes and re-sampling them, showing the mixed efficacy of static light models for predicting individualised exposure. I then additionally use this data to evaluate the differences in light exposure caused by different head positioning behaviours. Finally, I present egocentric lightprobes which quantify directional variation at different times of day for individuals within a space.

7.1.5 Chapter 6: Hemispheric differences in chromatic discrimination thresholds.

Throughout this thesis I have analysed the way that light around us varies. An important reciprocal action of this is visual adaptation to the environment. Particularly in daytime light levels, the human visual system adapts quickly to prolonged stimulation, and this has significant implications for our ability to discern between similar colours due to changes in our sensitivity

to colour difference from a background. A simple example of a spatial regularity in typical lighting environments is that the upper hemifield of the visual scene is likely different from the lower hemifield as a result of persistent differences between the sky – an illuminator with well-documented regularities in its spectral composition – and the lower half of the visual field, which is generally dominated by light reflected from surfaces that have a much broader range of colorimetric properties. I demonstrated the veracity of this heuristic by analysing a set of lightprobes, showing significant differences in the chromatic distribution of upper hemifields compared to lower hemifields - particularly in how they distribute across the black body locus and in their mean chromaticity. Then, using the chromatic properties of these different regions as an adapting stimulus, I conducted a 2-AFC chromatic discrimination experiment probing the upper and lower retina at an eccentricity of 10° from the fovea. This experiment found that under these adaptation conditions chromatic discrimination thresholds in the lower visual field are lower than those in the upper visual field. This shows that exposure to chromatic distributions typical of upper and lower visual fields in everyday life are sufficient to drive differences in chromatic discrimination performance.

7.2 Future directions

The work presented in this thesis represents a snapshot in an exciting time for psychophysical research and circadian science. While our knowledge of the non-visual effects of light is growing, developments in cutting-edge statistical techniques foreshadow an explosion in our understanding and capacity to study and analyse these effects. In this section, I explore the key directions of future development that the fields of visual and circadian science will take in the coming years, and propose specific research aims that will further the work presented throughout this thesis. Additionally, I will contextualise this forecast alongside the industrial lighting design context and discuss the extent to which research developments will translate into practice.

7.2.1 Further use of custom light measurement devices

Chapters 3, 4 and 5 of this thesis presented a range of physiologically-relevant light measurements from an open-plan office space in Oxford, UK. A significant focus of these chapters is the novel measurement methods used – the single-image circadian lightprobe imaging system and custom-built individualised light dosimeters with heading direction logging – which are used to carry out

data collection. Specific results identified from their use in the Oxford office may be limited in their scope, however, as there is large variation in light that is available across different spaces. The results presented here sit within a large existing literature of light measurement research, and as this literature builds, the variation between different spaces is increasingly quantified. Future research will continue to evolve methods of capturing light in physiologically-relevant ways across the full span of human-inhabited spaces. As discussed throughout this thesis, one of the most important targets for this is mixed-illuminant indoor environments, as the literature exploring these spaces is still limited compared to the literature exploring natural environments, yet these spaces prove to be more important as humans spend more time inside and more is understood about the wellbeing benefits of a healthy spectral diet [169, 325]. The global south has also been identified as a region of the world that is underrepresented in light quantification literature [326]. The devices presented here also represent steps forward for both allocentric and egocentric light measurement strands and pose methods for synthesising insights using both methods in tandem. Increasing the temporal, spatial and spectral resolution of allocentric measurements will continue to be a useful research endeavour. The use of specialised lenses or motorised tripods could facilitate automated capture schemes without the introduction of un-removable mirrors into the lightprobe image, and would be transformative for our ability to quantify light. The advancement of deep learning models in vision science will also provide opportunities for accurate post-hoc removal of mirrors from lightprobe images. Additionally, future work may seek to construct an analytic framework for interpreting higher-order angular harmonics to develop an understanding of the physiological effects of light across a broad range of angular frequencies. Future egocentric light measurements, with the ability to account for fixation, will further develop our ability to precisely calculate the light to which someone is exposed by accounting for gaze and head shape, as well as incorporating participant tracking and human mesh recovery to build a full 3D irradiance exposure map of an indoor space while accounting for ambulatory behaviours.

7.2.2 Efficient psychophysical research using machine learning

Statistical modelling and machine learning approaches promise to be of significant value to the world of human psychophysics. Perhaps the most pertinent example of this is the potential for the efficient measurement of psychophysical thresholds. Chapter 6 of this thesis presented a psychophysical experiment that found psychophysical thresholds using a 2-AFC paradigm. While an adaptive staircase was used to quickly approach discrimination thresholds based on

participants' responses, data collection still took a long time – data collection for each participant took around six hours. This is a significant challenge for many researchers aiming to conduct psychophysical research – putting off prospective participants from taking part, increasing the financial burden of data collection, and introducing strain on participants' focus during data collection, potentially compromising results. Furthermore, it limits the variables that are possible to examine in psychophysical study, as each additional increment of a variable introduces many more trials to a participant's experimental session. Recent work [323, 327] has leveraged the statistical modelling of perceptual space to efficiently sample thresholds across multiple perceptual dimensions in order to characterise discrimination thresholds across the visible gamut around 100 times faster than classical methods. This opens up the possibility for investigating more finely graded variables. For example, measuring discrimination threshold differences across a range of vertical eccentricities rather than a binary upper-lower split. The development of these techniques would enable wide-ranging parameterisation of the way in which light exposure affects local adaptation and could be used to expand the experimental paradigm presented in Chapter 6 across the entire retina while also measuring, for example, discrimination against a range of chromatic backgrounds, or across a range of temporal scales.

7.2.3 A full model of circadian effects of light

While our knowledge of the non-visual effects of light is rapidly growing, we are currently far from a fully parameterised model of circadian responses to light. The non-visual impact of light is modulated by the intensity, wavelength, duration, time of onset, sleep pressure and potentially many more factors, hence the widely-used analogy of a 'spectral diet' – light exposure is high-dimensional and likely depends on currently not understood contextual factors. Future experiments in the circadian science space – particularly those silencing other photoreceptors using specialised displays built for circadian metrology [196, 197] and circadian oscillator models built from sleep lab experiments with highly controlled lighting [159, 160] – can begin to piece together dose-response relationships across the many dimensions of circadian responses to light. ISETBio [328] is an example of a physiological model that has sought to unify data and knowledge from a range of investigations into individual components of the visual system. While it is an on-going project, it serves as an example of what models of non-visual effects may strive for in the coming years. Machine learning models, as described above, will likely play a significant role in unifying data from disparate experimental paradigms to form future models of circadian

responses to light. The methodologies to quantify an individual's light exposure over extended time periods and in complex environments developed in this thesis can provide input data that could drive such models.

7.3 Application of circadian science to lighting design

This project was a part of the Oxford Medical Research Council (MRC) iCASE programme and was funded by the MRC, Hoare Lea and Pembroke College Oxford, with Hoare Lea acting as the industrial partner throughout the project. This partnership involved frequent collaboration between myself, my supervisors and industrial advisors at Hoare Lea, including several work visits to the Hoare Lea offices in London. During this time, I spoke with lighting designers about their experiences and understanding of the non-visual effects of light, and the extent to which this is applied to their practice. I also attended several product demonstration sessions where lighting manufacturers showcased their products to the lighting design team. In this section, I will briefly discuss my experiences from these visits and consider the current state of the application of circadian science in lighting design practice. I will also speculate on the future trajectory of the application of physiological knowledge to real-world lighting design.

The role of lighting designers at companies such as Hoare Lea is to plan and create lighting solutions that fit a client's brief in terms of functional and aesthetic aims. These briefs come with additional constraints, such as financial budgets, environmental regulations [329], and lighting health standards [265]. Lighting designers receive computer-aided design (CAD) models of the project structure and analyse them using software such as ReLux or Rhinoceros 3D. These are software packages that use radiosity or reverse ray-tracing models to estimate the intensity of light cast by individual light sources. Lighting designers will import light sources from light manufacturers that are likely to suit the budget and aesthetic goals of the client and assess whether the desired light levels are reached throughout the design. Lighting manufacturers provide digital photometric data such as correlated colour temperature, intensity and uniformity measurements which allow lighting designers to model their products in CAD projects.

There is a range of standards and certification systems in the UK that can be applied to lighting design projects and focus on different aspects of lighting design. The British Standards Institute document 12464-1 is concerned with the lighting of indoor workplaces and lays out guidance for good practice in terms of quantity and quality of light in typical spaces [330]. Other

standards for building projects, such as the BREEAM framework, focus more on sustainability and energy consumption [331]. The only standard that explicitly points lighting designers towards circadian considerations is the WELL Building Standard [265]. The WELL Building Institute provides guidance on 10 aspects of building design, including lighting design, air quality, thermal comfort, and ease of movement. These concepts each contain preconditions – requirements for certification; and optimisations – additional standards for which additional credit is awarded within the certification score. Projects are awarded bronze, silver, gold or platinum certifications depending on the number of additional points scored. Standards and guidance for other aspects of building design are also considered, such as 'part L' regulations, based on the building regulation Approved Document L provided by the British government, which covers the conservation of heat and power, including the energy consumption density due to lighting throughout buildings. Compared to other standards, it is regarded as mandatory due to its connection to building regulations; though in reality, it is also not a statutory instrument, but rather guidance to help designers implement the regulations. Other documentation used to guide design decisions covers considerations such as fire safety and ventilation.

Within the lighting feature of the WELL standard, there are two preconditions. First, for natural light exposure, a minimum of 200 lux should be achieved for more than 30% of the floor space for at least 50% of daylit hours within a year. For areas where occupants are seated, the target is 300 lux for at least 30% of the floor space, and 100 lux for 95% of the floor space for at least 50% of the daylit hours in a year. Second, in consideration of visual performance, spaces should comply with one of a number of standards, including the CIE recommendations for indoor lighting [332]. This requires 200 lux at the desk level for anywhere where continuous work is being carried out. The WELL Building Institute itself provides an alternate quantification of this precondition requiring 320 lux in office areas. The optimisations in the lighting feature cover a range of topics such as visual balance and light directionality, glare control and occupant controllability. In order to score points for maintaining circadian health, melanopic EDI needs to be measured or modelled at above 136 lux or more, depending on the compliance with other optimisations. To score maximum points without qualifying via other optimisations, 250 lux melanopic EDI is required for at least four hours, starting before midday, each day. The utilisation of circadian-centric standards such as WELL is important as it is known that in many cases buildings do not meet standards for optimal circadian light [205, 333]. This emphasises the need for wider adoption of this standard, better techniques for quantifying light in the

built environment, as well as increased incentives for those who manufacture, design and build widely-used indoor spaces to ensure circadian health for the general public.

