Partial independence of brightness induction and brown inducton suggests a two-stage model for brightness induction

Joris Vincent

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2017

Reading Committee:

Steven L. Buck, Chair

Geoffrey M. Boyton

Ione Fine

Program Authorized to Offer Degree:
Psychology

University of Washington

Abstract

Partial independence of brightness induction and brown inducton
suggests a two-stage model for brightness induction

Joris Vincent

Chair of the Supervisory Committee:
Dr. Steven L. Buck
Psychology

Bright contexts surrounding achromatic stimuli generally induce perceptual darkness in those targets. In yellowish targets, such contexts also induce a qualitative brownness. This raises the question whether brownness induction is mediated by the same mechanisms as darkness and brightness induction in achromatic targets - whether brown is just dark yellow. While darkness induction by simple, uniform bright surrounds is usually explained by contrast mechanisms such a lateral inhibition, several brightness/darkness induction effects cannot solely be explained by these mechanisms. White's (1979) illusion consists of a high-contrast square-wave grating in which one target replace part of the white bars and another part target replaces of the black bars. The target replacing the white bars shares a larger border with the flanking black bars, but is usually perceived as brighter than the other target - contrary to the predictions of contrast mechanisms. Instead, this induction has been explained separately by low-level (e.g., multiscale spatial filtering accounts), mid-level (e.g., based on T-junctions, or Gestalt grouping), and higher-level processing (e.g., scission of the scene into different depth planes). This complex illusion is well-suited for studying the association between brightness induction and brownness induction. Participants judged both the brightness and brownness induction in White's classic illusion, a checkerboard variant, and a radial variant. Across participants, both brightness and brown induction were in the direction of assimilation more

often than in the direction of contrast, although the frequencies differed among the different stimuli. On the classic illusion, brownness induction and brightness induction were independent in their direction of effect. Achromatic brightness induction was almost exclusively observed in the assimilation direction, while brownness induction in either direction was observed across participants. This would suggest that brownness induction and brightness induction are mediated by independent mechanisms, producing uncorrelated effects. On the radial and checkerboard variants, participants showed congruent direction of brightness and brown induction more often than expected by chance. In contrast to the classic illusion, these results would suggest that brownness induction and brightness induction are mediated by shared mechanisms, producing the correlated effects. To resolve these conflicting results, a framework is suggested with at least two stages (lower- and higher-level) of processing underlying induction. A shared initial stage of visual processing is involved in both brightness and brownness induction, but the second stage consists of independent mechanisms for brightness and brownness. Different aspects of a stimulus display drive the stages differentially. The radial and checkerboard variants are suggested to drive the initial stage more strongly than the second stage, explaining the lack of independent brightness and brownness induction effects for these stimuli. The classic illusion is suggested to drive the second stage more strongly (than the radial and checkerboard variants), producing independent induction effects. Specifically, the second stage mechanisms would bias the brightness induction, but not the brownness induction, towards the assimilation direction. Several previous explanations of Whites illusion are discussed in relation to these two stages, providing a framework for integrating theories of brightness processing at different levels. The two-stage framework also generates several hypotheses about how well the brightness or brownness induction in a given stimulus predicts performance on another stimulus, or the other induction.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

This dissertation and the underlying research has been a project of love and frustration, stress and joy. For all the hours, effort and pain I put into, I could not have done it without the help and support of a great amount of people.

First of all, Steve, who has taught me more about conducting research than I could have hoped for. Thanks to your mentoring, I feel very confident in my abilities to design and conduct psychophysical experiments, and in general to ask exactly the questions necessary to get the answers I am looking for. Beyond that, you have been an inspiration for me in how to think about scientific research. While it is easy to get lost in the details, you set the example of never forgetting the bigger picture, and in vision science, that bigger picture should be "but how does it affect our perception? What do you see?". You have also always given me the room to explore my own scientific interests, even when those are not directly aligned with yours: you fully supported me devoting time and energy to developing my computational skills and interests, and we found an excellent way to turn that into a collaboration on the natural scenes paper. The greatest compliment this dissertation and research received, is you saying that there is no way anyone could confuse it for being your work - which I think is also a testament to your phenomenal mentoring.

The research project and dissertation would also not have been the same without the input and support from my committee: Geoff, Ione, John and Jay. All of you have had a major impact on how I view my own research, my academic career, and vision science in general. My committee has been a great support in branching out to more computational approaches, whether that meant recommending me for the Cold Spring Harbor Labs course, or sitting down with me to write some psychometric function fitting code. The questions

you posed for both my general and final exams also helped me think about my research and the field in new and exciting ways, and certainly some of those questions I would love to get back to at later stages in my career. I would like to thank all four of you for so excellently complementing Steve's mentoring, and shaping my career such that I can now move on to be a post-doc at Penn.

Throughout conducting the research in this dissertation, I have been lucky to share my thoughts and ideas with many other colleagues. In turn, their insightful comments have helped me shape this project better - whether it be drastic new ways of thinking about the problem, or smaller but no less important suggestions about wording. In particular, I have had very helpful conversations with Tanner, Hohjin, the Vision & Cognition group at the UW (particularly Kelly and Alex), and too many colleagues to list at the OSA Fall Vision Meeting, VSS, and ARVO. Whether you know it or not, there are parts of this dissertation that would not have been written without your input. And that has taught me one of the most important lessons about science: you should not try to conduct it alone.

In the five years that I spent on getting here, I have relied on other people not just for input and help with my research and academics. I also moved here from the other side of the world, and adjusting to that has not always been easy. Luckily, I am surrounded (both near and far) by amazing people who helped made it a lot easier. The Vision & Cognition group has helped make me feel at home in Seattle and at the UW. You made these 5 years in the basement of Guthrie slightly less grating, and sometimes even fun. In particular, I once again want to thank Steve, and Jeannette. You have shown a patience and support for me that has helped me thrive, such as letting me stay at your place while getting settled. I cannot thank you enough for providing that kind of foundation for my life in Seattle, which allowed me to focus more on developing myself.

During these five years, I am also lucky to have made some great friends in Seattle, who helped me relax after long days, celebrate victories, drown sorrows, and explore everything

this city (and some other places in the US) has to offer. My fellow cohortmembers in the psychology graduate program, in particular Frank, of course, Dave and Alli, Casey, everyone at Russell's, and everyone at Agrodolce - thank you for putting up with me for all these years. While I am very excited to start the next chapter of my life outside of Seattle, I will miss you all dearly.

And the hard part about that is, that I already have to miss many great friends back in Europe. While going home twice a year and trying to see as many people as I can certainly helps me keep in touch with most, Maarten, Manuel, Yosi and Colin have gone above and beyond that to remind me that I might be gone but not lost. It is certainly bittersweet to keep one foot grounded in Europe and the other in the US, but regularly talking, Skyping, protobowling, songguessing and gaming with you has been an excellent way of remaining a resident of both.

While presumably my European friends have missed me as much as I missed them, I am sure that Marlous, Gwen, mom and dad, have had to miss me more than anyone. I do not think that I can ever thank you enough. You not just supported me moving to, and staying, in Seattle for five years - you actually encouraged me. It takes a special kind of love to put your own interests aside like that. Beyond even that, you have also always had my back, listened when I needed to vent, and believed in me even when I myself did not. These five years, and this dissertation, would truly not have been possible without your help. Thank you.

DEDICATION

To my parents, who started me down this path all those years ago.

LIST OF FIGURES

LIST OF TABLES

Chapter 1

BROWN INDUCTION

Brown is one of the 11 basic color terms identified cross-culturally (Berlin & Kay, 1969): in languages across the world, a word for brown is the first to develop after the six Hering (1878) primary color categories of black, white, red, green, blue and yellow. These basic color terms are not only of linguistic significance. Most of them have been extensively studied and shown to have psychophysical properties that set them apart from other colors (Boynton & Olson, 1990). The six Hering (1878) primaries have long been appreciated to form 3 opponent pairs: black opposes white, red opposes green and blue opposes yellow. Of each opponent pair, only one end can be perceived in any given color. No color can look reddish and greenish at the same time, or blueish and yellowish. Mixtures between pairs are possible though: blueness and redness can combine to form purple, while yellowness and redness create orange. These six primaries thus provide the basis for most of modern color vision theory (Jameson & Hurvich, 1955; Hurvich & Jameson, 1955; De Valois & De Valois, 1993)

The notable exception to intense investigation has been brown, which up until a few years ago had largely been ignored, save a handful of studies (Kiesow, 1930; Bartleson, 1976; Uchikawa, Uchikawa, & Boynton, 1989; Boynton & Olson, 1990; Fuld, Wooten, & Whalen, 1981; Fuld, Werner, & Wooten, 1983; Quinn, Rosano, & Wooten, 1988). Yet brown has several unique and fascinating properties, that cannot yet be fully explained (Buck, 2015). What sets brown apart from the other 11 basic color terms (other than black), is that it is a dark color: it cannot be perceived as a bright light (Bartleson, 1976; DeLawyer et al., 2012), and is notably absent from the rainbow. Bartleson (1976) identified that to identify a surface as brown, it must be at low lightness (dim), in a specific range of hues (roughly orange), and ideally of intermediate strength (somewhat desaturated). However, he also noted that the

optimum for each of these dimensions depends on the others. He concluded that brown is not a hue category, but a color category: it requires definition by all three dimensions (hue, saturation and brightness) of color (Bartleson, 1976).

Thus, brown takes it place among the dark colors, where it is the counterpart to the bright color yellow. Brown and yellow, however, are not opponent. Both brownness and yellowness can be perceived in the same stimulus, as a caramel or butterscotch color (DeLawyer et al., 2012; Vincent, Buck, Armer, Delawyer, & Wilson, 2012) In fact, brown and yellow form a continuum: as brown stimulus is made lighter, it will transition from brown, through this range of butterscotch color combinations of yellow and brown, until all brownness disappears and only yellowness remains. While perceptually distinct from yellow, brown shares several key properties (DeLawyer et al., 2012; Buck, 2015).

First, like yellow, brown is opponent to blue (DeLawyer et al., 2012; Padgham & Saunders, 1975). A yellow percept can be canceled by adding blue light to it. Adding this blue light desaturates the yellow stimulus, until it looks achromatic, before adding perceptual blueness. Inversely, adding yellow light to a blue percept will cancel out the blueness. Again, this will first desaturate the blue stimulus until it appears achromatic, before adding perceptual yellowness. Brown shares the same relationship with blue. Adding blue light to a brown stimulus desaturates it, before adding perceptual blueness (DeLawyer et al., 2012). As a result, both yellowness and brownness are never perceived in the same stimulus as blueness; the blueness would have to cancel or be canceled out.

Second, both yellow and brown are red-green balance points (Kiesow, 1930; Padgham & Saunders, 1975; Buck & DeLawyer, 2012). A yellow stimulus can look reddish (peach or orange), it can look greenish, or it can be perceived without either redness or greenness. The redness in an orange percept can be canceled out by adding greenish, and conversely, any greenness in a yellow stimulus can be canceled by adding redness. Thus, an equilibrium yellow that appears neither reddish nor greenish represents a balance between redness and greenish, a red-green balance. Similarly, a brown percept can contain greenness (olive brown), or redness (maroon) (Bartleson, 1976; Buck & DeLawyer, 2012; Buck, 2015). The greenness in a

brown can be canceled out by adding redness, and any redness can be canceled out by adding greenness. Thus, an equilibrium brown, like an equilibrium yellow, represents a red-green balance. However, equilibrium brown and equilibrium yellow are perceived at different physical red-green balances: equilibrium brown requires more long wavelength light to be added to the stimulus than does equilibrium yellow (Buck & DeLawyer, 2012; DeLawyer et al., 2012; Buck, 2015). As a result, darkening an equilibrium yellow stimulus (without adjusting the red-green balance) creates a greenish brown: the red-green balance of the equilibrium yellow does not contain enough red for an equilibrium brown. Conversely, brightening an equilibrium brown (without adjusting the red-green balance) creates an orange percept, because the red-green balance contains too much red for an equilibrium yellow. These different balance points might be result of luminance dependent adaptation of the red-green opponent process (Vincent, Kale, & Buck, 2016).

Consequently, both yellow and brown can exist in absence of any other hue (Buck & DeLawyer, 2012; Bartleson, 1976; Uchikawa et al., 1989; Boynton & Olson, 1990). An equilibrium yellow will contain no redness, nor greenness, nor blueness; neither will an equilibrium brown. Thus, brown shares the crucial perceptual properties that make yellow one of the four opponent colors (Bartleson, 1976).

Bartelson on brown as a separate color and not just orange (or yellow) with a different label

Uchikawa et al. (1989) identified that brown is rarely identified by observers when presented in isolation on a black background. Under those conditions, it is often perceived as orange, because of its reddish shift compared to an equilibrium yellow (Bartleson, 1976). However, when surrounded by a sufficiently bright achromatic context, the same surfaces are much more likely to be identified as brown. Even a surround of one-sixteenth the size of the test stimulus is enough to yield substantial brownness in the percept (Uchikawa et al., 1989). Thus, brown is a context color: it can only be seen when presented alongside a sufficiently bright context. Figure 1.1 nicely illustrates this context dependence of brown. The center square on the topside of the cube is of identical chromaticity as the center square on the

front side. However, the former is surrounded by a bright context, and perceived as brown, while the latter is surrounded by a darker context and perceived as orange. Note that this categorical shift in color percept happens only for the orange and brown squares, but not for the other colored squares.

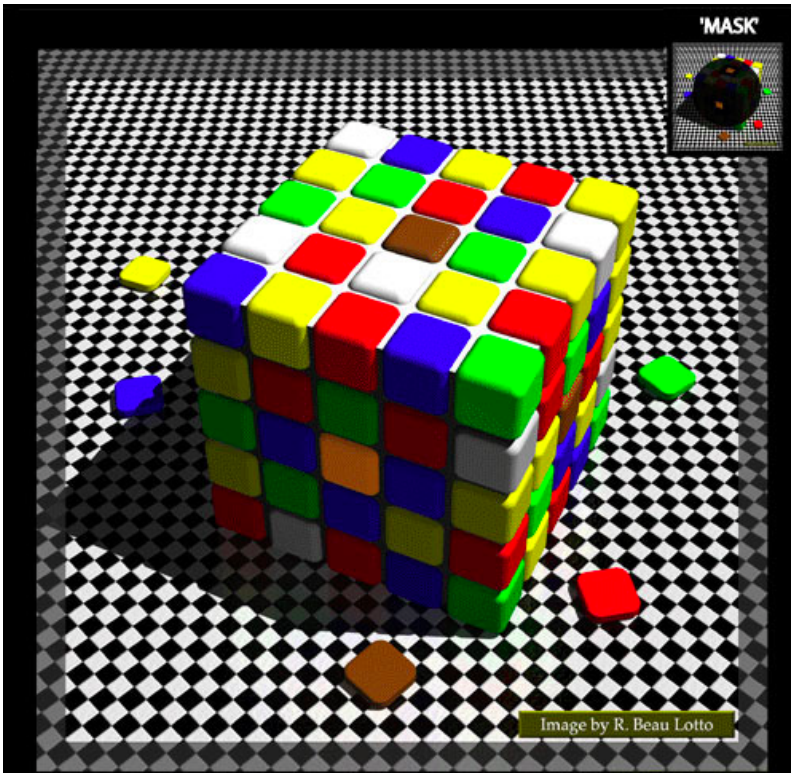

Figure 1.1: **Color cube illusion**, where the three-dimensional and lighting interpretations influence the perceptions of color. The middle tile on the upper surface appears brown, and the middle tile on the lower ("shadowed") surface of the cube appears orange; both are the same physical chromaticity. (Lotto, Purves, 2006©)

Bright contexts are thus able to induce brownness in otherwise yellowish (and orangish) stimuli. This makes brown unique among the basic color terms, and indeed among all colors. The conditions under which this brownness induction happens, as well as the mechanisms that mediate it, are as of yet not fully known. The rest of this chapter will present several recent investigations and findings concerning brownness induction, as well as tentatively

discuss mechanisms that might mediate brownness induction.

1.1 Measuring brownness

Brownness can be perceived over a range of light levels, either in combination with yellow, or without a trace of yellow. As a brown stimulus is made darker, the contribution of yellow decreases before disappearing entirely, leaving just brownness (with possible contributions from either green or red). This transition is marked by a perceptual brown boundary: the highest light level at which a stimulus will appear brown without a trace of yellow (DeLawyer et al., 2012, 2013; DeLawyer, Morimoto, & Buck, 2016; Buck et al., 2016). This brown boundary is a measure of how brown is induced into a target: the more brownness is induced, the lighter the target can be made while still looking exclusively brown. The brown boundary is not static; as outlined above, the lightness at which a surface will appear brown depends on, among others, the hue and saturation of the target (Bartleson, 1976). Moreover, different observers indicate different light levels as the brown boundary. Thus, the brown boundary can be used to express the amount of brownness induced into a target by a surround, and this measure of brownness induction can be compared within observers across stimuli (Vincent et al., 2012; DeLawyer et al., 2013; Buck et al., 2016). The following sections identify the influence (or lack thereof) of several factors on brownness induction.

1.2 Surround luminance effect on brownness induction

Brownness induction by a bright surround depends on the luminance of that surround. Previous research mentioned earlier established that for a surface to appear brown, it must be dark (Bartleson, 1976), and ideally darker than a bright surround (Uchikawa et al., 1989). The latter study noted that brighter surrounds increase brightness induction compared to darker surrounds, but did not parametrically investigate the relationship between surround luminance and brownness. This relationship can be established by measuring the brown boundary for a two degree target surrounded by a full-field (32 ° x 24 °) homogeneous surround at varying luminances (DeLawyer et al., 2012). Participants were able to reliably set

brown boundaries for each surround luminance. As expected, brownness induction strengthens with increasing surround luminance (DeLawyer et al., 2012). This is not a linear relationship, as the brown boundary is not found at a fixed ratio of target to surround luminance. Instead, the ratio is greatest (.35) for low luminance surrounds ($6\,cd/m^2$), but quickly drops off with increasing surround luminance, and levels at .16 for the highest surround luminance measured $110\,cd/m^2$ (Buck, unpublished). Thus, brownness induction can be measured reliably through brown boundaries, and is non-linearly related to the luminance of a contiguous full-field surround.

1.3 Spatial dependence of brownness induction

Brownness induction also depends on the distance between the target and the surround, and is strongest if the target and surround are contiguous (Vincent et al., 2012). A two degree test stimulus was presented with either a full-field black or white surround, or surrounded by a contiguous 16 degree diameter white disk on a black background. Brownness could be perceived in the presence of the full-field black surround, but not enough for a target to be devoid of yellow at any light level; participants were unable to set brown boundaries for this surround condition. The full-field white condition and the white contiguous disk both strongly induced brownness, with comparable brown boundaries (Vincent et al., 2012). However, when introducing a gap between the target and the white disk, the brownness induction was greatly weakened. Even a gap as small as 0.25° decreased the brown boundary by 50%. Increasing the size of this gap further reduced the brownness induction, but more gradually (Vincent et al., 2012). This suggests that proximity of the inducing surround to the brown target strengthens brownness induction, and that contiguity of target and surround has an especially strong effect. DeLawyer et al. (2016) replicated this finding for a four degree stimulus, surrounded by a two degree white or black annulus, on a white or black background. Brownness induction was greatest for the condition where the target was surrounded only by white (both annulus and background), and only slightly decreased when only the annulus was white (on a black background) (DeLawyer et al., 2016). When the

annulus was black, but the background white so that target was separated from the white surround by a two degree black gap, brownness induction was greatly reduced (DeLawyer et al., 2016). Again, no brown boundaries could be set when the target was surrounded only by black (both annulus and background). Buck et al. (2016) further investigated the spatial dependence of brownness induction, at two different time courses (27 ms and 1000 ms). For both durations, brownness induction was strongest when surround and target were contiguous, and induction attenuated with a distance between them (Buck et al., 2016). However, the previously reported steep drop off when target and surround are not contiguous, was only present for the 27 ms duration, and not for the 1000 ms duration (Buck et al., 2016, see Section 1.5 below). These conditions allow for a distinction between a contiguity induction effect and a context induction effect. The contiguity induction effect induces brownness when a target shares a border with a sufficiently bright surround. The context induction effect still induces substantial brownness for surrounds not contiguous with the target, at least when separated up to three degrees from the target. A similar distinction between contiguity effect and a context effect has previously been made for chromatic induction (Wesner & Shevell, 1992), and subsequently for chromatic contrast induction (Singer & D'Zmura, 1994). This spatial dependence of the contiguity induction effect, which is not found for the context induction effect, suggest that this type of brownness induction is mediated by two mechanisms, separable by their spatial constraints.

1.4 Complexity and brownness induction

Brownness induction does not seem to depend on the complexity of the inducing surround. Inhomogeneous surrounds of the same overall light level, area and proximity to the target stimulus, can differ from one another in their spatial distribution of the light. While laboratory experiments typically involve highly simplified displays with minimal complexity, naturalistic scenes have great complexity caused by many different surfaces and illumination conditions creating regions of interest. To be able to extrapolate from laboratory stimuli to naturalistic viewing conditions thus requires to be able to account for the effect of this

stimulus complexity. For brownness induction, Morimoto, Slezak, and Buck (2016) explored the effects of complexity in two experiments. All experiments measured brownness in a one degree test stimulus, separated from an inducing surround by a gap; the experiments differed in their surround configurations. The first experiment surrounded the stimulus with white homogeneous wedges extending from one to seven degrees away from the stimulus and equally spaced around. The number of wedges and the total area covered by these wedges were varied independently. While increasing the area covered by the surround strengthened brownness induction, increasing the number of wedges (for a given total surround area) did not. Decreasing the gap between the stimulus and the surrounding wedges from one to 0.25 ° strengthened the brownness induction, but did not lead to different results concerning surround complexity. In a second experiment, the stimulus was surround either by a homogeneous full-field surround, or a bubble surround that covered the display area. The bubble surrounds consisted of overlapping two degree circles that varied in luminance from each other. Three different mean surround luminances were tested; for each of these the bubble surrounds had the same mean luminance as the homogeneous surround. Despite the large change in complexity between the homogeneous and bubble surrounds, no changes in brownness induction were observed with complexity. Taken together, these experiments appear to indicate that beyond contiguity with, or distance from, the target, the complexity of the inducing surround does not affect brownness induction. There might yet be relevant, more specific, forms of complexity, e.g. that rely on the specific geometry or anisotropy of the surround. Certainly inferred lighting conditions, such is as in Figure 1.1 appears to induce brown, although it remains unclear whether such apparent shading has an effect beyond the brownness induction by the local (bright) surfaces.

1.5 Dynamics of brownness induction

Buck et al. (2016) investigated the temporal dynamics of brownness induction by flashing the target and inducing surround for either 1000 ms, or 27 ms. A two degree test stimulus was presented on a black background, either with or without a white disk surround 16 ° degrees

in diameter. The disk, when present, was either contiguous with the target, or separated by a gap of either 0.25 ° or 1 °. The 27 ms duration strongly induced brownness, although the amount of brownness induced was slightly less than in the 1000 ms condition (Buck et al., 2016). This suggests that strong brownness induction can happen very fast. At this short duration, the introduction of a gap between the target and surround led to the previously reported reduction of brownness induction. However, for this duration, increasing the size of the gap did not further weaken the induction effect (Buck et al., 2016). There was still an overall context induction effect for the 27 ms duration: brown boundaries could be set when stimulus was presented with the remote bright surround (1 ° gap between stimulus and surround annulus) brown boundaries could be set, but not when the stimulus was presented without any bright surround (full-field black) (Buck et al., 2016). For the 1000 ms duration, contiguity of the target and surround mattered less for the induction effect; a more linear relationship between surround distance and brownness induction was observed. Thus, the aforementioned contiguity effect on brownness induction seems to happen quickly, while the context effect of more remote surrounds happens more slowly.

1.6 Dichoptic brownness induction

To investigate the possible locus of brownness induction, DeLawyer et al. (2016) studied the dichoptic perception of brownness. Using a haploscope, the target and surround were either presented monocularly (both in the same eye), binocularly (both in both eyes), or dichoptically (target and surround in different eyes). In all conditions, the target and surround were perceptually aligned. The four degree circular target was surround by either a contiguous white ring two degrees in width on black background, or a two degree black ring on a white background. The former surround conditions provide a contiguity brownness induction effect, while the latter provides a context brownness induction effect. When presented dichoptically, the contiguous surround is perceptually contiguous, though not physically contiguous with the target on either retina.

No difference in brownness was observed between monocular and binocular viewing con-

ditions (DeLawyer et al., 2016). In these conditions, the contiguity effect provided stronger brownness induction than the context effect, in line with the previously reported studies on spatial dependence of brownness induction. In the dichoptic presentation, brownness induction was greatly reduced but not eliminated. Moreover, the spatial dependence appears to invert; in the dichoptic presentation, the context effect lead to greater brownness induction than the contiguity effect (DeLawyer et al., 2016). This suggests that while a remote context can induce brownness regardless of whether presented in the same or different eye as the target, the contiguity effect requires physical contiguity and not perceptual contiguity (DeLawyer et al., 2016). Thus, the contiguity effect appears to be mediated by a monocular mechanism operating on a physically contiguous surround . A separate mechanism is sensitive to the overall bright context surround the target, and this mechanism could operate after convergence of the signals from the two eyes (DeLawyer et al., 2016).

1.7 Mechanisms of brownness induction

A common thread through the previously presented studies investigating the properties of brownness induction appears to be a separation into two distinct effects. The spatial dependence of brownness induction suggests that contiguity effects are separable from larger context effects. The dynamics of brown induction suggest that the contiguity effect is fast, while the context effects take longer to develop. The dichoptic display of inducing surrounds suggest that the contiguity effect is sensitive to physical contiguity, while the context effects do not.

Overall, this paints a picture of two distinct mechanisms of brownness induction. Strong brownness induction is mediated by a fast acting mechanism, that is localized in space (sensitive to contiguity and proximity of surround and target), and operates in monocular pathways. A second mechanism is less localized in space (it is not sensitive to contiguity or proximity of surround and target), is processed more slowly, and could operate after convergence of signals from the two eyes.

A distinction between mechanisms on the basis of these features is not unique to the

processing of brown. Similar proposals have been put forward concerning both achromatic (D'Zmura & Singer, 1999) and chromatic induction (Wesner & Shevell, 1992; Singer & D'Zmura, 1994). Wesner and Shevell (1992) demonstrated that an achromatic ring surrounding a chromatic target can shift the target chromaticity; they also consider a distinction between contiguous and noncontiguous surround. However, they find that the effect of a noncontiguous surround context on chromaticity is weak when presented only with the target, but much stronger when presented with a chromatic adapting field (Wesner & Shevell, 1992). This is in contrast to the investigations on brown, which find strong effects of remote achromatic surrounds without chromatic adapting fields (Vincent et al., 2012; DeLawyer et al., 2016; Buck et al., 2016)). The question thus arises, how much does brownness have in common with these other forms of induction, and could brownness induction be completely explained by these other forms of induction?

Chapter 2

BRIGHTNESS INDUCTION IN ACHROMATIC STIMULI

Contexts similar to those presented in the previous chapter were already known to induce darkness in achromatic stimuli before being used to investigate the properties of brown induction. Generally, bright white surrounds make a target appear darker, while dark gray or black surrounds make a target appear brighter. One explanation for brownness induction then could be that brownness induction is simply darkness induction in yellow targets, and that the same mechanisms mediate both. The current investigation tests this explanation, by directly comparing brownness induction and brightness induction in the same displays.

The current chapter will provide a background on achromatic induction to put the current investigation into context. First, several key definitional issues on brightness and brightness induction will be addressed. Then, brightness contrast in displays comparable to those previously used to study brown induction will be discussed. Brightness assimilation, another form of brightness induction, will be introduced by presenting the much studied Whites illusion (White, 1979, 2010). This illusion has evaded a single accepted explanation, but several alternative explanations have previously been proposed and will be reviewed.

2.1 Luminance, lightness, brightness, darkness

Throughout the current work, the terms luminance, lightness and brightness will be used according to standard definitions (e.g., Kaiser, 1971; Kingdom, 2011). The photometric luminance is the physical intensity of a stimulus, scaled by human spectral sensitivity. Perceptually, this intensity is represented as brightness: the higher the intensity of a stimulus, all else being equal, the brighter it appears. However, the human visual system does not represent the intensity faithfully; depending on context, two equiluminant stimuli can ap-

pear to differ in brightness. Thus, brightness is not synonymous with luminance. Brightness differs from lightness in that the former is the perceived amount light received, such as from self-luminant sources, and the latter is the perceived reflectance of surfaces. Since all stimuli in the current work are presented on computer displays, and are thus self-luminant, all effects will be described in terms of brightness.

Brightness induction, in this work, will refer to a change in perceived brightness in a target attributable to a surrounding context. There is substantial discussion in the literature on whether brightness induction and darkness induction (or blackness perception, e.g., Cicerone, Volbrecht, Donnelly, & Werner, 1986; Shinomori, Schefrin, & Werner, 1997) should be considered symmetrical effects mediated by the same mechanism, or whether they are two distinct phenomena (e.g., Shevell, 1989; Rudd & Arrington, 2001; Rudd, 2013). This question is certainly pertinent to the relationship between brown induction and brightness induction, but the current investigation is not expected either to critically depend on assumptions about darkness and brightness induction or to reveal their differences. Throughout the rest of the current work, only the term brightness induction will be used to refer to both brightness and darkness induction.

2.2 Brightness contrast

Simultaneous brightness contrast is the induction of brightness in a target away from the brightness of surround context, i.e. a bright surround inducing darkness in a target, or a dim surround inducing brightness in a target. This can cause the two identical gray target to appear to differ in brightness (Figure 2.1).

The two targets in Figure 2.1 can be defined by the polarity of their physical *luminance contrast*. The gray target on the left in Figure 2.1A is more luminant than its surround, and thus has positive contrast with the surround. When moving from the target to the surround, the physical luminance decreases, and the difference in luminance is positive. The gray target on the right in Figure 2.1B is less luminant than its surround, and thus has negative contrast with the surround. When moving from the target to the surround, the physical luminance

Figure 2.1: **Brightness contrast** demonstration. The two gray center squares on the left and right are of the same physical luminance. Yet, the center square on the left is perceived as brighter than the center square of on the right. The dark surround of on the left induces brightness in its center square, while the bright surround of on the right induces darkness in its center square (adapted from Yeonan-Kim & Bertalmío, 2016).

increases, and the difference in luminance is positive. Generally, positive luminance contrast enhances the perceived brightness of the target that is more luminant than its surround, while negative luminance contrast reduces the perceived brightness (i.e. induces darkness) in the target that is less luminant than its surround.

