Technische Universität Berlin

Institute of Computer Engineering and Microelectronics

FG Computational Psychology

Fakultät IV
Marchstr. 23
10587 Berlin
https://www.psyco.tu-berlin.de/

Bachelor's Thesis in Computer Science

Testing a unified framework for perceived magnitude and integrated sensitivity in White's Illusion

Ori Toledo Cohen

Matriculation Number: 457185
Date of submission: 26.05.2025

First supervisor: Dr. Guillermo Andreas Aguilar Cornejo
Second supervisor: Prof. Guillermo Gallego

I hereby decalre that the thesis submitted is my own, unaided work, completed without any external help. Only the sources and resources listed were used. All passages taken from the sources and aids used, either unchanged or paraphrased, have been mark as such. Where generative AI tools where used, I have indicated the product name, manufacturer, the software version used, as well as the perspective purpose (e.g, checking and implementing language in the texts, systematic research). I am fully responsible for the selection, adoptation and all results of the AI-generated output I use. I have taken note of the Principles for Ensuring Good Research Practice at TU Berlin dated March 8th, 2017. https://www.tu.berlin/em/working-at-tu-berlin/important-documents/guidelinesdirectives/principles-for-ensuring-good-research-practice
I further declare that I have not submitted the thesis in the same or similar form to any other examination authority

Berlin, 26.05.2025

....
(Signature)

Abstract

This thesis examines how perceived magnitude and discriminability relate in the context of White's Illusion, using the framework proposed by Zhou, Duong, and Simoncelli (2024). White's illusion is a brightness illusion in which traget patches with same luminance, embedded in alternating black and white gratings, appear different depending on their spatial context. According to Zhou et al. (2024), both intensity judgments and sensitivity arise from a shared internal representation, linking the mean response to stimulus strength with its variability via Fisher Information. To test this, participants took part in magnitude estimation task- they viewed luminance-defined target patches embedded in black or white gratings and rated their brightness. From these ratings, mean responses and standard deviations were extracted. Sensitivity was derived by dividing the derivative of the mean response by the standard deviation. Integrating sensitivity yielded intensity functions, which were fitted to perceptual scales from a prior MLCM study on the same illusion. Each context was assigned its own integration constant, and a shared scale factor aligned units. Most fits showed strong alignment (average RMSE 6.9%). While the mean responses followed the expected power-law form, standard deviations-interpreted as internal noise, systematically deviated from the framework's multiplicative noise assumption. LOWESS smoothing revealed inverted U-shapes in many noise profiles, suggesting non-monotonic effects potentially linked to anchoring. These deviations challenge the framework's assumption of proportional noise and highlight complications when applying it to context-dependent illusions. Still, the consistent match between ME-derived and MLCM scales suggests the model robustly captures core aspects of contextual modulation. Future work should minimize task-related biases and collect ME and MLCM data within the same observers to clarify the sources of deviation.

Zusammenfassung

Diese Arbeit untersucht, wie wahrgenommene Helligkeit und Diskriminierbarkeit im Kontext der White'schen Illusion zusammenhängen, unter Verwendung des von Zhou et al. (2024) vorgeschlagenen theoretischen Rahmens. Die White'sche Illusion ist eine Helligkeitsillusion, bei der Zielbereiche mit identischer Luminanz, eingebettet in abwechselnd schwarze und weiße Gitter, je nach räumlichem Kontext unterschiedlich hell erscheinen. Laut Zhou et al. (2024) gehen sowohl Intensitätseinschätzungen als auch Sensitivität aus einer gemeinsamen internen Repräsentation hervor, die die mittlere Reaktion auf die Reizstärke mit ihrer Variabilität über die Fisher-Information verknüpft.

Zur Überprüfung dieser Annahme führten Teilnehmende eine Magnitudenschätzaufgabe (ME) durch: Sie betrachteten luminanzdefinierte Zielbereiche, eingebettet in schwarze oder weiße Gitter, und bewerteten deren Helligkeit. Aus diesen Bewertungen wurden Mittelwerte und Standardabweichungen extrahiert. Die Sensitivität wurde berechnet, indem die Ableitung der mittleren Reaktion durch die Standardabweichung geteilt wurde. Durch Integration der Sensitivität wurden Intensitätsfunktionen erhalten, die an perzeptuelle Skalen aus einer früheren MLCM-Studie zur gleichen Illusion angepasst wurden. Jedem Kontext wurde eine eigene Integrationskonstante zugewiesen, und ein gemeinsamer Skalierungsfaktor sorgte für die Angleichung der Einheiten. Die meisten Anpassungen zeigten eine hohe Übereinstimmung (durchschnittlicher RMSE 6,9%).

Während die Mittelwerte der Bewertungen der erwarteten Potenzfunktion folgten, wichen die Standardabweichungen- interpretiert als internes Rauschen- systematisch von der Annahme des Modells ab, dass das Rauschen proportional sei (multiplikatives Rauschen). Eine LOWESS-Glättung offenbarte in vielen Rauschprofilen umgekehrt U-förmige Verläufe, was auf nicht-monotone Effekte hindeutet, die möglicherweise mit Ankermechanismen zusammenhängen. Diese Abweichungen stellen die Annahme proportionalen Rauschens in Frage und verdeutlichen die Herausforderungen bei der Anwendung des Modells auf kontextabhängige Illusionen. Dennoch deutet die konsistente Übereinstimmung zwischen den aus Magnitudenschätzungen abgeleiteten Skalen und den MLCM-Skalen darauf hin, dass das Modell zentrale Aspekte kontextueller Modulation robust abbildet. Zukünftige Arbeiten sollten aufgabenspezifische Verzerrungen minimieren und ME- sowie MLCM-Daten innerhalb derselben Stichprobe erheben, um die Ursachen der Abweichungen besser zu klären.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Guillermo Aguilar, for his continuous support, guidance, and encouragement throughout the course of this thesis. His constructive feedback and insightful suggestions were invaluable to the development of this work. Working under his supervision was a truly rewarding academic experience, for which I am deeply grateful.

Statement on the Use of AI Tools

During the preparation of this thesis, I used generative AI tools (ChatGPT 4) exclusively to support spelling correction, grammar checking, and stylistic improvements. No content, ideas, or analyses were generated by AI. All scientific content, arguments, and conclusions are entirely my own.

Berlin, 26.05.2025   Signature: _____

Contents

Contents

List of Figures

List of Figures

List of Figures

List of Tables

1 Introduction

For centuries, people have been fascinated by how we perceive the world around us-how patterns of light, sound, and touch translate into coherent experiences. This long-standing curiosity led to the emergence of psychophysics, a scientific discipline devoted to uncovering the relationship between physical stimuli and perceptual experience. As the name suggests, "psycho" refers to mind, while "physics" refers to the measurable properties of the external world- capturing the core aim of the field: to quantify how objective input is transformed into subjective sensation. This endeavor faces a fundamental challenge: perceptual experiences are inherently private, accessible only through introspection and verbal report. Psychophysics addressed this by developing experimental methods that infer internal experience from behavioral responses, offering a way to systematically measure perception.

Through carefully designed experiments, psychophysics introduced methods to quantify subjective sensations in response to objective changes in stimuli, laying the groundwork for much of modern perceptual science. In the domain of visual perception, two central dimensions have shaped this line of inquiry- sensitivity and perceived magnitude. While traditionally examined through separate theoretical lenses, they capture complementary aspects of how sensory information is encoded and experienced.