The predominance of photopic illuminance measurements in the WELL standard speaks to the broader picture of the way light is modelled in industry. Light manufacturers provide data about their products, but this can be limited to intensity and CCT measurements, and doesn't often include EDI or full spectrum specifications. In my conversations with a number of lighting designers at Hoare Lea, we discussed that, while the team members are aware of the circadian and well-being implications of poorly lit indoor spaces, and have the technical knowledge to calculate mEDI values, circadian considerations can sometimes be de-prioritised in projects where sustainability goals and budgetary constraints are more limiting. However, designers in industry are increasingly seeing clients that specifically ask them to consider circadian-supportive lighting on their projects. This shows that it is becoming more of a consideration for clients, which circumvents the need for other drivers such as standards.

With increased media focus on the benefits of human centric lighting and the increased understanding of metrics to model these benefits, I speculate that the future promises much more integration of circadian science into the world of lighting design. Public-facing explainers of circadian science have been produced throughout the last ten years [334, 335], but recent efforts have focused on alerting the general public to the potential benefits of light health awareness [336–339]. Several software packages have also been developed to enable the calculation of melanopic quantities of light, many of which are compatible with currently-used software - such as the Grasshopper 3D plug-in for Rhinoceros 3D [340–345]. Undoubtedly, one of the main drivers for improvements in this space is improved scientific methodology to demonstrate measurable benefits of circadian lighting solutions. With lighting designers, such as those at Hoare Lea, being better informed about circadian effects of lighting and able to make melanopic calculations using modelling software, further development of these tools will enable integration across the different kinds of measurements that designers use to model spaces. Stronger modelling tools and increased client awareness would significantly increase the value proposition for lighting designers trying to attain the highest WELL certification. As circadian optimisation becomes more common, it will increase pressure on lighting manufacturers to produce human centric products, and clients to specify human centric designs, creating a virtuous cycle that ultimately absorbs circadian parameterisation into the standard set of lighting design parameters. While we are, at the time of writing, a long way off reaching this point, the increased research focus on the

non-visual effects of light, and the potential for a unified model of circadian responses to light mean that the future of the use of circadian science in industry is bright.

7.4 Final summary

As a final summary, this thesis sought to contribute to the characterisation of illumination in the built environment and to our understanding of the physiological changes and behaviours that occur in response. It did so across five experimental chapters. In the first, I propose a framework for analysing the non-visual effects of short-wavelength blocking filters, analysing a dataset of commercially-available filters and comparing their visual and non-visual attributes. In the second, I present an experiment using a specialised camera to create a dataset of physiologically-relevant images of a typical work environment and analysed the spatial and temporal dynamics of visual and non-visual components of light. I demonstrate that key changes that occur in natural environments are reproduced in mixed-illuminant environments through surface interreflections. In Chapter 3 I build on this work by proposing SICLIS, an imaging system that enables the capture of a full panoramic view of a scene in a single image. I conduct two experiments using SICLIS, analysing the directional structure of light across space and time. In the fourth, I focus on individualised light exposure, presenting custom-built light dosimetry devices with compass-heading logging. I demonstrate the importance of window-proximity to light exposure, as well as seasonal and diurnal dynamics. I utilise the lightprobes captured by SICLIS to assess the contribution of gaze behaviours on light exposure. Finally, in Chapter 5 I present a psychophysical experiment investigating hemispheric differences in chromatic discrimination thresholds. I show that there are regular statistical differences between the chromatic distributions of the upper and lower hemifields of natural lighting environments. In a psychophysical experiment, I show differences in the chromatic discrimination thresholds recorded in retinal regions above and below the fovea, with lower thresholds being found in the lower visual field. Finally, I have contextualised this work in the contemporary world of lighting design and circadian science and advocated for a push to further the use of physiological insights into the design of the indoor spaces we use every day. I believe that the work presented in this thesis provides an incremental step forward in the understanding and characterisation of light in the built environment and provides useful contributions from which future research will benefit.

References

[1] A. Ryer. *The Light Measurement Handbook*. International Light Technologies, 1997.

[2] CIE. CIE standard illuminant A - 1nm. In *International Commission on Illumination (CIE), Vienna, AT*, 2018.

[3] G. J. Ward. Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2):265–72, 1992.

[4] H. Kolb. Simple anatomy of the retina. Webvision. Moran Eye Center.

[5] R. L. De Valois and K. De Valois. A multi-stage colour model. *Vision Research*, 33(8):1053–65, 1993.

[6] M. S. Livingston and D. H. Hubel. Anatomy and physiology of a color system in the primate visual cortex. *Journal of Neuroscience*, 4(1):309–56, 1984.

[7] Y. Xiao, Y. Wang, and D. J. Felleman. A spatially organised representation of colour in the macaque cortical area v2. *Nature*, 421(6922):535–9, 2003.

[8] K. R. Gegenfurtner. Cortical mechanisms of colour vision. *Nature Reviews*, 4(7):563–72, 2003.

[9] B. R. Conway. Color vision, cones and color-coding in the cortex. *Neuroscientist*, 15(3):274–90, 2009.

[10] A. Stockman, L. T. Sharpe, and C. Fach. The spectral sensitivity of the human short-wavelength sensitive cones derived from thresholds and color matches. *Vision Research*, 39(17):2901–27, 1999.

[11] A. Stockman and L. T. Sharpe. The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40(13):1711–37, 2000.

[12] G. Wyszecki and W.S. Stiles. *Color science: concepts and methods, quantitative data and formulae*. (2nd edition), Wiley, 1982.

[13] R. J. Lucas, S. N. Peirson, D. M. Berson, T. M. Brown, H. M. Cooper, C. A. Czeisler, M. G. Figueiro, P. D. Gamlin, S. W. Lockley, J. B. O'Hagan, L. L. A. Price, I. Provencio, D. J. Skene, and G. C. Brainard. Measuring and using light in the melanopsin age. *Trends in Neurosciences*, 37(1):1–9, 2014.

[14] W. A. H. Rushton. Review lecture. pigments and signals in colour vision. *The Journal of Physiology*, 220(3):1–31, 1972.

[15] H. Kolb. Cone pathways through the retina. Webvision. Moran Eye Center.

[16] R. W. Williams. The human retina has a cone-enriched rim. *Visual Neuroscience*, 6(4):403–6, 1991.

[17] J. Mollon, B. Regan, and J. Bowmaker. What is the function of the cone-rich rim of the retina? *Eye*, 12:548–52, 1998.

[18] H. Kolb. Photoreceptors. Webvision. Moran Eye Center.

[19] C. A. Curcio, K. A. Allen, K. R. Sloan, C. L. Lerea, J. B. Hurley, I. B. Klock, and A. H. Milam. Distribution and morphology of human cone photoreceptors stained with anti-blue opsin. *Journal of Comparative Neurology*, 312(4):610–24, 1991.

[20] J. A. Castańo and H. G. Sperling. Sensitivity of the blue-sensitive cones across the central retina. *Vision Research*, 22(6):661–73, 1982.

[21] R. M. Boynton. *Human Color Vision*. Holt, Rinehart and Winston, 1979.

[22] B. C. Regan, C. Julliot, B. Simmen, F. Viénot, P. Charles-Dominique, and J. Mollon. Fruits, foliage and the evolution of primate vision. *Philosophical Transactions of The Royal Society B: Biological Sciences*, 356(1407):229–83, 2001.

[23] G. Jordan and J. D. Mollon. Sons and mothers: classification of colour-deficient and heterozygous subjects by counterphase modulation photometry. *Colour Vision Deficiencies XIII*, pages 385–92, 1997.

[24] M. P. Simunovic. Colour vision deficiency. *Eye*, 24(5):747–55, 2010.

[25] Jennifer Birch. Worldwide prevalence of red-green color deficiency. *Journal of the Optical Society of America A*, 29(3):313–20, 2012.

[26] M. A. Webster and D. I. A. Macleod. Factors underlying individual differences in the color matches of normal observers. *Journal of the Optical Society of America A*, 5(10):1722–35, 1988.

[27] Y. Asano, M. D. Fairchild, and L. Blondé. Individual colorimetric observer model. *PLOS One*, 11(2):e0145671, 2016.

[28] K. J. Emery and M. A. Webster. Individual differences and their implications for color perception. *Current Opinion in Behavioral Sciences*, 30:28–33, 2019.

[29] Y. H. Wang, J. Zhong, and X. M. Li. Age-related changes of lens thickness and density in different age phases. *International Journal of Ophthalmology*, 15(10):1591–7, 2022.

[30] J. Pokorny and V. C. Smith. Effect of field size on red-green color mixture equations. *Journal of the Optical Society of America A*, 66(7):705–8, 1976.

[31] D. M. Dacey. Physiology, morphology and spatial densities of identified ganglion cell types in primate retina. *Ciba Found Symposium*, 184:12–28, 1994.

[32] R. H. Masland. The fundamental plan of the retina. *Nature Neuroscience*, 4(9):877–86, 2001.

[33] D. M. Dacey and B. B. Lee. The 'blue-on' opponent pathway in primate retina originates from a distinct bistratified ganglion cell type. *Nature*, 367(6465):731–5, 1994.

[34] D. M. Dacey. Circuitry for color coding in the primate retina. *Proceedings of the National Academy of Sciences of the United States of America*, 93(2):582–8, 1996.

[35] V. A. Casagrande and X. Xu. *Parallel visual pathways: a comparative perspective*, pages 494–506. MIT Press, Cambridge, 2004.

[36] J. Martinovic. *Magno-, Parvo-, Koniocellular Pathways*, pages 1–5. Springer Science, 2015.

[37] S. G. Solomon. Retinal ganglion cells and the magnocellular, parvocellular, and koniocellular subcortical visual pathways from the eye to the brain. In *Neurology of Vision and Visual Disorders*, volume 178 of *Handbook of Clinical Neurology*, pages 31–50. Elsevier, 2021.

[38] E. S. Yamada, L. C. Silveira, and V. H. Perry. Morphology, dendritic field size, somal size, density, and coverage of m and p retinal ganglion cells of dichromatic cebus monkeys. *Visual Neuroscience*, 13(6):1011–29, 1996.

[39] D. M. Dacey. The mosaic of midget ganglion cells in the human retina. *Journal of Neuroscience*, 13(12):5334–55, 1993.

[40] A.M. Derrington, J. Krauskopf, and P. Lennie. Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of Physiology*, 357(1):241–65, 1984.

[41] T. N. Wiesel and D. H. Hubel. Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *Journal of Neurophysiology*, 29(6):1115–56, 1966.