A common neurophysiological explanation of simultaneous brightness contrast such as in Figure 2.1, is that it results from lateral inhibition between neurons (in the retina as well as in later stages such as early visual cortex) signaling the local luminance (Jameson & Hurvich, 1975; Moulden & Kingdom, 1989; Kingdom, 2011). Each of these neurons responds to the luminance in the small portion of the display that is in the cells receptive field center, but their response is also inhibited by the response of neighboring units (which respond to the luminance of neighboring portions of the display) in the cells receptive field surround. This

helps to rapidly adapt the retina to light, by normalizing the response of each neuron to the average of nearby neurons. The result is that these neurons, such as retinal ganglion cells, respond less to areas of uniform luminance, but strongly to luminance contrasts.

A similar direction of brightness induction can also be seen in displays where the surround is not physically contiguous with the target (Shevell, Holliday, & Whittle, 1992). These kinds of effects have been termed context effects (instead of contrast effects), and are considered to be mediated by different mechanisms than simultaneous brightness contrast. Together, these contrast and context effects are most like the displays used to study brownness induction presented in the previous chapter, and thus might represent similar mechanisms. These are not the kinds of displays and effects that will be used in the current investigation, and therefore the exact nature of contrast and context effects are not directly relevant. In the rest of the current work, the term contrast will be used to mean a shift in brightness induced in a target, in the direction away from the surround (whether it is physically contiguous with the target or not). The term simultaneous brightness contrast will be used to refer to the specific case of a brightness induction away from a surround that is physically contiguous with the target, and thought to be mediated by lateral inhibition.

2.3 Brightness assimilation

Brightness assimilation is the induction of brightness in the opposite direction from brightness contrast, i.e. a target in a bright context is made to look brighter than an identical target in a dark context. Brightness assimilation is most often due to context with high spatial frequency components, such as the high spatial frequency gratings in Figure 2.2. In the left-hand side of the display, the gray bars assimilate to the high luminance white interspersed bars and thus appear to be brighter than the same gray bars in the right-hand side of the display, where they assimilate to the black interspersed bars. This assimilation effect is limited to high spatial frequency contexts, as e.g. the bars in Figure 2.2 increase in width, the assimilation effect gradually decreases, and reverts to contrast at low enough frequencies (Helson & Rohles, 1959; Helson, 1963).

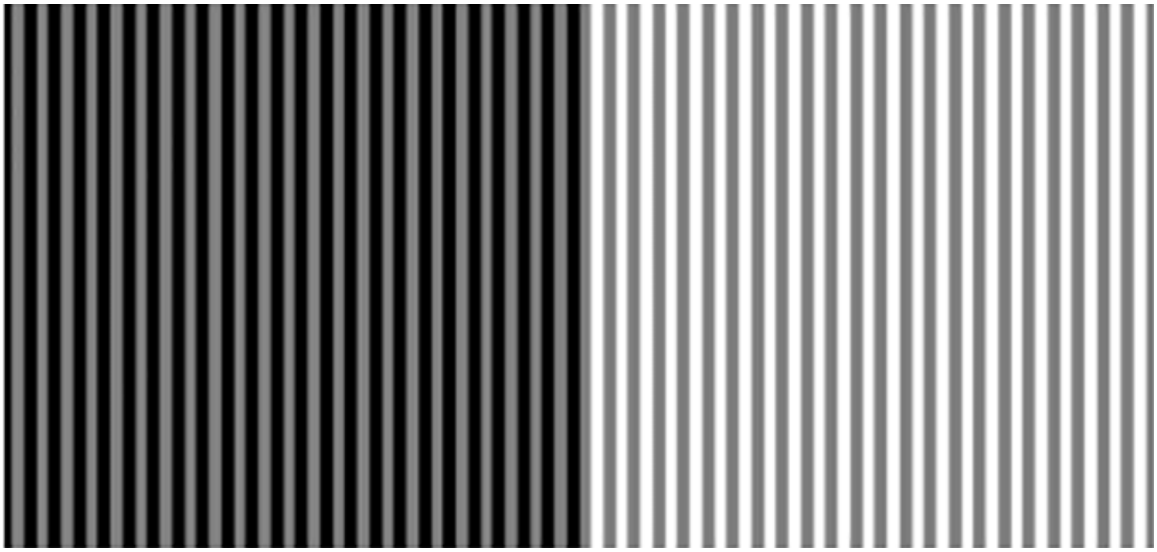

Figure 2.2: **Brightness assimilation** demonstration. All gray bars, on both the left and right side are of the same physical luminance. Yet, the gray bars on the left are perceived as darker than the gray bars on the right. The gray bars on the right assimilate towards the black surrounding features, while the gray bars on the right assimilate towards the white surrounding features (adapted from Yeonan-Kim & Bertalmío, 2016).

Assimilation has typically been explained as blending of the luminance values of the target regions with the luminance values of the surrounding context. This blending, however, is not due to optical blurring of the target and surround. The boundaries between the target element and the surround elements can be clearly identified in displays such as Figure 2.2, and are within the spatial resolution of the visual system. While the assimilation effect is strongest for high spatial frequencies, it is still pronounced for lower spatial frequencies that are well above the resolution threshold (Helson & Rohles, 1959).

However, the visual system does not just detect patterns at the highest spatial resolutions. Different neural units, specifically those in early visual cortex, are sensitive to different spatial frequency bands (Graham, 1981). Units with small receptive field sizes have high spatial frequency sensitivity: the small receptive field center and surround can accurately resolve the individual bars in fine gratings (such as in Figure 2.2). Units with larger receptive fields are less sensitive to such fine gratings; multiple bars of the grating will fall in the center,

as well as in the surround. The units with larger receptive fields end up responding to a blending of the different phases of the grating, averaging out their luminances, as a result of this lower spatial frequency sensitivity (Jameson & Hurvich, 1975).

The distributions of these receptive field sizes, and their resulting spatial frequency sensitivities and resolutions, can be observed as the absolute contrast sensitivity to different spatial frequencies (Graham, 1981). The sensitivity to luminance border contrast increases with spatial frequency for frequencies less than roughly 4 cycles per degree of visual angle, after which sensitivity steadily decreases (Davidson, 1968; Walker, 1978). For suprathreshold gratings, the apparent contrast of the grating shows a similar increase up to 4 cycles per degree, followed by a similar decrease in apparent contrast thereafter (Kulikowski, 1976). In other words, the different contrast sensitivities of different spatial frequency selective units translate well to apparent luminance border contrast for suprathreshold stimuli (Davidson, 1968; Kulikowski, 1976; Walker, 1978). Through these units, the luminance values of the target and surround elements would blend (Hurvich & Jameson, 1957; Jameson & Hurvich, 1975).

2.4 Whites illusion

White (1979) introduced a strong brightness illusion, which has since come to bear his name, that defies explanation in terms of either simultaneous brightness contrast or brightness assimilation. This illusion has been extensively studied since, and numerous explanations have been proposed. After a brief description of the illusion, and the features of the display that seem important for the effect to manifest, several of these explanations will be presented.

2.4.1 Description

White (1979) classic illusion consists of a square-wave grating with one (set of) target elements replacing parts of the white bars of the grating (i.e in phase with the grating) and another (set of) target elements replacing parts of the black bars of the grating (i.e. in counter-phase with the grating), as depicted in Figure 2.3. The bars that the target ele-

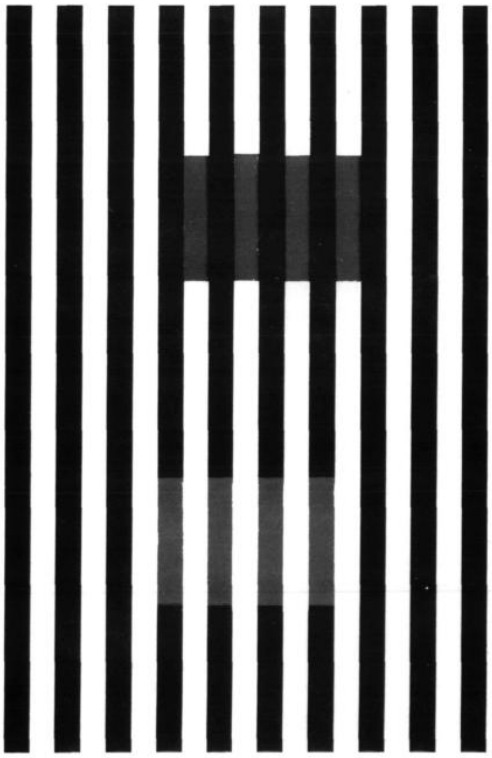

Figure 2.3: **White's illusion**, from White (1979). The two sets of gray target elements are physically identical, yet the upper set (coaxial with the white bars) appears darker than the lower set (coaxial with the black bars).

ments partially replace are dubbed collinear with the target, while the other bars of the grating are dubbed the flanking bars.

In the original illusion, presented in Figure 2.3, the target elements share a larger border with the flanking bars than with the collinear bars. Simultaneous brightness contrast would predict that the brightness of the target elements thus move away from the flanking bar. However, to most observers, the elements collinear with the black bars are perceived as significantly brighter than those collinear with the bright bars. This suggests that Whites illusion is an example of a brightness assimilation effect. This brightness assimilation explanation turns out to be insufficient for all versions of the illusion, as will be explained later in this chapter. For the purposes of clear labeling of the direction of effect, Whites

classic illusion will be considered to go in the *assimilation direction*, since the brightness of the target elements is shifted towards that of the flanking bars with which they share their larger border. Any variation that inverts this direction of the effect (i.e. the brightness of the target elements is shifted away from the flanking bars) will be considered to go in the *contrast direction*. Throughout the rest of current investigation, this labeling of the direction of effect is not meant to invoke any particular mechanism for the effect (except where otherwise indicated as assimilation or contrast *mechanisms*).

2.4.2 Important features

Across many studies of Whites illusion, using many variations, the (un)importance of several dimensions and features in the illusion have been identified.

The overall orientation of the display does not seem to affect the illusion: White (1979) originally used a vertically oriented grating, but horizontally oriented gratings have also been used extensively, with no reported decrease of the effect (White & White, 1985; Spehar, Gilchrist, & Arend, 1995; Adelson, 2000; Spehar, Clifford, & Agostini, 2002; Yazdanbakhsh, Arabzadeh, Babadi, & Fazl, 2002; Anstis, 2006; Barkan, Spitzer, & Einav, 2008; Betz, Shapley, Wichmann, & Maertens, 2015).

The number of target elements does not seem to affect the illusion. White (1979, 1981) presented not only the illusion in Figure 2.3, but also a variation with just a single element. Studies have used different numbers of target elements, with minimal reported changes in effect size, even going so far as to replace parts of all the bars in the grating (Anderson, 1997).

The size of the target elements is not crucial to the effect (White, 1979, 1981). When shortening the target elements to be half the width of a grating element, the assimilation effect remained. Other studies have similarly reported that shortening of the target elements to the point where they share a longer border with the collinear bars than the flanking bars does not remove the effect of the illusion.

The spatial properties of the grating do seem to affect the illusion. Higher spatial fre-

quency gratings enhance the effect (White, 1979), but at lower spatial frequencies the effect does not disappear. (White, 1979) and others (Howe, 2005) have also argued that the extent of the grating is important for the size of the effect, though this is not always reported (e.g., Zaidi, Spehar, & Shy, 1997). Thus, as (White, 1981) summarizes, the illusion does not seem to depend on the properties of the target elements, either their number of size. Instead, the extent, spatial frequency, and orientation of the grating (relative to the target) seem to be important for the magnitude of the effect (White, 1979, 1981). The grating can be described as having the effect of reducing contrast across borders parallel to it and/or of enhancing contrast across borders orthogonal to it (White, 1979, p. 415). Most explanations of Whites illusion can be considered attempts to explain why the grating has this mediating effect.

The luminance relationship between the target elements and the grating appears to be crucial to the effect. Brightness induction is only seen in Whites illusion if both sets of target elements are in the luminance range of the grating, i.e. all target elements must be more luminant than the white phase of the grating and less luminant than the black phase of the grating (Spehar et al., 1995).

Color, of particular interest to the current investigation, also preserves the illusion. In fact, Munker (1970) created an illusion similar to White (1979) in the color domain, which is why the illusion is occasionally referred to as the Munker-White illusion. The Munker-illusion consisted of chromatic targets, in a chromatic square-wave grating. In addition to a difference in brightness, this also leads to a difference in hue between the targets. For example, consider two green targets on a blue-yellow square-wave grating. The green target elements collinear with the yellow grating bars will appear more blueish green, while the green target elements collinear with the blue grating bars will appear more yellowish green. Thus, the target elements assimilate not only towards the brightness of the flanking bars, but also to the hue. This variant, although often referred to, is relatively little studied (White, 2010).

Other variations of Whites illusion have removed certain features from either the target elements or the gratings, to test specific explanations for the effect. Some such variations

will be discussed in relation to these explanations.

2.5 Explanations of Whites illusion

2.5.1 Assimilation through spatial averaging

White (1979) himself noted that neither simultaneous brightness contrast, nor brightness assimilation mechanisms through spatial averaging could explain the illusion. Both of these accounts would hypothesize the effect to depend on the ratio of the target area to the surrounding area, yet this hypothesis fails. For a given spatial frequency of the grating, this ratio of target element area to surround area is determined solely by the length of the borders that the target element shares with the collinear and flanking bars. In the original display, the rectangular target elements shared a longer border with the flanking elements of the grating than with the collinear elements. As mentioned above, White (1979) also reduced the length of the target elements to half squares, so that their border with the collinear bars is now twice the length of the border with the flanking bars. This variation demonstrates the insufficiency of explanations in terms of simultaneous brightness contrast and brightness assimilation mechanisms. In the original version of the illusion, simultaneous brightness contrast mechanisms would shift the brightness of the target elements away from the flanking bars. Yet If the target elements are shortened to have a longer border with the collinear bars than the flanking bars, simultaneous brightness contrast mechanisms shift the brightness of the target elements away from the collinear bars. Since this manipulation does not change the observed direction of effect, simultaneous brightness contrast mechanisms fail to capture both these versions of the illusion. Similarly, brightness assimilation mechanisms would shift the brightness of the target elements towards that of the flanking bars in the original illusion, but towards that of the collinear bars when the target elements are shortened to have a longer border with the collinear bars than the flanking bars. Thus, brightness assimilation mechanisms also fail to capture both these version of the illusion.

2.5.2 Pattern-specific inhibition

White (1979) initial attempt to explain the illusion is combination of three processes: brightness contrast mechanisms, brightness assimilation through spatial averaging, and assimilation from pattern-specific inhibition (White, 1981). The last of these processes is hypothesized to explain the mediating effect of the grating on the brightness assimilation. White (1981) isolated the effects of simultaneous brightness contrast mechanisms in simplified versions of the classic illusion comparing the gray elements surrounded by a homogeneous white or black surround, and interspersed with either white or black elements analogous to the flanking bars in the classic illusions. Across these variations, simultaneous brightness contrast mechanisms contribute their usual effect of inducing brightness in the target elements surrounded by the black, and inducing darkness in the elements surrounded by the white. He also isolated the effect of the brightness assimilation mechanisms, by interspersing gray target elements with white or black elements analogous to the flanking bars in the classic illusion. These reduced versions of the grating lead to the expected assimilation effect: interspersing the target elements with white induced brightness, while interspersing the target elements with black induced perceived darkness.

White (1981) also isolated what he dubbed pattern-specific inhibition. It has long been appreciated that neural units inhibit each other, and specifically that units with similar selectivity inhibit each other more so than those that differ in their selectivity. For the present discussion, this has the effect gratings of similar spatial frequency, which stimulate neural units that are tuned to this spatial frequencies, inhibit one another. Specifically, the perceived brightness contrast between the phases of a grating is reduced by similarly oriented gratings of similar spatial frequency nearby. By logical extension, any grating inhibits activation of itself most, within the same and adjacent regions of the display. In Whites classic illusion, the neural units that respond to the background grating will be strongly stimulated, and thus will also strongly self-inhibit. This has the effect the effect of reducing the simultaneous brightness contrast from the phases of the grating, or enhancing

the assimilation within the grating. If the gray target elements are experienced as part of the grating, this would also reduce the brightness contrast between the target elements and the flanking bars, and thus enhance the brightness assimilation effect. White (1981) provided evidence for this pattern-specific inhibition, by taking advantage of the orientation specificity of the inhibition. When the target elements are embedded in a grating that is orthogonal to the background grating, the assimilation effect is reduced, compared to when the target elements are oriented along the background grating. Moreover, the phase relation between the grating in which the target elements are embedded, and the larger surround grating, might influence this pattern-specific inhibition as well (White, 1981). Thus, White (1981) suggests, it is not merely brightness assimilation through spatial averaging that generates the effect in the classic illusion, but there is also a strong influence of pattern-specific inhibition by a grating with a large spatial extent that reduces the effect of simultaneous brightness contrast.

2.5.3 Multiscale spatial filtering

One of the attractive qualities of White (1981) explanation, is that it relies on relatively simple operators (contrast, assimilation, and inhibition) that have widely accepted neurophysiological mechanisms as their basis. The basis of the contrast and assimilation mechanisms lie in the center-surround structure of receptive fields of visual neurons (Walker, 1978; Kulikowski, 1976; Davidson, 1968; Jameson & Hurvich, 1975). These receptive field structures cannot only be psychophysically demonstrated, but electrophysiological recordings have provided further evidence for them as well (Graham, 1981). Moreover, these receptive field structures, and thus their effect as contrast and assimilation mechanisms, can be modeled as linear filters, operating on a display formalized as a matrix of luminance. Thus, explanations in terms of receptive field structures are relatively straightforward to implement computationally. Since they operate on a matrix of luminance values, without any further knowledge of the stimulus, these explanations and their computational formalizations are also generalizable to any display, including naturalistic scenes.

Moulden and Kingdom (1989) provided an initial computational formalization of mechanisms that might underlie Whites illusion. They argue for two mechanisms, that differ in their spatial range. A local mechanism would operate on the corners of the target elements, while a second mechanism with a larger spatial extent would sum the luminance over a larger range. The local mechanism consists of elongated cortical filters, made up of circularly symmetric center-surround antagonistic filter like the receptive fields of retinal ganglion cells. The corners of the target elements would excite the circular filters strongly, which in turn would drive the elongated cortical filters strongly. Through an unspecified mechanism, this would be disproportionately strong for the elongated filters that line on the edge between the target elements and the collinear bars. This disproportionate excitation would result in a contrast gain for this edge. The second, spatially extensive mechanism would involve a different kind of filter with a relatively small, circular receptive field center, and a surround that is elongated along one axis. Such a filter would have the effect of averaging the luminance values across a large area (along the axis of the surround elongation). Together, these two mechanisms could account for some variations in Whites illusion that Moulden and Kingdom (1989) tested.

Note that these two neurocomputational mechanisms would have the effect of simultaneous brightness contrast, and brightness assimilation, respectively. The contrast gain in the local mechanism would function as strong simultaneous brightness contrast across the edge between the target elements and the collinear bars: hypothetical cells whose receptive field captures the edge would emphasize the contrast across this edge, inducing brightness in the higher luminance region, and inducing darkness in the lower luminance region. The more spatially extensive mechanism has the effect of averaging luminance values that has previously been suggested as the basis for brightness assimilation.

By and far the most successful spatial filtering models, are the Oriented Difference-of-Gaussian (ODOG) family of models, first introduced by Blakeslee and McCourt (1999). These models suggest just one general shape of filter, and do not rely on disproportionate weighting of the filters in the corners of the target elements. Hence, these models are both

simpler and more plausible, as well as more generalizable (as they do not require specification of which region is the target). The ODOG family of models employ 42 filters in total, all of which have circular excitatory center, with a inhibitory surround elongated along one axis (usually twice the size of the center). These filters are arranged along 6 orientations (over 180 degrees, at 30 degree intervals), and 7 spatial scales (octave interval). Again, this allows the small filters (which represent small receptive fields, and thus high spatial frequency sensitivity), to act as contrast operators. These filters contribute simultaneous brightness contrast along edges in the display. The larger filters (representing larger receptive fields with corresponding lower spatial frequency sensitivity) will average luminance values for gratings which are too fine for them to detect, thus contributing brightness assimilation. Crucially, the ODOG family of models also introduce normalization rules: the output of each filter is divided by the output of all filters with the same orientation. Thus, the filters with fine spatial scales are normalized by those with larger spatial scales but the same orientation, and vice versa. Note that this normalization corresponds to the pattern-specific inhibition that White (1981) argued is the crucial mechanism mediating the illusions effect. Since each filter is normalized only by the other filters in the same orientation, this corresponds to inhibiting each filter-like unit by all other units that respond to similar patterns. Indeed, when no normalization rule is employed, the ODOG model fails to capture Whites illusion (Robinson, Hammon, & de Sa, 2007). In the ODOG family of models, different versions differ mainly in their normalization rules. The original ODOG normalizes each filter in its entirety by the average of all other filters of the same orientation, weighted according to the approximate contribution of their spatial frequency to the human contrast response function. Other versions of the model normalize each filter more locally; only the average in Gaussian window around each location is used to normalize that location (LODOG; Robinson et al., 2007). An even further refined model limits the normalization not only to nearby locations, but also to nearby spatial frequencies; each filter is normalized most by itself and similar spatial frequencies, while more extreme spatial frequencies have a smaller normalizing effect (FLODOG; Robinson et al., 2007). Note again how this implements White (1981) pattern-

specific inhibition mechanism: filter-like units that respond to a grating of a given orientation and spatial frequency are inhibited most strongly by themselves, and by units that respond strongly to similar patterns.

The ODOG-family of models, as well as several other filtering models, have enjoyed widespread success in predicting not only Whites illusion, but several other brightness phenomena (Robinson et al., 2007). Certainly spatial filtering models seem to be a good candidate for a universal brightness perception mechanism, but as of yet there has been little experimental evidence for the importance of normalization in many of these illusions. Moreover, there are other brightness illusions and nuances that these models have had less success with.

2.5.4 Junction analysis

Several authors have emphasized the importance of the precise geometry of the borders between target and grating regions, such as the precise alignment of the target elements with the grating bars. When the target elements are shifted laterally across the grating, the effect gradually decreases, and it goes away entirely when the target elements overlap equally with both the black and the white phase of the grating (White & White, 1985). Thus, the alignment of the target elements with the grating bars has an important influence on the original effect.

In case of a perfect alignment, a specific set of junctions defines the target. Each element in the display can be defined in terms of lines, the straight boundaries of elements, and junctions, the intersections where these lines meet. These junctions are labeled according to the number of lines that meet at that junction, and their relative orientations. In the Whites illusion, the target is defined by four lines: two lines where the gray target borders the flanking elements, and two lines where the target borders the collinear elements. These four lines form four junctions, one at each corner of the target element, where they meet additional lines, namely the borders between the flanking and collinear grating bars. Thus, each corner of the target forms a T-junction: the stem of the T is the border between the

target and the collinear grating element, one of the branches of the top of the T is the border between target and the flanking grating element, and the other branch of the top of the T is the border between the flanking and collinear bars (Zaidi et al., 1997). It is the top of these T-junctions that appear to be important for the illusion: the lines that form the top of the junction block simultaneous brightness contrast across them. Thus, when applying junction analysis it is the question of why the T-junction prevents brightness contrast across the top that is crucial to explaining the illusion.

Occlusion

One important problem for the human visual system is that the information from a three-dimensional world is projected onto the retina as a two-dimensional image. The visual system has to solve for the inverse of this projection, and attempt to determine what set of set of surfaces and illumination conditions is the most likely cause of what was detected. This necessarily involves guesswork of the best estimates, as an infinite number of states of the three-dimensional world can be projected to form the same two-dimensional retinal image. A feature of such projection is that lines in an image represent projections of edges: straight boundaries of a planar surfaces. Surfaces with depth relations can create edges through occlusion. A surface in front of another, such that it partially occludes the farther surface, will have an edge, and this edge projects as a line onto the retina. In more complicated displays, with multiple surfaces at multiple depth-planes, several lines can form junctions, such as T-junctions. Typically, such T-junctions have been considered to signal occlusion: the top of the T could indicate the edge of surface that occludes whatever surface has formed the stem of the T (Zaidi et al., 1997; Todorović, 1997, 2010; Anderson, 2001, 2003).

In Whites illusion, the tops of the T-junctions at the corners of the target elements are formed by the flanking bars. Thus, the visual system could be making the unconscious inference that the flanking bars form a set of surfaces that occlude a continuous surface that forms the target elements, on a homogeneous background. In other words, the target elements are interpreted as a surface on a homogeneous white or black background, that

is occluded by the bars that form the other phase of the grating. For the targets collinear with the white bars, this reduces the target elements to a surface on a homogeneous white background, which forms a simultaneous contrast display similar to Figure 2.3 (behind the occluding black bars). Since the target elements collinear with the white bars have darkness induced into them, this is consistent with the prediction from such a simultaneous contrast display. Similarly, the target elements collinear with the black bars are reduced to a surface on a homogeneous black background, behind the occluding white bars. Again, these target elements have brightness induced into them, which is consistent with a surface forming simultaneous contrast with a black surround. Thus, the T-junctions might serve as a depth cue, where the stem of the T is inferred to be behind the top of the T. The resulting interpretation as an occluded simultaneous contrast display predicts the correct direction of effect.

One limitation of this occlusion explanation is that the induction effect in White's illusion is often much stronger than the simultaneous brightness contrast display that this inference would generate (Anderson, 1997). Thus, an occlusion explanation does not seem to explain the full effect.

Scission

The T-junctions in White's illusion are of a specific form, with respect to the contrasts across the the top of the T (Anderson, 1997). The two luminance contrasts that form the branches of the top of the T are of different magnitude: one of the branches is the full contrast of the white and black bars of the grating, while the other is the contrast between the flanking bar and the target. Yet both branches of the top have the same contrast polarity: both branches have the higher luminance area on the same side of the branch. As a result, the two branches seem to form an aligned contour, with a consistent contrast polarity but a change in contrast magnitude.

X-junctions of similar form, where there is a change in contrast magnitude but a preservation of contrast polarity along one contour of the junction, are consistent with patterns that

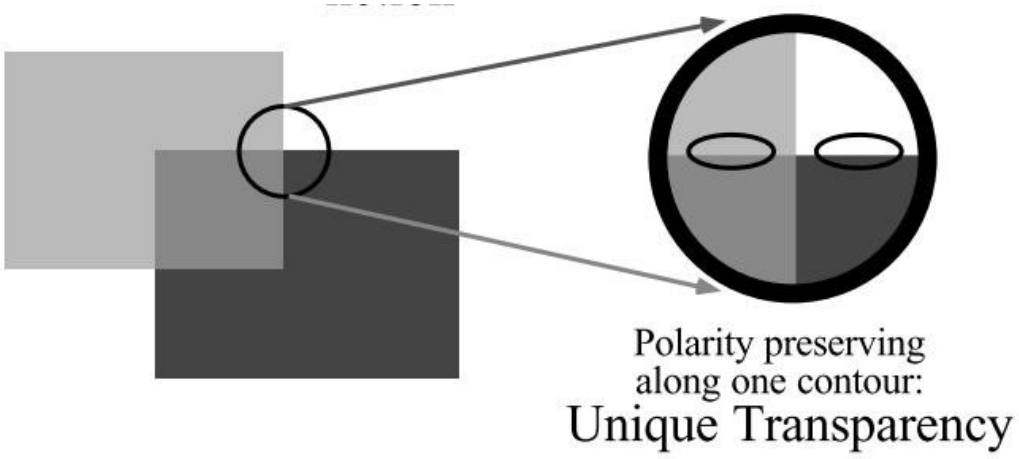

Figure 2.4: **X-junctions signaling transparency**, from Anderson (1997)

generate a percept of transparency (Anderson, 1997). Figure 2.4 shows an X-junction where the horizontal contour undergoes a change in contrast magnitude (low luminance contrast on the left side of the junction; high luminance contrast on the right side of the junction), but retains contrast polarity on both sides of the junctions (both top quadrants are more luminant than their respective bottom quadrant). As can be readily seen in the display that generates the junction, this leads to a percept of transparency across the horizontal contour to the left of the X-junction. In this case, Anderson (1997) argues that the visual system makes a *scission* inference: it decomposes the lower contrast region (the bottom-left quadrant of the X-junction in Figure 2.4) into multiple *causal layers*, or hypothetical surfaces, that give rise to the image. In the case of Figure 2.4, the multiple layers are two rectangular surfaces. Part of this inference is the simplifying assumption that each surface is of uniform lightness, which means that the lightness of the region of overlap has to be the product of the two lightnesses of the surfaces (Anderson, 1997, 2003; Anderson & Winawer, 2005). Some portion of the lightness of this region of overlap is assigned to the one surface, and the remainder is assigned to the other surface. For example, the bottom right surface can be interpreted as a semi-transparent layer: the lighter top-left surface is partially occluded

by a dark filter that lies on top of it. The inference can also go the other way, that the top-left surface is semi-transparent, since the physical luminance of a transparent surface is a function of both the transmittance of a surface and its reflection. If the top-left surface has a high surface reflectance that makes it appear light, but it is also semi-transparent, its superposition on the darker bottom-right surface would make the region of overlap darker than the rest of the surface.

Transparency is not the only possible inferred cause of this pattern of contrasts across such a junction either (Anderson, 1997, 2003; Anderson & Winawer, 2005). The same pattern could be the due to changes in illumination, e.g., overlapping shadows. In that case, the top-left area is a brightly illuminated surface, while the bottom-right area is a cast shadow. The region of overlap is then caused by the shadow cast on the top-left surface. The lightness of this area is attributed to the top-left surface, which the shadow dims. In the absence of more information, inferences of transparency and inferences of illumination changes cannot be disentangled in this display, and this is not necessary for application of this scission theory to Whites illusion. Conscious awareness of scission is presumably not required for its effect on lightness judgments (Anderson, 1997, 2003; Anderson & Winawer, 2005).

A T-junction with a similar contour as outlined above, could be formed if the top quadrants of the X-junction in Figure 2.4 were identical (Anderson, 1997, 2003). The horizontal lines would form the top of the T, and the line separating the bottom two quadrants would be the stem of the T. A T-junction with these luminance relationships would presumably still signal transparency. This is the kind of T-junction that is present in the Whites illusion as well (Figure 2.5, left hand side). Thus, the T-junctions in the Whites illusion are of the kind that would lead to a scission or decomposition of the lower contrast region into multiple causal layers (Anderson, 1997, 2003). In the Whites illusion, this lower contrast region is always the target, since the contrast between the target and the flanking bars is less than the contrast between the flanking and collinear bars. Perceptual scission would decompose the target regions in Whites illusion into multiple layers.