1.1 Sensitivity and Weber's law

Imagine the following scenario: you sit in a completely dark room. Someone lights a single candle, and the change is immediate, the flame cuts through the darkness, and the space is filled with light. Mere single flame made a dramatic difference in comparison to the former pitch black room, it is much easier to see your surroundings. A second candle is added to the room, and again, the difference is noticeable. The more lit candles are added- the room grows steadily brighter, but each new flame seems to matter a little less. By the time fifty candles are burning, lighting one more barely changes the room's appearance.

This intuitive experience is a perfect demonstration of the concept of sensitivity, which is the ability to detect small differences between stimulus intensities. It is commonly measured using the just-noticeable difference (JND), the smallest change in a stimulus that leads to a perceptible difference. The JND increases with stimulus intensity, meaning it is harder to notice a difference in a stronger stimulus- a relationship known as Weber's law (Sowden, 2012). In the candle example, the brighter the room becomes, the greater the increase in light needed to produce a noticeable change. Weber's law formalizes this observation, stating that the ratio between the JND and the overall stimulus

magnitude remains roughly constant, as shown in Equation 1.1:

$$\Delta I = k \cdot I \tag{1.1}$$

JND (ΔI) between two stimulus intensities is proportional to the baseline intensity (I), where k is a constant. This implies that as stimulus intensity increases, a larger difference is needed for a change to be perceptible.

This proportional relationship has been validated across a wide range of sensory modalities, including vision, audition, and touch. Sensitivity is highest at low stimulus levels and diminishes as physical intensity increases, making it harder to distinguish changes at higher magnitudes.

Weber's law expresses a constant ratio between the JND and stimulus intensity. If sensitivity is interpreted as the inverse of the JND, then it can be concluded that:

$$D(I) = \frac{1}{\Delta I} \propto \frac{1}{I} \tag{1.2}$$

This means, sensitivity ($D(I)$) decreases as intensity increases, following an inverse proportional trend. Figure 1.1 illustrates this relationship: sensitivity is highest at low intensities and declines as intensity increases. This inverse trend implies that equal physical increments produce diminishing perceptual effects.
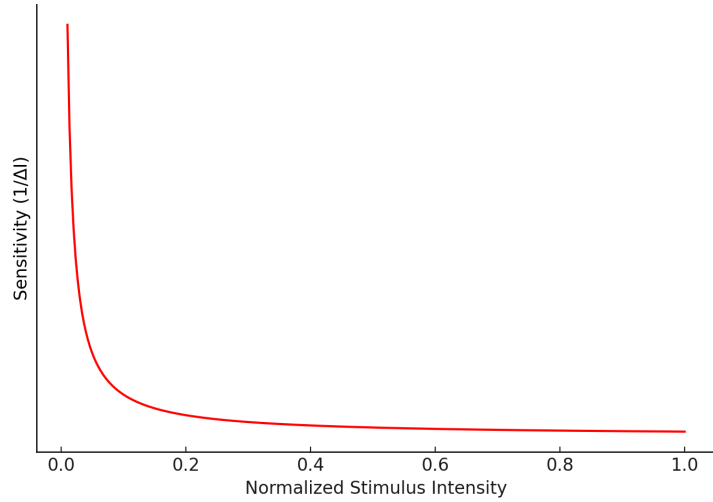
Figure 1.1: **Illustration of Weber's law.** The plot depicts the inverse relationship between sensitivity (y-axis), and normalized stimulus intensity (x-axis). As stimulus intensity increases, sensitivity decreases, indicating that larger changes in intensity are required to notice a difference. This reflects Weber's principle that the just-noticeable difference (JND) is proportional to the baseline intensity, leading to diminishing perceptual sensitivity at higher stimulus levels.

Figure 1.2: **Illustration of Fechner's law.** The plot shows the logarithmic relationship between perceived intensity (y-axis) and normalized stimulus intensity (x-axis), as proposed by Fechner. As stimulus intensity increases, the rate of change in perceived intensity diminishes, resulting in a conca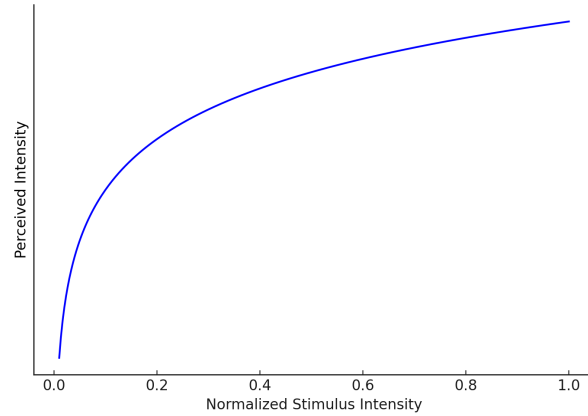ve upward curve. This reflects the principle that equal increments in physical intensity produce progressively smaller perceptual effects.

1.2 Perceived intensity and Stevens' law

Returning to the candlelit room: the first few candles dramatically transform the space, while later additions make increasingly smaller perceptual differences. Eventually, the light seems to plateau, even though more candles continue to be added. This shift not only illustrates a decline in sensitivity but also raises the question of how brightness is perceived as intensity accumulates. While physical intensity of a stimulus can be measured directly, perceived magnitude refers to how strong a stimulus feels- a subjective judgment that does not always align linearly with the physical stimulus.

Fechner's law (Fechner, 1860) was the first attempt to formalize this relationship. Building on Weber's law, it proposed that perceived intensity increases logarithmically with physical stimulus strength. That is, each just-noticeable difference adds an equal step to perceived magnitude, leading to a cumulative internal scale, as seen in Figure 1.2. This approach offered a foundational link between sensory thresholds and perception. However, it assumed a fixed logarithmic form across all modalities.

Stevens' law (Stevens, 1957) later refined this view. Stevens proposed that the relationship between physical and perceived intensity follows a power law:

$$s = k \cdot I^{\alpha} \tag{1.3}$$

where α is a modality-specific exponent, and k is a constant. The value of α determines the shape of the perceptual response: compressive ($\alpha < 1$), linear ($\alpha = 1$), or expansive ($\alpha > 1$).
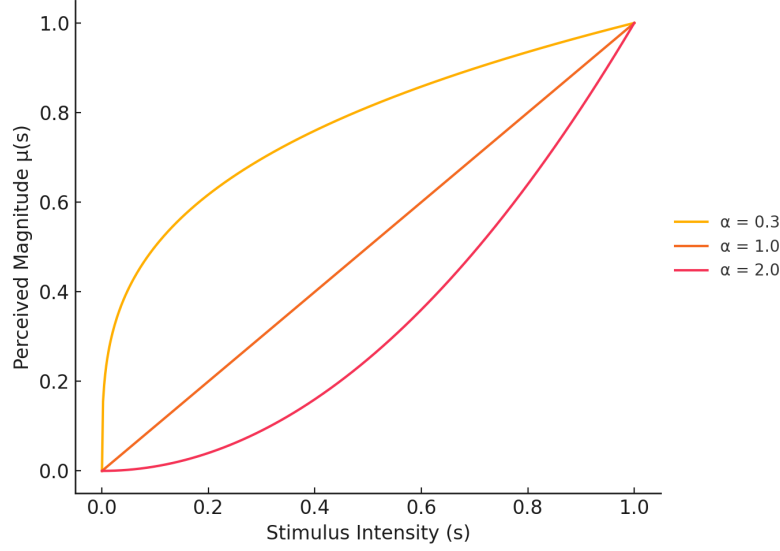
Figure 1.3: **Illustration of Stevens' power law.** The plot shows perceived magnitude ($\mu(s)$) as a function of stimulus intensity (s), for different exponent values (α). Each curve represents a power-law relationship: $\alpha = 0.3$ (compressive, yellow), $\alpha = 1.0$ (linear, orange), and $\alpha = 2.0$ (expansive, red). The x-axis represents physical stimulus intensity, while the y-axis reflects perceived magnitude. Lower exponents yield diminished perceptual growth, while higher exponents exaggerate it, capturing modality-specific variations in sensory scaling.