[42] P. H. Schiller. Parallel information processing channels created in the retina. *Proceedings of the National Academy of Science of The United States of America*, 107(40):17087–94, 2010.

[43] A. Kling, G. D. Field, D. H. Brainard, and E. J. Chichilnisky. Probing computation in the primate visual system at single-cone resolution. *Annual Review of Neuroscience*, 42:169–86, 2019.

[44] L. E. Wool, J. D. Crook, J. B. Troy, O. S. Packer, Q. Zaidi, and D. M. Dacey. Nonselective wiring accounts for red-green opponency in midget ganglion cells of the primate retina. *Journal of Neuroscience*, 38(6):1520–40, 2018.

[45] L. M. Hurvich and D. Jameson. An opponent-process theory of colour vision. *Psychological Review*, 64(6):384–404, 1957.

[46] V. A. Casagrande, F. Yazar, K. D. Jones, and Y. Ding. The Morphology of the Koniocellular Axon Pathway in the Macaque Monkey. *Cerebral Cortex*, 17(10):2334–45, 2007.

[47] Y. Tsukamoto and N. Omi. Off bipolar cells in macaque retina: type-specific connectivity in the outer and inner synaptic layers. *Frontiers in Neuroanatomy*, 9:122, 2015.

[48] L. E. Wool, O. S. Packer, Q. Zaidi, and D. M. Dacey. Connectomic identification and three-dimensional color tuning of s-off midget ganglion cells in the primate retina. *Journal of Neuroscience*, 39(40):7893–909, 2019.

[49] L. T. Sharpe, A. Stockman, W. Jagla, and H. Jägle. A luminous efficiency function, V*(λ), for daylight adaptation. *Journal of Vision*, 5(11):3, 2005.

[50] J. Carroll, J. Neitz, and M. Netiz. Estimates of L:M cone ratio from ERG flicker photometry and genetics. *Journal of Vision*, 2(8):531–42, 2002.

[51] H. Hofer, J. Carroll, J. Neitz, M. Neitz, and D. R. Williams. Organization of the human trichromatic cone mosaic. *Journal of Neuroscience*, 25(42):9669–79, 2005.

[52] U. S. Kim, O. A. Mahroo, J. D. Mollon, and P. Yu-Wai-Man. Retinal ganglion cells-diversity of cell types and clinical relevance. *Frontiers in Neurology*, 12:635, 2021.

[53] J. D. Crook, B. B. Peterson, O. S. Packer, F. R. Robinson, P. D. Gamlin, J. B. Troy, and D. M. Dacey. The smooth monostratified ganglion cell: evidence for spatial diversity in the y-cell pathway to the lateral geniculate nucleus and superior colliculus in the macaque monkey. *Journal of Neuroscience*, 28(48):12654–71, 2008.

[54] C. E. Rhoades, N. P. Shah, M. B. Manookin, N. Brackbill, A. Kling, and G. et al. Goetz. Unusual physiological properties of smooth monostratified ganglion cell types in primate retina. *Neuron*, 103:658–82, 2019.

[55] D.M. Dacey, Y. J. Kim, O. S. Packer, and P. B. Detwiler. On-off direction selective ganglion cells in macaque monkey retina are tracer-coupled to an on-off direction selective amacrine cell type. *Investigations in Ophthalmology & Vision Science*, 60:5280, 2019.

[56] J. Krauskopf, D. R. Williams, and D. W. Heeley. Cardinal directions of color space. *Vision Research*, 22(9):1123–31, 1982.

[57] D. I. A. MacLeod and R. M. Boynton. Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America A*, 69(8):1183–6, 1979.

[58] CIE. Commission internationale de l'eclairage proceedings, 1931. *Cambridge: Cambridge University Press*, 1932.

[59] W. D. Wright. A re-determination of the trichromatic coefficients of the spectral colours. *Transactions of the Optical Society*, 30(4):141, 1929.

[60] W. D. Wright. A re-determination of the mixture curves of the spectrum. *Transactions of the Optical Society*, 31(4):201, 1930.

[61] J. Guild. The colorimetric properties of the spectrum. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 230(681):149–87, 1931.

[62] CIE. Colorimetry. Part 5: CIE 1976 $L^*u^*v^*$ colour space and $u'v'$ uniform chromaticity scale diagram. *(Vienna: Central Bureau of the Commission Internationale de l' Éclairage)*, 1976.

[63] CIE. Colorimetry. Part 4: CIE 1976 $L^*a^*b^*$ colour space. *(Vienna: Central Bureau of the Commission Internationale de l' Éclairage)*, 1976.

[64] R. Pastilha and A. Hurlbert. Seeing and sensing temporal variations in natural daylight. *Progress in brain research*, 273(1):275–301, 2022.

[65] A. Panorgias, J. J. Kulikowski, N. R. A. Parry, D. J. McKeefry, and I. J. Murray. Phases of daylight and the stability of color perception in the near peripheral human retina. *Journal of Vision*, 12(3), 2012.

[66] M. Spitschan, G. K. Aguirre, D. H. Brainard, and A. M. Sweeney. Variation of outdoor illumination as a function of solar elevation and light pollution. *Scientific Reports*, 6:26756, 2016.

[67] CIE. Spatial Distribution of Daylight - CIE Standard General Sky. *(Vienna: Central Bureau of the Commission Internationale de l' Éclairage)*, 2004.

[68] D. B. Judd, D. L. MacAdam, G. Wyszecki, H. W. Budde, H. R. Condit, S. T. Henderson, and J. L. Simonds. Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America A*, 54(8):1031–40, 1964.

[69] CIE. CIE standard illuminant D65. In *International Commission on Illumination (CIE), Vienna, AT*, 2019.

[70] T. Esposito and K. Houser. Correlated color temperature is not a suitable proxy for the biological potency of light. *Scientific Reports*, 12:20223, 2022.

[71] C. Li, G. Cui, M. Melgosa, X. Ruan, Y. Zhang, L. Ma, K. Xiao, and M. R. Luo. Accurate method for computing correlated color temperature. *Optics Express*, 24:14066–78, 2016.

[72] J. Hernández-Andrés, R. L. Lee, and J. Romero. Calculating correlated color temperatures across the entire gamut of daylight and skylight chromaticities. *Applied Optics*, 38:5703–9, 1999.

[73] D. Jameson and L. M. Hurvich. Perceived color and its dependence on focal, surrounding, and preceding stimulus variables. *Journal of the Optical Society of America*, 49(9):890–8, 1959.

[74] S. K. Shevell. Dual role of chromatic backgrounds. *Vision Research*, 18(1971):1649–61, 1978.

[75] K. T. Blackwell and G. Buchsbaum. The effect of spatial and chromatic parameters on chromatic induction. *Color Research & Application*, 13(3):166–73, 1988.

[76] O. Creutzfeldt, B. Lange-Malecki, and E. Dreyer. Chromatic induction and brightness contrast: a relativistic colour model. *Journal of the Optical Society of America A*, 7(9):1644–53, 1990.

[77] R. O. Brown and D. I. A. MacLeod. Color appearance depends on the variance of surround colors. *Vision Research*, 34(3):844–9, 1997.

[78] S. L. Guth. Model for color vision and light adaptation. *Journal of the Optical Society of America A*, 8(6):976–93, 1991.

[79] M. D. Bownds and V. Y. Arshavsky. What are the mechanisms of photoreceptor adaptation? *Behavioural and Brain Sciences*, 18(3):415–24, 1995.

[80] G. L. Fain, H. R. Matthews, M. C. Cornwall, and Y. Koutalos. Adaptation in vertebrate photoreceptors. *Physiological Reviews*, 81(1):117–51, 2001.

[81] M. A. Webster and D. Leonard. Adaptation and perceptual norms in color vision. *Journal of the Optical Society of America A*, 25(11):2817–25, 2008.

[82] M. A. Webster. Visual adaptation. *Annual Review of Vision Science*, 1(1):547–67, 2015.

[83] G. Jordan and J. Mollon. Adaptation of colour vision to sunlight. *Nature*, 386:135–6, 1997.

[84] O. Rinner and K. R. Gegenfurtner. Time course of chromatic adaptation for color appearance and discrimination. *Vision Research*, 40(14):1813–26, 2000.

[85] D. H. Foster. Color constancy. *Vision Research*, 51(7):674–700, 2011.

[86] B. J. Craven and D. H. Foster. An operational approach to colour constancy. *Vision Research*, 32(7):1359–66, 1992.

[87] D. H. Foster and S. M. Nascimento. Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 257(1349):115–21, 1994.

[88] J. M. Kraft and D. H. Brainard. Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Science of the United States of America*, 96(1):307–12, 1999.

[89] J. M. Kraft, S. I. Maloney, and D. H. Brainard. Surface-illuminant ambiguity and color constancy: Effects of scene complexity and depth cues. *Perception*, 31(2):247–63, 2002.

[90] H. E. Smithson and Q. Zaidi. Colour constancy in context: Roles for local adaptation and levels of reference. *Journal of Vision*, 4(9):3, 2004.

[91] H. E. Smithson. Sensory, computational and cognitive components of human colour constancy. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 360(1458):1329–46, 2005.

[92] H. J. A. Dartnall, J. K. Bowmaker, J. D. Mollon, and H. B. Barlow. Human visual pigments: microspectrophotometric results from the eyes of seven persons. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 220(1218):115–30, 1983.

[93] S. L. Polyak. The retina: The anatomy and the histology of the retina in man, ape and monkey, including the consideration of visual functions, the history of physiological optics and the histological laboratory technique. *Journal of the American Medical Association*, 118(15):1337, 1942.

[94] S. Hecht, S. Shlaer, and M. H. Pirenne. Energy, quanta and vision. *Journal of General Physiology*, 25(6):819–40, 1942.

[95] T. D. Lamb. Why rods and cones? *Eye*, 30(2):179–85, 2016.

[96] L. T. Sharpe and A. Stockman. Rod pathways: the importance of seeing nothing. *Trends in Neurosciences*, 22(11):497–504, 1999.

[97] H. Delange. Experiments on flicker and some calculations on an electrical analogue of the foveal systems. *Physica*, 18(11):935–50, 1952.

[98] H. Sun, J. Pokorny, and V. C. Smith. Rod-cone interactions assessed in inferred magnocellular and parvocellular postreceptoral pathways. *Journal of Vision*, 1(1):5, 2001.

[99] J. Pokorny, M. Lutze, D. Cao, and A. J. Zele. The color of night: surface color categorization by color defective observers under dim illuminations. *Visual Neuroscience*, 25(3):475–80, 2008.