Since the collinear bars and the flanking bars are assumed to each be of uniform lightness,

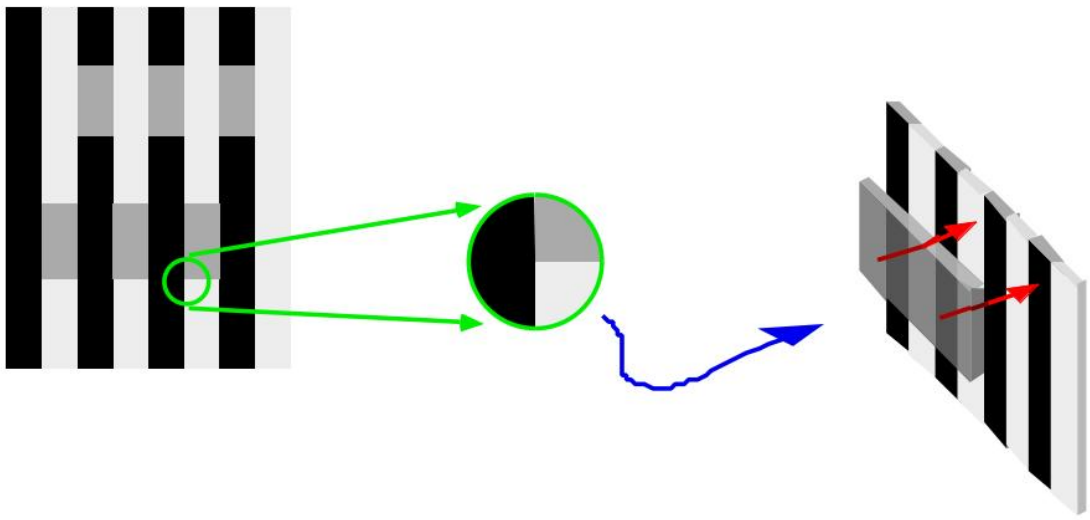

Figure 2.5: **Scission of White's illusion**, from Anderson (1997)

the inference is that the target elements are formed by a semi-transparent surface (or a change in illumination) on top of the grating (Figure 2.5). The brightness of the target, then, is the product of brightness of the collinear bars, and brightness of the semi-transparent surface. For the target elements in phase with the grating, the inferred transparent layer has to have no surface reflectance since the black bars of the grating that are seen through it still appear uniformly black. The white bars are visible through the layer, but the luminance in the target regions is lower than in the white bars, so the layer is inferred to have some opacity. The brightness of the target elements, then, is assumed to come from the white bars that they partially occlude. As a result, some brightness is taken out of the inferred layer that forms the target elements. This layer is thus perceived as darker than just a veridical interpretation of its luminance. For the target elements in counterphase with the grating, the black collinear bars do not contribute any brightness because they are black. Yet in the target regions there is a higher luminance, so the target layer is inferred to have some nonzero reflectance. Thus, the luminance of the target elements is assigned to just the target

layer. Presumably, the light from the white flanking bars shining through the target layer is of higher luminance than that reflected by the target layer, so that the flanking bars are not impeded by the transparent target layer. The brightness of these target elements then, is determined entirely by the luminance of the target elements on the black bars (Anderson, 1997, 2003).

Scission, just like occlusion, provides an explanation of why T-junctions in Whites illusion appear to prevent brightness contrast across the top of the T. Note that as mentioned earlier, scission does not require an inference of transparency, but could also be explained in terms of changes in illumination. Also, a conscious percept of decomposition of the Whites illusion into layers is presumed not to be necessary for scission to affect the brightness of the target elements.

Versions without T-junctions

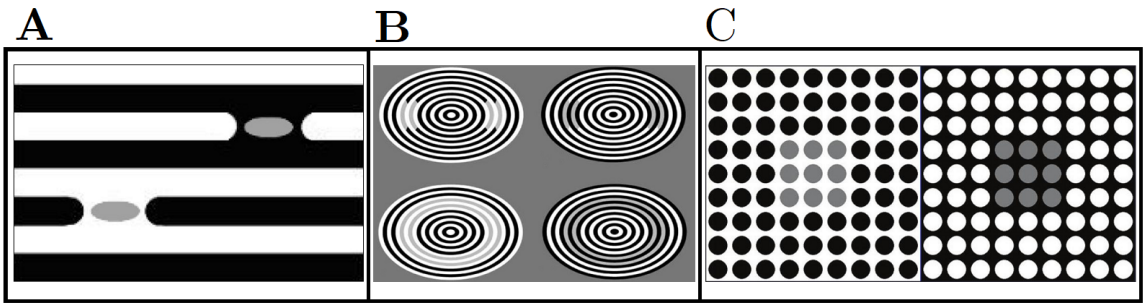

Figure 2.6: **White's Illusion variants without T-junctions**. **A** Yazdanbakhsh et al. (2002) variant, maintaining the oriented grating but removing T-junctions with the target elements. **B** Howe (2005) radial variants, substitute the oriented grating for a radial grating. In the top row, T-junctions are still present at the corners of the target elements; in the bottom row, T-junctions have been removed entirely. **C** White (2010) groundal (left) and figural (right) dotted illusions.

In response to explanation of White's illusion using junction analysis, several versions of the illusion without T-junctions have been produced. When the edges of the target elements are rounded, so that now the target elements are ellipsoids, the Whites effect is reduced, but

still present (Figure 2.6A) (Yazdanbakhsh et al., 2002). Since the area immediately bordering the target elements is not the luminance of the collinear bar, but that of the flanking bars, this variant of the illusion cannot be explained by simultaneous brightness contrast. In radial versions of Whites illusion, where the grating is now isotropic, the effect is comparable to the original illusion (Figure 2.6B) (Howe, 2005). The variants in the top row of Figure 2.6B still contain T-junctions similar to those in the original illusion. In the bottom row, the target elements have been elongated to form a complete ring, which is analogous to lengthening a target element in the original illusion to replace an entire bar in the grating. In these variants, which contain no T-junctions at all, the effect is still as strong as in the original illusion (Howe, 2005). Similar illusions had previously been used to study the induction effects of near and far contexts (Hong & Shevell, 2004a, 2004b).

White (2010) himself created several version of the illusion without T-junctions that have been received surprisingly little attention (Figure 2.6C). In what White calls the figural dotted illusion, the grating is replaced by a regular grid of dots (either black on white background, or white on black background), some of which have been replaced by the intermediate gray target dots. The target dots assimilate to, rather than contrast with, the background that they are placed on. Again, this illusion contains no T-junctions at all, yet a very strong illusory brightness induction effect is present. As White notes, if all versions are to be explained in the same general terms, the role of T-junctions must be incidental to the effects (White, 2010, p. 4).

2.6 Gestalt grouping and anchoring

A final class of explanations that will be presented in this chapter, attempts to group the target elements with their flanking bars, based on the Gestalt principles of perception (A. L. Gilchrist et al., 1999; A. L. Gilchrist, 2008; A. Gilchrist, 2015; A. L. Gilchrist, 2015). Grouping perceived surfaces in a display is argued to be important to solve the anchoring problem. The retina, through adaptation of neural responses, does not veridically represent the luminance at each retinal location. Rather, it encodes the relative luminances of different

areas. Adaptation of retinal responses is a necessary process, since the response range of neurons is far smaller than the range of possible luminance values in the world. Luminances encountered in the world can differ by as much as 6 log units across illumination conditions, where the firing rate of the average neuron spans maybe 2 log units. Thus, retinal neurons need to adapt their gain to keep the encountered luminance values within their dynamic range. However, a side-effect of this adaptation is that information is lost about the absolute luminances. A given range of relative luminance values encoded by the retina can represent very low to less low luminances, but the same range of relative luminance values could also be the retinal response to a range of high to very high luminances. To solve for this ambiguity, the visual anchors the relative luminances to the highest relative light level, and that light level is perceived as white (A. L. Gilchrist et al., 1999; A. Gilchrist, 2015). All other relative luminances are perceived as shades darker than that white. In the simplest displays, consisting only of two surfaces under the same illuminant, this anchoring rule works well to explain perceived brightness. For more complex displays, which can contain any number of surfaces, illuminations, and their combination, the simple highest luminance anchoring rule fails.

Instead, A. L. Gilchrist et al. (1999) argue that anchoring happens separately for groupings of regions within the display. The image is parsed into at least two frameworks: a global framework that includes the entire visual field, and any number of subordinate local frameworks. Anchoring then takes places within each framework; the highest relative luminance value in the framework is assigned the highest brightness (white), the lowest luminance value is assigned the lowest brightness (black), and all other relative luminance values in the framework are spaced accordingly. Thus, the perceived brightness of a region depends on which local framework it is deemed to belong to (A. L. Gilchrist et al., 1999; A. Gilchrist, 2015; A. L. Gilchrist, 2015).

These local frameworks are defined by grouping of regions, on the basis of several grouping principles, borrowed from Gestalt psychology. Most importantly, regions perceived as coplanar, moving together, or lying in the same shadow, are each grouped together into frame-

works. T-junctions have also been presented as a grouping factor; the tops of T-junctions would prevent grouping across them, while the stems connect different regions within one framework, analogous to how these junctions could signal occlusion (A. L. Gilchrist et al., 1999; A. Gilchrist, 2015). Note that this invokes the Gestalt principle of good continuation: all regions on one side of the of the top of T are perceived as a continuation of the same contour, distinct from the regions on the other side of the T of the T. It is also these T-junctions that form the primary grouping in Whites illusion. However, since other grouping principles also could define frameworks, the anchoring theory is not easily disproven by variants of the illusion without T-junctions.

In Whites classic illusion, the T-junctions at each corner of each target element form the contour across the top of the T, namely the edge of the flanking bar. All the regions on the other side of top, i.e. the target element and the collinear bar, are grouped together in the same framework. For the target elements collinear with the white bars, this means that they are grouped with these white bars into one local framework. In this framework, the collinear white bars are assigned as white, and the only other relative luminance value left is that of the gray target elements, whose perceived brightness is thus shifted down. As a result, darkness is induced into these target elements. For the target elements collinear with the black bars, and thus grouped with them, the target elements get assigned the higher brightness. Thus, brightness is induced into these target elements. Both frameworks are at least slightly adjusted by the global framework, so that the white bars of the grating still appear brighter than the target elements collinear with black bars (since the former have a higher relative luminance). As a result, the two sets of target elements still look intermediate to the black and white bars of the grating, but because of the anchoring within the local framework, brightness has been induced in the direction of the observed effect in the illusion.

2.7 Integration of explanations

The explanations of Whites illusion presented above are not a comprehensive review of all explanations that have been given for the illusion, or of all the details of, evidence for, and

evidence against, each of these explanations. Rather, this discussion has attempted two main goals.

First, it presents an overview of different classes of brightness processing theories. Generally, these classes are ordered along different levels of analysis, usually with regards to whether and how they incorporate depth information (White, 2010). The low-level explanations do not invoke any inferences about depth relations, or any interpretation of the scene in general. The original explanation in terms of pattern-specific inhibition, as well as the multiscale spatial filtering models, attempt to explain the illusion at this level. At the highest level, explanations of the illusion rely on inferences about the scene, depth relations between assumed surfaces, and related processes such as scission. The scission theory, and most other explanations on the basis of junction analyses, are argued to operate at this level. In between, at the mid-level, make some inferences about the scene, but do not require a full understanding of the underlying surfaces causing the display. The Gestalt grouping theory, as well as certain explanations based on T-junctions that are not presented here, generally operate at this level. These different levels of analysis will prove useful in the discussion of the results of the current investigation, and are therefore included here.

Second, the discussion of these explanations is also meant to show the range of different explanations for the very same illusion. This reflects the current state of understanding of brightness and lightness processing, where there are many competing theories (Kingdom, 2011). For each of the presented explanations, there is evidence in favor, but also evidence against that some would consider a crucial invalidation of the explanation. A full review of this evidence, for each of these explanations, is beyond the scope of the current investigation, and not essential to it. Moreover, there have been few attempts to integrate these different explanations. While most authors concede that a full account of brightness processing in all displays probably requires multiple explanations, providing a unified framework has proven to be difficult. Kingdom (2011) notes that the problem [of combining filtering models with junction analysis and scission] is a profound one however, because the languages of layer decomposition and filtering are so very different.

The current work will not be testing these different explanations, nor competing theories of brightness in general. It is not an attempt to resolve any of these debates. Rather, it will investigate the relationship between brown induction and brightness induction, and how that might inform theories on either. Specifically, the question guiding the current investigation is whether brownness induction in White's classic illusion, and two variants, can be explained by the achromatic brightness induction in these stimuli.

Chapter 3

PSYCHOPHYSICAL METHOD

3.1 Stimuli

Three stimuli were used to investigate the relationship between brownness induction and brightness induction: a version of Whites (1979, 1981) classic illusion, a checkerboard variant, and set of concentric rings. Each stimulus consists of an articulated high-contrast surround context that induces a brightness difference in two targets. The three stimuli were tested in separate experiments.

3.1.1 Whites classic illusion

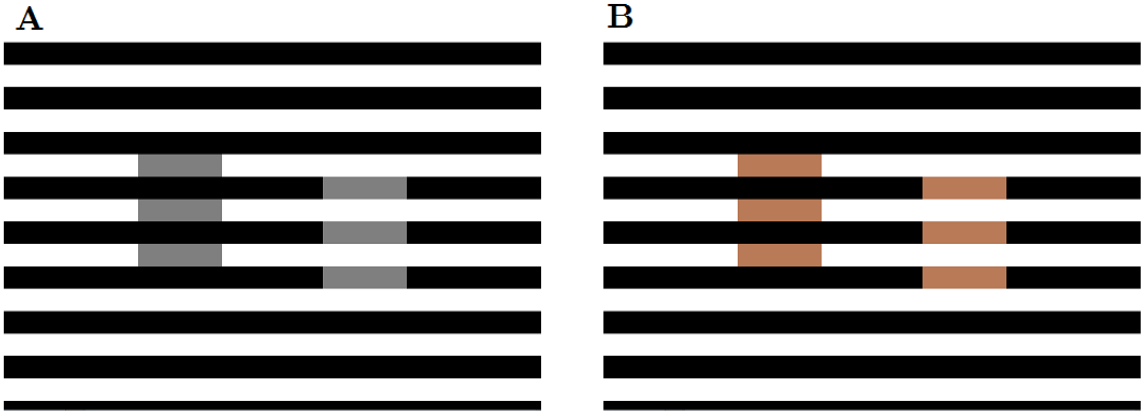

Figure 3.1: **White's classic illusion stimulus**

Whites (1979) classic illusion consists of a square-wave grating with one (set of) target elements in phase with the grating (i.e. partly replacing the white bars of the grating) and another (set of) target elements in counterphase with the grating (i.e. partly replacing of

the black bars of the grating). The number, size and nature of the target elements, as well as the orientation of the surround grating, have been varied by many authors (e.g., White, 1979; White & White, 1985; Blakeslee & McCourt, 1999; Anderson, 2001, 2003; Howe, 2005), usually to very little change in the induction effect.

In the current investigation, a horizontally oriented square-wave grating of approximately 0.40 cpd surrounded two sets of 3 target elements each. The white phase of the grating was achromatic (1931 CIE x = .33, y = .33; L/(L+M) cone ratio = .66, S/(L+M) cone ratio = .0032) at a constant 95 cd/m^2. The black phase of the grating was <0.001 cd/m^2. The total display subtended 32° x 24° visual angle.

On the left hand side of the display, the target elements were in phase with the surround grating, partly replacing the white bars. On the right hand side, the target elements were in complete counterphase with the surround grating, partly replacing the black bars. The target elements were centered at 5.5° horizontal eccentricity, 5° in width, and spanning approximately 7.5° degrees vertically (3 full cycles).

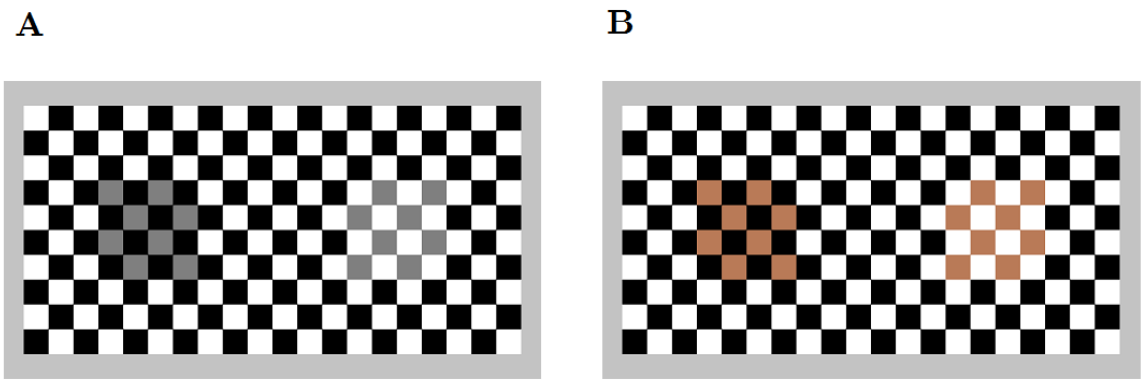

Figure 3.2: **Checkerboard variant stimulus**

3.1.2 Checkerboard variant

Several variants of Whites illusion have sought to eliminate T-junctions, to investigate explanations of the induction effect based on these junctions (Howe, 2005; White, 2010). One such variant is to replace the surround grating with a black and white checkerboard, where the target elements substitute for one or more white or black checks. Such a display substitutes the T-junctions with X-junctions, for which junction analysis provides no prediction for. Blakeslee and McCourt (2004) found that in a checkerboard variant, the direction of effect depends on the spatial frequency: at high spatial frequencies the target brightness is shifted towards that of the majority of the adjacent checks (*assimilation*), while at low spatial frequencies the target brightness is shifted away from that of the majority of the adjacent checks (*contrast*).

In the current investigation, a 20 x 10 checkerboard of 0.5 ° square checks subtended 10 ° x 5 ° of the visual field. The white checks were achromatic (CIE 1964 10 ° x = .33, y = .33; L/(L+M) cone ratio = .66, S/(L+M) cone ratio = .0032) at a constant 95 cd/m^2, the black checks at <0.001 cd/m^2. The checkerboard was placed on a 24 ° x 16 ° achromatic background at 48 cd/m^2. On the left hand side of the display, the target replaced 8 white checks of the display in a 2 ° x 2 ° region. On the right hand side of the display, the target replaced 8 black checks of the display. Both targets were centered at 2.5 ° horizontal eccentricity.

3.1.3 Radial variant

Two radial variants of Whites illusion were introduced by Howe (2005) to test explanations of Whites effect based on T-junctions. In both variants, the context grating is no longer oriented vertically or horizontally, but radially outward from a center point. This forms a set of concentric rings, of equal width but not diameters, alternating between black and white. In Howes (2005) first variant, the targets replace part of the intermediate rings, either the intermediate white or intermediate black rings (see Figure 2.6B). This is a closer approximation of Whites illusion, in that the target elements only partly replace elements

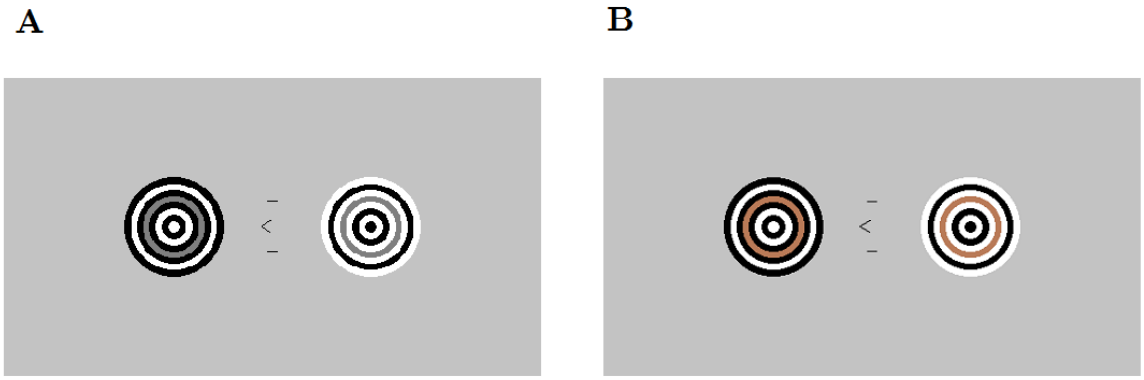

Figure 3.3: **Radial variant stimulus**

of the grating. As a result, T-junctions are maintained, as the corners of the gray test elements adjoin both black and white background elements. The second variant eliminates these T-junctions entirely, by extending the gray test elements to replace the intermediate rings entirely (see Figure 2.6B Howe, 2005). Each test ring is now bordered exclusively by either the two white background rings or the two black background rings adjacent. This variant is analogous to lengthening the target elements in Whites original illusion to replace entire bars of the background grating. Contrary to what explanations based on T-junctions would predict, the illusory brightness difference remains. Hong and Shevell (2004a, 2004b) found that in this second radial variant, the more remote rings significantly influenced the target brightness, even changing the relation between contiguous surround ring luminance and target brightness. This, too, is not predicted by some explanations based on T-junctions, since the T-junctions depend only on the contiguous rings. Moreover, since this variant is isotropic (i.e it has contrast power in all orientations) multiscale spatial filtering models have trouble accurately capturing the direction of brightness induction in this illusion (Robinson et al., 2007).

In the current investigation, two sets of 7 concentric rings were placed on a 24° x 16° achromatic background at 48°, centered at 2° horizontal eccentricity each. The width of

each ring was 0.125 °, including the inner circle of the display. Each entire set of rings was 2 ° in diameter. Alternatively, each set of concentric rings can be considered to be a radial square-wave grating with a frequency of 8 cpd. The white rings were achromatic (CIE 1964 10 ° x = .33, y = .33; L/(L+M) cone ratio = .66, S/(L+M) cone ratio = .0032) at a constant 95 cd/m^2, the black rings at <0.001 cd/m^2.

The target elements were the middle (4th) ring in each set. Unlike the other two stimuli, this radial variant of Whites illusion requires two separate gratings, each containing one full target ring. The two displays differ in their phase: the inner ring of the left display is black, while the inner ring of the right display is white. The target ring in the left display replaces a white ring, and is contiguous with two black rings; the target ring in the right display replaces a black ring, and is contiguous with two white rings.

3.2 Tasks

3.2.1 Brightness matching

An asymmetric brightness matching task was employed to measure the brightness induction effect. To measure brightness in isolation of any chromatic effects, both targets were kept achromatic (CIE 1964 10 ° x = .33, y = .33; L/(L+M) cone ratio = .66, S/(L+M) cone ratio = .0032). On every trial, one target, the reference, was kept at a fixed luminance, pseudo-randomly drawn from the uniform integer distribution between 1 cd/m^2 and 90 cd/m^2. Participants adjusted the physical luminance of the other target, the match, until it matched the reference in brightness. Participants could not adjust the chromaticity of the match target. Each session of the brightness matching task consisted of 50 trials, during which the side of the reference was counterbalanced, so that 25 trials had the reference on the left side.

Brightness matching does not provide a measure of the brightness of each target; instead, the difference in physical luminance at match is a measure of the brightness difference between the targets as a result of the surrounding context. An outstanding question in in-

vestigations of White's illusion is whether this difference is due to darkening of one target, or brightening of the other target, or both. The present investigation will not attempt to answer this question, but attempts only to measure the brightness induction effect generated by the different stimuli.

Even at a brightness match, the two targets might differ in their perceived lightness, albedo, shading or transparency. Participants were instructed to match the two targets on brightness, and told that even at the best possible setting, the match and reference might not look identical. To find the best possible match, participants were to find an initial setting that they perceived as a match, then move the light level away from this setting by a few steps, and try to find their original setting again.

3.2.2 Brown boundary setting

To measure the brownness induction effect, an analogous brownness matching task could not be used. The perceptual differences between the two targets were judged to be too great to overcome in a matching task. The brown targets, by virtue of being chromatic, can differ from one another in three perceptual dimensions (hue, brightness and saturation), as well as the aforementioned properties of lightness, albedo, shading or transparency. Matching the two targets on brightness, or brownness, would still leave significant perceptual differences, which made any judgment of perceptual equality difficult.

Instead, brownness was measured through an appearance judgment that uses the quality of brown being the dark complimentary of yellow. A bright yellow stimulus will look pure yellow, without a trace of brown. As this stimulus is made darker, brownness will perceptually mix in, to form a color that has both yellowness and brownness in it (e.g., caramel or butterscotch). When made sufficiently dark, the yellowness will disappear, and the target will appear as a pure brown (without a trace of yellow). A perceptual brown boundary exists between this pure brown percept, and the yellow-brown color mixture. The brown boundary must map onto a physical luminance of the stimulus, which can be found by adjusting the luminance of a stimulus to the highest light level at which it appears exclusively

brown. The degree of brownness induction is inversely related to this brown boundary: the stronger the brown induction by a surrounding context, the more luminant it can be made before yellowness starts to mix in. This property allows for quantitative measurement and comparison of the perceptual brownness of different stimuli, with a task that relies solely on the percept of brownness itself. An analogous task is not possible for brightness induction, as the achromatic brightness dimension lacks such a boundary between percepts.

Observers can reliably identify their brown boundary for a given stimulus, but it can vary widely from stimulus to stimulus, context to context, and observer to observer. Participants were instructed to adjust the physical luminance of each target to the highest value at which it looks exclusively brown, without a trace of yellow. Additionally, they were instructed to maintain a red-green balance for each target as they adjusted the luminance. A yellow stimulus can appear reddish (orange), or greenish, or as a red-green equilibrium that appears exclusively yellow. Similarly, a brown stimulus can appear as a either reddish brown, a greenish brown, or a red-green equilibrium brown. However, previous work indicates that the physical red-green equilibrium point for yellow is not the physical red-green equilibrium point for brown. An equilibrium yellow, when made darker, will look like a greenish brown and requires additional redness to become an equilibrium brown. Conversely, an equilibrium brown, when made brighter, will appear orangish, rather an equilibrium yellow. As participants adjust the luminance of the target to identify their brown boundary, the perceptual red-green balance will shift accordingly. To counteract this shift, participants adjusted the relative amounts of redness and greenness until the stimulus appeared as approximately equilibrium brown.

Like with the achromatic matches, participants were to find an initial setting that they perceived as a brown boundary, then move the away from this setting by a few steps, and try to find their original setting again. On every trial participants first adjusted one target to the brown boundary, followed by the second target; the order was counterbalanced across trials. Both targets were displayed simultaneously throughout the trial. Participants completed 16 trials per session, thus making 32 brown boundary settings. Since the brown

boundary settings span a smaller luminance range than the achromatic matches, fewer trials are necessary to provide a reasonable sampling of the stimulus space.

Brightness induction and brownness induction are measured in notably different ways. By comparing the luminance difference between the two targets at matching brightness, the inducing effect of the context is measured as a difference in physical luminance between the two targets at matching brightness. The brown boundary settings, on the other hand, measure the brownness of each target as an effect of the inducing context. By comparing the brown boundaries of the targets, the relative brownness induction, analogous to the brightness induction from achromatic matching, can be inferred. Thus, while luminance difference between the two targets arises from different kinds of perceptual judgements for achromatic matching and brown boundary setting, they both express the magnitude of the induction effect of the surrounding context.

3.3 Procedure

All stimuli were presented on a ViewSonic G90fB (Model No. VS 10794) 19 inch (48.3 cm) CRT monitor, using custom routines written for MATLAB (Mathworks, 2016) and Psychtoolbox (Brainard, 1997). The gamma lookup table in the system was replaced with a linearized gamma table, calculated from luminance and chromaticity measurements by a Spectrascan PR 705 photometer. To reduce adaptation effects, the stimulus display cycled continuously on (1 s) and off (3 s), replaced by a full 32 ° x 24 ° achromatic field at 48 cd/m^2. When the stimulus was not presented, a fixation cross appeared in the center of the display. In an otherwise dark room, participants viewed the stimulus binocularly through natural pupils from a chin rest at approximately 63 cm distance from the screen.

A total of 25 participants participated in either 1 experiment (N = 17), 2 experiments (N = 6) or all three experiments (N = 2). All participants were undergraduate Psychology students (XX male, XX female) at the University of Washington, compensated for their participation with course credit. The sample is representative of the UW undergraduate population, which is largely of Caucasian and East-Asian heritage. All participants were

screened for color vision deficiency using the Ishihara pseudo-isochromatic plates; one participant (initials AGP) was a deuteranope, all other participants were color normal. All participants had normal or corrected-to-normal spatial vision.

Participants adjusted the luminance using two keys on a keyboard. When satisfied with their setting, participants would accept their setting with another keypress, and a new trial would appear. Fixation was not enforced, but participants were encouraged to accept settings only while fixating on the center of the display, as indicated by the fixation cross during the 3s off-cycle. Participants completed one session of achromatic matching at a time (approx. 45 minutes), and a maximum of one session per day.

Chapter 4
PSYCHOPHYSICAL RESULTS

4.1 Number of sessions completed

Participants typically completed 8 sessions of both achromatic matching (= 400 trials), and brown boundary setting (= 128 trials), although this varied between participants and stimuli due to scheduling reasons. Since the brown boundary settings span a smaller luminance range than the achromatic matches, fewer trials are necessary to provide a reasonable sampling of the stimulus space. This, combined with the longer duration for each trial in the brown boundary setting, leads to a lower number of trials completed for that task. A complete overview of the number of trials each participant completed per stimulus can be found in Appendix A.

Two participants were not included in subsequent analysis for specific stimuli: one completed fewer than 2 sessions of achromatic matching (<100 trials), the other fewer than 4 sessions of brown boundary setting (<64).

This chapter will first outline the analysis procedure for the achromatic matches, and the brown boundaries, for each observer using an exemplar. After outlining the analysis of individual participants results, the results for all participants on the classic illusion are presented. This is followed by the checkerboard and radial variations, and each stimulus is analyzed separately. The chapter finished by comparing the findings across the three stimuli.