1.3 Unified framework

At first glance, Weber's and Stevens' laws appear to address complementary aspects of perception: Weber's law describes how small a stimulus change must be in order to be noticeable, while Stevens' law quantifies how strong a stimulus is subjectively perceived. Despite their shared goal of modeling perceptual experience, the two laws have historically been treated as separate, leading to several theoretical and empirical issues (Krueger, 1989). First, maintaining two independent models suggests that the human brain uses distinct mechanisms to encode sensitivity and perceived magnitude- despite both being derived from the same sensory input. This complicates efforts to develop unified, efficient models of perception. Second, the two laws impose an empirical conflict between the logarithmic scaling implied by Weber's law and the power-law scaling observed by Stevens (1957). Treating them as unrelated prevents a coherent explanation of how perceived intensity and sensitivity interact or co-vary across conditions. Finally, maintaining this separation limits the ability to connect the psychophysical with the physiological. Without a unified framework, the link between neural encoding and behavioral responses remains unclear, reducing our ability to describe and develop generalizable models of perception.

   A recent paper by Zhou et al. (2024) proposes a mathematical framework that resolves this longstanding divide by anchoring both Stevens' and Weber's laws to shared properties of internal sensory representations. The framework posits that sensory stimuli are

transformed into an internal representation- the brain's internal encoding of stimulus intensity, which combines neural response strength (mean response, denoted as $\mu(s)$) and variability of that neural response which is interpreted as noise (denoted as $\sigma(s)$). The mean response is modeled as:

$$\mu(s) = k \cdot s^{\alpha} \tag{1.4}$$

where s is the stimulus intensity, k is a scaling constant, and α determines the nonlinearity of the response. Noise ($\sigma(s)$), representing the deviation in the internal representation, is assumed to scale proportionally with the mean response itself:

$$\sigma(s) = c \cdot \mu(s) + b \tag{1.5}$$

where c serves as proportionality constant and b is the intercept. This formulation captures different models of internal noise depending on the values of c and b. Additive noise corresponds to the case $c = 0$ and $b \neq 0$, resulting in constant noise independent of stimulus intensity. Multiplicative noise arises when $c \neq 0$ and $b = 0$, where $\sigma(s)$ scales proportionally with the mean response. This model supports both Weber's and Stevens' laws for any exponent α according do the framework proposed by Zhou et al. (2024). When both c and b have nonzero values, the resulting mixed noise model includes both stimulus-dependent and independent components. The framework uses Fisher Information (FI) to quantify the precision of internal representations in noisy systems. While FI provides a general and powerful description, it is computationally intractable. To address this, the framework employs Fisher sensitivity- a simpler approximation that serves as a lower bound on FI, and is defined as:

$$D(s) = \frac{\mu'(s)}{\sigma(s)} \tag{1.6}$$

where $\mu'(s)$ is the derivative of the mean response with respect to stimulus intensity, and $\sigma(s)$ is the internal noise.

This unification explains Weber's law under various noise configurations, with proportional noise allowing both Weber's and Stevens' laws to coexist. As illustrated in Figure 1.4, the left and middle panels demonstrate how Weber's law is compatible with Stevens' law for any exponent, as long as internal noise scales proportionally with the mean response. The right panel extends this idea by assuming that perceptual distances between clearly visible (suprathreshold) stimuli correspond to integrated sensitivity. Under this assumption, these distances follow a logarithmic mapping, reconciling all three laws within a unified framework.

While Zhou et al. (2024) framework demonstrates strong potential for unifying classical laws across sensory modalities, its application to context-dependent phenomena
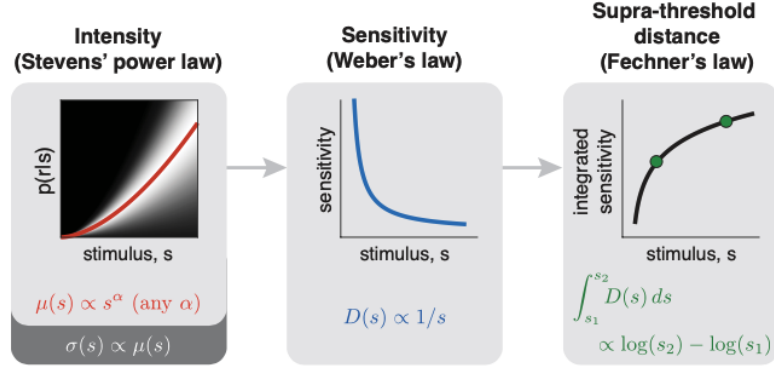
Figure 1.4: **Unified framework linking Weber, Stevens', and Fechner.** This figure illustrates how Zhou et al. (2024) model reconciles three classical psychophysical laws. The left and middle panels demonstrate that Weber's law and Stevens' law are compatible when internal noise scales proportionally with the mean response. The right panel shows that, assuming above threshold (suprathreshold) perceptual distances align with integrated sensitivity, the framework predicts a logarithmic internal mapping, consistent with Fechner's law.

like visual illusions remains untested. Visual illusions pose a unique challenge to the framework's core assumption- that perceived magnitude and sensitivity derive from a fixed internal representation of the stimulus. In illusions, the appearance of a stimulus is systematically altered by its spatial context, potentially challenging the assumption of direct stimulus-response mapping. Context effects like contrast (where the stimulus appear less like its surround) and assimilation (where stimulus appear more like its surround) demonstrate that intensity perception depends not only on absolute stimulus properties but also on relative spatial structure. If the framework can account for such phenomena, it would mark a significant step toward a general model of perception that accommodates both basic scaling laws and more complex contextual modulations.

1.4 Research question

The framework proposed by Zhou et al. (2024) offers a unified explanation for perceived magnitude and sensitivity by linking both to shared properties of internal sensory representations. It addresses the long-standing separation between two central aspects of perception- sensitivity and perceived intensity. The framework is supported by experimental findings across multiple sensory domains and grounded in a mathematically consistent structure. Yet, its validity under context-dependent phenomena like visual illusions remains untested. Context effects, such as contrast or assimilation, systematically alter perception without changes in the physical stimulus itself. These effects present a challenge to models based on direct mappings between stimulus intensity and internal representation.