[100] S. L. Elliott and D. Cao. Scotopic hue percepts in natural scenes. *Journal of Vision*, 13(13):15–, 2013.

[101] W. Nagel. Appendix: adaptation, twilight vision and the duplicity theory. *Helmholtz's Treatise on Physiological Optics Translated from the Third German Edition by JPC Southall*, pages 313–43, 1924.

[102] C. Sung, C. M. Davenport, J. C. Hennessey, I. H. Maumenee, S. G. Jacobson, J. R. Heckenlively, R. Nowakowski, G. Fishman, P. Gouras, and J. Nathans. Rhodopsin mutations in autosomal dominant retinitis pigmentosa. *Proceedings of the National Academy of Sciences of the United States of America*, 88(15):6481–5, 1991.

[103] A. J. Zele, J. Kremers, and B. Feigl. Mesopic rod and s-cone interactions revealed by modulation thresholds. *Journal of the Optical Society of America A*, 29(2):A19–26, 2012.

[104] A. Stockman and L. T. Sharpe. Into the twilight zone: the complexities of mesopic vision and luminous efficiency. *Ophthalmic and Physiological Optics*, 26(3):225–39, 2006.

[105] M. Aguilar and W.S. Stiles. Saturation of the Rod Mechanism of the Retina at High Levels of Stimulation. *Optica Acta: International Journal of Optics*, 1(1):59–65, 1954.

[106] D. C. Hood and M. A. Finkelstein. Sensitivity to light. In: K. Boff, L. Kaufman, & J. Thomas (Eds.). *Handbook of Perception and Human Performance*, 1, 1986.

[107] A. J. Zele and D. Cao. Vision under mesopic and scotopic illumination. *Frontiers in Psychology*, 5:1594, 2015.

[108] G. S. Lall, V. L. Revell, H. Momiji, J. A. Enezi, C. M. Altimus, A. D. Güler, C. Aguilar, M. A. Cameron, S. Allender, M. W. Hankins, and R. J. Lucas. Distinct contributions of rod, cone, and melanopsin photoreceptors to encoding irradiance. *Neuron*, 66:417–28, 2010.

[109] R. Frederiksen, A. Morshedian, S. A. Tripathy, T. Xu, G. H. Travis, G. L. Fain, and A. P. Sampath. Rod photoreceptors avoid saturation in bright light by the movement of the g protein transducin. *The Journal of neuroscience*, 41(15):3320–30, 2021.

[110] B. B. Lee, V. C. Smith, J. Pokorny, and J. Kremers. Rod inputs to macaque ganglion cells. *Vision Research*, 37(20):2813–28, 1997.

[111] G. D. Field, M. Greschner, J. L. Gauthier, C. Rangel, J. Shlens, A. Sher, D. W. Marshak, A. M. Litke, and E. J. Chichilnisky. High-sensitivity rod photoreceptor input to the blue-yellow color opponent pathway in macaque retina. *Nature Neuroscience*, 12(9):1159–64, 2009.

[112] D. Cao, B. B. Lee, and H. Sun. Combination of rod and cone inputs in parasol ganglion cells of the magnocellular pathway. *Journal of Vision*, 10(11):4, 2010.

[113] J. L. Ecker, O. N. Dumitrescu, K. Y. Wong, N. M. Alam, S. Chen, T. LeGates, J. M. Renna, G. T. Prusky, D. M. Berson, and S. Hattar. Melanopsin-expressing retinal ganglion-cell photoreceptors: Cellular diversity and role in pattern vision. *Neuron*, 67(1):49–60, 2010.

[114] I. Provencio, I. R. Rodriguez, G. Jiang, W. P. Hayes, E. F. Moreira, and M. D. Rollag. A novel human opsin in the inner retina. *Journal of Neuroscience*, 20(2):600–605, 2000.

[115] M. Münch and A. Kawasaki. Intrinsically photosensitive retinal ganglion cells: classification, function and clinical implications. *Current Opinion in Neurology*, 26(1):45–51, 2013.

[116] L. S. Mure. Intrinsically photosensitive retinal ganglion cells of the human retina. *Frontiers in Neurology*, 12, 2021.

[117] C. A. Curcio and K. A. Allen. Topography of ganglion cells in human retina. *The Journal of Comparative Neurology*, 300(1):5–25, 1990.

[118] M. T. H. Do and K. Yau. Intrinsically photosensitive retinal ganglion cells. *Physiological Reviews*, 90(4):1547–81, 2010.

[119] T. Woelders, A. E. Allen, and R. J. Lucas. Melanopsin enhances image persistence. *Current Biology*, 33(23):5048–56, 2023.

[120] G. K. Aguirre. Vision: a prolonged and colorless experience. *Current Biology*, 34(3):89–91, 2024.

[121] D. M. Dacey, H. Liao, B. B. Peterson, F. R. Robinson, V. C. Smith, J. Pokorny, K. Yau, and P. D. Gamlin. Melanopsin-expressing ganglion cells in primate retina signal colour and irradiance and project to the LGN. *Nature*, 433(7027):749–54, 2005.

[122] F. Zaidi, J. Hull, S. Peirson, K. Wulff, D. Aeschbach, J. Gooley, G. Brainard, K. Gregory-Evans, J. Rizzo, C. Czeisler, R. Foster, M. Moseley, and S. Lockley. Short-wavelength light sensitivity of circadian, pupillary, and visual awareness in humans lacking an outer retina. *Current Biology*, 17(24):2122–28, 2007.

[123] T. M. Brown, S. I. Tsujimura, A. E. Allen, J. Wynne, R. Bedford, G. Vickery, A. Vugler, and R. J. Lucas. Melanopsin-based brightness discrimination in mice and humans. *Current Biology*, 22(12):1134–41, 2012.

[124] M. Spitschan, A. S. Bock, J. Ryan, G. Frazzetta, D. H. Brainard, and G. K. Aguirre. The human visual cortex response to melanopsin-directed stimulation is accompanied by a distinct perceptual experience. *Proceedings of the National Academy of Science of the United States of America*, 114(46):(46): 12291–96, 2017.

[125] A. J. Zele, P. Adhikari, B. Feigl, and D. Cao. Cone and melanopsin contributions to human brightness estimation. *J. Opt. Soc. Am. A*, 35(4):B19–25, 2018.

[126] A. J. Zele, B. Feigl, P. Adhikari, M. Maynard, and D. Cao. Melanopsin photoreception contributes to human visual detection, temporal and colour processing. *Scientific Reports*, 8:3842, 2018.

[127] D. Cao, A. Chang, and S. Gai. Evidence for an impact of melanopsin activation on unique white perception. *Journal of the Optical Society of America A*, 35(4):B287–91, 2018.

[128] L. Tähkämö, T. Partonen, and A. K. Pesonen. Systematic review of light exposure impact on human circadian rhythm. *Chronobiology International*, 36(2):151–70, 2018.

[129] M. Spitschan. Melanopsin contributions to non-visual and visual function. *Current Opinion in Behavioral Sciences*, 30:67–72, 2019.

[130] C. Blume, C. Garbazza, and M. Spitschan. Effects of light on human circadian rhythms, sleep and mood. *Somnologie*, 23:147–56, 2019.

[131] T. Bedrosian and R. Nelson. Timing of light exposure affects mood and brain circuits. *Translational Psychiatry*, 7:e1017, 2017.

[132] S. R. Pandi-Perumal, I. Trakht, V. Srinivasan, D. W. Spence, G. J. M. Maestroni, N. Zisapel, and D. P. Cardinali. Physiological effects of melatonin: Role of melatonin receptors and signal transduction pathways. *Progress in Neurobiology*, 85:335–53, 2008.

[133] M. A. Aranda and T. M. Schmidt. Diversity of intrinsically photosensitive retinal ganglion cells: circuits and functions. *Cellular and Molecular Life Sciences*, 78:889–907, 2021.

[134] A. C. Rupp, M. Ren, D. C. Altimus, C. M. Fernandez, M. Richardson, F. Turek, S. Hatter, and T. M. Schmidt. Distinct iprgc subpopulations mediate light's acute and circadian effects on body temperature and sleep. *eLife*, 8:e44358, 2019.

[135] M. E. Stabio, S. Sabbah, L. E. Quattorchi, M. C. Ilardi, P. M. Fogerson, M. L. Leyrer, M. T. Kim, I. Kim, M. Schiel, J. M. Renna, K. L. Briggman, and D. M. Berson. The m5

cell: A color-opponent intrinsically photosensitive retinal ganglion cell. *Neuron*, 97(1):251, 2018.

[136] M. E. Estevez, P. M. Fogerson, M. C. Ilardi, B. G. Borghuis, E. Chan, S. Weng, O. N. Auferkorte, J. B. Demb, and D. M. Berson. Form and function of the m4 cell, an intrinsically photosensitive retinal ganglion cell type contributing to geniculocortical vision. *Journal of Neuroscience*, 32(39):13608–20, 2012.

[137] C. Beier, Z. Zhang, M. Yurgel, and S. Hattar. The projections of iprgcs and conventional rgcs to retinorecipient brain nuclei. *Journal of Comparative Neurology*, 529(8):1763–75, 2020.

[138] M. L. Aranda, J. D. Bhoi, O. A. Payán Parra, S. Ki Lee, T. Yamada, Y. Yang, and T. M. Schmidt. Genetic tuning of retinal ganglion subtype identity to drive visual behaviour. *Nature Communications*, 16(1):8678, 2025.

[139] Liao, H. and Ren, X. and Peterson, B. B. and Marshak, D. W. and Yau, K., Gamlin, P. D., et al. Melanopsin-expressing ganglion cells on macaque and human retinas form two morphologically distinct populations. *Journal of Comparative Neurology*, 524:2845–72, 2016.

[140] L. S. Mure, F. Vinberg, A. Hanneken, and S. Panda. Functional diversity of human intrinsically photosensitive retinal ganglion cells. *Science*, 366:1251–5, 2019.

[141] J. Hannibal, A. T. Christiansen, S. Heegaard, J. Fahrenkrug, and J. F. Kiilgaard. Melanopsin expressing human retinal ganglion cells: Subtypes, distribution, and intraretinal connectivity. *Journal of Comparative Neurology*, 525(8):1934–61, 2017.

[142] G. Esquiva, P. Lax, J. J. Pérez-Santonja, J. M. García-Fernández, and N. Cuenca. Loss of melanopsin-expressing ganglion cell subtypes and dendritic degeneration in the aging human retina. *Frontiers in Aging Neuroscience*, 9:79, 2017.