4.2 Achromatic matches

Figure 4.1 shows all achromatic brightness matches by one participant (AGP) on the Whites illusion. Each trial plots the physical luminances of the two target patches at perceived equal brightness. If perceived brightness was solely a function of target luminance, and not

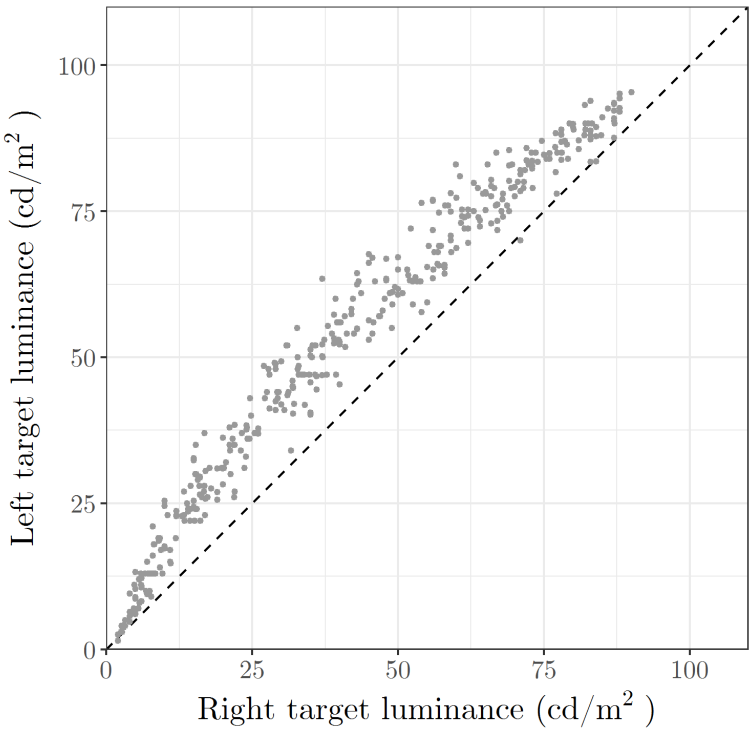

Figure 4.1: **Example achromatic brightness matching data** for participant AGP on the classic illusion.

context, the two target patches should be perceived as equally bright only when both have have the same physical luminance, regardless of what that luminance is. Each trial would then fall on the unity line (dashed in Figure 4.1), indicating equiluminance.

The trials do not fall on the equiluminant line; perceived equal brightness is reached when the target patches have different luminances. This difference in luminance necessary to create matching brightness in the target patches is smallest when both target patches are very dim or very bright. In particular, the difference in luminance is minimal when the luminance of the reference target approaches the luminance of either the black or white elements of the background grating. This bounding of the brightness difference could be due to either the reference or the target patch being indistinguishable from the elements it is embedded in, and has been previously reported (Spehar et al., 1995).

The effect of the context on target patch brightness is neither purely additive nor a pure scaling. If the effect of the induction effect of the context was purely additive, the difference in perceived brightness between the target patches would be constant. The trials would lie on a line parallel to the unity line, either above or below it. If the effect of the induction effect was a pure scaling, the difference in perceived brightness would increase with reference patch luminance. The trials would lie on a line of greater or smaller slope than the unity line. The bounding of the matching function to the unity line at extreme luminances, and the resulting curving away from the unity line at intermediate luminances indicates that neither of these two cases apply. Instead, the luminance difference at brightness match follows a nonlinear function of the reference luminance, and is smallest for extreme luminances and greatest for intermediate luminances. This achromatic brightness matching data can be accurately described using nonlinear models, as outlined in Appendix B. However, such a description of the data is not possible for the brown boundary data, and provides no clear advantages over the analyses outlined in the rest of this chapter.

4.2.1 Magnitude of the induction effect

The magnitude of the induction effect on any trial can be expressed as the distance of the point from the unity line. This is proportional to the signed difference between the target patch luminances at perceived equal brightness. The left target patch luminance is the ordinate, so a positive induction magnitude indicates that the left target patch needed to be at higher physical luminance than the right target patch to appear equally bright. (Throughout this chapter, the left target patch luminance will be the ordinate, and the right target patch luminance will be the abscissa.) A positive induction magnitude means that the left target appeared darker than the right target patch. The sign of the induction magnitude is chosen arbitrarily; switching the abscissa and ordinate will also switch the sign of the induction magnitude.

Figure 4.2 shows a histogram of the brightness induction magnitudes for the Whites illusion data from participant AGP in Figure 4.1. In this plot, the equiluminant line is

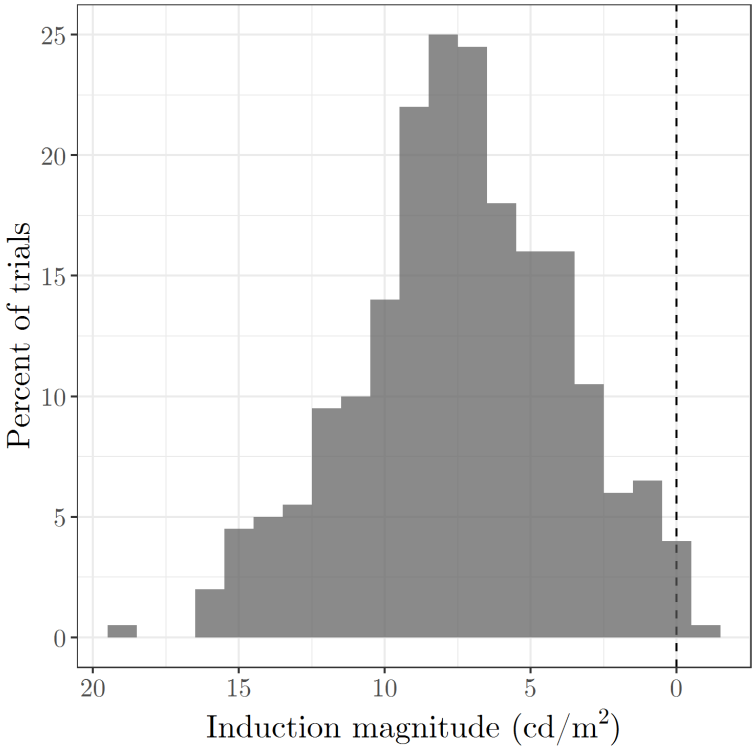

Figure 4.2: **Histogram of induction magnitudes**

indicated by the dashed vertical line at zero induction magnitude. Note that positive values lie left of the equiluminant line, and negative values lie to the right of it. The histogram can be interpreted as a projection of the scatterplot onto a line orthogonal to the unity line. The majority of the data lie to the left of the equiluminant line; on average the left target appeared significantly darker than the right target (M = 8.21, SD = 3.43, t(290) = 40.82, p <.0001) for this participant.

4.3 Brown boundaries

Figure 4.3 overlays the brown boundary (brown) settings of the same participant (AGP) onto the achromatic matches in Figure 4.1 for Whites illusion. Unlike the brightness matches, the brown boundary settings do not span the whole range of luminances. Participants set both

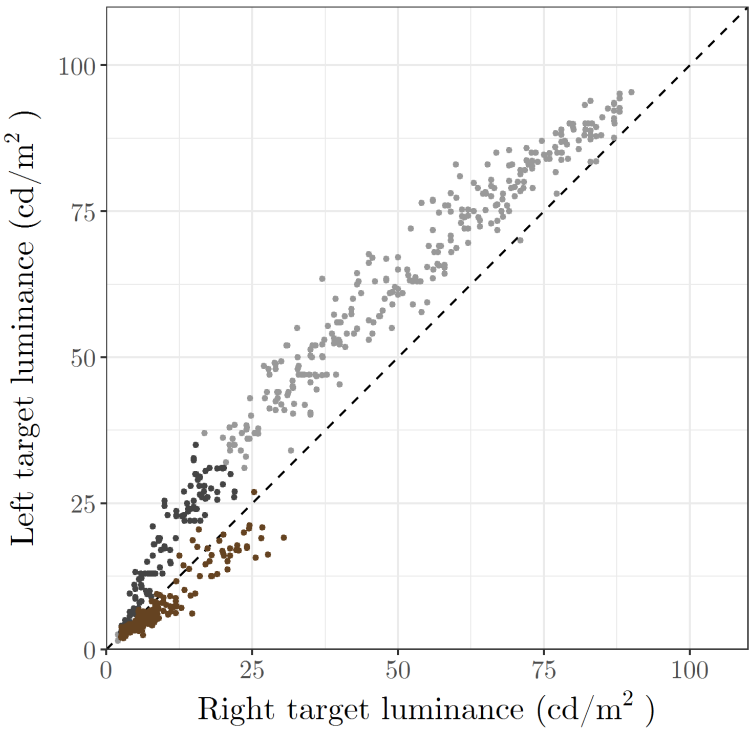

Figure 4.3: **Example of brown boundary data**, (brown) of participant AGP, compared to their achromatic brightness matching data (gray), on the classic illusion. The brightness matches in the same luminance range as the brown boundaries are plot in dark gray.

target patches to the highest luminance at which they perceived the patches as uniquely brown (without a trace of yellow). Thus, the brown boundary task necessarily limits the luminance range for the brown boundaries to lower luminances.

To make a meaningful comparison between the achromatic brightness induction magnitude and the brown induction magnitude, only trials with comparable patch luminances are taken into account. The achromatic brightness matches are bounded at the high and low luminance, and show a nonlinear relationship between patch luminance and induction magnitude. The projection of the brown boundary data onto the unity line defines a range of patch luminances over which brown boundaries are set. The achromatic trials that project onto the unity line within the range defined by the brown boundaries (dark gray in Figure

4.3) are considered for comparison of brown and brightness induction. All other achromatic trials (light gray in Figure 4.3) are not considered for comparison of induction magnitudes for brown and brightness.

4.3.1 Magnitude of induction effect

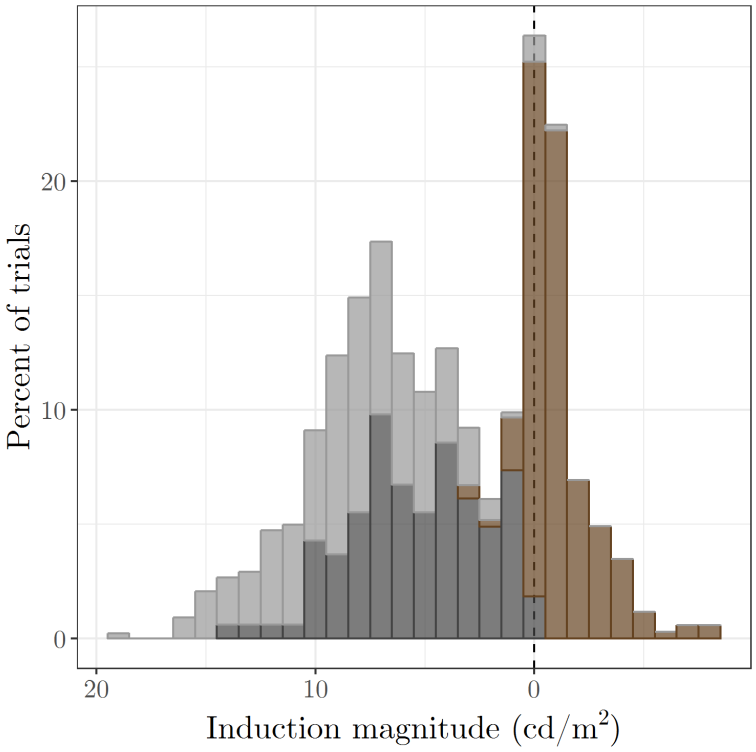

Figure 4.4: **Comparing induction magnitudes** of brown boundary data (brown) and achromatic brightness matching (gray), of participant AGP, on the classic illusion. The brightness matches in the same luminance range as the brown boundaries are plot in dark gray.

Figure 4.4 shows the distribution of the brown boundary induction magnitudes (brown), all achromatic matching induction magnitudes (light gray), and the induction magnitude for the selection of achromatic matches with the same luminances as the brown boundaries, as described above (dark gray) for participant AGP on Whites illusion. The heights of

the histograms are normalized to each subset individually, i.e., the percentages shown were calculated separately for each of the three categories.

For this participant, the brown boundary settings predominately go in the opposite direction from the achromatic matches. The overall achromatic matches indicate that the left target appeared brighter than the right target. The subset of achromatic matches in the luminance range of the brown boundaries show an effect in that same direction (M = 5.32, SD = 3.11, t(108) = 17.90, p <.0001) but of smaller magnitude (t(212.87) = 8.05, p <.0001). Note that this reduced magnitude is expected from the bounding of the achromatic matches. The brown boundaries, however, lie to the right of the equiluminant line, indicating that the right target looks significantly more brown that the left target (M = -1.19, SD = 1.60, t(230) = -11.31, p <.0001). In addition, for this participant, the achromatic matches and brown boundaries are significantly different in signed magnitude, t(135.68) = -20.65, p <.0001, over their shared luminance range. All these magnitudes and comparisons can also be found in the first row of Table 4.1.

4.4 Whites classic illusion

Twelve participants completed sufficient sessions of both achromatic matching and brown boundary setting on Whites classic illusion. Figure 4.5 shows the achromatic matches, brown boundaries, and their respective magnitudes for each participant. Note the range of induction magnitudes (i.e. the horizontal axis on the histograms) differs between individual participants.

Magnitudes Table 4.1 contains the corresponding mean induction magnitudes for the achromatic matches, for the subset in the luminance range of the brown boundaries, and for the brown boundaries, along with standard deviations and null-hypothesis significance test results. Each individual set of magnitudes is compared to a null hypothesis that the mean magnitude is 0, using one-sample t-tests; significant results are bolded. The magnitudes of the subset of achromatic matches is also compared to the overall set of achromatic matches, as well as to the brown boundaries, using two separate Welchs independent samples t-tests.

Figure 4.5: **White's classic illusion**, achromatic brightness matches (gray) and brown boundaries (brown) of all participants)

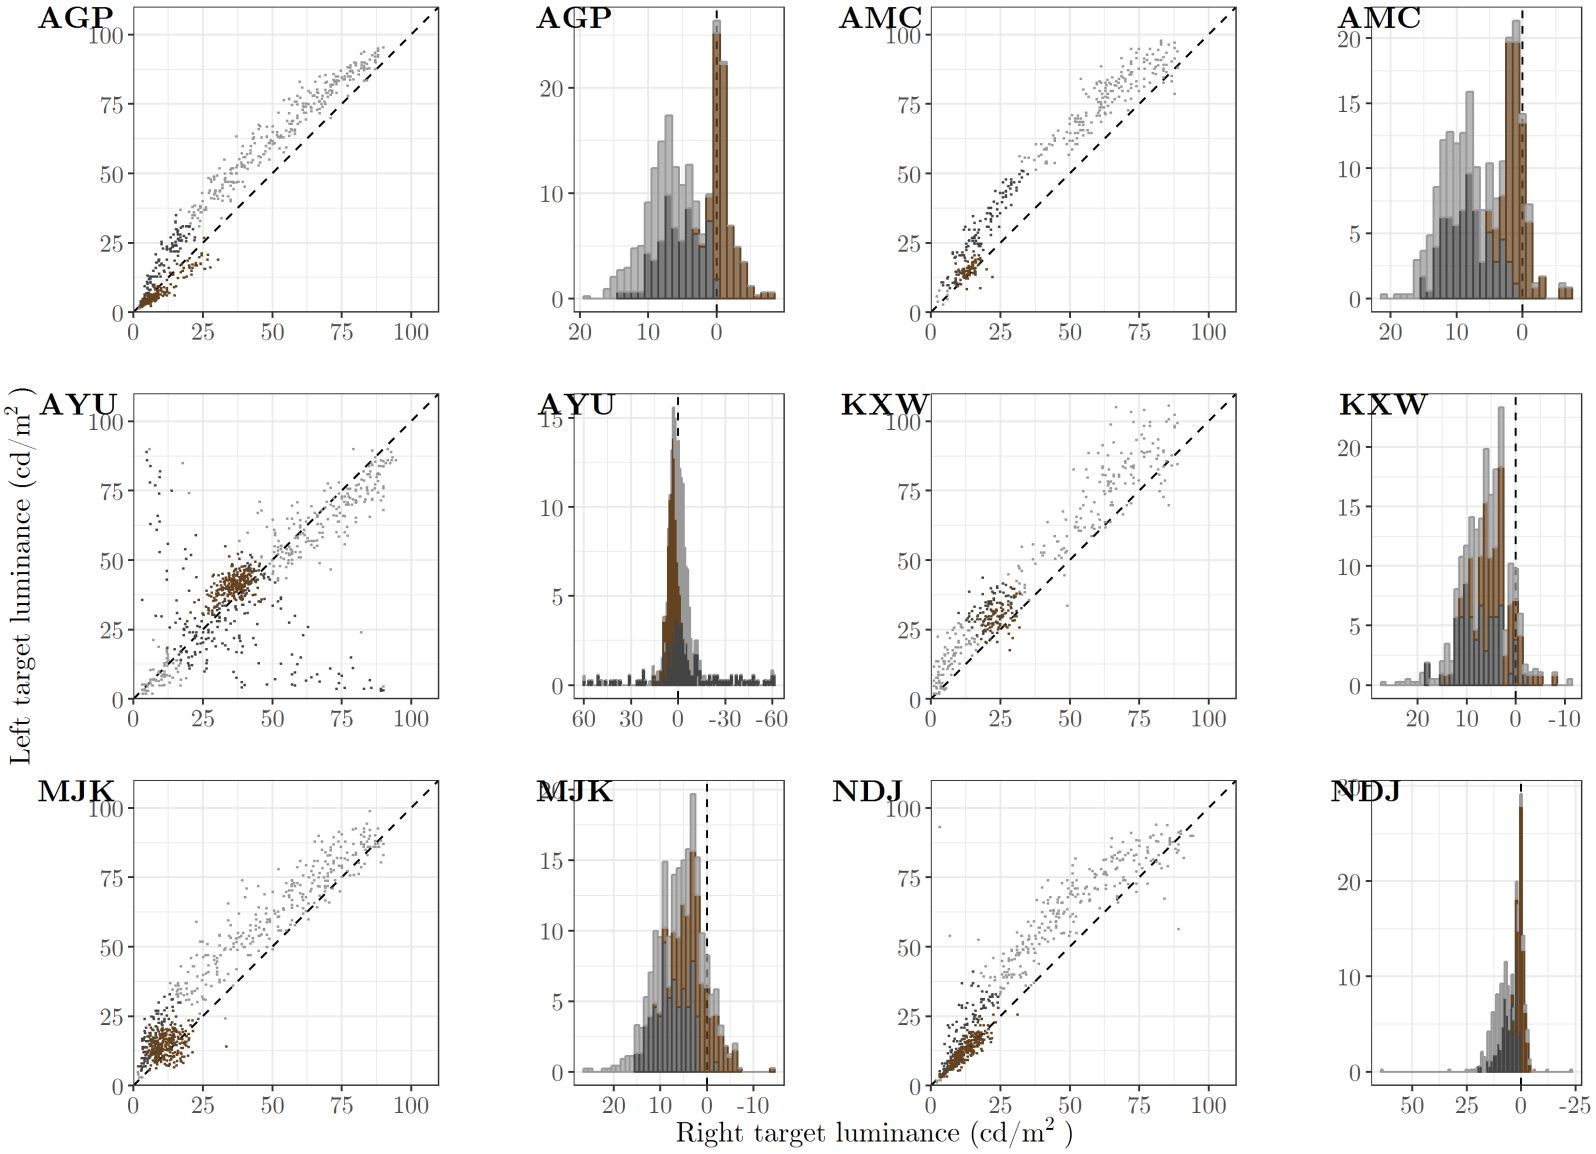

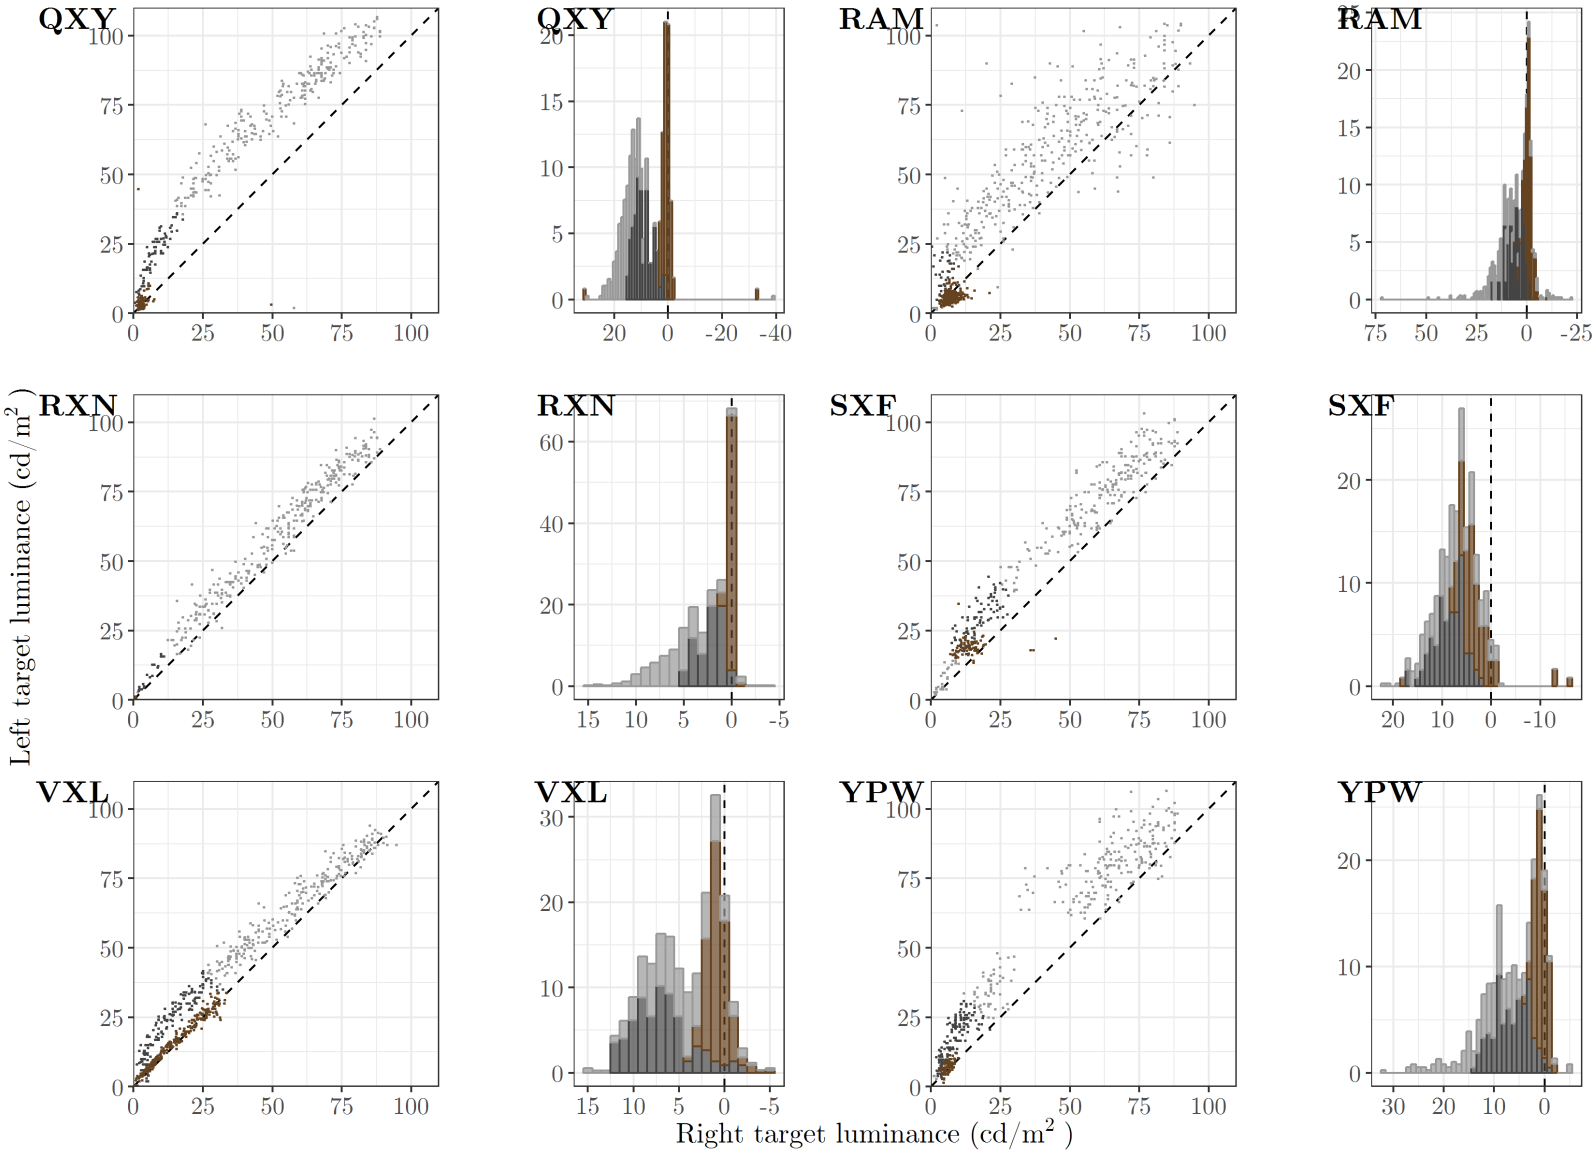
QXY
QXY
RAM
RAM
RXN
RXN
SXF
SXF
VXL
VXL
YPW
YPW
Left target luminance (cd/m^2)
Right target luminance (cd/m^2)

A Bonferroni correction for multiple comparisons is applied to all tests per participant.

Eleven out of the twelve participants show a significant positive mean achromatic induction magnitude, indicating they perceived the right target as brighter than the left target. The remaining participant shows a significant negative mean achromatic induction magnitude, indicating they perceived the left target as brighter than the right target. Eight out of twelve participants show a significant positive mean brownness induction magnitude, indicating that they perceive the left target as more brown than the right target. Two participants show a significant negative mean brownness induction magnitude, indicating that they perceive the right target as more brown than the left target. The remaining two participants do not show a significant difference in brownness between the two targets.

Seven out of twelve participants show a significantly reduced mean induction magnitude for brown compared to the achromatic matches in the same luminance range, with both effects going in the same direction. While the mean induction magnitudes may be significantly different, this is not necessarily of theoretically significance; different induction magnitudes could arise from a shared underlying mechanism with different response gain for brightness induction than for brown induction. Four of the remaining participants also show a difference in mean induction magnitudes between brown and the subset of achromatic matches, but for these participants, the effects also went in different directions.

Table 4.1: Induction magnitudes and corresponding significance tests for White's classic illusion

Participant	Induction magnitudes						Comparisons	
	Achromatic match		Subset of achromatic matches		Brown boundaries		Achromatic − Subset	Subset − Brown boundaries
	Mean (SD)	t(df); p	Mean (SD)	t(df); p	Mean (SD)	t(df); p	t(df); p	t(df); p
AGP	8.21 (3.43)	**(290) = 40.82; p<.0001**	5.32 (3.11)	**t(108) = 17.90; p<.0001**	-1.19 (1.60)	**t(230) = -11.31; p<.0001**	**t(212.87) = 8.05; p<.0001**	**t(135.68) = -20.65; p<.0001**
AMC	9.18 (4.44)	**t(201) = 29.36; p<.0001**	8.11 (3.43)	**t(117) = 25.72; p<.0001**	0.88 (1.90)	**t(78) = 4.12; p=0.0001**	**t(294.50) = 2.42; p=0.0163**	**t(189.31) = -18.97; p<.0001**
AYU	-1.90 (8.46)	**t(274) = -3.72; p=0.0002**	-3.59 (21.65)	t(224) = -2.48; p=0.0137	3.49 (2.82)	**t(255) = 19.77; p=0.0000**	t(279.88) = 1.10; p=0.2713	**t(230.69) = 4.86; p=0.0000**
KXW	6.77 (5.65)	**t(252) = 19.06; p<.0001**	7.37 (4.15)	**t(69) = 14.87; p<.0001**	3.89 (3.67)	**t(79) = 9.49; p<.0001**	**t(147.61) = -0.97; p=0.3325**	**t(138.91) = -5.41; p<.0001**
MJK	7.41 (5.58)	**t(297) = 22.91; p<.0001**	7.40 (3.56)	**t(101) = 20.96; p<.0001**	1.97 (3.56)	**t(255) = 8.83; p=0.0000**	t(275.83) = 0.03; p=0.9768	**t(185.88) = -13.01; p<.0001**
NDJ	9.05 (7.02)	**t(261) = 20.85; p<.0001**	5.57 (3.99)	**(111) = 14.76; p<.0001**	-0.10 (1.48)	t(255) = -1.09; p=0.2747	**t(343.34) = 6.04; p<.0001**	**t(124.62) = -14.60; p<.0001**
QXY	15.09 (5.11)	**t(255) = 47.25; p<.0001**	9.20 (3.64)	**t(71) = 21.47; p<.0001**	0.63 (5.18)	t(79) = 1.09; p=0.2777	**t(158.11) = 11.01; p<.0001**	**t(141.91) = -11.90; p<.0001**
RAM	10.13 (9.80)	**t(358) = 19.60; p<.0001**	5.77 (4.35)	**t(40) = 8.49; p<.0001**	-0.92 (1.75)	**t(255) = -8.38; p<.0001**	**t(96.17) = 5.11; p<.0001**	**t(42.09) = -9.72; p<.0001**
RXN	5.37 (3.07)	**t(302) = 30.47; p<.0001**	2.29 (1.37)	**t(16) = 6.90; p<.0001**	0.14 (0.24)	**t(79) = 5.37; p<.0001**	**t(26.17) = 8.19; p<.0001**	**t(16.21) = -6.44; p<.0001**
SXF	7.84 (4.61)	**t(235) = 26.10; p<.0001**	8.49 (3.10)	**t(83) = 25.10; p<.0001**	3.42 (4.37)	**t(79) = 7.00; p<.0001**	**t(217.68) = -1.44; p=0.1511**	**t(141.92) = -8.54; p<.0001**
VXL	5.07 (3.52)	**t(248) = 22.73; p<.0001**	6.81 (3.10)	**t(150) = 27.03; p<.0001**	0.72 (1.22)	**t(223) = 8.80; p<.0001**	**t(348.07) = -5.16; p<.0001**	**t(181.57) = -23.02; p<.0001**
YPW	9.92 (6.31)	**t(255) = 25.16; p<.0001**	6.40 (3.39)	**t(143) = 22.63; p<.0001**	0.75 (1.17)	**t(79) = 5.79; p<.0001**	**t(397.35) = 7.25; p<.0001**	**t(194.31) = -18.13; p<.0001**

4.4.1 Congruence

That the achromatic induction effect goes in the positive direction for almost all participants, while the brown boundaries show less consistency, seems to suggest that the directions of these effects might be independent. To test for such independence, the (in)congruence of direction of effect within individual participants is crucial. If brightness and brownness induction are mediated by the same mechanism(s), the induction effects should go in the same direction for an individual participant. This does not require the induction effects to be of the same magnitude; there could be trivial difference that cause the two measurements to yield different magnitudes, e.g., different response gains for brightness and brownness, or differences in how sensitive the tasks are to detect induction effects. Yet the direction of effects should only be different if the two induction effects are independent: finding, e.g., a significant positive effect for brightness and a significant negative effect for brownness, for the same participant, cannot be explained by differences in response gain or task differences.