Assimilation effects provide a meaningful extension in this direction. In assimilation, the appearance of a target is influenced by its surround, becoming more similar to
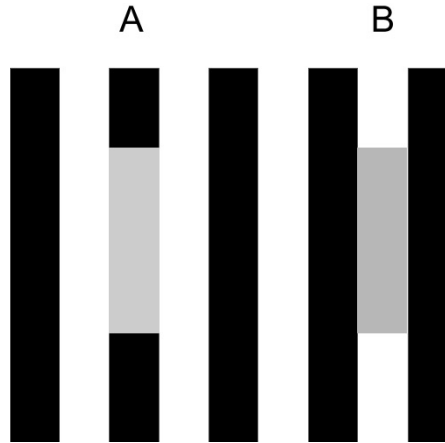
Figure 1.5: **Contextual modulation demonstrated in White's illusion.** Two identical gray rectangles are placed over alternating black and white grating. Despite their physical equivalence, the one embedded in black appears brighter than the one embedded in white. This classic illusion exemplifies an assimilation effect, where the target adopts the visual characteristics of its surrounding context.

it. White's Illusion (White, 1979) exemplifies this: two identical patches, placed over alternating black and white grating, appear different in brightness (see Figure 1.5). The patch on the black stripe appears lighter, and the one on the white stripe darker, despite being physically identical. This illusion demonstrates how perceived intensity can be modulated by spatial context, challenging models that assume a direct mapping from physical to perceptual magnitude.

White's Illusion is especially suited for testing the framework proposed by Zhou et al. (2024), not only because of its strong contextual modulation, but also because perceptual intensity data for this illusion have already been collected. These data were obtained through a Maximum Likelihood Conjoint Measurement (MLCM) experiment conducted by (Vincent, Maertens, & Aguilar, 2024) at the Computational Psychology Department of TU Berlin. MLCM provides perceptual scales based on comparative judgments, offering an independent and methodologically distinct estimate of perceived intensity. Since Zhou et al. (2024)'s framework yields intensity functions by integrating sensitivity derived from magnitude estimation, the availability of MLCM data creates a unique opportunity: it allows a direct comparison between the analytically derived intensity functions and independently measured perceptual scales. This makes it possible to assess whether the framework can recover established perceptual patterns in a context-dependent illusion like White's.

This leads to the research question of this thesis: **How do perceived magnitude and integrated sensitivity relate to each other in White's Illusion, within the framework proposed by Zhou et al. (2024)?**

1.5 Magnitude estimation and MLCM

Magnitude estimation (ME) is a psychophysical method for quantifying perceived intensity. Observers assign numerical values to stimuli based on how strong or intense they perceive them (Kingdom & Prins, 2016).

ME offers several advantages. First, it is highly sensitive to fine perceptual differences, allowing participants to report subtle variations in perceived intensity using continuous numerical values (Stevens, 1957). Second, it is straightforward to implement, requiring minimal instruction and no complex stimulus pairings, which makes it efficient for both researchers and participants (Gescheider, 1997). Third, it provides direct access to perceived magnitude on an interval scale, enabling immediate modeling of perceptual responses (Stevens & Galanter, 1957).

However, ME also has notable drawbacks. Small shifts in stimulus range can lead to range effects, a bias where participants' responses are influenced by the range of stimuli presented, rather than by their absolute properties. Another problematic aspect is anchoring effects (Garcia-Marques & Fernandes, 2023), where participants' estimates are influenced by initial or previously seen values, even when those values are irrelevant to the current judgment. Additionally, participants often interpret the numerical scale differently, leading to inconsistencies (Mertens, Mertens, & Lerche, 2021). These issues make ME less reliable for between-subject comparisons or for capturing stable perceptual metrics.

To address these limitations, alternative methods such as partition scaling and maximum likelihood conjoint measurement (MLCM) have been developed. MLCM, in particular, is a comparative judgment method in which participants evaluate which of two stimuli appears more intense, across a systematically varied stimulus grid. By modeling the frequency of chosen comparisons under varying conditions, MLCM estimates underlying perceptual scales for each stimulus dimension (Knoblauch, Maloney, Knoblauch, & Maloney, 2012). Unlike ME, MLCM yields relative rather than absolute scale values, but with higher internal consistency and resistance to contextual biases.

2 Methods

2.1 Participants

To investigate the relationship between perceived brightness and integrated sensitivity within the context of White's Illusion, a magnitude estimation experiment was conducted with nine participants. Eight participants were naïve observers, while one participant ('OTC99') was the author of this thesis- who also designed the experiment and was highly familiar with its theoretical background and objectives. All participants took part exclusively in this experiment. No data were reused from previous studies.

2.2 Stimulus

The stimuli consisted of gray target patches embedded in a high-contrast square-wave grating, as can be seen in Figure 2.1. The grating spanned 16° × 12° (width × height) at 0.5 cycles per degree, resulting in 8 full cycles (16 bars). Luminance ranged from 5.25 cd/m² (black bars) to 525 cd/m² (white bars), overlaid on a uniform gray background of 95 cd/m². Two target patches, each 4° high, were vertically centered within the grating. Eighteen target luminance levels were used, evenly distributed across the display's available luminance range.

2.3 Experiment Design

The experiment was conducted in a darkened room using a JVC monitor (376 mm × 301 mm, 1024 × 768 px), operating at 60 Hz with 16-bit grayscale resolution (65536 levels). The monitor received an analog signal and supported a measured luminance range of 0.76 to 550 cd/m². Participants were seated 76 cm from the screen, resulting in approximately 34 pixels per visual degree. A chinrest was used to ensure a stable viewing position. Stimulus presentation was controlled using the Python library HRL (`https://github.com/computational-psychology/hrl`), and responses were recorded using a keypad.

Each trial presented a single target target patch embedded in a high-contrast square-wave grating (16° × 12°, 0.5 cycles/deg, 8 cycles total). The grating consisted of alternating black and white bars on a uniform gray background. Two vertically centered target patches (each 4° high) were placed either on a black or white stripe. Eighteen luminance levels were sampled across the display's usable range.

Participants completed five blocks of 36 trials (18 per context), totaling 180 trials. Trial order was randomized. Brightness ratings were entered using a keypad, on a scale
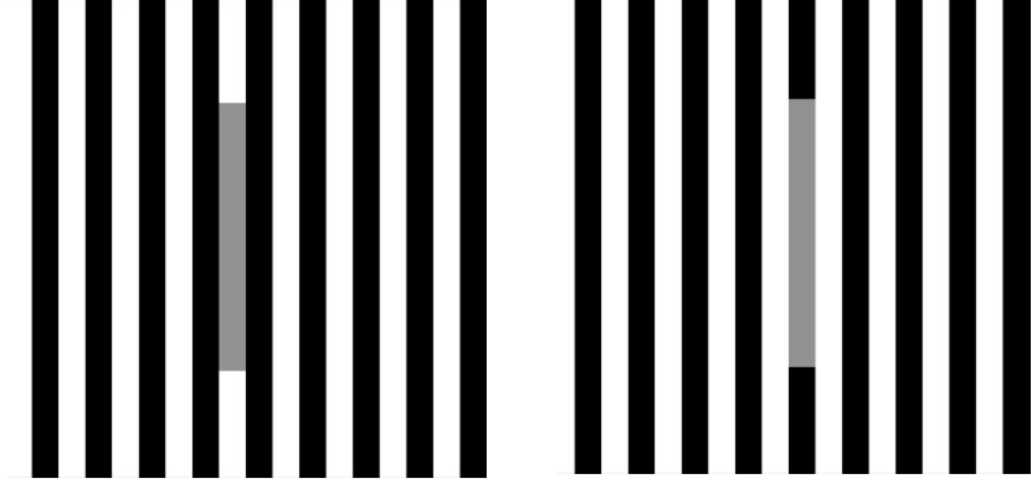
Figure 2.1: **Stimuli used in the magnitude estimation experiment.** Each stimulus consisted of a gray target patch embedded in a square-wave grating of black and white bars. The target's luminance varied across trials, and its position alternated between black and white phases of the grating. This design enables measurement of context-dependent changes in perceived brightness, following the structure of White's Illusion.

from 0 (black) to 100 (white). Breaks were provided every 18 trials.