[143] S. Nasir-Ahmad, S. C. S. Lee, P. R. Martin, and U. Grünert. Melanopsin-expressing ganglion cells in human retina: Morphology, distribution, and synaptic connections. *Journal of Comparative Neurology*, 527(1):312–27, 2019.

[144] M. T. H. Do. Melanopsin and the intrinsically photosensitive retinal ganglion cells: Biophysics to behavior. *Neuron Review*, 104:209–26, 2019.

[145] A. S. Bordt, S. S. Patterson, J. A. Kuchenbecker, M. A. Mazzaferri, J. N. Yearick, E. R. Yang, J. M. Ogilvie, J. Neitz, and D. W. Marshak. Synaptic inputs to displaced intrinsically-photosensitive ganglion cells in macaque retina. *Scientific Reports*, 12:15160, 2022.

[146] J. Rabin, J. Lovell, N. Tahvilian, A. Ku, T. Contreras, J. Carachure, L. Geabou, J. Sies, and D. Nguyen. Putative retinal and cortical potentials from melanopsin sensitive ganglion cells. *Eye*, 35:2651–2, 2020.

[147] T. M. Schmidt, S. Chen, and S. Hattar. Intrinsically photosensitive retinal ganglion cells: many subtypes, diverse functions. *Trends in Neuroscience*, 32(11):572–80, 2011.

[148] S. S. Patterson, J. A. Kuchenbecker, J. R. Anderson, M. Neitz, and J. Neitz. A color vision circuit for non-image-forming vision in the primate retina. *Current Biology*, 30(7):1269–74, 2020.

[149] S. S. Patterson, M. A. Mazzaferri, A. S. Bordt, J. Chang, M. Neitz, and J. Neitz. Another blue-on ganglion cell in the primate retina. *Current Biology*, 30(23):R1409–10, 2020.

[150] P. A. Barrionuevo, M. L. Sandoval Salinas, and J. M. Fanchini. Are iprgcs involved in human color vision? hints from physiology, psychophysics, and natural image statistics. *Vision Research*, 217, 2024.

[151] A. P. Patton and M. H. Hastings. The suprachiasmatic nucleus. *Current Biology*, 28(15):R816–22, 2018.

[152] R. Y. Moore. Suprachiasmatic nucleus in sleep–wake regulation. *Sleep Medicine*, 8:27–33, 2007.

[153] K. Choi, C. Shin, T. Kim, H.J. Chung, and H. J. Suk. Awakening effects of blue-enriched morning light exposure on university students' physiological and subjective responses. *Scientific Reports*, page 345, 2019.

[154] H. R. Wright and L. C. Lack. Effect of light wavelength on suppression and phase delay of the melatonin rhythm. *Chronobiology International*, 18(5):801–8, 2001.

[155] H. R. Wright, L. C. Lack, and D. J. Kennaway. Differential effects of light wavelength in phase advancing the melatonin rhythm. *Journal of Pineal Research*, 36(2):140–4, 2004.

[156] J. F. Duffy and K. P. Wright. Entrainment of the human circadian system by light. *Journal of Biological Rhythms*, 20(4):326–38, 2005.

[157] T. M. Brown. Melanopic illuminance defines the magnitude of human circadian light responses under a wide range of conditions. *Journal of Pineal Research*, 69:e12655, 2020.

[158] M. Cyr, D. Z. Artenie, A. Al Bikaii, D. Borsook, and J. A. Olson. The effect of evening light on circadian-related outcomes: A systematic review. *Sleep Medicine Reviews*, 64:101660, 2022.

[159] A. C. Skeldon, T. Rodriguez Garcia, S. F. Cleator, C. Della Monica, K. K. G. Ravindran, V. L. Revell, and D. J. Dijk. Method to determine whether sleep phenotypes are driven by endogenous circadian rhythms or environmental light by combining longitudinal data and personalised mathematical models. *PLoS computational biology*, 19(12):e1011743, 2023.

[160] A. C. Skeldon and D. J. Dijk. The complexity and commonness of the two-process model of sleep regulation from a mathematical perspective. *NPJ Biological Timing and Sleep*, 2(1):24, 2025.

[161] C. Ribelayga, Y. Cao, and S. C. Mangel. The circadian clock in the retina controls rod-cone coupling. *Neuron*, 59(5):790–801, 2008.

[162] A. J. Lewy, T. A. Wehr, F. K. Goodwin, D. A. Newsome, and S. P. Markey. Light suppresses melatonin secretion in humans. *Science*, 210(4475):1267–9, 1980.

[163] C. Cajochen, K. Kräuchi, and A. Wirz-Justice. Role of melatonin in the regulation of human circadian rhythms and sleep. *Journal of Neuroendocrinology*, 15(4):432–7, 2003.

[164] I. N. Karatsoreos, S. Bhagat, E. B. Bloss, J. H. Morrison, and B. S. McEwen. Disruption of circadian clocks has ramifications for metabolism, brain, and behavior. *Proceedings of the National Academy of Science of the United States of America*, 108(4):1657–62, 2011.

[165] J. A. Evans and A. J. Davidson. Chapter ten - health consequences of circadian disruption in humans and animal models. *Progress in Molecular Biology and Translational Science*, 119:283–323, 2013.

[166] W. H. Walker, J. C. Walton, A. C. DeVries, and R. J. Nelson. Circadian rhythm disruption and mental health. *Translational Psychiatry*, 10(28), 2020.

[167] C. Vetter. Circadian disruption: What do we actually mean? *European Journal of Neuroscience*, 51(1):531–50, 2020.

[168] A. B. Fishbein, K. L. Knutson, and P. C. Zee. Circadian disruption and human health. *Journal of Clinical Investigation*, 131(19):e148286, 2021.

[169] F. S. Webler, M. Spitschan, R. G. Foster, M. Andersen, and S. N. Peirson. What is the 'spectral diet' of humans? *Current Opinion in Behavioral Sciences*, 30:80–6, 2019.

[170] T. M. Brown, G. C. Brainard, C. Cajochen, C. A. Czeisler, J. P. Hanifin, and S. W. et al. Lockley. Recommendations for daytime, evening, and nighttime indoor light exposure to best support physiology, sleep, and wakefulness in healthy adults. *PLoS Biology*, 20(3):e3001571, 2022.

[171] H. Burkhalter, A. Wirz-Justice, K. Denhaerynck, T. Fehr, J. Steiger, R. M. Venzin, C. Cajochen, T. E. Weaver, and S. De Geest. The effect of bright light therapy on sleep and circadian rhythms in renal transplant recipients: a pilot randomized, multicentre wait-list controlled trial. *Transplant International*, 28(1):59–70, 2015.

[172] K. Wulff and R. G. Foster. Insight into the role of photoreception and light intervention for sleep and neuropsychiatric behaviour in the elderly. *Current Alzheimers Research*, 14(10):1022–9, 2017.

[173] S. Niu, Y Lu, Y. Lin, and K. Liu. Effectiveness of light intervention on daytime sleepiness, fatigue, and sleep quality among female shift-working nurses - a randomized control trail. *Applied Nursing Research*, 86:152017, 2025.

[174] J. J. Gooley, S. M. W. Rajaratnam, G. C. Brainard, R. E. Kronauer, C. A. Czeisler, and S. W. Lockley. Spectral responses of the human circadian system depend on the irradiance

and duration of exposure to light. *Science Translational Medicine*, 2(31):31ra33–31ra33, 2010.

[175] G. V. Vartanian, B. Y. Li, A. P. Chervenak, O. J. Walch, W. Pack, P. Ala-Laurila, and K. Y. Wong. Melatonin suppression by light in humans is more sensitive than previously reported. *Journal of Biological Rhythms*, 30(4):351–4, 2015.

[176] C. Nowozin, A. Wahnschaffe, A. Rodenbeck, J. de Zeeuw, S. Hadel, R. Kozakov, H. Schopp, M. Munch, and D. Kunz. Applying melanopic lux to measure biological light effects on melatonin suppression and subjective sleepiness. *Current Alzheimer Research*, 14(10):1042–52, 2017.

[177] A. S. Prayag, R. P. Najjar, and C. Gronfier. Melatonin suppression is exquisitely sensitive to light and primarily driven by melanopsin in humans. *Journal of Pineal Research*, 66(4):e12562, 2019.

[178] J. Hong, S. J. Choi, S. H. Park, H. Hong, V. Booth, and E. Y. et al. Joo. Personalized sleep-wake patterns aligned with circadian rhythm relieve daytime sleepiness. *iScience*, 24:103129, 2021.

[179] Y. M. Song, S. J. Choi, S. H. Park, S. J. Lee, E. Y. Joo, and J. K. Kim. A real-time, personalized sleep intervention using mathematical modeling and wearable devices. *Sleep*, 46:1–13, 2023.

[180] J. Klooster, G. F. J. M. Vrensen, L. J. Müller, and J. J. L. Van der Want. Efferent projections of the olivary pretectal nucleus in the albino rat subserving the pupillary light reflex and related reflexes a light microscopic tracing study. *Brain Research*, 688(2):34–46, 1995.

[181] M. Spitschan. Photoreceptor inputs to pupil control. *Journal of Vision*, 19(9):1–5, 2019.

[182] P. D. Gamlin, D. H. McDougal, J. Pokorny, V. C. Smith, K. W. Yau, and D. M. Dacey. Human and macaque pupil responses driven by melanopsin-containing retinal ganglion cells. *Vision Research*, 47(7):946–954, 2007.

[183] S. Tsujimura, K. Ukai, D. Ohama, A. Nuruki, and K. Yunokuchi. Contribution of human melanopsin retinal ganglion cells to steady- state pupil responses. *Proceedings of the Royal Society B – Biological Sciences*, 277:2485–92, 2010.

[184] P. A. Barrionuevo and D. Cao. Luminance and chromatic signals interact differently with melanopsin activation to control the pupil light response. *Journal of Vision*, 16(11):29, 2016.

[185] L. A. Ostrin, K. S. Abbott, and H. M. Queener. Attenuation of short wavelengths alters sleep and the ip rgc pupil response. *Ophthalmic and Physiological Optics*, 37(4):440–50, 2017.

[186] M. Spitschan, S. Jain, D. H. Brainard, and G. K. Aguirre. Opponent melanopsin and S-cone signals in the human pupillary light response. *Proceedings of the National Academy of Science of the United States of America*, 111(43):15568–15572, 2014.

[187] H. McAdams, A. Igdalova, M. Spitschan, D. H. Brainard, and G. K. Aguirre. Pulses of melanopsin-directed contrast produce highly reproducible pupil responses that are insensitive to a change in background radiance. *Investigative Ophthalmology and Vision Science*, 59:5615–26, 2018.