A negative achromatic induction magnitude indicates that right the target is perceived as darker than the left target. Since the right target shares a larger border with the darker elements of the surround, this negative direction of effect is in line with the predictions made by simultaneous brightness contrast. The negative direction of effect will therefore be referred to as the *contrast direction*. Similarly, the positive direction of effect will be labeled the *assimilation direction*. Throughout the current investigation, this labeling will be used for the direction of effect for both the achromatic brightness matching and the brown boundary setting, although this is certainly more arbitrary for the latter than the former. It is not meant to provide any explanatory value or invoke any specific mechanism, but simply as shorthand for the direction of effect.

Table 4.2 provides the observed number of participants for each combination of direction of effect (assimilation, none, contrast) for achromatic matching and brown boundary setting. The marginal totals indicate the total number of participants who show a specific direction of effect for the achromatic matching (right) and brown boundary setting (bottom). If the

achromatic matching results are independent of the brown boundary setting, any individual participant represent a random draw from the distributions defined by these marginal totals. This is the null hypothesis that achromatic matching and brown boundary setting are independent results. Since eleven out of twelve (= 92%) participants show an achromatic effect in the assimilation direction, under the null hypothesis the number of participants who show an achromatic assimilation effect and a brown assimilation effect should be 92% of the first column total (eight) in Table 4.2, which would be 7.33 participants. These expected numbers of participants for each cell under the null hypothesis, the product of the marginal probabilities for that row and column times the total number of participants, are presented in parentheses.

A chi-squared test of independence reveals that this null hypothesis cannot be rejected, $\chi^2(df = 2, N = 12) = .55, p = .76$. Note that because of the low number of observations in each cell, a chi-squared test approximation of the expected number of participants in each cell is not necessarily accurate. Nevertheless, on Whites classic illusion, the directions of effect for brightness induction and brownness induction appear to be independent of one another, for individual observers. Seven of the twelve participants show congruent directions of effect for both tasks, compared to the 7.5 expected.

Table 4.2: Observed (and expected) number of participants for each combination of brownness and brightness induction directions, for White's classic illusion

		Brown			
		Assimilation	None	Contrast	Total
Achromatic	Assimilation	7 (7.33)	2 (1.83)	2 (1.83)	11
	None	0 (0.0)	0 (0.0)	0 (0.0)	0
	Contrast	1 (.67)	0 (.17)	0 (.17)	1
	Total	8	2	2	12

4.5 Checkerboard variation

Eleven participants completed sufficient sessions of both achromatic matching and brown boundary setting on the checkerboard variation. Figure 4.6 shows the achromatic matches, brown boundaries, and their respective magnitudes for each participant. The corresponding descriptive and inferential statistics can be found in magnitudes table 4.3.

Five out of the eleven participants show a significant contrast effect for the achromatic matching, three participants show a significant assimilation effect, and the remaining three participants show no significant effect. Five out of eleven participants show a significant assimilation effect for the brown boundary setting, five participants show a significant contrast effect, and one participant shows no significant effect. Of the six participants who showed congruent direction of effect, three show a significantly reduced mean induction magnitude for brown boundary setting compared to achromatic matching.

Figure 4.6: **Checkerboard variant**, achromatic brightness matches (gray) and brown boundaries (brown) of all participants)

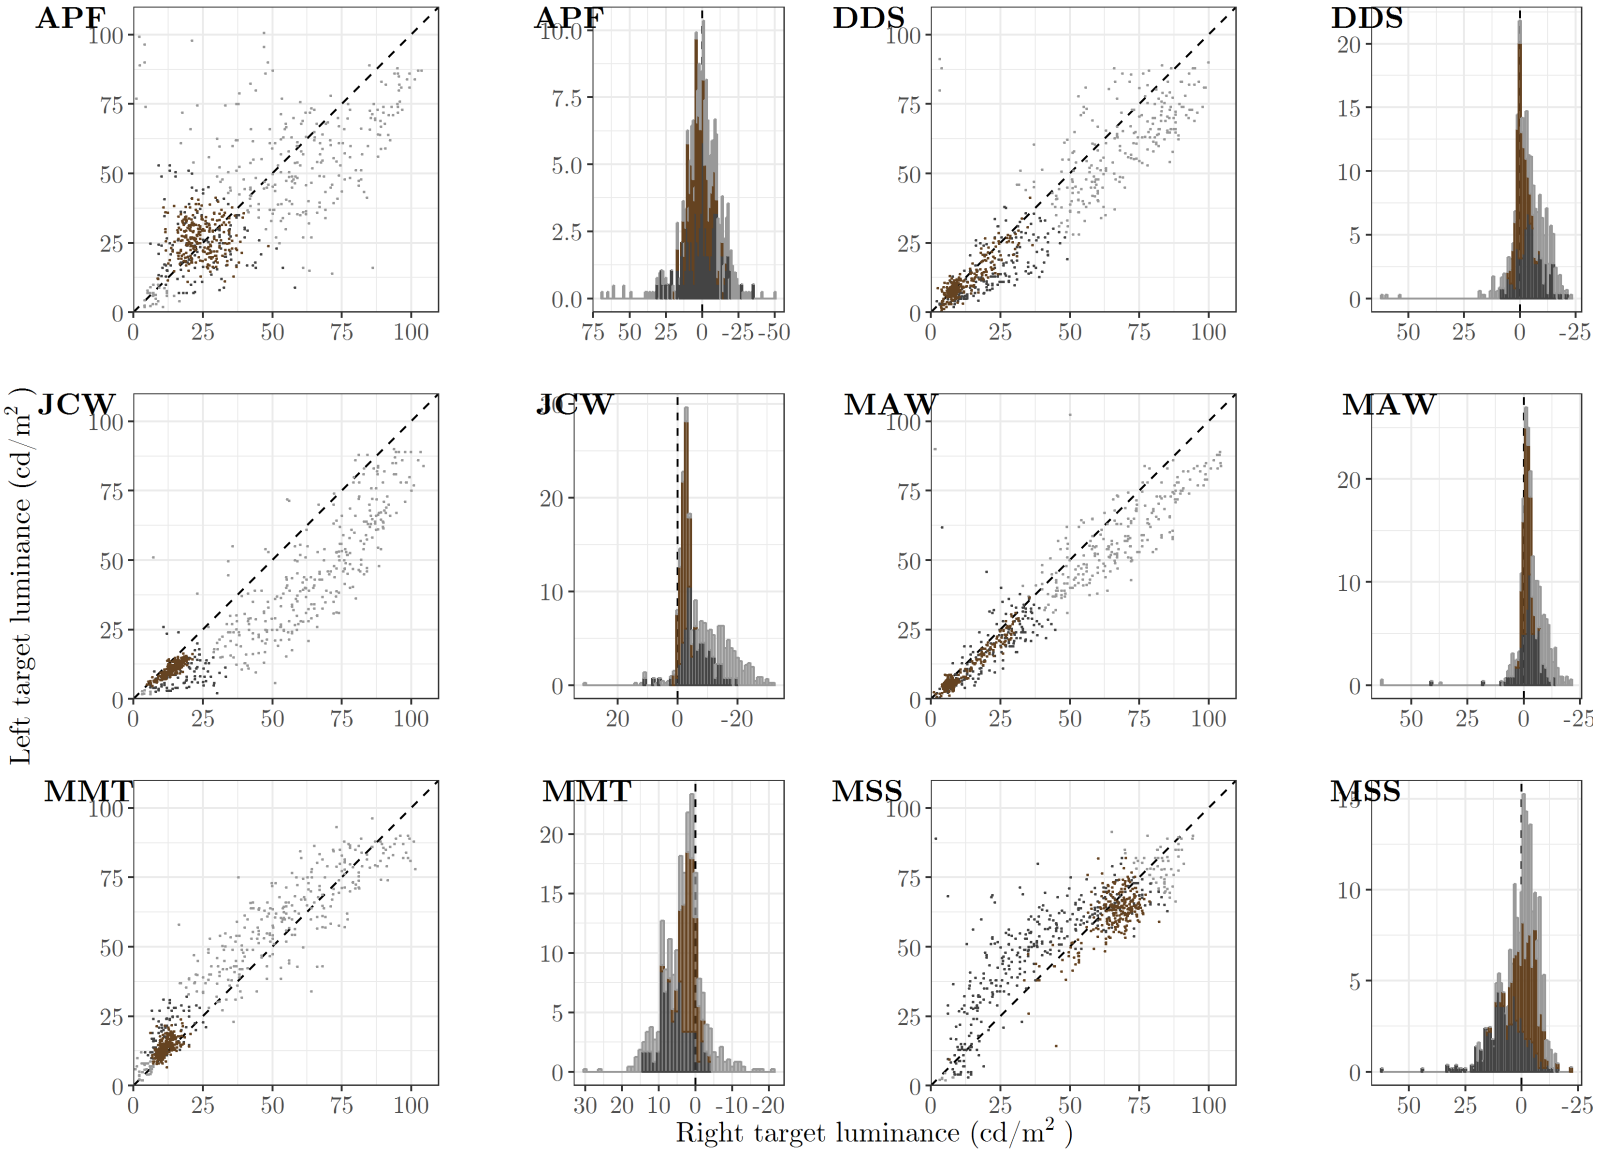

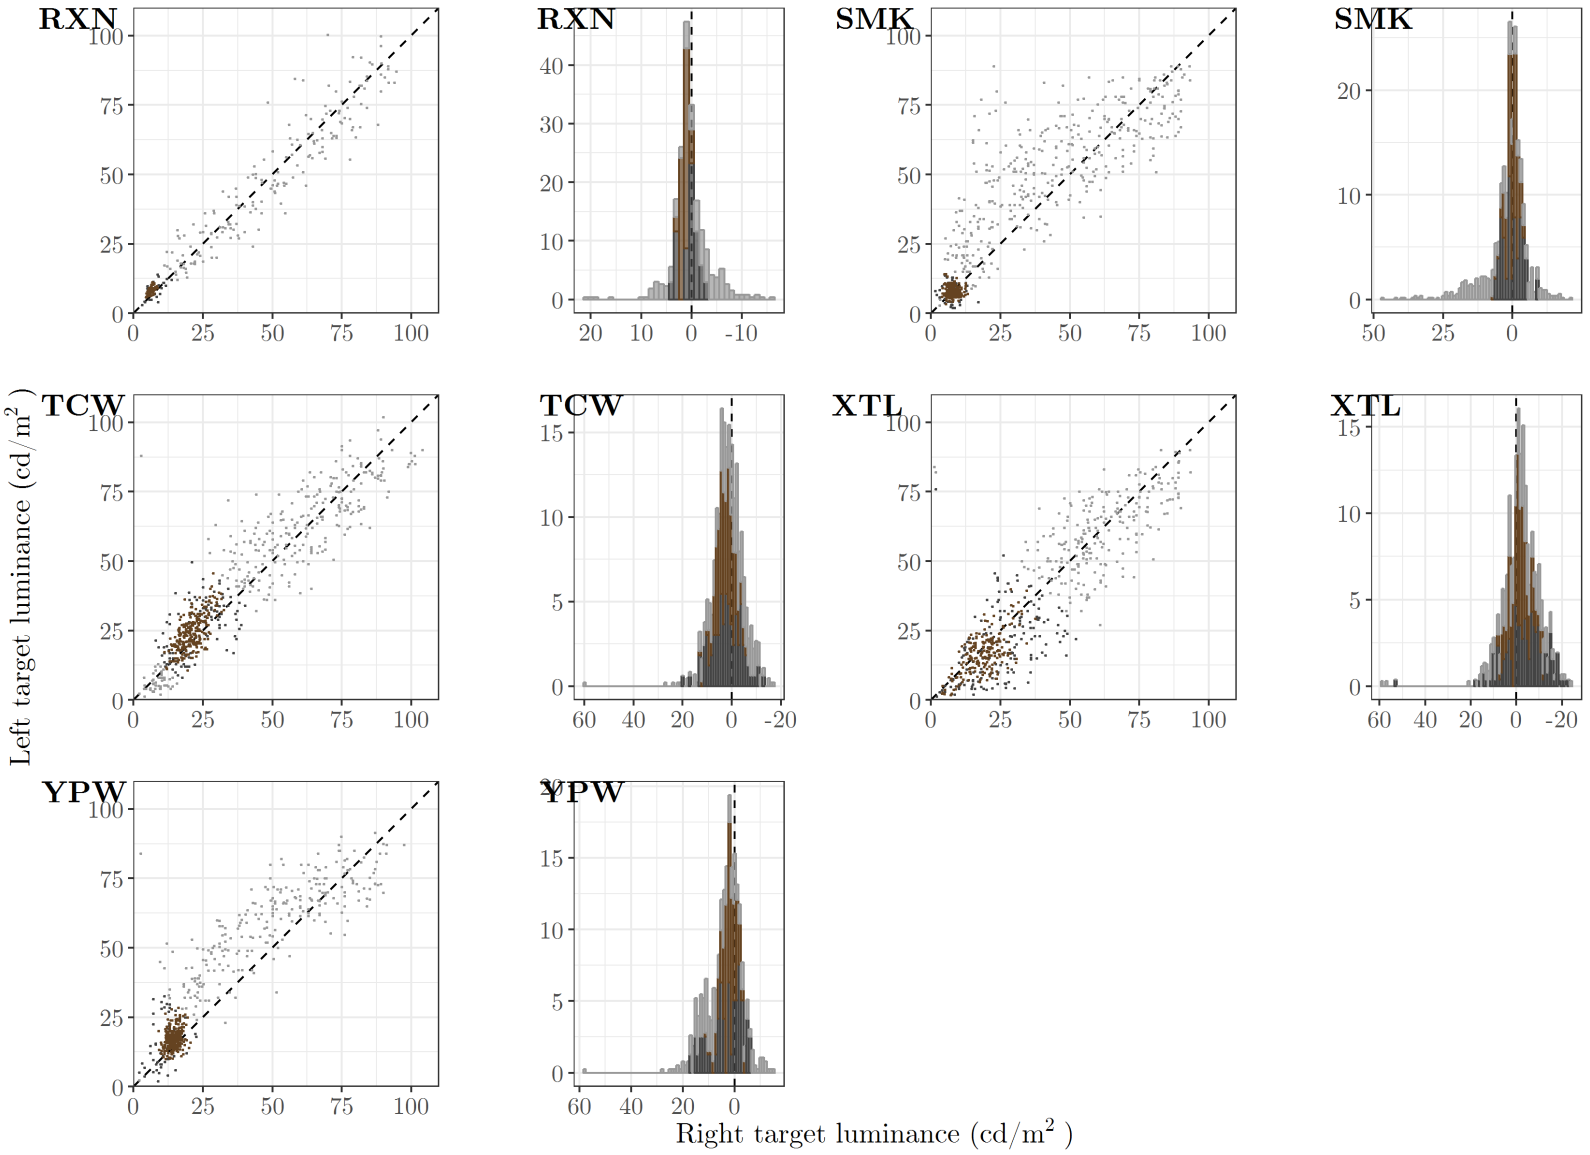
RXN
RXN
SMK
SMK
TCW
TCW
XTL
XTL
YPW
YPW
Left target luminance (cd/m^2)
Right target luminance (cd/m^2)

Table 4.3: Induction magnitudes and corresponding significance tests for the checkerboard variation

Participant	Induction magnitudes						Comparisons	
	Achromatic match		Subset of achromatic matches		Brown boundaries		Achromatic − Subset	Subset − Brown boundaries
	Mean (SD)	t(df); p	Mean (SD)	t(df); p	Mean (SD)	t(df); p	t(df); p	t(df); p
APF	-2.42 (17.03)	t(273) = -2.35; p=0.0193	0.66 (11.58)	t(125) = 0.64; p=0.5262	1.62 (6.47)	**t(255) = 4.00; p=0.0001**	t(342.16) = -2.11; p=0.0354	t(164.40) = 0.87; p=0.3876
DDS	-5.25 (11.12)	**t(229) = -7.16; p<.0001**	-4.98 (5.59)	**t(169) = -11.61; p<.0001**	-0.71 (2.28)	**t(255) = -4.96; p¡.0001**	t(356.14) = -0.31; p=0.7532	**t(206.69) = 9.45; p<.0001**
JCW	-13.60 (8.68)	**t(310) = -27.65; p<.0001**	-6.58 (5.27)	**t(88) = -11.77; p<.0001**	-2.22 (1.24)	**t(243) = -27.98; p<.0001**	**t(237.00) = -9.44; p<.0001**	**t(91.58) = 7.72; p<.0001**
MAW	-6.04 (9.20)	**t(237) = -10.12; p<.0001**	-2.53 (5.58)	**t(161) = -5.76; p<.0001**	-1.40 (1.45)	**t(255) = -15.48; p<.0001**	**t(393.31) = -4.74; p<.0001**	t(174.81) = 2.52; p=0.0127
MMT	3.01 (7.49)	**t(320) = 7.20; p<.0001**	5.39 (4.09)	**t(78) = 11.72; p<.0001**	1.68 (1.86)	**t(255) = 14.45; p<.0001**	**t(223.13) = -3.82; p=0.0002**	**t(88.18) = -7.82; p<.0001**
MSS	-2.36 (5.90)	**t(65) = -3.24; p=0.0019**	6.86 (9.02)	**t(333) = 13.91; p<.0001**	-2.57 (4.84)	**t(255) = -8.50; p<.0001**	**t(133.39) = -10.50; p<.0001**	**t(532.05) = -16.30; p<.0001**
RXN	-0.78 (5.90)	t(176) = -1.76; p=0.0805	0.33 (1.83)	t(22) = 0.87; p=0.3924	1.34 (0.71)	**t(79) = 16.95; p<.0001**	t(99.03) = -1.90; p=0.0601	t(23.91) = 2.57; p=0.0167
SMK	5.07 (11.17)	**t(365) = 8.68; p<.0001**	-0.15 (3.59)	t(33) = -0.24; p=0.8135	-0.06 (1.84)	t(216) = -0.47; p=0.6392	**t(111.05) = 6.15; p<.0001**	t(35.76) = 0.14; p=0.8896
TCW	0.68 (8.18)	t(289) = 1.41; p=0.1605	1.93 (6.07)	**t(109) = 3.33; p=0.0012**	2.28 (3.35)	**t(255) = 10.91; p<.0001**	t(263.43) = -1.67; p=0.0968	t(138.33) = 0.58; p=0.5653
XTL	-0.46 (9.78)	t(228) = -0.72; p=0.4745	-3.73 (9.21)	**t(170) = -5.29; p<.0001**	-1.63 (4.05)	**t(207) = -5.80; p<.0001**	**t(377.37) = 3.41; p=0.0007**	**t(223.73) = 2.77; p=0.0061**
YPW	6.21 (9.01)	**t(246) = 10.82; p<.0001**	3.19 (6.17)	**t(52) = 3.77; p=0.0004**	1.96 (2.52)	**t(255) = 12.43; p<.0001**	**t(105.91) = 2.94; p=0.0040**	t(55.66) = -1.43; p=0.1582

4.5.1 Congruence

Table 4.4 provides the observed number of participants for each combination of direction of effect (assimilation, none, contrast) for achromatic matching and brown boundary setting. A chi-squared test of independence reveals that the null hypothesis of independence can be rejected, $\chi^2(df = 4, N = 11) = 9.90, p = .04$. On the checkerboard variation, the directions of effect for brightness induction and brownness induction appear to be dependent on one another, for individual observers. Six of the eleven participants show congruent directions of effect for both tasks, compared to the 3.54 expected. Note that because on the checkerboard variation, at least one no effect was observed both for brightness induction and for brown induction, this contingency table has more degrees of freedom associated with it than the classic Whites illusion.

Table 4.4: Observed (and expected) number of participants for each combination of brownness and brightness induction directions, for the checkerboard variation

		Brown			
		Assimilation	None	Contrast	Total
Achromatic	Assimilation	2 (1.36)	1 (.27)	0 (1.36)	3
	None	3 (1.82)	0 (.36)	1 (1.82)	4
	Contrast	0 (1.82)	0 (.36)	4 (1.82)	4
	Total	5	1	5	11

4.6 Radial variation

Twelve participants completed sufficient sessions of both achromatic matching and brown boundary setting on the radial variation. Figure 4.7 shows the achromatic matches, brown boundaries, and their respective magnitudes for each participant. The corresponding descriptive and inferential statistics can be found in magnitudes table 4.5.

Figure 4.7: **Radial variant**, achromatic brightness matches (gray) and brown boundaries (brown) of all participants)

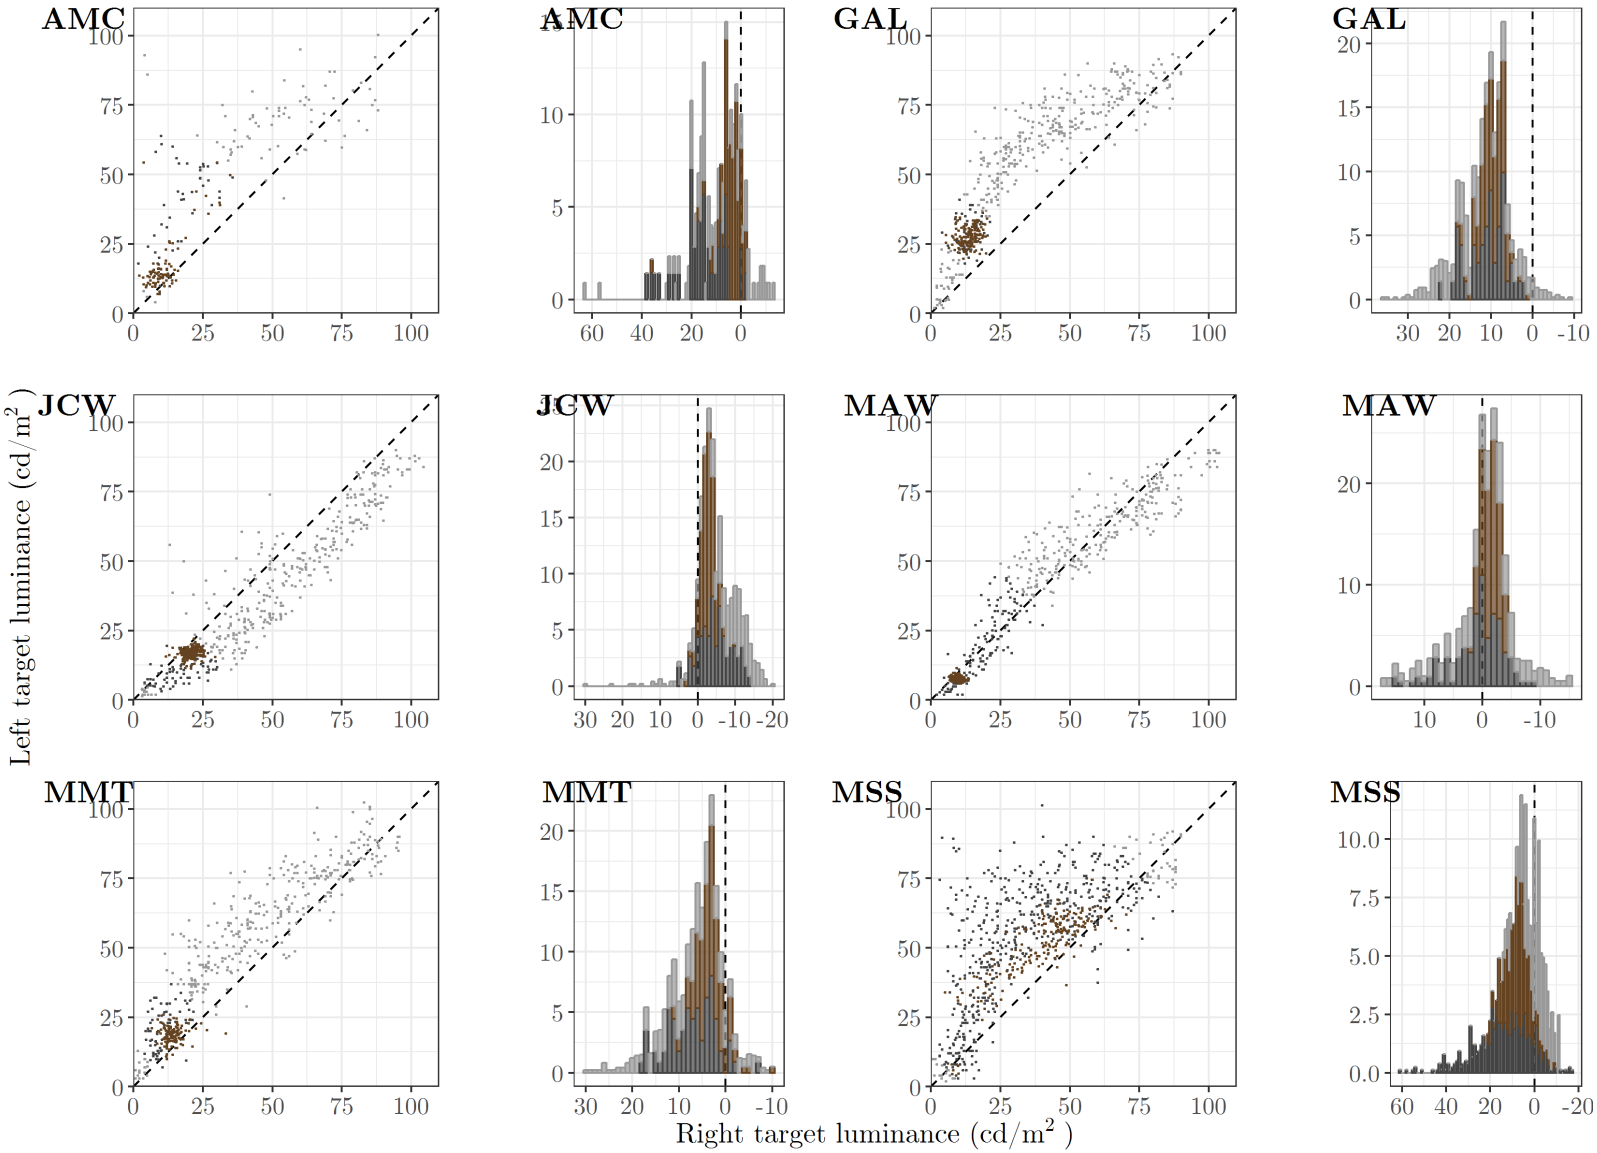

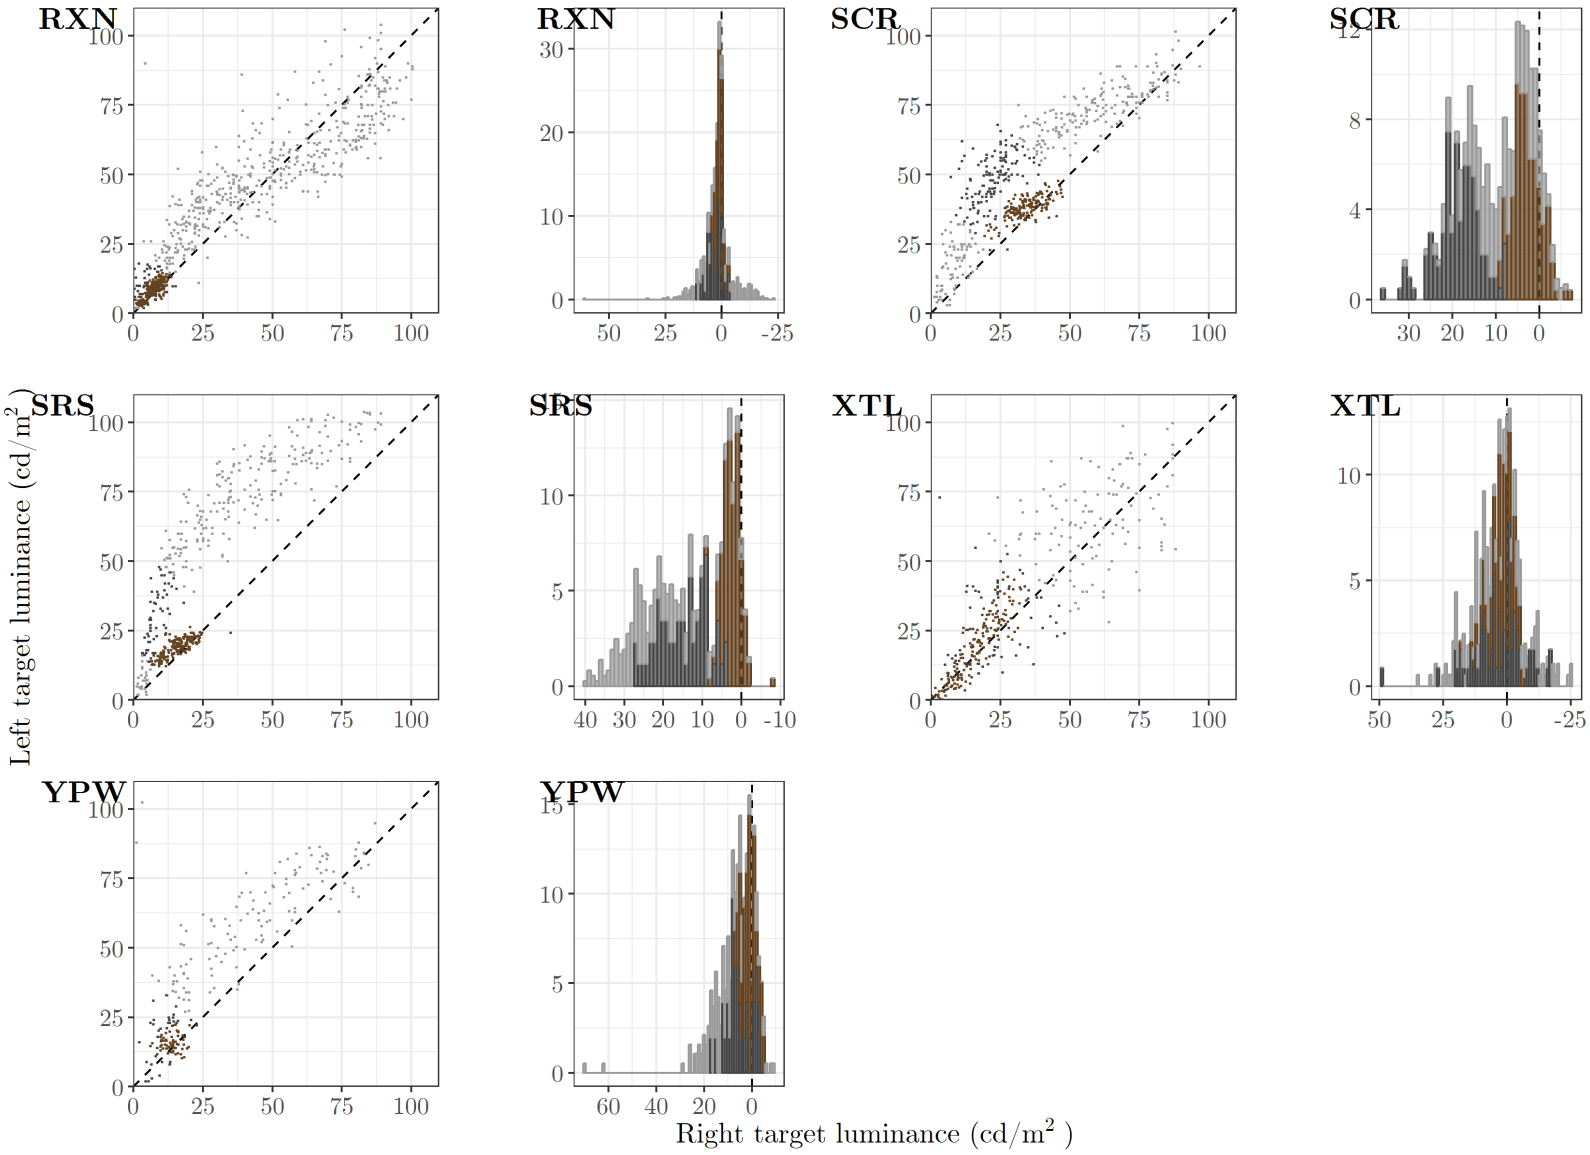
RXN
RXN
SCR
SCR
SRS
SRS
XTL
XTL
YPW
YPW
Left target luminance (cd/m^2)
Right target luminance (cd/m^2)

Eight out of the twelve participants show a significant assimilation effect for the achromatic matching, two participants show a significant contrast effect, and the remaining 2 participants show no significant effect. Nine out of twelve participants show a significant assimilation effect for the brown boundary setting, the remaining three participants show a significant contrast effect. Of the nine participants who showed congruent direction of effect, eight show a significantly reduced mean induction magnitude for brown boundary setting compared to achromatic matching.