2.4 Experiment procedure

The experiment consisted of five blocks, each containing 36 trials: 18 with the target on a white bar and 18 on a black bar. All blocks were completed within a single session. Trial order was randomized. Each participant completed 180 trials in total. In every trial, participants assigned a brightness value to the target patch using a scale from 0 (completely black) to 100 (completely white). Break screens appeared automatically after every 18 trials, and participants were free to take additional breaks as needed.

2.5 Data analysis

The analysis focused on the analytical approach suggested in the framework by Zhou et al. (2024). First, the analysis examined how different levels of luminance were perceived across the two contextual conditions. For each luminance level, perceived brightness ratings were averaged to estimate the mean perceptual response. To characterize the relationship between physical stimulus intensity and perceived brightness, a power function following Equation 1.4 was fitted to the data using nonlinear regression, implemented through Python's `scipy.curve_fit` method.

Next, response variability- interpreted as internal noise according to the framework, was measured at each luminance level by calculating the standard deviation of brightness estimates. As with the mean response, Python's `scipy.curve_fit` method was used to find the best-fitting function of the form in Equation 1.5.

To ensure that the fitted functions for both mean response and standard deviation accurately captured the data, residual analyses were performed. For each fit, the root mean squared error (RMSE) and bias were computed. RMSE quantifies the average magnitude of residuals- the differences between predicted and observed values-and reflects overall fit quality. It is defined as:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2} \tag{2.1}$$

where y_i is the observed value, $\hat{y}_i$ is the predicted value, and N is the number of data points. Lower RMSE values indicate closer alignment between model and data.

Bias, in contrast, captures the average signed error across all points and reveals systematic over- or underestimation. It is computed as:

$$\text{Bias} = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i) \tag{2.2}$$

A bias close to zero suggests that residuals are symmetrically distributed around the true values, supporting the adequacy of the functional form used.

Using the estimated mean response and noise functions, a sensitivity function was derived to quantify how distinguishable closely spaced luminance levels are. Theoretically, this reflects the precision of the internal sensory representation and is linked to the concept of Fisher Information. Since the mean response is modeled as a power function as seen in Equation 1.4, its derivative is expressed analytically as:

$$\mu'(s) = \alpha \cdot k \cdot s^{\alpha - 1} \tag{2.3}$$

Taken together the definition of Fisher sensitivity (1.6) and the Equations 1.4 and 1.5:

$$D(s) = \frac{\mu'(s)}{\sigma(s)} = \frac{\alpha k s^{\alpha - 1}}{c k s^{\alpha} + b} \tag{2.4}$$

Then, taking Equation 2.4 and integrating it yields the intensity function:

$$I(s) = \int \frac{\alpha k s^{\alpha - 1}}{c k s^{\alpha} + b} dx = \frac{\log(\alpha(c k s^{\alpha} + b))}{c} + C \tag{2.5}$$

where C is an unknown constant. Equation 2.5 can be rewritten into a more compact and readable form:

$$I(s) = \frac{1}{c} \log(\alpha(cks^\alpha + b)) + C \tag{2.6}$$

In practice, it is impossible to compute the value of the constant C. Since indefinite integration always introduces an additive constant, its value cannot be recovered from the derivative alone- all functions that differ only by a constant have the same derivative. Thus, the integration step inherently yields an undetermined constant C. This makes the comparison between the data obtained through MLCM and the analitically computed intensity functions non-trivial, as C determines the vertical shift in the resulting intensity functions. Without knowing C, the computed functions lack an important aspect of their absolute relationship, preventing direct comparison beyond their relative shapes.

To address this, the ME-derived intensity functions were fitted to the perceptual scales obtained through MLCM. Each context (target on black or white grating) was assigned an independent integration constant, C_w (for white) and C_b (for black), and both functions were jointly scaled by a common factor g. This common scaling factor preserves the internal unit of perceptual magnitude across contexts, reflecting the assumption of a shared internal axis. The parameters α, c, k, and b were fixed to the values obtained previously. Thus, the fitting was performed across three degrees of freedom: C_w, C_b, and g.

The integrated sensitivity functions were defined separately for each context. These represent the result of analytically integrating the Fisher sensitivity function:

$$\begin{aligned}
I_w(s) &= \frac{1}{c} \log\left(\alpha(cks^\alpha + b)\right) + C_w \\
I_b(s) &= \frac{1}{c} \log\left(\alpha(cks^\alpha + b)\right) + C_b
\end{aligned} \tag{2.7}$$

Here, $I_w(s)$ and $I_b(s)$ are the context-specific intensity functions derived from the sensitivity function. The constants C_w and C_b account for the ambiguity introduced by indefinite integration.

To enable comparison with perceptual scales from MLCM, these functions were jointly scaled by a global factor g:

$$\hat{I}_w(s) = g \cdot I_w(s), \quad \hat{I}_b(s) = g \cdot I_b(s) \tag{2.8}$$

The scaling preserves the shape of the functions while aligning them to the units of the perceptual scales. Only the constants C_w, C_b, and g were fitted, while the rest of the parameters were fixed from the earlier model fits.

The intensity functions of each participant from the ME experiment was then fitted independently to each participant from the MLCM dataset, computing the root mean squared error (RMSE) of the fit in each case. The RMSE provides a quantitative measure of the goodness of fit between the analytical prediction and the observed perceptual scales. This procedure resulted in a total of 72 comparisons across all participant pairs.

3 Results

3.1 Mean response

The first step in applying the framework suggested by Zhou et al. (2024) is computing the mean response functions of each participant per luminance level, for each context. Figure 3.1 shows the mean response functions obtained from all nine participants through the ME experiment. The x-axis depicts normalized luminance levels, while the y-axis represents the mean response for each luminance. In each panel, the black markers depict data for targets embedded in the black grating, whereas the light gray markers depict data for targets embedded in the white grating.

The plots show that for each participant, the mean response functions obtained under the two context conditions are relatively similar in shape and progression. This pattern holds across participants, resulting in a high degree of consistency in the overall structure of the functions. In other words, not only are the black and white context curves similar within each participant, but the response profiles also closely resemble one another across the entire participant group. These trends are reflected in the estimated α and k parameters, summarized in Table 3.1.

Table 3.1: **Estimated power-law parameters α and k for all participants.** The table lists the fitted parameters for each participant's perceived magnitude function, for the white and black context conditions. The exponent α determines the nonlinearity of the perceptual response, while the scaling factor k adjusts overall magnitude. Most participants exhibit higher α values in the white context, indicating more expansive scaling relative to the black context.