[188] M. Spitschan, M. Gardasevic, F. P. Martial, R. J. Lucas, and A. Allen. Pupil responses to hidden photoreceptor-specific modulations in movies. *PLOS One*, 14(5):e0216307, 2019.

[189] S. Mathôt, H. Berberyan, P. Büchel, V. Ruuskanen, A. Vilotijević, and W. Kruijne. Effects of pupil size as manipulated through iprgc activation on visual processing. *NeuroImage*, 283:120420, 2023.

[190] CIE. CIE S 026/E:2018: CIE System for Metrology of Optical Radiation for ipRGC-Influenced Responses to Light. In *International Commission on Illumination (CIE), Vienna, AT*, 2018.

[191] T. Young. II. The Bakerian Lecture. on the theory of light and colours. *Philosophical Ttransactions of the Royal Society of London*, 92:12–48, 1802.

[192] D. H. Brainard and A. Stockman. *The Optical Society of America Handbook of Optics, 3rd edition, Volume III: Vision and Vision Optics.*, chapter Colorimetry. New York: McGraw Hill, 2010.

[193] O. Estévez and H. Spekreijse. The "silent substitution" method in visual research,. *Vision Research*, 22:681–91, 1982.

[194] M. Spitschan and T. Woelders. The Method of Silent Substitution for Examining Melanopsin Contributions to Pupil Control. *Frontiers of Neurology*, 9:941, 2018.

[195] D. H. Foster, K. Amano, S. M. Nascimento, and M. J. Foster. Frequency of metamerism in natural scenes. *Journal of the Optical Society of America A*, 23(10):2359–72, 2006.

[196] A. C. Hexley, A. O. Yöntem, M. Spitschan, H. E. Smithson, and R. Mantiuk. Demonstrating a multi-primary high dynamic range display system for vision experiments. *Journal of the Optical Society of America A*, 37(4):A271–84, 2020.

[197] T. W. Nugent and A. J. Zele. A five-primary maxwellian-view display for independent control of melanopsin, rhodopsin, and three-cone opsins on a fine spatial scale. *Journal of Vision*, 22(12):20, 2022.

[198] M. Spitschan, G. K. Aguirre, and D. H. Brainard. Selective stimulation of penumbral cones reveals perception in the shadow of retinal blood vessels. *PLoS One*, 10(4):e0124328, 2015.

[199] M. A. Webster and J. D. Mollon. Adaptation and the color statistics of natural scenes. *Vision Research*, 37(23):3283–98, 1997.

[200] G. Tkačik, P. Garrigan, C. Ratliff, G. Milcinski, J. Klein, L. Seyfarth, P. Sterling, D. Brainard, and V. Balasubramanian. Natural images from the birthplace of the human eye. *PLoS One*, 6(6):e20409, 2011.

[201] T. Morimoto, S. Kishigami, J. M. M. Linhares, S. M. C. Nascimento, and H. E. Smithson. Hyperspectral environment illumination maps: characterising directional spectral variation in natural environments. *Optics Express*, 27(22):32277–93, 2019.

[202] S. M. C. Nascimento, F. Vio, P. Ferreira, and D. H. Foster. Statistics of spatial cone-excitation ratios in natural scenes. *Journal of the Optical Society of America A*, 19(8), 2002.

[203] Specim. Specim iq portable hyperspectral imaging camera. `https://www.specim.com/iq/`, 2025.

[204] Westboro Photonics. Imaging circadiometers. `https://wphotonics.com/products/imaging-circadiometers/`, 2025.

[205] S. Ezpeleta, E. Orduna-Hospital, T. Solana, J. Aporta, I. Pinilla, and A. Sánchez-Cano. Analysis of photopic and melanopic lighting in teaching environments. *Buildings*, 11:439, 2021.

[206] P. Lalande, C. M. H. Demers, F. J. Lalonde, A. Potvin, and M. Hérbert. patial representations of melanopic light in architecture. *Architectural Science Review*, 64(6):522–33, 2021.

[207] J. Hraška. Nonvisual aspects of daylight in the built environment. *IOP Conference Series: Materials Science and Engineering*, 1252:012063, 2022.

[208] H. J. Bailes and R. J. Lucas. Human melanopsin forms a pigment maximally sensitive to blue light ($\lambda \approx 479nm$) supporting activation of $G_{q/11}$ and $G_{i/o}$ signalling cascades. *Proceedings of the Royal Soceity B Biological Sciences*, 280:1759, 2013.

[209] R. J. Lucas, M. S. Freedman, M. Muñoz, J. M. Garcia-Hernández, and R. G. Foster. Regulation of the Mamillian Pineal by Non-rod, Non-cone, Ocular Photoreceptors. *Science*, 284(5413):505–7, 1999.

[210] M. von Schantz, I. Provencio, and R. G. Foster. Recent developments in circadian photoreception: More than meets the eye. *Investigations in Ophthalmology & Vision Science*, 41(7):1605–7, 2000.

[211] S. Hattar, H. Liao, M. Takao, D. M. Berson, and K. Yau. Melanopsin-containing retinal ganglion cells: Architecture, projections, and intrinsic photosensitivity. *Science*, 295:1065–70, 2002.

[212] D. M. Berson, F. A. Dunn, and M. Takao. Phototransduction by retinal ganglion cells that set the circadian clock. *Science*, 295:1070–3, 2002.

[213] C. Cajochen. Alerting effects of light. *Sleep Medicine Reviews*, 11(6):453–464, 2007.

[214] J. L. Souman, A. M. Tinga, S. F. te Pas, R. van Ee, and B. N. S. Vlaskamp. Acute alerting effects of light: A systematic literature review. *Behavioural Brain Research*, 337:228–39, 2018.

[215] J. M. Zeitzer, D. J. Dijk, R. Kronauer, E. Brown, and C. Czeisler. Sensitivity of the human circadian pacemaker to nocturnal light: melatonin phase setting and suppression. *Journal of Endocrinology & Metabolism*, 526:695–702, 2000.

[216] J. P. Hanifin, S. W. Lockley, K. Cecil, K. West, M. Jablonski, and et al. Warfield, B. Randomized trial of polychromatic blue-enriched light for circadian phase setting, melatonin suppression, and alerting responses. *Physiology and Behaviour*, 198:57–66, 2019.

[217] I. Ho Mien, E. C. Chua, P. Lau, L. C. Tan, I. T. Lee, and et al. Yeo, S. C. Effects of exposure to intermittent versus continuous red light on human circadian rhythms, melatonin suppression, and pupillary constriction. *PLoS One*, 9(5):e96532, 2014.

[218] M. I. Silvani, R. Werder, and C. Perret. The influence of blue light on sleep, performance and wellbeing in young adults: A systematic review. *Frontiers in Physiology*, 13:943108, 2022.

[219] W. T. Ham, H. A. Mueller, and D. H. Sliney. Retinal sensitivity to damage from short wavelength light. *Nature*, 260:153–5, 1976.

[220] M. G. Figueiro, B. Wood, B. Plitnick, and M. S. Rea. The impact of light from computer monitors on melatonin levels in college students. *Neuro Endocrinology Letters*, 32:158–63, 2011.

[221] J. G. Lawrenson, C. C. Hull, and L. E. Downie. The effect of blue-light blocking spectacle lenses on visual performance, macular health and the sleep-wake cycle. *Ophthalmic and Physiological Optics*, 37:644–54, 2017.

[222] R. Kazemi, N. Alighanbari, and Z. Zamanian. The effects of screen light filtering software on cognitive performance and sleep among night workers. *Health Promotion Perspectives*, 9(3):233–40, 2019.

[223] C. Formentin, S. Carraro, M. Turco, L. Zarantonello, P. Angeli, and S. Montagnese. Effect of morning light glasses and night short-wavelength filter glasses on sleep-wake rhythmicity in medical inpatients. *Frontiers in Physiology*, 11(5), 2020.

[224] A. Partyka, M. Dyrek, P. Sikorski, and E. Sikorska. The benefits and limitations of using blue-blocking filters – a review of results of clinical trials. *European Journal of Clinical and Experimental Medicine*, 22(1):194–200, 2024.

[225] M. Spitschan, R. Lazar, and C. Cajochen. Visual and non-visual properties of filters manipulating short-wavelength light. *Ophthalmic and Physiological Optics*, 39:459–68, 2019.

[226] R. Lazar, F. Fazlali, M. Dourte, C. Epple, O. Stefani, M. Spitschan, and C. Cajochen. Afternoon to early evening bright light exposure reduces later melatonin production in adolescents. *NPJ Biological Timing and Sleep*, 2(1):1–19, 2024.

[227] A. E. Brøndsted, J. H. Lundeman, and L. Kessel. Short wavelength light filtering by the natural human lens and iols – implications for entrainment of circadian rhythm. *Acta ophthalmologica*, 91(1):52–7, 2013.

[228] J. Pokorny, V. C. Smith, and M. Lutze. Ageing of the human lens. *Applied Optics*, 26(8):1437–40, 1987.

[229] CIE. Light as a true visual quantity: Principles of measurement. *CIE*, 1994.

[230] M. Ashraf, R. K. Mantiuk, A. Chapiro, and S. Wuerger. castlecsf — a contrast sensitivity function of color, area, spatiotemporal frequency, luminance and eccentricity. *Journal of Vision*, 24(4), 2024.

[231] G. Glickman, E. M. Harrison, M. Herf, L. Herf, and T. M. Brown. Optimizing the potential utility of blue-blocking glasses for sleep and circadian health. *Translational Vision Science and Technology*, 14(25), 2025.

[232] J. Neitz, J. Carroll, Y. Yamauchi, M. Neitz, and D. R. Williams. Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, 35:783–92, 2002.

[233] B. Wozniak, J. Maule, A. Franklin, and J. Bosten. Color discrimination and chromatic balance perception after adaptation to natural and color-reflected scenes. *Journal of Vision*, 23(9), 2023.

[234] A. E. Skelton, J. Maule, S. Floyd, B. Wozniak, A. Majid, J. M. Bosten, and A. Franklin. Effects of visual diet on colour discrimination and preference. *Proceedings of the Royal Society B*, 291, 2024.

[235] T. Morimoto and H. E. Smithson. Discrimination of spectral reflectance under environmental illumination. *Journal of the Optical Society of America A*, 35(4):B244–55, 2018.

[236] A. E. Skelton, A. Franklin, and J. M. Bosten. Color vision is aligned with natural scene statistics at 4 months of age. *Developmental Science*, 26:e13402, 2023.

[237] T. Morimoto, C. Zhang, K. Fukada, and K. Uchikawa. Spectral measurement of daylights and surface properties of natural objects in japan. *Optics Express*, 30(3):3183–204, 2022.