Table 4.5: Induction magnitudes and corresponding significance tests for the radial variation

Participant	Induction magnitudes						Comparisons	
	Achromatic match		Subset of achromatic matches		Brown boundaries		Achromatic − Subset	Subset − Brown boundaries
	Mean (SD)	t(df); p	Mean (SD)	t(df); p	Mean (SD)	t(df); p	t(df); p	t(df); p
AMC	10.29 (13.43)	**t(72) = 6.55; p<0.0001**	14.92 (9.31)	**t(46) = 10.99; p<0.0001**	4.10 (4.96)	**t(86) = 7.70; p<0.0001**	**t(117.30) = -2.23; p=0.0277**	**t(60.44) = -7.42; p<0.0001**
GAL	13.07 (8.49)	**t(352) = 28.91; p<0.0001**	10.90 (4.61)	**t(46) = 16.20; p<0.0001**	9.81 (2.78)	**t(159) = 44.64; p<0.0001**	**t(94.40) = 2.68; p=0.0087**	t(56.15) = -1.54; p=0.1294
JCW	-7.08 (6.88)	**t(324) = -18.56; p<0.0001**	-4.91 (4.42)	**t(74) = -9.64; p<0.0001**	-2.62 (1.70)	**t(255) = -24.62; p<0.0001**	**t(168.12) = -3.41; p=0.0008**	**t(80.52) = 4.41; p<0.0001**
MAW	0.09 (7.20)	**t(259) = 0.21; p<0.0001**	0.63 (4.52)	**t(139) = 1.64; p<0.0001**	-1.47 (1.41)	**t(127) = -11.82; p<0.0001**	t(388.84) = -0.90; p=0.3660	**t(168.06) = -5.22; p<0.0001**
MMT	8.06 (7.57)	**t(324) = 19.19; p<0.0001**	7.03 (5.11)	**t(74) = 11.91; p<0.0001**	3.40 (3.00)	**t(127) = 12.84; p<0.0001**	t(158.66) = 1.42; p=0.1567	**t(104.36) = -5.61; p<0.0001**
MRM	-2.75 (7.15)	**t(140) = -4.56; p<0.0001**	4.56 (3.63)	**t(18) = 5.47; p<0.0001**	-2.75 (2.66)	**t(63) = -8.24; p<0.0001**	**t(40.29) = -7.10; p<0.0001**	**t(24.03) = -8.14; p<0.0001**
MSS	0.29 (5.91)	t(53) = 0.36; p=0.7213	13.98 (12.67)	**t(495) = 24.57; p<0.0001**	7.60 (5.45)	**t(192) = 19.38; p<0.0001**	**t(116.19) = -13.90; p<0.0001**	**t(680.92) = -9.24; p<0.0001**
RXN	1.36 (9.75)	**t(493) = 3.10; p=0.0020**	2.93 (3.36)	**t(65) = 7.08; p<0.0001**	1.02 (1.37)	**t(223) = 11.10; p<0.0001**	**t(251.64) = -2.60; p=0.0099**	**t(71.51) = -4.51; p<0.0001**
SCR	8.74 (6.88)	**t(265) = 20.72; p<0.0001**	18.65 (5.30)	**t(133) = 40.69; p<0.0001**	2.75 (3.03)	**t(159) = 11.48; p<0.0001**	**t(333.72) = -15.90; p<0.0001**	**t(202.90) = -30.75; p<0.0001**
SRS	21.01 (9.88)	**t(241) = 33.07; p<0.0001**	14.87 (6.16)	**t(57) = 18.37; p<0.0001**	2.27 (2.11)	**t(159) = 13.61; p<0.0001**	**t(136.57) = 5.97; p<0.0001**	**t(61.89) = -15.25; p<0.0001**
XTL	4.26 (12.88)	**t(123) = 3.69; p=0.0003**	3.46 (10.65)	**t(75) = 2.83; p=0.0059**	1.96 (4.35)	**t(159) = 5.71; p<0.0001**	t(181.08) = 0.48; p=0.6326	t(87.11) = -1.18; p=0.2419
YPW	12.02 (10.28)	**t(251) = 18.56; p<0.0001**	5.25 (5.36)	**t(67) = 8.07; p<0.0001**	1.36 (2.92)	**t(127) = 5.25; p<0.0001**	**t(210.51) = 7.38; p<0.0001**	**t(88.66) = -5.56; p<0.0001**

4.6.1 Congruence

Table 4.6 provides the observed number of participants for each combination of direction of effect (assimilation, none, contrast) for achromatic matching and brown boundary setting. A chi-squared test of independence reveals that the null hypothesis of independence should be be rejected, $\chi^2(df = 2, N = 12) = 9.33, p = .009$. On the radial variation, the directions of effect for brightness induction and brownness induction appear to be dependent on one another, for individual observers. Ten of the twelve participants show congruent directions of effect for both tasks, compared to the 6.50 expected.

Table 4.6: Observed (and expected) number of participants for each combination of brownness and brightness induction directions, for the radial variation

		Brown			
		Assimilation	None	Contrast	Total
Achromatic	Assimilation	8 (6.0)	0 (0.0)	0 (2.0)	8
	None	1 (1.5)	0 (0.0)	1 (0.5)	2
	Contrast	0 (1.5)	0 (0.0)	2 (0.5)	2
	Total	9	0	3	12

4.7 Comparison across stimuli

Figure 4.8A plots the mean achromatic induction magnitude (abscissa) and the mean brown induction magnitude (ordinate) for each observer for each stimulus (colors). Again, a positive induction magnitude (plot to the right for achromatic, and up for brown) is labeled an assimilation effect; a negative induction (to the left and down) is labeled a contrast effect. The upper-right and lower-left quadrants contain participants with congruent directions of effect; the other two quadrants contain participants with incongruent directions of effect.

Most, but not all, points lie in either the upper-right, or lower-left quadrant; most ob-

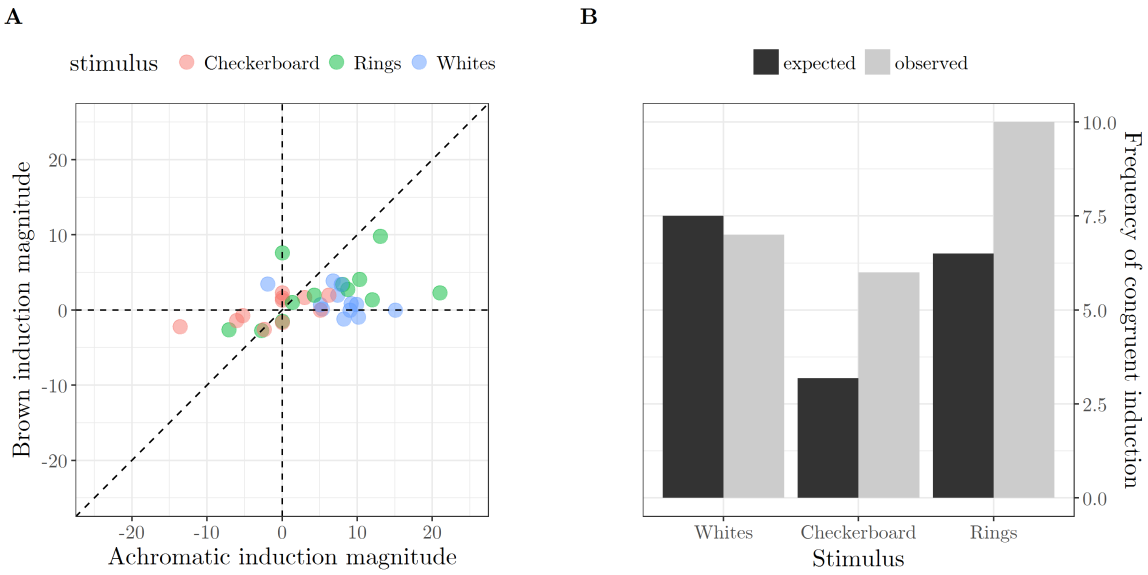

Figure 4.8: **Brightness induction and brownness induction compared across stimuli. A** Mean brightness induction magnitude (abscissa) and mean brownness induction magnitude (ordinate), for each observers, for each stimulus. **B** Expected and observed number of participants who show congruent directions of induction effect, for each stimulus

servers show congruent direction of achromatic brightness and brownness induction. Figure 4.8B further visualizes this relationship for each stimulus, by plotting the observed and expected number of participants with congruent directions of effect. However, there are also clear differences in the distribution of effects between stimuli. On the classic Whites illusion, almost all participants show an assimilation effect for brightness (right half side of Figure 4.8A). In line with the independence of brightness and brown induction, such a shift is not seen for brown boundary setting, where participants remain roughly equally distributed between contrast and assimilation. On the checkerboard variation, this does not seem to be the case, with all participants distributed around the origin. Participants are almost equally split between contrast and assimilation, for both achromatic matching and brown boundary setting. On the radial variation, the distribution of participants is most pronounced; some participants show strong assimilation effects for both achromatic matching and brown

boundary setting, while others show strong contrast effects for both.

In summary, the three stimuli differ in the distribution of achromatic induction effect directions, in the distribution of brownness induction effect direction, and in the congruence between achromatic and brownness induction effect directions. On the classic Whites illusion, all but one participant show a significant achromatic assimilation effect, but the brownness induction effects are not as strongly biased towards assimilation. As a result, the number of participants who show congruent directions of effect is not significantly different from what would be expected under the null hypothesis that the achromatic and brownness induction are independent.

On the radial variant, more participants show an assimilation effect than a contrast effect, for both brownness and brightness induction. Significantly more participants showed congruent directions of effect than would be expected if brightness and brownness induction were independent.

On the checkerboard variant, participants are almost evenly split between assimilation and contrast, for both brightness and brownness induction. Moreover, the number of participants who show congruent directions of effect for brightness and brownness appears to be significantly more than would be expected if these were independent. However, the evidence for this is not strong, as one participant showing a different direction of effect for either brownness or brightness could change the statistical significance.

Thus, on Whites classic illusion, brightness and brownness induction appear to be independent, while on the radial variant, they are not independent.

Chapter 5

TWO-STAGE FRAMEWORK FOR INDUCTION

5.1 Independence of brightness and brownness

The psychophysical results presented in the previous chapter do not lend themselves to straightforward interpretation of the relationship between brightness induction and brownness induction. The most direct evidence for brownness and brightness induction mediated by a single underlying mechanism, would come from identical direction and magnitude of effect. The three stimuli paint different pictures when it comes to this congruence.

On the radial variation, the direction of brightness and brownness are not independent. While there are individual differences in the direction of both brightness and brownness induction, more observers tended to show congruent induction directions than expected. This suggests that achromatic brightness induction and brownness induction are mediated by the same underlying mechanisms for this stimulus (Figure 5.1A). This shared set of mechanisms process the stimulus in the same way, regardless of whether the target elements are achromatic. The reduced magnitude of brownness induction compared to brightness induction could be explained as an attenuating response gain, or any other transformation that would reduce the difference between the two targets when mapping from the output of these mechanisms to the psychophysical brown boundaries.

On Whites classic illusion, the direction of brightness and brownness induction are independent. Moreover, brightness assimilation is overwhelmingly more likely than achromatic brightness induction in the contrast direction, while brown induction does not share this strong bias. This suggests that achromatic brightness induction and brownness induction are mediated by independent mechanisms (for this stimulus). If brownness and brightness induction would share an underlying mechanism, they would have to be correlated for any

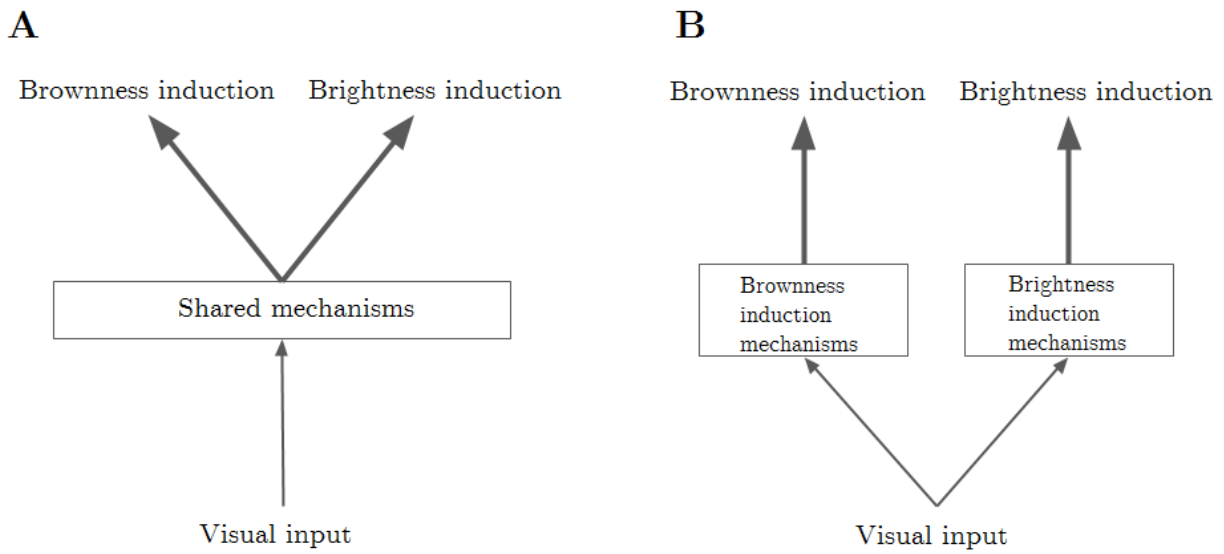

Figure 5.1: **Possible relationships between brightness and brownness mechanisms. A** A set of shared mechanisms underlies both brightness induction and brownness induction. **B** Brightness induction and brownness induction are mediated by independent sets of mechanisms.

stimulus. However, the direction of achromatic brightness induction does not predict the direction of brownness induction on the classic illusion. Thus, some factor has to influence either brownness, or brightness, without influencing the other. Two separate (sets of) mechanisms for brightness induction and brownness induction could operate in parallel, on the same visual input, and produce uncorrelated induction effects (Figure 5.1B).

The checkerboard variation fails to provide a case for either model. Brightness and brownness induction appear not to be independent, at least in that the directions of effect are significantly more congruent than expected. However, the evidence for this is not very strong, as one participant could change the statistical significance. The checkerboard variation is also interesting in that it shows no bias for either contrast or assimilation direction effects; for both brightness and brownness induction, participants are roughly equally split between the two directions. The checkerboard appears to support the parallel mechanisms interpretation

like the radial variant, but due to the weak significance of the data, no conclusions will be drawn from this variant.

5.2 Introducing two-stage framework

A parsimonious account of the data requires a schema in which brownness and brightness induction both parallel and shared (in series). An attempt to provide a single set of mechanisms that can explain all three tested variants of Whites illusion (as well possible additional stimuli) should allow for independence of brightness and brownness under some conditions but not other conditions. This can be achieved by considering at least two stages of induction processing (Figure 5.2).

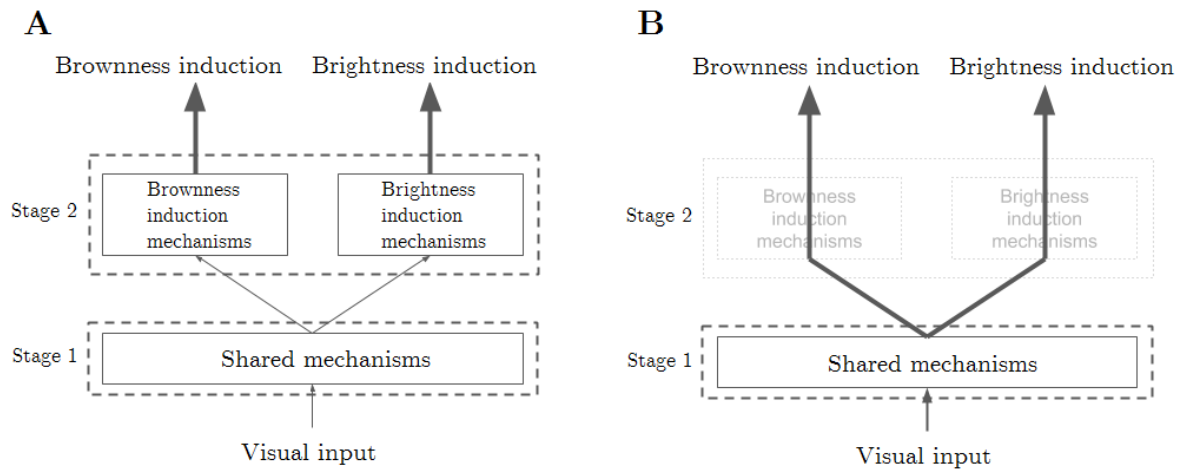

Figure 5.2: **Two stages of processing A**, can explain the effect of independent mechanisms mediating brownness and brightness induction, as a weak contribution of the shared first stage mechanisms, compared to the contributions of the independent second stage mechanisms. **B** The same two stage framework can also explain the lack of independence of brightness and brownness induction, as a strong contribution from the shared first stage mechanisms, compared to the contributions of the independent second stage mechanisms.

A first stage of processing is shared between brightness and brownness induction, after

which additional brightness induction mechanisms operate independently from additional brownness mechanisms. Since brightness processing (however minimal) must occur for any stimulus, and brownness processing is unique to sufficiently long wavelength stimuli in certain contexts, the first stage of this framework must be composed of mechanisms that are general enough to apply to any stimulus that undergoes brightness processing. In the second stage, mechanisms can be highly specialized to either brightness on brownness. Different stimuli are hypothesized to drive these two stages of processing differentially, which allows for their effects to be teased apart (Figure 5.2A and B).

All three stimuli currently investigated would drive the initial stage, which would induce some brightness into the targets in each stimulus. The output of this initial stage feeds into the brownness induction mechanism, which thus receives input that has already undergone induction. Any brownness induction would compound the initial brightness induction, but this is hypothesized to be minimal, as explained later in this chapter. The brownness induction would be strongly associated with the initial stage of processing; certainly the direction of effect would have to be consistent (or consistently inverted) between the first stage and the brownness processing in the second stage (Figure 5.2A).

The initial stage of shared processing would also feed into a second stage of additional brightness processing. Similar to the relationship with brownness induction, the second stage of brightness processing alters the output of the first stage to further adjust the perceived brightness of the target. Different properties of a display would drive this second stage, compared to the initial stage. As a result, this second stage could alter the overall brightness induction effect, including the direction effect. The second stage does not feed into the brownness mechanism, or only minimally so. It is this second stage that allows brightness induction and brownness induction to be independent of each other.

Note that when attempting to explain, predict or fit just brightness induction data for a single stimulus, a model would be able to capture the effect of multiple stages of processing as a single stage. Especially, the multiscale spatial-filtering models of brightness perception reviewed in chapter 2, which transform physical luminances into perceived brightness,

generate output in arbitrary units. These arbitrary units need to be scaled to fit with psychophysical measurements of induction effects to be meaningful, and thus can be made to fit any magnitude of effect, for a single stimulus. Certainly this is part of the appeal of such models: they can operate on any stimulus, whether artificial and simple or naturalistic and complex, and be scaled to fit with psychophysical results.

However, when comparing across different stimuli, presumably one method of scaling the model output to psychophysical results would be used. If these stimuli differentially involve the two stages of processing, and if these differences are not adequately captured in a model, a single scaling factor cannot make the model fit all the results with high accuracy. Indeed, when comparing predictions with a given set of parameters, filtering models often get right the direction of effect in different illusions, but not the magnitude of effect (Robinson et al., 2007). Moreover, comparing brightness induction effects to brownness induction, for multiple stimuli, can provide patterns of results that require different sets of parameters for different effects. Since participants can show different effects for these two types of induction, it is not readily clear how one set of parameters would predict both directions, unless taking into account that certain mechanisms are involved in brightness processing but not brownness processing (or vice versa). It is the comparison across these stimuli and tasks that help tease apart how different features of the stimuli drive the stages of processing.

The main evidence in the current investigation for a two-stage framework of induction comes from the independence of brightness and brownness induction on Whites classic illusion, compared to the lack of independence on the radial variant. The radial variant of Whites illusion thus appears to strongly drive the initial stage of brightness processing, but not the second stage (Figure 5.2B). Within individual participants, brownness induction seems to follow brightness induction quite strictly. Whites classic illusion, seems to drive both stages of brightness processing: some brightness induction from the initial stage is carried over into brownness induction, but the contribution of the second brightness induction stage independently induces more brightness into the target (Figure 5.2A). Presumably it is this second stage that strongly biases the brightness induction towards the assimilation direc-

tion on the classical illusion, which is not seen for brownness. The checkerboard variation, in many ways, seems to lie in between the other two stimuli. More participants show congruent directions of effect than on the classic illusion, but not enough to convincingly show a lack of independence between brightness and brownness induction. The checkerboard variation also does not show a strong bias towards either direction, neither for brightness nor brownness induction. Based on these results, the checkerboard variation represents a stimulus that drives both stages of brightness processing. For some participants it drives the initial stage more, leading to congruent directions of effect, for other participants a stronger drive of the second stage of brightness processing leads to incongruent directions of effect.

5.3 Brightness in the two-stage framework

While the goal of current investigation is not to outline a complete theory of brightness or brownness perception, or to argue for or against any individual mechanism previously identified in the literature, it is worthwhile to consider which previously proposed mechanisms could be associated with each of the proposed two stages. Many more explanations of brightness induction than brownness induction are available, so a discussion of these explanations and their possible place in the two-stage framework provides a more substantial base of evidence. After discussion of these brightness mechanisms and their place in the two-stage framework, predictions for mechanisms involved in brownness processing will be discussed.

Since White's classic illusion and the radial variant are hypothesized to differ in how much they drive the second stage, what features of these stimuli set them apart, and thus might indicate the kind of brightness processing that takes place in the second stage?

An immediate candidate would be the anisotropy of the classic illusion. While the radial variation has equal power in all orientations, the classic illusion is defined by the horizontal and vertical lines in the display. Junction analysis explicitly uses these edges to explain the effects in this stimulus. In the classic illusion, the top edge of T-junctions is argued to prevent contrast effects from inducing darkness and brightness into the target that looks brighter, and the target that looks darker, respectively. Based on junction analysis, Anderson (1997, 2001)

argued that perceptual scission of the classic illusion into separate perceived layers, or image sources, explains the effect. The target elements collinear with the white bars are perceived as a semi-transparent layer on top of the background grating. Some of the brightness of the luminance of this layer is perceptually taken out and assigned to the underlying layer (the white elements of the grating that have been replaced). This reassigning of brightness from one layer to another could be due to an inference of transparency (e.g., the target layer is a semi-transparent surface partially occluding the background grating), or of differences in illumination (e.g., the target layer is a shadow cast onto the background grating). In either case, the grating is the true surface that contributes the brightness of the target area, which the target layer attenuates. This makes that these target elements are perceived as darker, to reflect the true set of sources contributing to the image. Indeed, many observers describe this set of target elements in the classic illusion to have some perceptual quality of transparency. This suggest image interpretation, depth relations, transparency and colour constancy as additional candidates for the second stage of brightness processing.

The radial variation has low complexity (both sides of the display contain just 7 elements) and is isotropic. Moreover, it contains edges between the rings, but no junctions. Howe (2005) argues that because a brightness induction effect is still observed in the absence of T-junctions, this variation suggests that junction interpretation is not the primary mechanism driven by all variants of Whites illusions. Yazdanbakhsh et al. (2002) made a similar argument, by selectively removing the T-junctions in the classic illusion. Image filtering models have predicted different directions of effect for this radial variation, depending on the model parameters. Two sets of parameters of the FLODOG seem to capture the assimilation direction of effect on a radial variation, while other parameters and models predict either no effect, or contrast effects (Robinson et al., 2007). Since the psychophysical results from the current investigation also show that this variation is susceptible to large individual differences in the direction and magnitude of induction, it is not clear which direction of effect a filtering model should predict. Instead, the brightness induction in the radial variation has been suggested to result from anchoring within frameworks based on perceptual grouping

(Howe, 2005; A. L. Gilchrist et al., 1999; A. Gilchrist, 2015).

Taking together, these interpretations of the features of the classic illusion and the radial variant, suggest an initial categorization of brightness mechanisms within the two-stage framework. The initial stage of brightness processing might perform some grouping, and a preliminary anchoring of lightness within these groupings. This would not require understanding of the scene, and could be based on Gestalt grouping of low-level features (luminance, contrast, orientations). Alternatively, spatial filtering models could mimic the mechanisms of this initial stage of processing (Figure 5.3) This would require some explanation for why some observers perceive the radial variant as going in the contrast direction, while other observers perceive it as going in the assimilation direction. The second stage of brightness processing would alter this initial anchoring by parsing the scene and making inferences about illumination, depth, transparency and other higher-level features (Figure 5.3).

This initial categorization leads to general predictions for the checkerboard variation of Whites illusion, and the congruence of brightness and brownness on this variant. This variant shares the anisotropy of the classic illusion, as well as straight edges. However, it replaces the T-junctions that are present in the classic illusion with X-junctions that previous junction analyses make no clear predictions for. Thus, it lacks certain key features that could drive brightness processing in the second stage. This would predict that most of the brightness induction in the checkerboard variant is mediated by the initial stage of processing. Since brownness is in general mediated mainly by the initial stage, the checkerboard variation is predicted to show a lack of independence of brightness induction and brownness induction, seen as more congruent directions of effect than predicted under the null hypothesis. This is indeed the pattern of results that is observed for the checkerboard variation; however, as pointed out earlier, this conclusion will only be made tentatively.

5.4 Brownness in the two-stage framework

Fewer factors thought to influence brownness induction have been identified than for brightness induction. Those factors that have been identified do seem to suggest that brownness induction is mediated by a spatially localized contrast mechanism, and less localized general context mechanism, both inducing brownness in yellow targets surrounded by a bright context. The two mechanisms are analogous to simultaneous brightness contrast and brightness context mechanisms. Simultaneous brightness contrast is considered to be an effect mediated by low-level mechanisms, such as lateral inhibition. Hence, it is considered to be part of the first stage of induction processing in the current framework. If the simultaneous brightness contrast effect is mediated by similar mechanisms, it too would be be considered part of the first stage. Brightness context effects, and the corresponding brownness context effects, are of less clear origin. However, these effects can be observed in very simple displays, with just a single target and a single remote context. It does not seem to invoke any of the mechanisms that have been proposed to depend on T-junctions, or other interpretations of the display. Context effects, then, can also be considered low-level effects, and are tentatively placed in the initial stage of processing. The two mechanisms proposed to mediate brownness in previously studied displays both are placed in the initial stage.

In the current investigation, presumably different mechanisms have to mediate brownness induction than previously identified, since brownness induction in the assimilation direction is observed (at least for some participants). Could the same explanations that have been proposed for brightness induction these stimuli also apply to brownness induction? On the radial variant, and possibly the checkerboard variant, brownness induction appears to follow brightness induction in the direction of effect. Presumably, the mechanisms identified as underlying this brightness induction processing in the initial stage also mediate brownness on these stimuli.

On the classic illusion, brightness induction is heavily biased towards the assimilation direction, while brownness is not. This suggests that the mechanism responsible for biasing

brightness induction in this direction is exclusive to brightness induction, and not applicable to brownness induction. In other words, on the classic illusion the second stage of brightness processing is driven strongly, and independently from brownness. From the previous section, higher-level processing would fit the role, as the classic illusion contains more features that have been argued to drive these higher-level mechanisms, e.g. T-junctions and their possible signalling of either occlusion, or transparency or changes in illumination. Informally, the contribution of these mechanisms to brightness induction could be gleaned from observers indicating that they perceive the achromatic target elements collinear with the white bars in the classic illusion to appear transparent. It is unclear whether observers have the same percept of transparency for the analogous brown targets. Most observers do informally indicate a difference in saturation between the two sets of brown target elements in the classic illusion. Physically, a semi-transparent layer on a bright white background would look desaturated compared to the same semi-transparent surface on a black background. The white light shining through the semi-transparent layer would perceptually mix with the light reflected by semi-transparent layer itself, thus desaturating the perceived color of this semi-transparent layer. Desaturation of these sets of target elements in the classic illusion thus might indicate a scission of inference, similar to the achromatic version. However, slight desaturation of a target has been found to enhance the perception of brownness (Bartleson, 1976; DeLawyer et al., 2016). Thus, it is unclear whether layer decomposition and perceived transparency would induce brownness in the same or opposite direction as brightness, if it affects brownness perception at all.