Participant	α		k	
	in white	in black	in white	in black
OTC99	0.616	0.480	107.405	99.425
MA97	0.675	0.458	104.980	100.419
JH02	0.810	0.561	105.320	95.164
ND99	0.472	0.543	90.650	96.389
LK99	0.439	0.367	101.238	94.256
ST05	0.421	0.285	93.467	85.551
JF00	0.372	0.298	105.652	103.206
GC99	0.474	0.263	104.371	98.258
JS00	0.810	0.561	105.320	95.164

Table 3.1 presents the estimated parameters for the mean response function fits. Across participants, the values of α (the exponent that determines the nonlinearity of
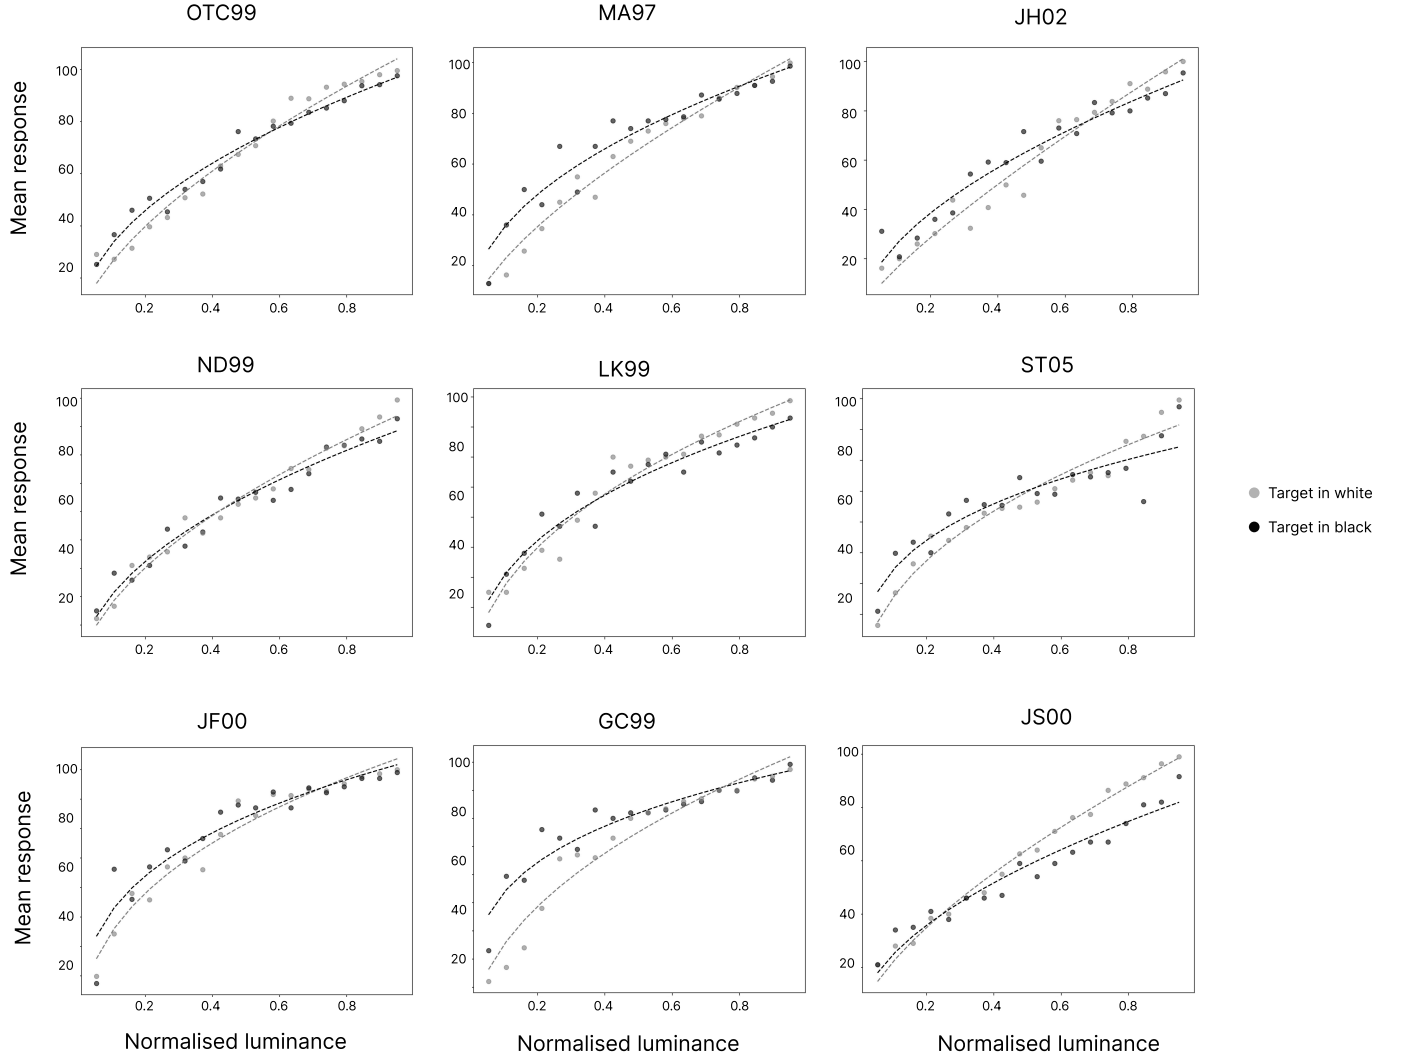
Figure 3.1: **Perceived brightness as a function of luminance.** For each participant, mean responses (y-axis) are plotted against stimulus luminance (x-axis) for both context conditions. The black curve corresponds to targets embedded in a black grating, and the gray curve to targets on white. Across all participants, the assimilation pattern is evident- the target embedded in black grating was perceived systematically brighter than the target embedded in white grating. This effect wasn't constant across the entire luminance range- all participants experienced reversal of the assimilation effect. On average, at higher luminances the target embedded in white grating was perceived as brighter than the one embedded in black.

the mean response function) and k (the scaling factor) are relatively consistent between the two context conditions, supporting the observation of structural similarity in the response functions. Most participants show slightly higher α values in the white context, indicating a more expansive scaling in that condition, but the overall differences remain moderate.

Residual analyses for the standard deviation fits are reported Figure 5.1, and their RMSE and bias values in Table 5.1. RMSE values varied between participants, with some showing higher deviation in the black context (e.g., JF00, ST05), while others showed larger errors in the white context (e.g., JH02, GC99). Bias values were generally small, ranging between -0.23 and 0.27, indicating no strong systematic over- or underestimation.

The mean response plots reveal a notable pattern: consistent with White's Illusion as an assimilation effect, targets embedded in the black grating are initially perceived as brighter than those in the white grating. However, this relationship reverses at higher luminance levels, with the white-context targets eventually appearing brighter. Most participants exhibit this perceptual shift around luminance values of 0.6 to 0.8, though some, such as LK99 and LS00, experience the reversal considerably earlier.

3.2 Standard deviation

Figure 3.2 shows the standard deviation functions for each participant per context. The x-axis represents normalized luminance, and the y-axis represents the standard deviation values.

The parameters estimated from the model (c which is a scaling constant and b which is the intercept) varied considerably across participants and contexts. Most participants showed a $c > 0$, indicating that noise decreased with increasing perceived intensity. This trend was consistent across both contexts for the majority of participants. Positive values of c, though marginal, were observed in participants ND99 and ST05 in at least one context. Intercept values also varied widely, ranging from approximately 2.0 to over 30.0. Some participants, such as JF00 and GC99, exhibited high baseline noise levels, while others, including ST05 and JS00, showed comparatively low values. A full summary of the parameters is provided in Table 3.2.

The corresponding noise functions, shown in Figure 3.2, reveal substantial variability in shape and trend, unlike the more consistent patterns observed in the mean response functions in Figure 3.1. In some participants (e.g., OTC99, LK99, GC99), the functions across contexts appear nearly parallel. Others (e.g., MA97, JH02, ND99, JF00) display clear intersections between the two context-specific curves. ST05 and JS00 exhibit overlapping starting points, with the functions gradually diverging across the stimulus range.