[238] Y. Nayatani and G. Wyszecki. Color of daylight from north sky. *Journal of the Optical Society of America A*, 53(5):626–9, 1963.

[239] G. T. Winch, M. C. Boshoff, C. J. Kok, and A. G. Du Toit. pectroradiometric and colorimetric characteristics of daylight in the southern hemisphere: Pretoria, south africa. *Journal of the Optical Society of America A*, 56(4):456–64, 1966.

[240] M. A. Webster, Y. Mizokami, and S. M. Webster. Seasonal variations in the color statistics of natural images. *Network: Computation in Neural Systems*, 18(3):213–33, 2007.

[241] L. E. Welbourne, A. B. Morland, and A. R. Wade. Human colour perception changes between seasons. *Current Biology*, 25(15):R646–7, 2015.

[242] J. Vincent, A. M. Kale, and S. L. Buck. Luminance-dependent long-term chromatic adaptation. *Journal of the Optical Society of America A*, 33(3):A164–9, 2016.

[243] S. L. Buck and T. DeLawyer. Dark versus bright equilibrium hues: rod and cone biases. *Journal of the Optical Society of America A*, 31(4):A75–81, 2014.

[244] N. E. Klepeis, W. C. Nelson, W. R. Ott, J. P. Robinson, A. M. Tsang, P. Switzer, J. V. Behar, S. C. Hern, and W. H. Engelmann. The National Human Activity Pattern Survey (NHAPS): A resource for assessing exposure to environmental pollutants. *Journal of Exposure Science and Environmental Epidemiology*, 11:231–52, 2001.

[245] Y. Zeng, H. Sun, B. Lin, and Q. Zhang. Non-visual effects of office light environment: Field evaluation, model comparison and spectral analysis. *Building and Environment*, 197:107859, 2021.

[246] O. Stefani and C. Cajochen. hould we re-think regulations and standards for lighting at workplaces? a practice review on existing lighting recommendations. *Frontiers in Psychiatry*, 12:652161, 2021.

[247] R. Zhang, C. Campanella, S. Aristizabal, A. Jamrozik, J. Zhao, P. Porter, S. Ly, and B. A. Bauer. Impacts of dynamic led lighting on the well-being and experience of office occupants. *International Journal of Environmental Research and Public Health*, 17(19):7217, 2020.

[248] A. B. Watson and J. I. Yellott. A unified formula for light-adapted pupil size. *Journal of Vision*, 12(10):1–16, 2012.

[249] R. F. Hess and K. Nordby. Spatial and temporal properties of human rod vision in the achromat. *Journal of Physiology*, 371(1):387–406, 1986.

[250] A. G. Shapiro, J. Pokorny, and V. C. Smith. Cone-rod receptor spaces with illustrations that use crt phosphor and light-emitting-diode spectra. *Journal of the Optical Society of America A*, 13(12):2319–28, 1996.

[251] Date and Time. Oxford, england, united kingdom.

[252] L. Shiwen, L. Steele, C. A. L. Dahlsjö, S. N. Pierson, A. Shenkin, T. Morimoto, H. E. Smithson, and M. Spitschan. Hyperspectral characterisation of natural illumination in woodland forest environments. *bioRxiv*, 2021.

[253] H. E. Smithson, P. S. Anderson, G. Dinkova-Bruun, R. A. E. Fosbury, G. E. M. Gasper, P. Laven, T. C. B. McLeish, C. Panti, and B. K. Tanner. Color-coordinate system from a 13th-century account of rainbows. *Journal of The Optical Society of America A*, 31(4):A341–9, 2014.

[254] P. A. Barrionuevo and D. Cao. Contributions of rhodopsin, cone opsins and melanopsin to postreceptoral pathways inferred from natural scene statistics. *Journal of The Optical Society of America A*, 31(4):A131–9, 2014.

[255] M. Englezhou and A. Michael. Investigation of the daylight spectrum in an indoor environment using the cie s 026 melanopic metrics. *Lighting Research & Technology*, 55(7):690–711, 2023.

[256] C. Yu, S. Pont, R. Pastilha, and A. Hurlbert. Unveiling the temporal dynamics of diurnal and crepuscular illumination. *Journal of Vision [Vision Sciences Society Annual Meeting Abstract]*, 24(10):1211, 2024.

[257] D. H. Foster, K. Amaro, and S. M. C. Nascimento. Time-lapse ratios of cone excitations in natural scenes. *Vision Research*, 120:45–60, 2016.

[258] S. Tir, R. G. Foster, and S. N. Peirson. Evaluation of the digital ventilated cage$^{\ddot{o}}$ system for circadian phenotyping. *Scientific Reports*, 15(1):3674, 2025.

[259] C. Yu, M. Wijntjes, E. Eisemann, and S. Pont. Quantifying the spatial, temporal, angular and spectral structure of effective daylight in perceptually meaningful ways. *Optics Express*, 31(5):8953–74, 2023.

[260] M. Spitschan, G. K. Aguirre, D. H. Brainard, and A. M. Sweeney. Variation of outdoor illumination as a function of solar elevation and light pollution. *Scientific Reports*, 6(1):26756, 2016.

[261] T. Machida, N. Yokoya, and H. Takemure. Surface reflectance modelling of real objects with interreflections. *Proceedings of IEEE International Conference on Computer Vision*, pages 170–7, 2003.

[262] D. H. Brainard and W. T. Freeman. Bayesian color constancy. *Journal of The Optical Society of America A*, 14(7):1393–411, 1997.

[263] CIE. Discomfort Glare in the Interior Working Environment, Publication CIE No. 55 (TC-3.4). *Commission Internationale de l'Eclairage (CIE), Vienna, Austria*, page 43, 1983.

[264] J. Wienold and J. Christoffersen. Evaluation methods and development of a new glare prediction model for daylight environments with the use of ccd cameras. *Energy and Buildings*, 2006.

[265] International WELL Building Institute. WELL Building Standard v2.

[266] K. Y. Wong, F. A. Dunn, and D. M. Berson. Photoreceptor adaptation in intrinsically photosensitive retinal ganglion cells. *Neuron*, 48(6):1001–10, 2005.

[267] A. Mury, S. Pont, and J. Koenderink. Representing the light field in finite three-dimensional spaces from sparse discrete samples. *Applied Optics*, 48(3):450–7, 2009.

[268] Y. Li, Y. Wang, W. Huang, and Z. Zhang. Automatic image stitching using SIFT. *Proceedings of the IEEE Conference on Audio, Language and Image Processing*, pages 568–71, 2008.

[269] Z. Wang, A. C. Bovick, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4), 2004.

[270] J. Nilsson and T. Akenine-Möller. Understanding ssim. *ArXiv*, 2020.

[271] A. Mury, S. Pont, and J. Koenderink. Structure of light fields in natural scenes. *Applied Optics*, 48(28):5386–95, 2009.

[272] L. Xia, S. C. Pont, and I. Heynderickx. Light diffuseness metric, pt. 1: theory. *Lighting Research and Technology*, 49(4):411–27, 2017.

[273] F. Zhang, H. de Ridder, P. Barla, and S. Pont. Effects of light map orientation and shape on the visual perception of canonical materials. *Journal of Vision*, 20(4):13, 2020.

[274] S. M. Berman and R. D. Clear. A practical metric for melanopic metrology. *Lighting Research & Technology*, 51(9):1178–91, 2019.

[275] P. Lalande, M. Hébert, A. Potvin, J. F. Lalonde, M. Watchman, and C. M. Demers. Representing the photobiological dimension of light in northern architecture. *Indoor and Built Environment*, 32(9):1875–90, 2023.

[276] P. Lalande, M. Hébert, J. F. Lalonde, C. Espinoza-Sanhueza, and C. M. Demers. Imagery dataset of interior and exterior northern architectural spaces for photobiological lighting analyses. *Data in Brief*, 55:110658, 2024.

[277] S. Ponting, R. Kelly Waskett, M. Spitschan, and H. E. Smithson. Quantifying cie alpha-opic signals in the indoor built environment. *Journal of the Optical Society of America A*, 42:B379–90, 2025.

[278] K. Scheuermaier, A. M. Laffan, and J. F. Duffy. Light exposure patterns in healthy older and young adults. *Journal of Biological Rhythms*, 25:113–22, 2010.

[279] S. L. Hartmeyer, A. Davies, and M. Andersen. Post-occupancy evaluation of office lighting – what can we learn from light- dosimetry? *LEUKOS*, 2024.

[280] S. W. de Vries, M. Gkaintatzi-Masouti, J. van Duijnhoven, J. Mardaljevic, and M. P. J. Aarts. Recommendations for light-dosimetry field studies based on a meta-analysis of personal light levels of office workers. *Lighting Research and Technolgy*, 57:47–70, 2025.

[281] M. Spitschan, J. Zauner, M. N. Tengelin, C. A. Bouroussis, P. Caspar, and F. on behalf of the MeLiDos consortium Eloi. Illuminating the future of wearable light metrology: Overview of the melidos project. *Measurement*, 235:114909, 2024.

[282] C. Wang, D. Yan, and X. Ren. Modeling individual's light switching behavior to understand lighting energy use of office building. *Energy Prodecia*, 88:781–7, 2016.

[283] K. J. Leoniak and W. Cwalina. The role of normative prompts and norm support cues in promoting light-switching behavior: A field study. *Journal of Environmental Psychology*, 64:1–11, 2019.

[284] A. M. Biller, P. Balakrishnan, and M Spitschan. Behavioural determinants of physiologically-relevant light exposure. *Communications Psychology*, 2, 2024.

[285] K. R. Bhandari, H. Mirhajianmoghadam, and L. A. Ostrin. Wearable sensors for measurement of viewing behavior, light exposure, and sleep. *Sensors*, 21(21):7096, 2021.

[286] S. L. Hartmeyer, F. S. Webler, and M. Andersen. Towards a framework for light-dosimetry studies: Methodological considerations. *Lighting Research and Technology*, 55(4-5):377–99, 2023.

[287] J. R. Stampfli, B. Schrader, C. di Battista, R. Häfliger, O. Schälli, G. Wichmann, C. Zumbühl, P. Blattner, C. Cajochen, R. Lazar, and M. Spitschan. The light-dosimeter: A new device to help advance research on the non-visual responses to light. *Lighting Research and Technology*, 55(4-5):474–86, 2023.

[288] S. L. Hartmeyer and M. Andersen. Towards a framework for light-dosimetry studies: Quantification metrics. *Lighting Research and Technology*, 56(4):337–65, 2024.