The current investigation suggests that brownness on the radial and checkerboard variation could possibly be mediated by the same mechanisms as brightness induction, as part of initial stage in the proposed two-stage framework. Proposed explanations such as spatial filtering accounts, and anchoring within frameworks defined by Gestalt groupings should thus also apply to brownness, and testing these mechanisms in isolation could be a future direction in the investigation of brownness induction. Oh Whites classic illusion, these same mechanisms would be at play, but additional mechanisms could be involved in brownness

induction. While scission and inferences about transparency or illumination are hypothesized to play a role in brightness perception on this stimulus, it is not clear whether these explanations also apply to brownness induction. If they do, additional mechanisms must shift either brightness induction or brownness induction, so that the two induction effects are independent of each other.

5.5 Additional considerations

Similar cases for at least two stages of brightness processing have also been made by Robinson and de Sa (2008), and expanded on by Kaneko and Murakami (2012). The former found that a fast-acting mechanism mediated simultaneous brightness contrast for very brief stimulus durations, while a slower mechanism is responsible for brightness assimilation. After this initial brightness induction, the strength of the illusion would change with longer exposure. The two separate time courses also seem to have different spatial properties (Kaneko & Murakami, 2012). The mechanism mediating brightness perception for short stimulus durations is spatially localized to produce fast simultaneous contrast, while a mechanism with a longer time course appears to take into account more remote context instead. These findings are compatible with the aforementioned ideas that initial stage of brightness processing would act on low-level features to provide initial anchoring, which is then augmented by the larger context in the second stage. This is not to say that the categorization of brightness mechanisms into two stages presented here, is a conclusive and comprehensive account of brightness processing. Firstly, the hypothesized two-stage theory of brightness processing cannot rule out any previously proposed mechanism. Several additional explanations brightness induction, as well as many variations on the explanations reviewed in Chapter 2, can be found in the (extensive) literature. The omission of these additional explanations from the present discussion does not rule out that these explanations could be included in either stage of the framework presented here. Secondly, there is certainly no reason presented here to assume that there are just two stages of brightness processing: each of the two hypothesized stages could include multiple mechanisms that operate sequentially or simultaneously. The current

investigation does not have enough paradigms and variations on the Whites illusion to differentiate between two or more stages. Perhaps a more detailed investigation would reveal a hierarchy of mechanisms mediating brightness perception, performing different computations on different features, from the lower-level features to the higher-level scene understanding. Lastly, both feedforward and feedback inhibition between different stages, or different points in the hierarchy of brightness processing, would create complex interactions between mechanisms. Ultimately, a combination of many stimulus features, driving many mechanisms, is the most likely origin of our rich world of brightness perceptions.

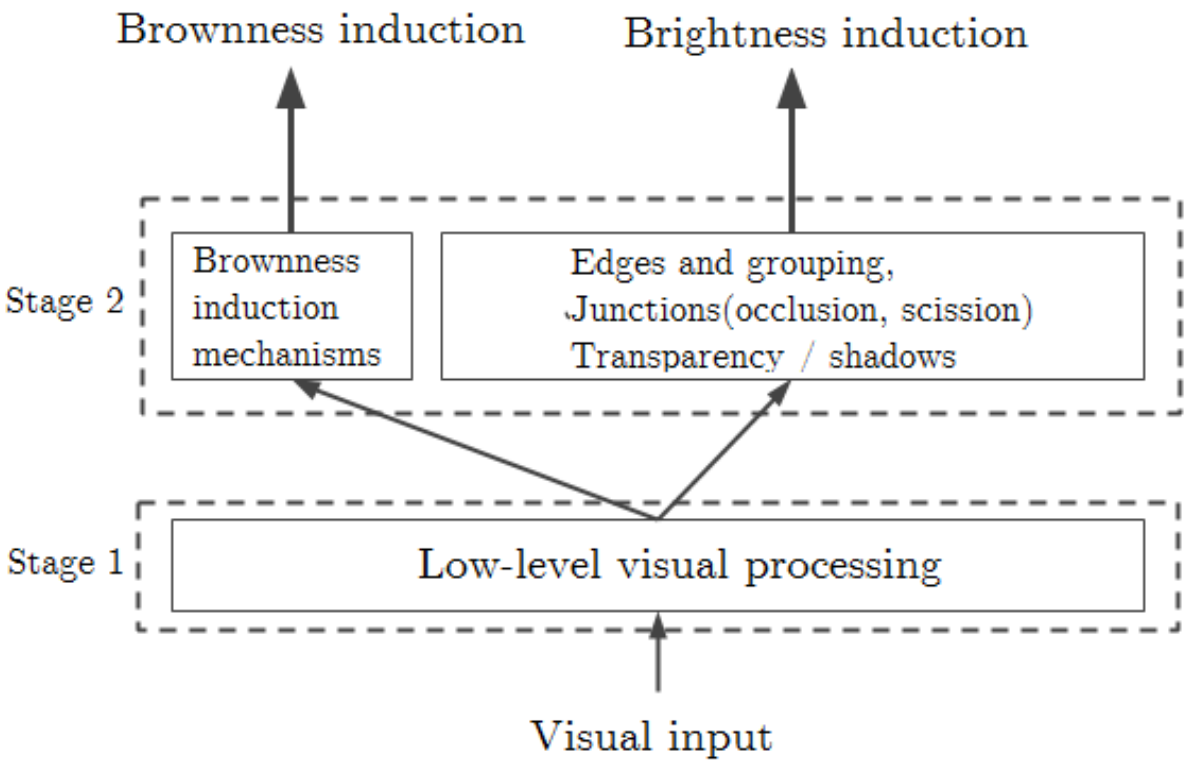

Figure 5.3: **The full two-stage framework of brightness and brownness induction** posits that the initial, shared, stage consists of lowlevel visual processing, such as the mechanisms modeled in spatial filtering models. The second stage could fit mid- and high-level visual mechanisms such as those that are sensitive to edge and junctions, grouping of regions on the basis of Gestalt principles, and inferences about depth and transparency.

Chapter 6

CONCLUSION

The current investigation aimed to further our understanding of the perception of brownness. This color category is remarkable in that it is perceived only in a sufficiently bright context, and thus is not explicitly incorporated into opponent process models of color vision that typically focus on bright colors. Opponent process theories of color vision are assumed to hold for both bright and dark colors, such that red and green are opponent when both are bright but also when both are dark. However, the bright color yellow does not seem to be represented among the dark colors. Previous work has argued that instead it is replaced by brown, which fulfills the same role in opponent process theory among dark colors that yellow fulfills among bright colors. Brown and yellow share the crucial properties for their place in opponent processing: both cancel, and are canceled by, blue, both are red-green balances, and both can appear independent of any other hue percept. An interesting and as of yet unexplained difference between yellow and brown is that while both are red-green balance points, equilibrium brown tends be redshifted, so that it is closer to a dark orange than a dark equilibrium yellow. The similarities between yellow and brown suggest that brown is the dark counterpart of yellow in classic theories of color vision; a unique role, since similar color category shifts between bright and dark colors do not happen for the other opponent colors. Why brown has this unique position is subject to mere speculation at this time.

Further understanding of the role of brown in color vision theories entails a better understanding of what mediates the induction of brown. Previous work has shown that brownness can be reliably induced in yellow targets by bright contexts. Several key manipulations of this surround, such as its luminance, distance from the target, complexity, temporal dynamics, and comparison of monocular and dichoptic presentations, all suggest that there are

at least two mechanisms mediating the induction of brownness. A fast, spatially localized mechanism strongly induces brownness, and presumably operates in monocular pathways. A slower mechanism takes into account the larger context around the stimulus, is not contiguity dependent, and could operate after convergence of signals from the two eyes.

The type of contexts that induce brownness, as well as the factors that appear to influence the effect of these contexts, bear striking resemblance to those that have been known to induce darkness in achromatic targets. The displays and effects presented in Chapter 1 are similar to simultaneous brightness contrast and context effects, as discussed in Chapter 2. This begs the question whether brownness is just the perception of darkness in yellow stimuli. This would not explain why a categorical shift takes place, but it would allow the perception of brownness to be explained by known and hypothesized mechanisms of brightness perception.

The goal of the current investigation was to explore what mechanisms might mediate the induction of brownness, specifically focusing on whether mechanisms of brightness perception might be sufficient to explain brownness. By comparing the effect of contexts that are known to induce darkness, and the extent to which these contexts also induce brownness, the case for shared underlying mechanisms can be evaluated. Specifically, if these context induce brownness and darkness in the same direction and with comparable magnitude, it is likely that these two percepts are mediated by the same mechanisms. This hypothesis was tested on Whites classic brightness illusion, a checkerboard variant, and a radial variant. These two variants eliminate some but not all of the features that have previously been hypothesized to mediate the effect in the classic illusion. For each of these stimuli, brownness induction and achromatic brightness induction were measured, and compared.

The three stimuli paint a different picture of the relationship between brightness and brownness induction. While brownness induction and brightness induction appear to be independent on the classic illusion, this is not the case on the radial variation. The results from the checkerboard variation suggest that brownness induction and brightness induction might not be independent, but the evidence is not overwhelming. From this comparison it is clear that brownness induction and brightness induction can behave similarly, or independently,

depending on the features of the stimulus.

In order for different stimuli to yield different results with regards to the independence of brightness and brownness, it is suggested that these two types of induction share some, but not all of the same underlying mechanisms. The shared mechanisms, which could generate congruent brownness and brightness induction are presumably low-level visual processing mechanisms. Simultaneous contrast seems to be involved in the perception of both brightness and brownness, based on previous literature. The accounts and explanations of brightness assimilation through spatial averaging, and pattern-specific inhibition, would certainly suggest that these mechanisms are not exclusive to the domain of brightness, and could also affect brownness. Comparing the different stimuli tested in the current investigation suggests that higher-level mechanisms involved in brightness processing, such as perceptual scission, could be isolated from brownness processing. While the current investigation suggests a framework of at least two stages of brightness processing, only the first of which is shared between brightness and brownness induction, a more complete theory likely involves a hierarchy of mechanisms, some of which also feed into brownness processing.

In introducing the two-stage framework, the current investigation provides initial speculation on which mechanisms previously invoked to explain Whites illusion might fit in which stage. More direct evidence for these divisions can be accrued by developing stimuli that allow for direct comparison of the different features necessary for the different explanations. For each such stimulus, comparing the brightness induction and brownness induction allows for assignment of the hypothesized mechanisms to either the shared first stage, or the isolated second stage. Such psychophysical tests of the multiple stages of processing can clarify and extend the framework, and can also be used to investigate the relationships between the different mechanisms and different stages.

Similarly, computational modeling of previously hypothesized explanations of Whites illusion can be used to explore how new stimuli are expected to involve each mechanism. The multiscale spatial filtering models provide a powerful tool, as they allow for comparison among different brightness illusions. The versatile models can make a prediction of the

brightness induction in any stimulus, and indeed any arbitrary visual input. By comparing the pattern of predictions across stimuli, to psychophysical measurements of the perceived brightness induction, different models and their variations can be evaluated. Such models allow for exploration of plausible neurophysiological mechanisms to give additional evidence for explanations brightness perception based on the image analysis. For instance, models expected to capture just the initial stage of brightness processing should be able to explain brightness on the radial variant in the current study, but fail (or at least underestimate) the classic illusion since the latter involves additional processing. Yet that same model should be able to explain brownness on both of these illusions, as that is limited to the initial stage.

The presented comparison of induction effects suggests that while brownness shares some properties and possibly mechanisms with brightness, brownness cannot just be considered darkness in yellow stimuli. Moreover, the results suggest that multiple mechanisms are likely involved in brightness processing. The multistage framework presented here is an initial suggestion of how such a combination of multiple mechanisms could be thought of. The framework also suggests specific hypotheses for how explanations should generalize from brightness to brownness, and from one stimulus to another. Thus, the multistage framework of induction effects is offered as a jumping off point for further investigation into the relationship between brightness and brownness, and into the mechanisms underlying both of these perceptual dimensions in general.

References

Adelson, E. H. (2000). Lightness Perception and Lightness Illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., Vol. 3, pp. 339–351). Cambridge, Massachusetts.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. *Perception*, *26*(4), 419–453. doi: 10.1068/p260419

Anderson, B. L. (2001). Contrasting theories of White's illusion. *Perception*, *30*(12), 1499–1507. doi: 10.1068/p3012ds

Anderson, B. L. (2003). Perceptual organization and White's illusion. *Perception*, *32*, 269–284.

Anderson, B. L., & Winawer, J. (2005). Image segmentation and lightness perception. *Nature*, *434*(March), 79–83. doi: 10.1038/nature03343.

Anstis, S. (2006). White's effect in lightness, color and motion. In *Seeing spatial form* (pp. 91–100).

Barkan, Y., Spitzer, H., & Einav, S. (2008). Brightness contrast-contrast induction model predicts assimilation and inverted assimilation effects. *Journal of Vision*, *8*(7), 1–26. doi: 10.1167/8.7.27

Bartleson, C. J. (1976). Brown. *Vision Research*, *1*(4), 181–191.

Berlin, B., & Kay, P. (1969). *Basic color terms: their university and evolution*. Berkeley, CA: California UP.

Betz, T., Shapley, R. M., Wichmann, F. A., & Maertens, M. (2015). Noise masking of White's illusion exposes the weakness of current spatial filtering models of lightness perception. *Journal of Vision*, *15*, 1–17. doi: 10.1167/15.14.1.doi

Blakeslee, B., & McCourt, M. E. (1999). A multiscale spatial filtering account of the White effect, simultaneous brightness contrast and grating induction. *Vision Research*, *39*, 4361–4377.

Blakeslee, B., & McCourt, M. E. (2004). A unified theory of brightness contrast and assimilation incorporating oriented multiscale spatial filtering and contrast normalization. *Vision Research*, *44*(21), 2483–2503. doi: 10.1016/j.visres.2004.05.015

Boynton, R. M., & Olson, C. X. (1990). Salience of Chromatic Basic Color Terms Confirmed by Three Measures. *Vision Research*, *30*(9), 1311–1317.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial vision*, 433–436. doi: 10.1163/156856897X00357

Bressan, P., & Kramer, P. (2008). Gating of remote effects on lightness. *Journal of Vision*, *8*(2), 1–8. doi: 10.1167/8.2.16

Buck, S. L. (2015). Brown. *Current Biology*, *25*(13), R536–R537. doi: 10.1016/j.cub.2015.05.029

Buck, S. L., & DeLawyer, T. (2012, dec). A new comparison of brown and yellow. *Journal of Vision*, *12*(14), 9. doi: 10.1167/12.14.9

Buck, S. L., & Delawyer, T. (2014). Dark vs. bright equilibrium hues: rod and cone biases. *Journal of the Optical Society of America A*, *31*(4), 1–15. doi: 10.1364/JOSAA.31.000A75

Buck, S. L., Shelton, A., Stoehr, B., Hadyanto, V., Tang, M., Morimoto, T., & DeLawyer, T. (2016). Influence of surround proximity on induction of brown and darkness. *Journal of the Optical Society of America A*, *33*(3), A12–21. doi: 10.1364/JOSAA.33.000A12

Cicerone, C. M., Volbrecht, V. J., Donnelly, S. K., & Werner, J. S. (1986). Perception of blackness. *Journal of the Optical Society of America. A, Optics and image science*, *3*(4), 432–436. doi: 10.1364/JOSAA.3.000432

Davidson, M. (1968). Perturbation approach to spatial brightness interaction in human vision. *J Opt Soc Am*, *58*(9), 1300–1308.

De Valois, R. L., & De Valois, K. K. (1993). A Multi-Stage Color Model. *Vision Research*, *33*(8), 1053–1065.

De Valois, R. L., & De Valois, K. K. (1996). On "a three-stage color model". *Vision Research*, *36*(6), 833–836. doi: 10.1016/0042-6989(96)84513-5

DeLawyer, T., Foote, K., Kwong, C., Lin, T., Short, W., Suh, E., & Buck, S. L. (2012, dec). The effects of luminance surrounds on the perception of the color brown. *Journal of Vision*, *12*(14), 36. doi: 10.1167/12.14.36

DeLawyer, T., Frederick, A., Kaplan, S., Lin, T., Shonka, T., & Buck, S. L. (2013). Dependence of the color brown on the spatial configuration of high luminance surrounds. In *presented at association for research in vision and ophthalmology.* Seattle, Washington.

DeLawyer, T., Morimoto, T., & Buck, S. L. (2016). Dichoptic perception of brown. *Journal of the Optical Society of America A*, *33*(3), A123–128. doi: 10.1364/JOSAA.33.00A123

D'Zmura, M., & Singer, B. (1999). Contrast gain control. In L. T. Sharpe & K. R. Gegenfurtner (Eds.), *Color vision: From genes to perception* (pp. 369–385). Cambridge, UK: Cambridge University Press.

Fuld, K., Werner, J. S., & Wooten, B. R. (1983). The possible elemental nature of brown. *Vision Research*, *23*(6), 631–637.

Fuld, K., Wooten, B. R., & Whalen, J. J. (1981). The elemental hues of short-wave and extraspectral lights. *Perception & Psychophysics*, *29*(4), 317–322. doi: 10.3758/BF03207340

Gilchrist, A. (2015). Theoretical approaches to lightness and perception. *Perception*, *44*(4), 339–358. doi: 10.1068/p7935

Gilchrist, A. L. (2008). Perceptual organization in lightness. In J. Wagemans (Ed.), *Oxford handbook of perceptual organization* (pp. 1–25). Oxford: Oxford University Press. doi: 10.1093/oxfordhb/9780199686858.013.031

Gilchrist, A. L. (2015). Response to Maniatis' Is a unified model of contrast and constancy possible? Reply to Gilchrist. *Vision Research*, *108*(January), 117. doi: 10.1016/j.visres.2014.12.016

Gilchrist, A. L., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., ... Economou, E. (1999). An anchoring theory of lightness perception. *Psychological review*, *106*(4), 795–834. doi: 10.1037/0033-295X.106.4.795

Graham, N. (1981). *Psychophysics of spatial-frequency channels.*

Helson, H. (1963). Studies of anomalous contrast and assimilation. *Journal of the Optical Society of America*, *53*(1), 179. doi: 10.1364/JOSA.53.000179

Helson, H., & Rohles, F. H. J. (1959). A Quantitative Study of Reversal of Classical Lightness-Contrast. *The American Journal of Psychology*, *72*(4), 530–538.

Hering, E. (1878). *Zur Lehre vom Lichtsinne*.

Hong, S. W., & Shevell, S. K. (2004a, jan). Brightness contrast and assimilation from patterned inducing backgrounds. *Vision Research*, *44*(1), 35–43. doi: 10.1016/j.visres.2003.07.010

Hong, S. W., & Shevell, S. K. (2004b). Brightness induction: unequal spatial integration with increments and decrements. *Visual Neuroscience*, *21*(3), 353–7.

Howe, P. D. L. (2005). White's effect: removing the junctions but preserving the strength of the illusion. *Perception*, *34*, 557–564. doi: 10.1068/p5414

Hurvich, L. M., & Jameson, D. (1955). Some quantitative aspects of an opponent-colors theory. II. Brightness, saturation, and hue in normal and dichromatic vision. *Journal of the Optical Society of America*, *45*(8), 602–616. doi: 10.1364/JOSA.45.000602

Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological review*, *64, Part 1*(6), 384–404. doi: 10.1037/h0041403

Jameson, D., & Hurvich, L. M. (1955). Some Quantitative Aspects of an Opponent-Colors Theory I Chromatic Responses and Spectral Saturation. *Journal of the Optical Society of America*, *45*(7), 546. doi: 10.1364/JOSA.45.000546

Jameson, D., & Hurvich, L. M. (1975). From Contrast to Assimilation: In Art and in the Eye. *Leonardo*, *8*(2), 125–131.

Kaiser, P. K. (1971). Luminance and brightness. *Applied optics*, *10*(12), 2768–2770. doi: 10.1364/AO.10.002768

Kaneko, S., & Murakami, I. (2012). Flashed stimulation produces strong simultaneous brightness and color contrast. *Journal of vision*, *12*(12), 1–18. doi: 10.1167/12.12.1

Kiesow, F. (1930). Uber die Entstehung der Braunempfindung. *Neue Psychologische Studien*, *6*, 119–30.

Kingdom, F. A. A. (2011). Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research*, *51*(7), 652–673. doi: 10.1016/j.visres.2010.09.012

Kulikowski, J. J. (1976). Effective contrast constancy and linearity of contrast sensation. *Vision Research*, *16*(12), 1419–1431. doi: 10.1016/0042-6989(76)90161-9

Mathworks, T. (2016). *MATLAB.*

Morimoto, T., Slezak, E., & Buck, S. L. (2016). No effects of surround complexity on brown induction. *Journal of the Optical Society of America A*, *33*(3), A45–52. doi: 10.1364/JOSAA.33.000A45

Moulden, B., & Kingdom, F. (1989). White's effect: A dual mechanism. *Vision Research*, *29*(9), 1245–1259. doi: 10.1016/0042-6989(89)90071-0

Moulden, B., & Kingdom, F. A. A. (1991). The local border mechanism in grating induction. *Vision Res*, *31*(11), 1999–2008.

Munker, H. (1970). *Abbildung auf der Netzhaut und übertragungstheoretische Beschreibung der Farbwahrnehmung* (Unpublished doctoral dissertation). Ludwig-Maximilians-Universität, München.

Naka, K. I., & Rushton, W. A. H. (1966). S-potentials from colour units in the retina of fish (Cyprinidae). *Journal of Physiology1*, *185*, 536–555.

Padgham, C. A., & Saunders, J. E. (1975). *The Perception of Light and Colour.* Academic press.

Quinn, P. C., Rosano, J. L., & Wooten, B. R. (1988, feb). Evidence that brown is not an elemental color. *Perception & Psychophysics*, *43*(2), 156–64.

Robinson, A. E., & de Sa, V. R. (2008). Brief presentations reveal the temporal dynamics of brightness induction and White's illusion. *Vision Research*, *48*(22), 2370–2381. doi: 10.1016/j.visres.2008.07.023

Robinson, A. E., Hammon, P. S., & de Sa, V. R. (2007). Explaining brightness illusions using spatial filtering and local response normalization. *Vision Research*, *47*(12), 1631–1644. doi: 10.1016/j.visres.2007.02.017

Rudd, M. E. (2013). Edge integration in achromatic color perception and the lightness darkness asymmetry. , *13*, 1–30. doi: 10.1167/13.14.18.doi

Rudd, M. E., & Arrington, K. F. (2001, dec). Darkness filling-in: a neural model of darkness induction. *Vision Research*, *41*(27), 3649–62.

Shevell, S. K. (1989). On neural signals that mediate induced blackness. *Vision Research*, *29*(7), 891–900.

Shevell, S. K., Holliday, I., & Whittle, P. (1992, dec). Two separate neural mechanisms of brightness induction. *Vision Research*, *32*(12), 2331–40.

Shinomori, K., Schefrin, B. E., & Werner, J. S. (1997). Spectral mechanisms of spatially induced blackness: data and quantitative model. *Journal of the Optical Society of America. A, Optics, image science, and vision*, *14*(2), 372–87. doi: 10.1364/JOSAA.14.000372

Singer, B., & D'Zmura, M. (1994). Color Contrast. *Vision Research*, *34*(23), 3111–3126.

Spehar, B., Clifford, C. W. G., & Agostini, T. (2002). Induction in variants of White's effect: Common or separate mechanisms? *Perception*, *31*(2), 189–196. doi: 10.1068/p10sp

Spehar, B., Gilchrist, A. L., & Arend, L. (1995). The critical role of relative luminance relations in White's effect and grating induction. *Vision Research*, *35*(18), 2603–2614. doi: 10.1016/0042-6989(95)00005-K

Todorović, D. (1997). Lightness and junctions. *Perception*, *26*, 379–394.

Todorović, D. (2010). Context effects in visual perception and their explanations. *Review of Psychology*, *17*(1), 17–32.

Uchikawa, H., Uchikawa, K., & Boynton, R. M. (1989). Influence of achromatic surrounds on categorical perception of surface colors. *Vision Research*, *29*(7), 881–890.

Vincent, J., Buck, S. L., Armer, J., Delawyer, T., & Wilson, L. (2012). Remote rings bias equilibrium brown and yellow. In *Icvs biannual meeting.* Winchester, UK.

Vincent, J., Kale, A. M., & Buck, S. L. (2016). Luminance-dependent long-term chromatic adaptation. *Journal of the Optical Society of America A*, *33*(3), A164–169. doi: 10.1364/JOSAA.33.00A164

Walker, J. T. (1978). Brightness enhancement and the Talbot level in stationary gratings. *Perception and Psychophysics*, *23*(4), 356–359. doi: 10.3758/BF03199722

Wallach, H. (1948). Brightness constancy and the nature of achromatic colors. *Journal of Experimental Psychology*, *38*(3), 310–324. doi: 10.1037/h0053804

Wesner, M. F., & Shevell, S. K. (1992). Color Perception Within a Chromatic Context: Changes in Red / Green Equilibria Caused by Noncontiguous light. *Vision Research*, *32*(9), 1623–1634.

White, M. (1979). A new effect of pattern on perceived lightness. *Perception*, *8*, 413–416.

White, M. (1981). The effect of the nature of the surround on the perceived lightness of grey bars within square-wave test gratings. *Perception*, *10*, 215–230.

White, M. (2010). The Early History of White's Illusion. *Colour: Design & Creativity*, *5*(7), 1–7.

White, M., & White, T. (1985). Counterphase lightness induction. *Vision research*, *25*(9), 1331–1335.

Yazdanbakhsh, A., Arabzadeh, E., Babadi, B., & Fazl, A. (2002). Munker-White-like illusions without T-junctions. *Perception*, *31*(6), 711–715. doi: 10.1068/p3348

Yeonan-Kim, J., & Bertalmío, M. (2016). Retinal lateral inhibition provides the biological basis of long-range spatial induction. *PLoS ONE*, *11*(12), 1–23. doi: 10.1371/journal.pone.0168963

Zaidi, Q., Spehar, B., & Shy, M. (1997). Induced effects of backgrounds and foregrounds in three-dimensional configurations: the role of T-junctions. *Perception*, *26*, 395–408.

Appendix A

NUMBER OF TRIALS COMPLETED

Table A.1: **Number of trials completed**, by each participant for each stimulus and task.

	Classic illusion		**Checkerboard variant**		**Radial variant**	
Participant	**Achromatic**	**Brown**	**Achromatic**	**Brown**	**Achromatic**	**Brown**
AGP	400	231	0	0	0	0
AMC	320	79	0	0	120	87
APF	0	0	400	256	0	0
AYU	500	256	0	0	0	0
DDS	0	0	400	256	0	0
GAL	0	0	0	0	400	160
JCW	0	0	400	244	400	256
KXW	323	80	0	0	0	0
MAW	0	0	400	256	400	128
MJK	400	256	0	0	0	0
MMT	0	0	400	256	400	128
MRM	0	0	0	0	160	64
MSS	0	0	400	256	550	193
NDJ	374	256	0	0	0	0
QXY	328	80	0	0	0	0
RAM	400	256	0	0	0	0
RXN	320	80	200	80	560	224
SCR	0	0	0	20	400	160
SMK	0	0	400	217	100	0
SRS	0	0	0	0	300	160
SXF	320	80	0	0	0	0
TCW	0	0	400	256	0	0
VXL	400	224	0	0	0	0
XTL	0	0	400	208	200	160
YPW	400	80	300	256	320	128

Appendix B

PSYCHOMETRIC FUNCTION FITTING OF ACHROMATIC BRIGHTNESS MATCHING DATA

B.1 Linear model

The simplest family of models that could fit the achromatic brightness matching data, would be linear models. In such a model, the left target luminance at match would be a linear transformation of the right target luminance at match. One such linear transformation would be an additive shift: the left target luminance at match would be a fixed (per participant, per stimulus) number of candelas higher than the right target at match, i.e., equation B.1. An additive model would be applicable if the surround context influenced the target brightnesses by a fixed amount, regardless of the luminances of the targets.

$$Lum_{left} = Lum_{right} + b \tag{B.1}$$

Another linear transformation that could be applied, is a scaling: the left target luminance at match would be scaled from the right target at match, i.e., equation B.2. B.2. A scaling model would be applicable if the surround context influenced the target brightnesses dependent on the luminance of the targets, but keeping their relative brightness constant.

$$Lum_{left} = a * Lum_{right} \tag{B.2}$$

These two transformations could also be combined, such that the left target luminance at match is both scaled and shifted compared to the right target luminance at match. i.e., equation B.3.

$$Lum_{left} = a * Lum_{right} + b \tag{B.3}$$

Figure B.1 shows the results of fitting each of these three models to the achromatic brightness matching data of participant AGP on the classic illusion. The scatterplots show the bright-

ness matching data, identical to Figure 4.1. The black line in each plot indicates the model prediction. Models were fit by finding the parameter values for a best-fit in the least-squares sense; the bar graph in Figure B.1D shows the goodness-of-fit of each model expressed as R^2. The variability inherently present in the data limit the performance of any model, so as a comparison, the split-half reliability of the data is presented in B.1D. The achromatic brightness matching data was randomly split into a "training" and "testing" half, and the left target luminances of each trial in the testing half was predicted as the mean left target luminance of all trials in the training half with the same right target luminance.