Residual analyses for the standard deviation fits are reported Figure 5.1, and their RMSE and bias values in Table 5.2. In some participants (e.g., JS00, LK99), RMSE values were low in both contexts, indicating that the noise function captured response variability well. Others, such as JH02 and GC99, showed high errors, suggesting a rather
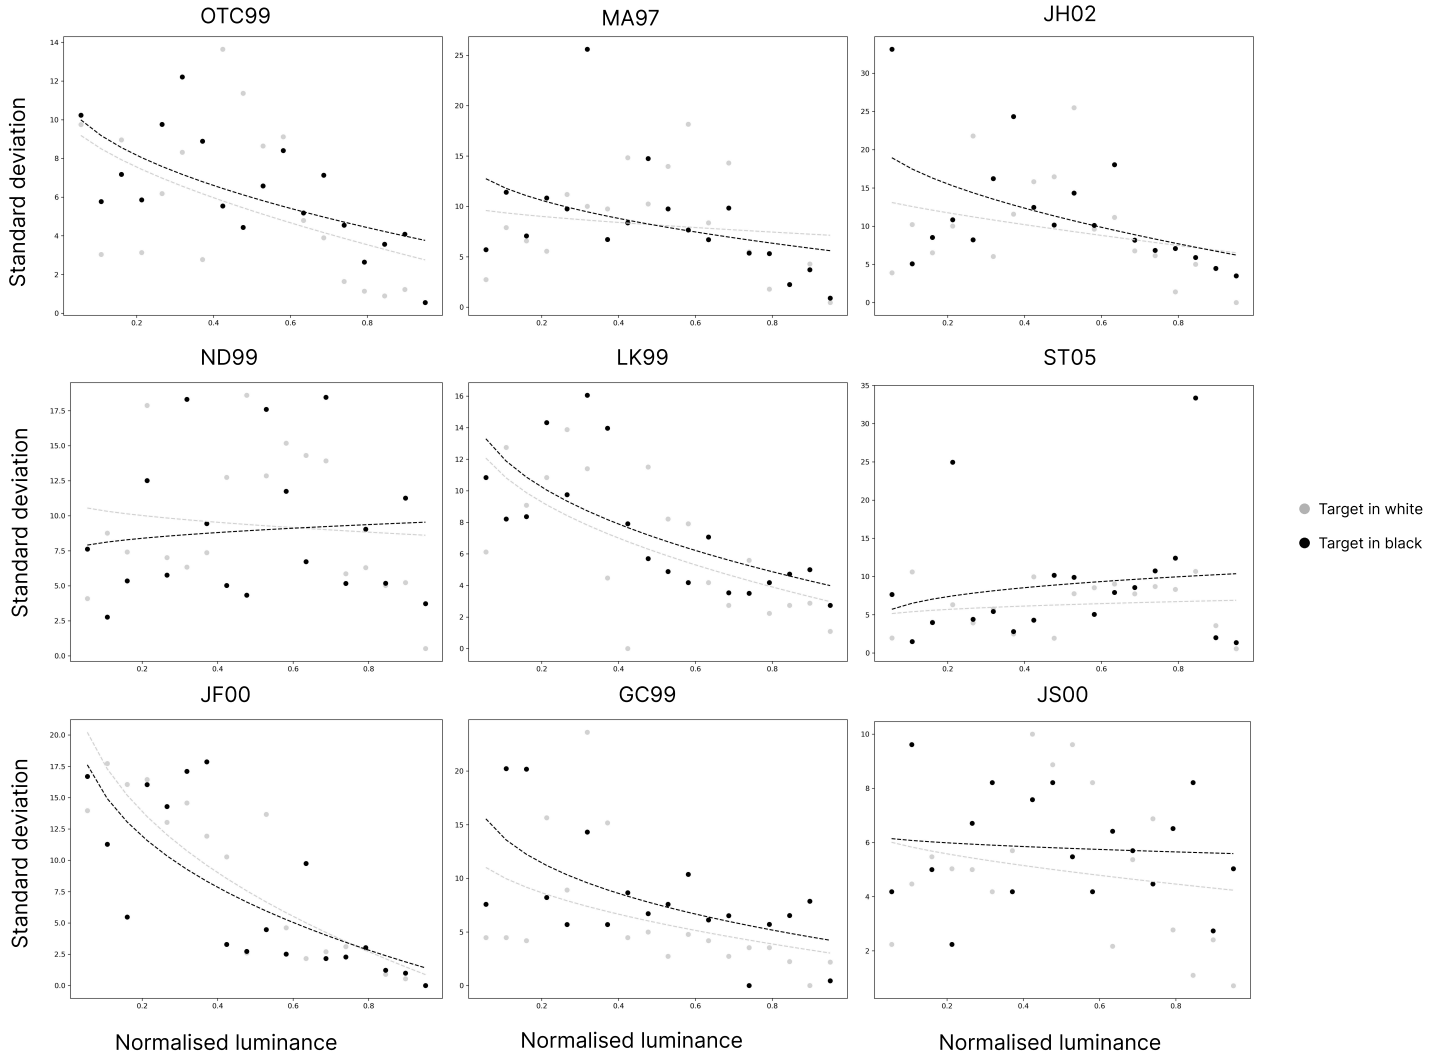
Figure 3.2: **Standard deviation functions across participants and context conditions.** Each subplot shows the variability of brightness estimates as a function of normalized luminance for one participant, separately for targets embedded in white (gray points and dashed line) and black (black points and dashed line) grating phases. The y-axis represents the standard deviation of reported brightness, interpreted as internal noise. Most participants exhibit decreasing noise with increasing luminance.

Table 3.2: **Estimated parameters of the standard deviation functions for all participants.**
The table lists the fitted c (proportionality constant) and b (intercept) values from the
noise model, for the white and black context conditions. The parameter c indicates how
noise scales with the mean response, while the intercept reflects baseline noise level. Most
participants show negative c values, implying decreasing variability with increasing intensity,
contrary to the multiplicative noise assumption of the framework.