[289] L. Price, M. Khazova, and J. O'Hagan. Performance assessment of commercial circadian personal exposure devices. *Lighting Research & Technology*, 44(1):17–26, 2012.

[290] M. P. J. Aarts, J. van Duijnhoven, M. B. C. Aries, and A. L. P. Rosemann. Performance of personally worn dosimeters to study non-image forming effects of light: Assessment methods. *Building and Environment*, 117:60–72, 2017.

[291] M. Münch, A. Wirz-Justice, S. A. Brown, T. Kantermann, K. Martiny, O. Stefani, C. Vetter, K. P. Wright, K. Jr., Wulff, and D. J. Skene. The role of daylight for humans: Gaps in current knowledge. *Clocks and Sleep*, 2(1):61–85, 2020.

[292] S. L. Hartmeyer, F. R. Rudawski, M. Knoop, and M. Andersen. Exploring spatiotemporal dynamics in light-dosimetry. *Journal of Physics: Conference Series*, 2600:112005, 2023.

[293] V. Grygorenko. Sensing - magnetic compass with tilt compensation. *Cypress Perform*, 001-32379*C, 2014.

[294] M. Adamsson, L. Thorbjörn, and M. Takeshi. Seasonal variation in bright daylight exposure, mood and behavior among a group of office workers in sweden. *Circadian Rhythms*, 16:2, 2018.

[295] I. Gilchrist, V. Brown, and J. Findlay. Saccades without eye movements. *Nature*, 390:130–1, 1997.

[296] I. D. Gilchrist, V. Brown, J. M. Findlay, and M. P. Clarke. Using the eye–movement system to control the head. *Proceedings of the Royal Society of London B*, 265(1408):1831–6, 1998.

[297] J. M. Samonds, W. S. Geisler, and N.J. Priebe. Natural image and receptive field statistics predict saccade sizes. *Nature Neuroscience*, 21:1591–9, 2018.

[298] N. Paraskevoudi and J. S. Pezaris. Full gaze contingency provides better reading performance than head steering alone in a simulation of prosthetic vision. *Scientific Reports*, 11:11121, 2021.

[299] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015.

[300] K. Baumann, C.and Dierkes. Neon accuracy test report. *Pupil Labs*, 2023.

[301] U. Nakade and M. Spitschan. Predicting visual field boundaries from head features. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 42(6):C19–26, 2025.

[302] B. Xiao, B. Hurst, L. MacIntyre, and D. H. Brainard. The color constancy of three-dimensional objects. *Journal of Vision*, 12(4):1–15, 2012.

[303] D. L. MacAdam. Visual sensitivities to color differences in daylight. *Journal of the Optical Society of America*, 32(5):247, 1942.

[304] J. Krauskopf and K. Gegenfurtner. Color discrimination and adaptation. *Vision Research*, 32(11):2165–75, 1992.

[305] D. L. MacAdam. *Sources of color science.* MIT Press, 1970.

[306] R. Lafer-Sousa, K. L. Hermann, and B. R. Conway. Striking individual differences in color perception uncovered by 'the dress' photograph. *Current Biology*, 25(13):R545–6, 2015.

[307] L. D. Daoudi, M. Kunchulia, and M. H. Herzog. The role of one-shot learning in #thedress. *Journal of Vision*, 17:1–7, 2017.

[308] S. Aston and A. Hurlbert. What #thedress reveals about the role of illumination priors in color perception and color constancy. *Journal of Vision*, 2017.

[309] C. Witzel, C. Racey, and J. K. O'Regan. The most reasonable explanation of "the dress": implicit assumptions about illumination. *Journal of Vision*, 17(2), 2017.

[310] R. A. Warwick, N. Kaushansky, N. Sarid, A. Golan, and M. Rivlin-Etzion. Inhomogeneous encoding of the visual field in the mouse retina. *Current Biology*, 28(5):655–65, 2018.

[311] Y. Qiu, Z. Zhao, D. Klindt, M. Kautzky, K. P. Szatko, F. Schaeffel, K. Rifai, K. Franke, L. Busse, and T. Euler. Natural environment statistics in the upper and lower visual field are reflected in mouse retinal specializations. *Current Biology*, 31(5):3233–47, 2021.

[312] D. Brainard. The psychophysics toolbox. *Spatial Vision*, 10(4):433–6, 1997.

[313] R. A. Bone and J. T. Landrum. Heterochromatic flicker photometry. *Archives of Biochemistry and Biophysics*, 430(2):137–42, 2004.

[314] T. Morimoto. *Colour constancy for glossy objects under complex lighting environments.* PhD thesis, University of Oxford, 2019.

[315] N. Prins. The psi-marginal adaptive method: How to give nuisance parameters the attention they deserve (no more, no less). *Journal of Vision*, 13(7):3, 2013.

[316] J. Krauskopf, D. R. Williams, M. B. Mandler, and A. M. Brown. Higher order color mechanisms. *Vision research*, 26(1):23–32, 1986.

[317] M. W. Levine and J. J. McAnany. The relative capabilities of the upper and lower visual hemifields. *Vision Researcher*, 45(21):2820–30, 2005.

[318] M. Danilova and J. Mollon. The symmetry of visual fields in chromatic discrimination. *Brain and Cognition*, 69:39–46, 2008.

[319] A. Skelton, S. Floyd, J. Maule, B. Wozniak, A. Majid, A. Franklin, and J. Bosten. Is color discrimination influenced by the chromatic statistics of different visual environments? *Journal of Vision*, 21(9):1945, 2021.

[320] F. H. Previc. Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. *Behavioural and Brain Sciences*, 13(3):519–42, 1990.

[321] M. D. Fairchild and L. Reniff. Time course of chromatic adaptation for color-appearance judgments. *Journal of the Optical Society of America A*, 12:824–33, 1995.

[322] A. Werner. Spatial and temporal aspects of chromatic adaptation and their functional significance for colour constancy. *Vision Research*, 104:80–9, 2014.

[323] F. Hong, R. Bouhassira, J. Chow, C. Sanders, M. Shvartsman, P. Guan, A. H. Williams, and D. H. Brainard. Comprehensive characterization of human color discrimination thresholds. *bioRxiv*, 2025.07.16.665219, 2025.

[324] W. R. J. Brown. The influence of luminance level on visual sensitivity to color differences. *Journal of the Optical Society of America A*, 41:684–8, 1951.

[325] J. Koenderink, A. van Doorn, and K. Gegenfurtner. Colors and things. *I-Perception*, 11(6), 2020.

[326] A. Gutiérrez, B. Silva, J. M. Fanchini, T. Morimoto, P. A. Barrionuevo, and M. L. Sandoval-Salinas. Spectral dataset of natural objects' reflectance from the southern cone of south america. *Scientific Data*, 12(1), 2025.

[327] L. Owen, J. Browder, B. Letham, G. Stocek, C Tymms, and M. Shvartsman. Adaptive nonparametric psychophysics. *arXiv*, 210409549, 2021.

[328] N. P. Cottaris, H. Jiang, X. Ding, B. A. Wandell, and D. H. Brainard. A computational-observer model of spatial contrast sensitivity: Effects of wave-front-based optics, cone-mosaic structure, and inference engine. *Journal of Vision*, 19, 2019.

[329] Ministry of Housing, Communities & Local Government and Department for Levelling Up, Housing & Communities. Conservation of fuel and power: Approved document L. *GOV.UK*, 2023.

[330] British Standards Institution (BSI). Bs en 12464-1:2021, light and lighting — lighting of work places — indoor work places. Technical report, British Standards Institution (BSI), 2021.

[331] BRE Group (n.d.) BREEAM Infrastructure. Available at: https://bre.ac/breeam-infrastructure/ (Accessed: 1 June 2026).

[332] CIE and ISO. CISO 8995-1:2002(E)/CIE S 008/E:2001 Lighting of work places - Part 1: Indoor. In *International Commission on Illumination (CIE), Vienna, AT*, 2001.

[333] E. Kolberg, S. Pallesen, G. Hjetland, I. Nordhus, E. Thun, and E. Flo-Groeneboom. Insufficient melanopic equivalent daylight illuminance in nursing home dementia units across seasons and gaze directions. *Lighting Research and Technology*, 54(2):163–77, 2022.

[334] R. Foster. Why do we sleep? In *TedGlobal2013*, 2013.

[335] BBC Sounds. *Just One Long Thing: Sleep Special – with Prof Russell Foster*. BBC Sounds, 2024.

[336] C. Weinzaepflen and M. Spitschan. Enlighten your clock: How your body tells time. (c. weinzaepflen, illus.)., 2021.

[337] C. Guidolin, M. Spitschan, A. Hahn, A. M. Biller, H. S. Heinrichs, I. Bednarek, L. Bickerstaff, M. Stegmann, M. Fernandez-Alonso, and N. Tabandeh. *Light O'Clock*. Translational Sensory and Circadian Neuroscience Unit, 2025.

[338] C. Weinz, A. M. (Ed.) Biller, A. (Ed.) Parreño, S. (Ed.) Stegbauer, V. (Ed.) Streibel, and M. (Ed.) The. *Lumi's Delight*. C. Weinz, Illus., 2025.

[339] Light Tribe by Phos. *In Light of it All*. Phos, 2025.

[340] M. Inanici, M. Brennan, and E. Clark. Multi-spectral lighting simulations: Computing circadian light. *International Building Performance Simulation Association (IBPSA) 2015 Conference, Hyderabad, India*, 2015.

[341] M. Gkaintatzi-Masouti, C. Pierson, J. Van Duijnhoven, M. Andersen, and Aarts M. A simulation tool for indoor lighting design considering iprgc-induced responses. *BuildSim Nordic Conference, Copenhagen, Denmark*, 2022.

[342] M. Spitschan, S. Nam, and J. A. Veitch. luox: Platform for calculating quantities related to light and lighting [software]. *Available from https://luox.app/*, 2022.

[343] M. Spitschan, J. Mead, C. Roos, C. Lowis, B. Griffiths, P. Mucur, M. Herf, S. Nam, and J. A. Veitch. luox: novel validated open-access and open-source web platform for calculating and sharing physiologically relevant quantities for light and lighting. *Wellcome Open Res*, 6(69), 2022.

[344] B. Jung, Z. Cheng, M. Brennan, and M. Inanici. Multispectral lighting simulation approaches for predicting opsin-driven metrics and their application in a neonatal intensive care unit. *IBPSA 2023 Conference, Shanghai, China*, 2023.

[345] M. T. Aguilar-Carrasco, I. Acosta, and S. Domínguez-Amarillo. Circalight, a new circadian light assessment tool for grasshopper environment: Development and reliability testing. *Lighting and Building Engineering*, 71:106527, 2023.