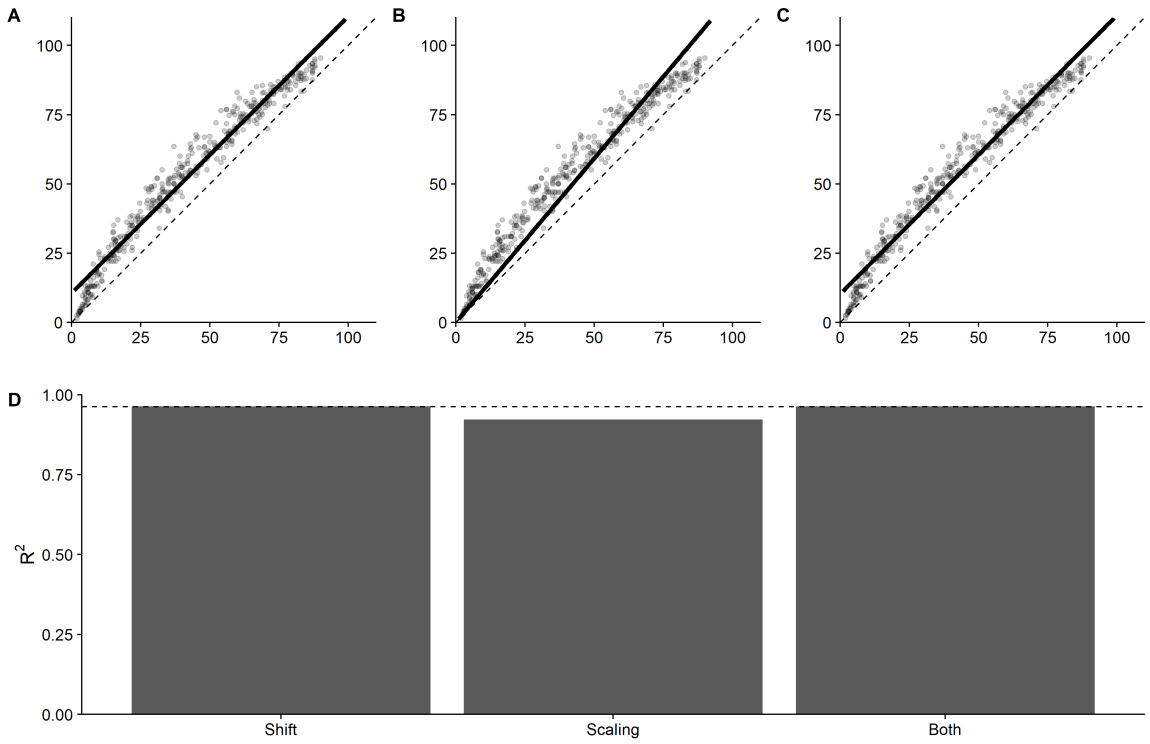

Figure B.1: **Linear model fits** to the example data from participant AGP on the Classic illusion. **A** Left target luminance is a linear translation (shift by additive constant) of the right target luminance. **B** Left target luminance is a linear scaling of the right target luminance. **C** Left target luminance is a linear transformation of right target luminance that consists of both a translation and a scaling. **D** Goodness-of-fit, expressed as R^2, of the three models to the exemplar data, compared to the split-half reliability (dotted line).

For this exemplar, an additive shift of 10.50 cd/m^2 accurately captures the data (R^2 = .96), as accurate as the split-half reliability (R^2 = .96). The best-fitting scaling of 1.18 does notably worse, R^2 = .92. The combined linear model performs as well as the additive shift (R^2 = .96). The best-fitting parameters of the combined linear model show that the model relies mainly on the additive shift (10.21 cd/m^2 in the combined model), and barely on the scalar (1.01 in the combined model). Since the combined model contains an extra parameter, compared to the other two models, a better fit for the combined model could be spurious; this model can always fit the data better than either of the two restricted models. An F-test can be used test if the additional parameter has significant explanatory value, based on the residual sum of squares (RSS) and number of parameters (p) of the unrestricted (2) and restricted model (1), i.e., equation B.4.

$$F = \frac{\left(\frac{RSS_1 - RSS_2}{p_2 - p_1}\right)}{\left(\frac{RSS_2}{n - p_2}\right)}, withdf = (p_2 - p_1, n - p_2) \tag{B.4}$$

$$F = \frac{RSS_1 - RSS_{Combined}}{\left(\frac{RSS_{Combined}}{n-2}\right)}, withdf = (1, n - 2) \tag{B.5}$$

For the linear models, this reduces to equation B.5. In the case of the exemplar data, the combined model is a significant improvement over the simple scaling (F(1,398) = 442.76, p <.0001), but not over the additive shift (F(1,398) = 442.76, p = .49). Together, this suggests that the exemplar data can be accurately described as linear additive shift in the luminances of the two targets at match, without the need for an additional scaling parameter.

Table B.1 reports the best-fitting parameters, in the least-squares sense, and corresponding goodness-of-fit, expressed as R^2 for each model and participant on the Classic illusion. The same pattern for the exemplar data can be seen for most participants: an additive shift and combined model describe the data accurately (mean R^2 = .86, and .88, respectively), and better than a scalar (mean R^2 = .81). The combined model relies mainly on an additive shift (mean = 12.24 cd/m^2) comparable to that of the best-fitting additive shift in isolation (mean = 10.25 cd/m^2), and relies very little on the scaling term (mean = 0.95). For most participants, however, the scalar term does have explanatory power, as the combined

Table B.1: **Linear model fits on the Classic illusion** for all participants. Scalar and shift terms in the combined models that provide a significant improvment are bolded.

	Reliability	Shift		Scaling		Both		
Participant	R^2	Shift	R^2	Scalar	R^2	Scalar	Shift	R^2
AGP	.96	10.50	.96	1.18	.92	1.01	10.21	.96
AMC	.95	12.43	.95	1.21	.91	1.02	11.57	.95
AYU	.04	-3.76	.24	0.84	.33	**0.61**	**14.04**	**.40**
KXW	.91	9.76	.93	1.18	.91	**1.03**	**8.48**	**.94**
MJK	.89	10.50	.92	1.15	.84	**0.92**	**13.92**	**.93**
NDJ	.82	11.33	.87	1.19	.80	**0.96**	**13.01**	**.87**
QXY	.94	19.51	.93	1.34	.82	**1.04**	**17.95**	**.93**
RAM	.60	13.70	.73	1.21	.58	**0.87**	**18.85**	**.74**
RXN	.94	7.36	.97	1.13	.96	**1.04**	**5.60**	**.97**
SXF	.94	11.32	.95	1.19	.91	1.01	10.69	**.95**
VXL	.95	8.10	.96	1.11	.91	**0.92**	**11.56**	**.97**
YPW	.93	12.23	.94	1.21	.89	**1.03**	**10.93**	**.94**
Mean	**.81**	**10.25**	**.86**	**1.16**	**.81**	**0.95**	**12.24**	**.88**

model predicts the matches significantly better than the translation term alone (all combined models outperform the scalar term alone).

Table B.2 reports the best-fitting parameters, in the least-squares sense, and corresponding goodness-of-fit, expressed as R^2 for each model and participant on the Checkerboard variant. The Checkerboard variant shows a subtly different pattern of model fits as the Classic illusion: an additive shift and combined model describe the data accurately (mean $R^2 = .69$, and .75, respectively), and better than a scalar (mean $R^2 = .68$). The combined model does rely on an additive term (mean = $8.36\,\mathrm{cd/m^2}$) for a significant improvement over the scaling, for all but one participant (DDS). However, this additive term is generally very different in direction and magnitude from the additive shift in the restricted model (mean = $-0.72\,\mathrm{cd/m^2}$). Instead, for all participants, the combined model relies also on the scaling

Table B.2: **Linear model fits on the Checkerboard variant** for all participants. Scalar and shift terms in the combined models that provide a significant improvment are bolded.

	Reliability	Shift		Scaling		Both		
Participant	R^2	Shift	R^2	Scalar	R^2	Scalar	Shift	R^2
APF	0.53	-2.05	0.12	0.85	0.22	**0.54**	**18.83**	**0.38**
DDS	0.87	-7.26	0.75	0.85	0.78	**0.84**	0.55	**0.78**
JCW	0.87	-17.03	0.76	0.71	0.79	**0.81**	**-6.68**	**0.80**
MAW	0.88	-6.53	0.77	0.85	0.82	**0.80**	**3.33**	**0.82**
MMT	0.91	4.93	0.85	1.04	0.82	**0.86**	**10.96**	**0.87**
MSS	0.80	7.55	0.66	1.06	0.57	**0.75**	**18.49**	**0.74**
RXN	0.92	-0.92	0.91	0.97	0.91	**0.94**	**2.07**	**0.91**
SMK	0.77	6.54	0.57	1.03	0.50	**0.73**	**17.99**	**0.67**
TCW	0.91	1.44	0.82	1.00	0.82	**0.89**	**6.46**	**0.83**
XTL	0.73	-2.63	0.68	0.93	0.69	**0.88**	**2.72**	**0.69**
YPW	0.79	8.03	0.71	1.09	0.62	**0.79**	**17.25**	**0.76**
Mean	**0.81**	**-0.72**	**0.69**	**0.94**	**0.68**	**0.80**	**8.36**	**0.75**

term (mean = 0.80) for a significant improvement over the additive shift model. Thus, the brightness matches on the checkerboard variant can be best explained by a linear model containing both an additive term and a scaling.

Table B.3 reports the best-fitting parameters, in the least-squares sense, and corresponding goodness-of-fit, expressed as R^2 for each model and participant on the Radial variant. The Radial variant shows a pattern that corresponds to that for that Classic illusion in broad strokes: an additive shift and combined model describe the data accurately (mean R^2 = .69, and .74, respectively), and better than a scalar (mean R^2 = .53). The combined model does rely on an additive term (mean = 17.78 $\mathrm{cd/m^2}$) for a significant improvement over the scaling, for all participants. Like on the Classic illusion, and in contrast to the Checkerboard variant, this additive term is generally in the same direction as the additive shift in the restricted model (mean = 10.03 $\mathrm{cd/m^2}$). In contrast to the Classic illusion, and

Table B.3: **Linear model fits on the Radial variant** for all participants. Scalar and shift terms in the combined models that provide a significant improvement are bolded.

	Reliability	Shift		Scaling		Both		
Participant	R^2	Shift	R^2	Scalar	R^2	Scalar	Shift	R^2
AMC	0.46	17.12	0.39	1.21	-0.02	**0.65**	**30.54**	**0.54**
GAL	0.92	18.13	0.76	1.31	0.50	**0.87**	**23.18**	**0.77**
JCW	0.89	-9.44	0.86	0.82	0.87	**0.87**	**-3.05**	**0.88**
MAW	0.92	0.40	0.86	0.97	0.86	**0.86**	**7.13**	**0.88**
MMT	0.92	11.12	0.83	1.15	0.72	**0.87**	**16.67**	**0.85**
MRM	0.93	-2.66	0.82	0.90	0.87	**0.78**	**8.11**	**0.90**
MSS	0.67	17.87	0.40	1.25	0.03	**0.69**	**29.06**	**0.49**
RXN	0.84	2.19	0.68	0.96	0.68	**0.73**	**14.44**	**0.80**
SCR	0.92	17.06	0.77	1.26	0.48	**0.83**	**23.21**	**0.80**
SRS	0.93	28.03	0.78	1.57	0.47	1.04	**26.61**	**0.78**
XTL	0.66	5.60	0.49	1.05	0.44	**0.76**	**16.31**	**0.54**
YPW	0.79	14.96	0.65	1.24	0.44	**0.83**	**21.17**	**0.68**
Mean	**0.82**	**10.03**	**0.69**	**1.14**	**0.53**	**0.81**	**17.78**	**0.74**

in line with Checkerboard variant, for all but one participant, the combined model relies also on the scaling term (mean = 0.81) for a significant improvement over the additive shift model. Thus, the brightness matches on the Radial variant can be best explained by a linear model containing both an additive term and a scaling.

B.2 Luminance Response Function (LRF)

While the linear models can capture the overall trend of the achromatic brightness matching data in terms of direction and magnitude of induction, Figure B.1A&C reveal a shortcoming of these models. Their linear nature requires the predicted luminance to fall along a straight line; the data, however, is better described as a curve. At extreme luminances, nearing the lowest or the highest luminance in the display, the induction effect is weaker than at

intermediate luminances. The linear models seem to capture the induction effect at these intermediate levels, but as a result overestimate the effect at the extreme luminances. This is not a trivial limitation of the linear models: while indeed most of the luminances tested fall in the intermediate range that these models can describe well, the fact that the induction effect seems to be bounded by the luminances of the inducing context is certainly of theoretical relevance to explanations of White's illusion.

To overcome this limitation of linear models, and attempt to capture the full curve of the brightness matching data, a nonlinear model must be applied. To construct a model of achromatic brightness matching, it is worthwhile to first consider the relationship between physical luminance and perceived brightness for a single stimulus. Naka and Rushton (1966) first characterized the responses of retinal cells to light contrast using the equation that has come to bear their name since, equation B.6. This sigmoidal function can be used to describe the relationship between some contrast C, and a neural response R (usually in terms of spike rate).

$$R(C) = R_{max}\frac{C^p}{C^p + \sigma^p} + b \tag{B.6}$$

The R_{max} parameter is a scaling factor that determines the maximal response: as the contrast increases, the R asymptotically approaches R_{max} and the response thus saturates. The parameter σ determines the contrast level at which R is at half its asymptotic maximum, and is thus referred to as the *semi-saturation constant*. While the Naka-Rushton equation is most often applied as a contrast response function, charactarizing a neurons or systems response R to contrast C, it can serve a purpose in the current investigation as the basis for a Luminance Response Function (LRF). Equation B.7 describes the relationship between the luminance of a stimulus, and its perceived brightness with the same shape as the Naka-Rushton equation.

$$Brightness(Lum) = \alpha\frac{Lum^{p+q}}{Lum^p + \sigma^p} + b \tag{B.7}$$

One addition to the LRF in B.7 is the parameter q, which prevents the LRF from saturating. There is no reason to assume that perceived brightness would saturate over the luminance

ranges tested in the current investigation. Another aspect of the LRF to note is that perceived brightness here is in arbitrary units, since there are no physical units to measure perceived brightness in (without comparison to another stimulus).

B.2.1 Matching LRFs

A single LRF describes the relationship between physical luminance and the perceived brightness of a single stimulus. By combining two LRFs, one can describe the relationship between the two targets. Defining an LRF for the right target allow for a prediction to be made about the perceived brightness of the right target at any physical luminance, e.g., $50\,\mathrm{cd/m^2}$ (Figure B.2, red arrow). This predicted brightness can then be applied to the LRF for the left target (Figure B.2, purple arrow). That LRF allows then gives a physical luminance value for the left target, e.g., $64.2\,\mathrm{cd/m^2}$ (Figure B.2, blue arrow), at which the LRF for the left target would predict that same perceived brightness. Thus, the combination of the two LRFs allows one to predict the a luminance of the left target, that would predict the same perceived brightness as a known luminance of the right target. Note that this calculation has the same logic as the psychophysical achromatic brightness matching paradigm: one target is fixed at a known luminance, and it generates some perceived brightness. This brightness is hard to quantify perceptually; in the model, it is in arbitrary units. Yet, adjusting another target until it has the same brightness, perceptually in the psychophysical task, or in the same arbitrary units in the model, leads to a clearly defined physical luminance of that other target that has the same brightness.

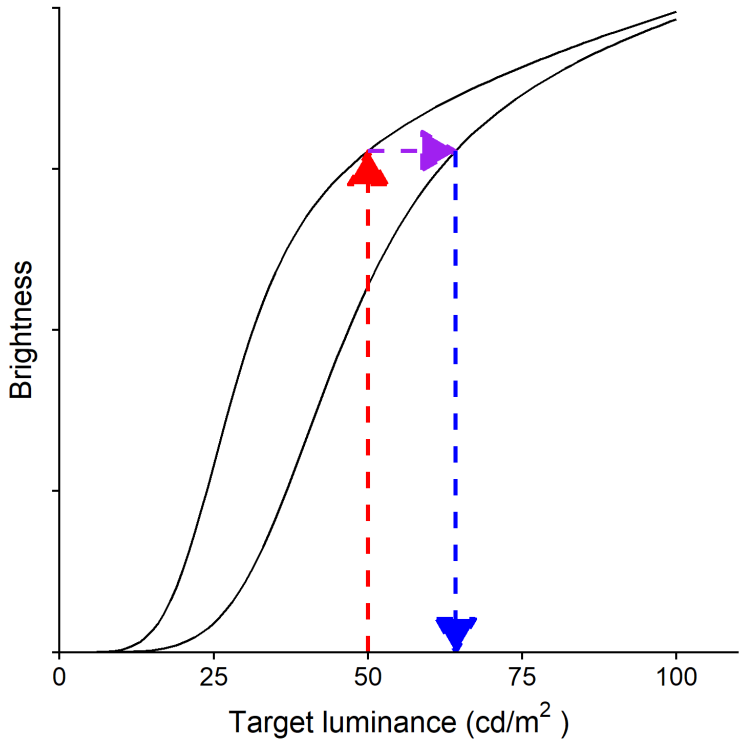

Figure B.2: **Predicting achromatic brightness matches from LRFs** After defining an LRF for the right target, the brightness response to any physical luminance can be predicted (red arrow). Finding an equal brightness response on the LRF for the left target (purple arrow), leads to a physical luminance of the left target (blue arrow) that should be produce the same perceived brightness.

B.2.2 Fitting double LRF models

A combination of two LRFs, then, can predict achromatic brightness matches. To fit this model, in the least-squares sense, requires feeding in a set of physical luminance of the right target, and adjusting the parameters until the error between the predicted and actual left target luminance is minimized. The double LRF model provides several parameters that can be adjusted to try and achieve a best fit: the translation parameter b, the scaling parameter α, the semi-saturation constant σ are the most straightforward candidates, and will be the focus of the current investigation.

Figure B.3 shows the LRF models and corresponding fits for changes in the three aforementioned parameters. The first model has two identical LRFs for the right and left targets, but the LRF for the left target has been linearly translated through the b term in equation B.7 (Figure B.3A). On the exemplar (participant AGP on the Classic illusion), a translation of -.15 (arb. units) fits the data best in the least-squares sense, $R^2 = .89$. However, the scatterplot in Figure B.3D, reveals that this model is far from adequate in capturing the data. The linear translation creates an asymptote at the lower luminance extreme, so that the left target is never predicted to be below $20\,\mathrm{cd/m^2}$. This suggest that the effect of the surround context is not to make the left target brighter by a fixed amount. The second model has the same right target LRF as the first model, but the left target LRF this time is a scaling of the right LRF, by change of the a term in equation B.7 (Figure B.3B). On the exemplar, a translation of .95 fits the data best in the least-squares sense, $R^2 = .87$. This is worse than the linear translation of the LRF, and the scatterplot in Figure B.3E, reveals that this model too fails to capture the shape of the brightness matching curve. The seems to have a negligible effect below $45\,\mathrm{cd/m^2}$, where it would predict that the two targets have to be nearly the same physical luminance to match in perceived brightness. This suggest that the effect of the surround context is also not to make the left target appear brighter by a fixed ratio. The third and last model under investigation here has the same right target LRF as the first and second model. The left target LRF differs from this right target LRF by its semi-saturation constant σ in equation B.7 (Figure B.3C). Thus, compared to the right target, the left target has a shallower slope and reaching half-max brightness later. Since the model is fairly sensitive to this semi-saturation constant, the σ for the right LRF has also been fit to the data simultaneous with the left LRF σ. To provide a fairer comparison with the other two models (which only fit 1 parameter), the same LRF has been used for the right target for all three models. This effectively reduces all models to a 1 parameter model, with the approximately optimal parameter values for the right target LRF. Across all three models, the parameters p and q in equation B.7 have been kept constant, at 5 and .3, respectively. On the exemplar, the semi-saturation constant for the right target LRF was

fit at $\sigma = 26.14$, which was the parameter value then used for all three models. A left target LRF semi-saturation constant of $\sigma = 42.04$ fits the data best in the least-squares sense, $R^2 = .98$. This is better than both the translation and the scaling of the LRF. Figure B.3F also reveals that this model accurately capture the full shape of the brightness matching function; it curves close to the unity line at the more extreme luminance, and bends farther away at the intermediate luminances. Figure B.3G presents the goodness-of-fit, expressed as R^2 for all three models. This suggests that the effect of the context is to change the shape of the nonlinear function relating physical luminance to perceived brightness of each target.

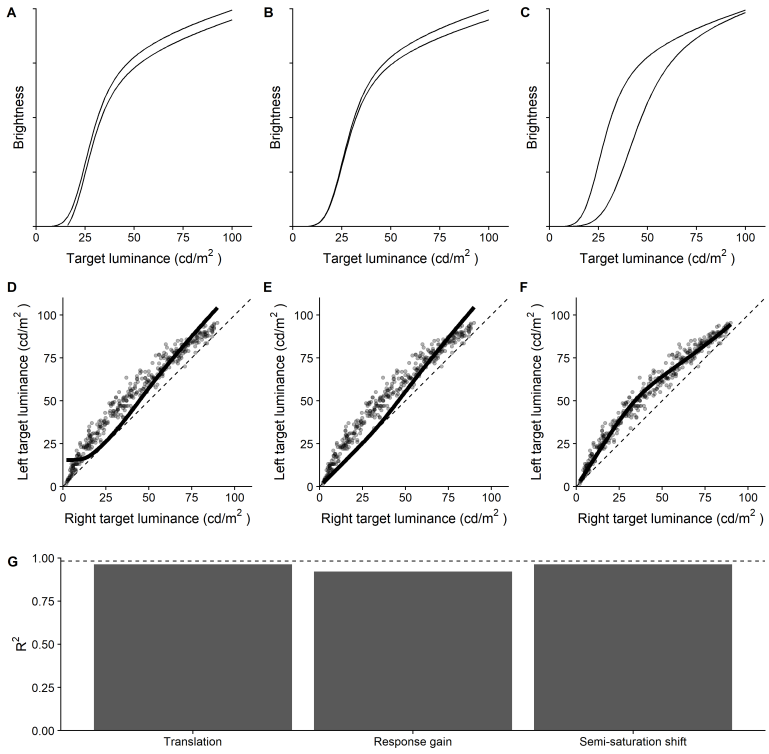

Figure B.3: **Luminance response function based model**, and their fits, for example data from participant AGP on the Classic illusion. **A&D** Two LRFs that only differ in the horizontal position, i.e., linear translation, and the fit of this model to the exemplar data. **B&E** Two LRFs that only differ by their output scaling, i.e., linear response gain, and the fit of this model to the exemplar data. **C&F** Two LRFs that differ in their semi-saturation constant, i.e., a difference in the shape of the nonlinearity, and the fit of this model to the exemplar data. **G** Goodness-of-fit, expressed as R^2, of the three models to the exemplar data, compared to the split-half reliability (dotted line).

Table B.4: **LRF model fits on the Classic illusion** for all participants.

	Reliability	Translation		Scaling		Semi-saturation		
Participant	R^2	Shift	R^2	Scalar	R^2	SS Left	SS Right	R^2
AGP	0.98	-0.15	0.89	0.95	0.87	42.04	26.14	0.98
AMC	0.98	-0.18	0.88	0.94	0.85	46.22	27.63	0.97
AYU	0.34	0.00	0.28	2.48	0.33	471.21	564.90	0.33
KXW	0.95	-0.15	0.89	0.95	0.87	46.31	31.70	0.93
MJK	0.95	-0.12	0.85	0.96	0.82	35.04	15.47	0.94
NDJ	0.95	-0.16	0.78	0.96	0.73	41.84	24.94	0.90
QXY	0.96	-0.27	0.79	0.92	0.72	51.09	22.28	0.94
RAM	0.79	-0.17	0.60	0.95	0.53	38.40	16.50	0.75
RXN	0.98	-0.15	0.93	0.95	0.93	49.81	40.85	0.97
SXF	0.97	-0.16	0.90	0.95	0.87	44.08	26.54	0.96
VXL	0.98	-0.09	0.91	0.97	0.89	31.24	16.61	0.97
YPW	0.97	-0.17	0.89	0.95	0.86	46.85	24.81	0.95
Mean	**0.90**	**-0.15**	**0.80**	**1.08**	**0.77**	**78.68**	**69.87**	**0.88**

Table B.4 reports the best-fitting parameters, in the least-squares sense, and corresponding goodness-of-fit, expressed as R^2 for each LRF model and participant on the Classic illusion. For all participants, the model with a change in semi-saturation constants between the two models vastly outperforms the other two LRF models. Since these are not nested models, unlike the linear models, there less to be gleaned from comparing the various parameters.

Rather, Table B.5 shows the goodness-of-fit, expressed as R^2, for each linear model and each LRF model, for each participant on the Classic illusion. With the exception of participants AYU, KXW and MJK, most achromatic brightness matching data is better explained by the LRF model based on shifted semi-saturation constants, than by any linear model.

This can also be seen in Figure B.4, which plots the mean goodness-of-fit, expressed as R^2

Table B.5: **All model fits on the Classic illusion** for all participants, in R^2 sense.

		Linear			LRF		
Participant	Reliability	Shift (+)	Scalar (x)	Both (+x)	Shift (LRF+)	Scalar (LRFx)	semi-saturation (LRFss)
AGP	0.98	0.96	0.92	0.96	0.89	0.87	0.98
AMC	0.98	0.95	0.91	0.95	0.88	0.85	0.97
AYU	0.34	0.24	0.33	0.40	0.28	0.33	0.33
KXW	0.95	0.93	0.91	0.94	0.89	0.87	0.93
MJK	0.95	0.92	0.84	0.93	0.85	0.82	0.94
NDJ	0.95	0.87	0.80	0.87	0.78	0.73	0.90
QXY	0.96	0.93	0.82	0.93	0.79	0.72	0.94
RAM	0.79	0.73	0.58	0.74	0.60	0.53	0.75
RXN	0.98	0.97	0.96	0.97	0.93	0.93	0.97
SXF	0.97	0.95	0.91	0.95	0.90	0.87	0.96
VXL	0.98	0.96	0.91	0.97	0.91	0.89	0.97
YPW	0.97	0.94	0.89	0.94	0.89	0.86	0.95
Mean	**0.90**	**0.86**	**0.81**	**0.88**	**0.80**	**0.77**	**0.88**

for each model, for the three stimuli. The dotted lines here indicate the ±1 standard error of the mean reliability across observers. Notable is that all models on average perform better on the Classic illusion than on either the Checkerboard or the Radial variant. Secondly, the LRF-based models with a change in semi-saturation constant (LRFss) outperforms the other models on the Classic illusion and the Radial variant. On the Checkerboard variant, it seems to be slightly inferior to the combined linear model (+x).

B.3 conclusion

The achromatic brightness matching data generally follows a distinct pattern, across participants and stimuli. This suggests that this pattern could be captured mathematically, and the brightness matches thus predicted. Simple linear models are initially successful at such

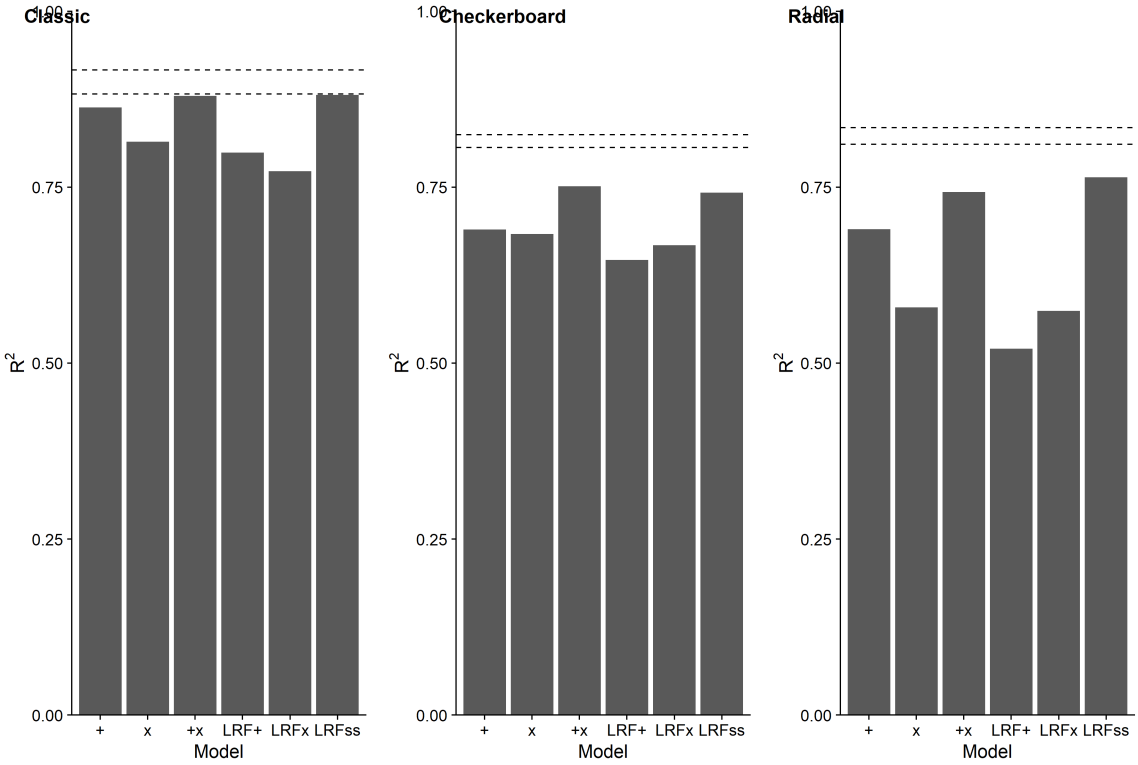

Figure B.4: **Comparison of models**, both linear and LRF based models, for all three stimuli. The linear models consist of translation by an additive term (+), scaling (x), and both translation and scalig (+x). The LRF-based models consist of translation by an additive term (LRF+), scaling (LRFx), and different semi-saturation constants for the two LRFs (LRFss). The dotted lines here indicate the ±1 standard error of the mean reliability across observers.

prediction. These linear models seem to mainly capture the approximately linear relationship between right target luminance and left target luminance at intermediate lightlevels. While these models work well across the three different stimuli, different linear transformations appeart to be important for the different stimuli; on the Classic illusion, a simple additive term translating the luminance is enough to capture most of the data. On the Checkerboard and ring variants, a scalar term provides additional predictive power. At more extreme luminance, the assumption of linear breaks down, and the linear models clearly fail to capture the induction effect at these lightlevels To overcome the limitations of the linear models, the

data can be fit with a family of models derived from Luminance Response Functions. Based on the Naka-Rushton equation, these sigmoidal functions describe the relationship between a target's physical luminance and its perceived brightness. By combining two of the LRFs, brightness matches can be predicted. These LRF-based models are better at capturing the psychophysical pattern of achromatic brightness matches. Specifically, by varying the shape of the nonlinear functions through a difference semi-saturation constants, the reduced induction effect at more extreme luminances can be accurately predicted. These LRF-based models are the most effective at capturing the sorts of achromatic brightness matching data presented in the current work.