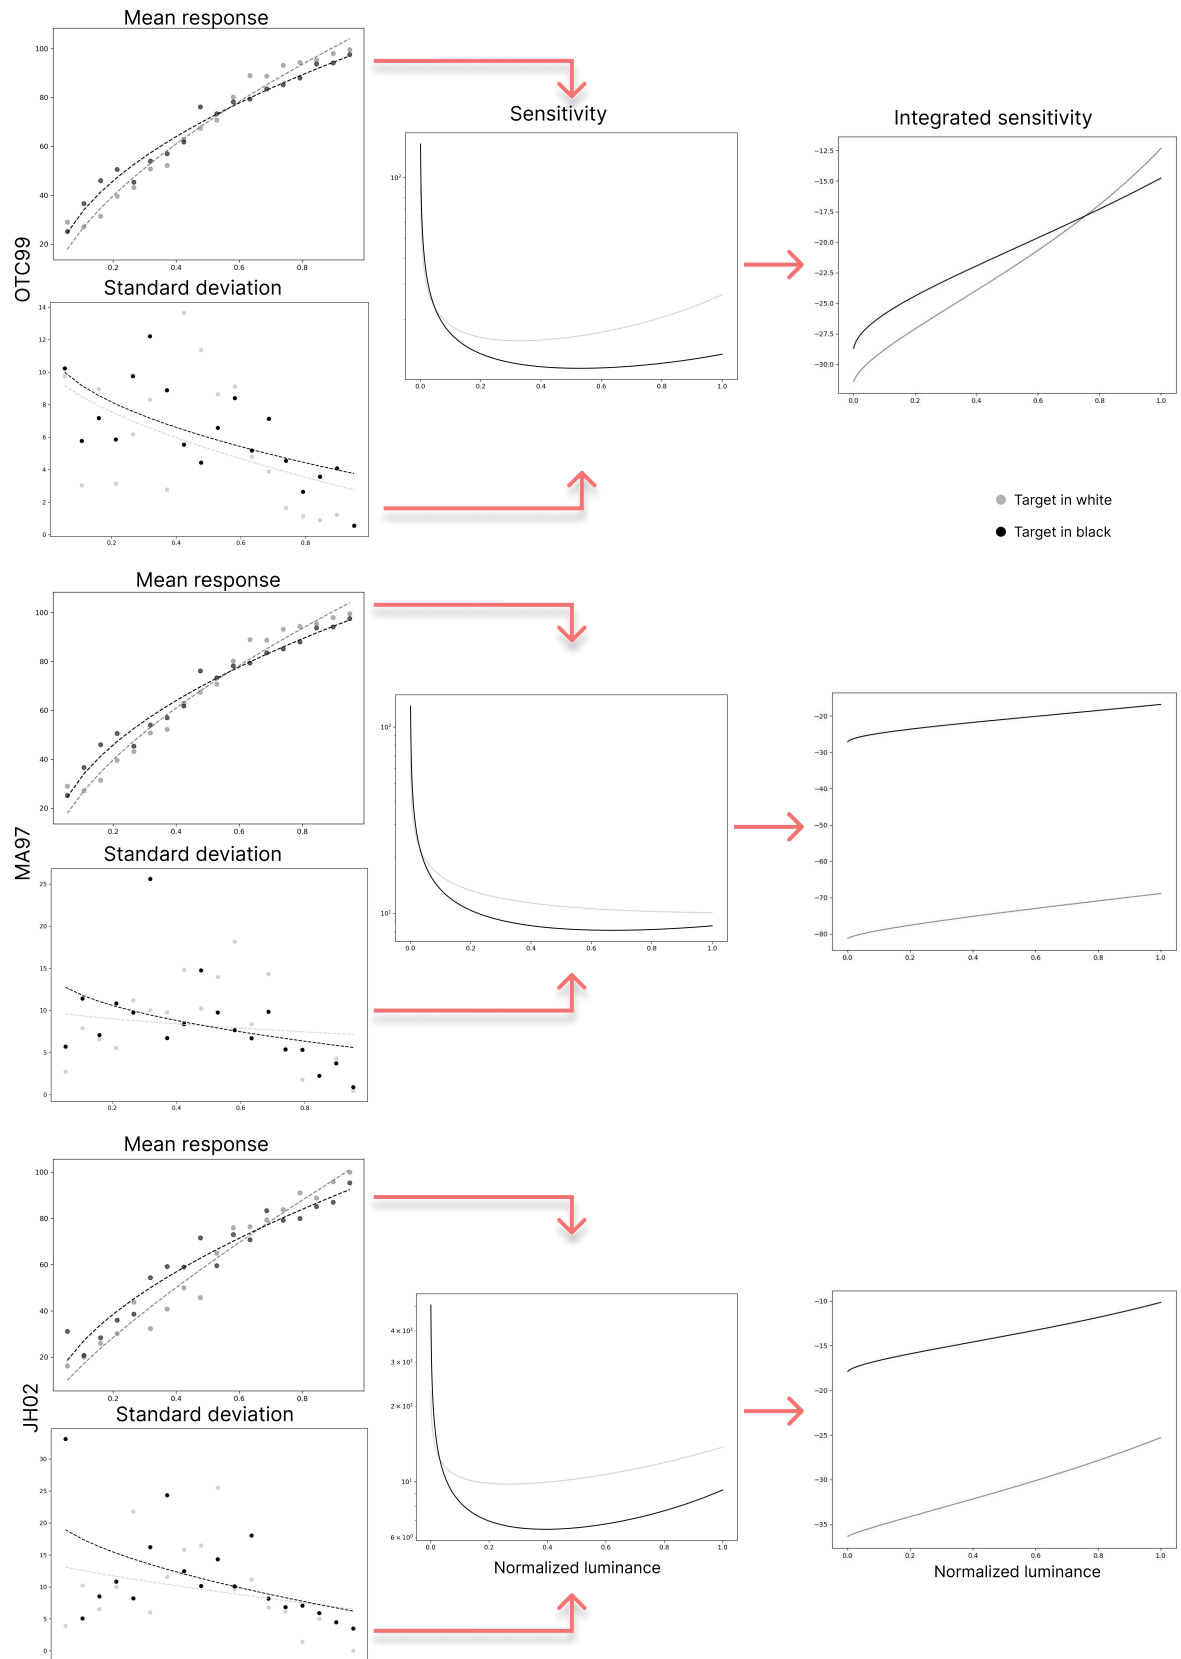| Participant	c		b	
	in white	in black	in white	in black
OTC99	-0.074	-0.086	10.518	12.106
MA97	-0.028	-0.099	10.005	15.388
JH02	-0.072	-0.172	13.809	22.147
ND99	0.024	-0.026	7.334	11.072
LK99	-0.128	-0.154	15.693	18.312
ST05	0.027	0.098	4.400	2.0214
JF00	-0.284	-0.278	30.411	29.694
GC99	-0.105	-0.220	13.783	25.640
JS00	-0.021	-0.008	6.315	6.296

poor fit. ST05 showed a large difference between contexts, with a particularly bad fit in
the black condition. Bias values were zero in all cases due to the symmetric structure
of residuals. Overall, the noise model suggested by the framework (Equation1.5) model
worked well for some participants but did not fit all data equally. Bias values of the stan-
dard deviations were exactly zero across all participants and conditions. This outcome
reflects the symmetric structure of residuals- overestimations and underestimations oc-
curred in balanced proportion across the luminance scale. While RMSE values captured
the overall deviation between predicted and observed noise, the zero bias indicates that
the model did not introduce systematic upward or downward shifts.

3.3 Intensity

After obtaining the values from the mean response and standard deviation plots, the sen-
sitivity functions (as shown in Equation 2.4) and the intensity functions (Equation 2.6)
were computed. The results are displayed in Figure 3.3. Figure 3.4 shows all the resulting
intensity functions on a single plot.

The intensity functions can be roughly divided into two groups. Group A includes
the participants OTC99, LK99, and JF00. This group is characterized by S-like, curvy
intensity functions, and the curves for these participants cross each other around a nor-
malized luminance of 0.8. Group B consists of the rest of the participants- MA97, JH02,
ND99, ST05, GC99, and JH00. Their intensity functions are more logarithmic in ap-
pearance and do not exhibit any crossing points. These groupings reflect differences
in the relative structure of the intensity functions across contexts: some participants
show consistent vertical separation between conditions, while others display a reversal
at higher luminance levels. However, these differences must not be interpreted in terms

Mean response

Standard deviation

Sensitivity

Integrated sensitivity

ND99

Target in white

Target in black

LK99

ST05

Normalized luminance

Normalized luminance
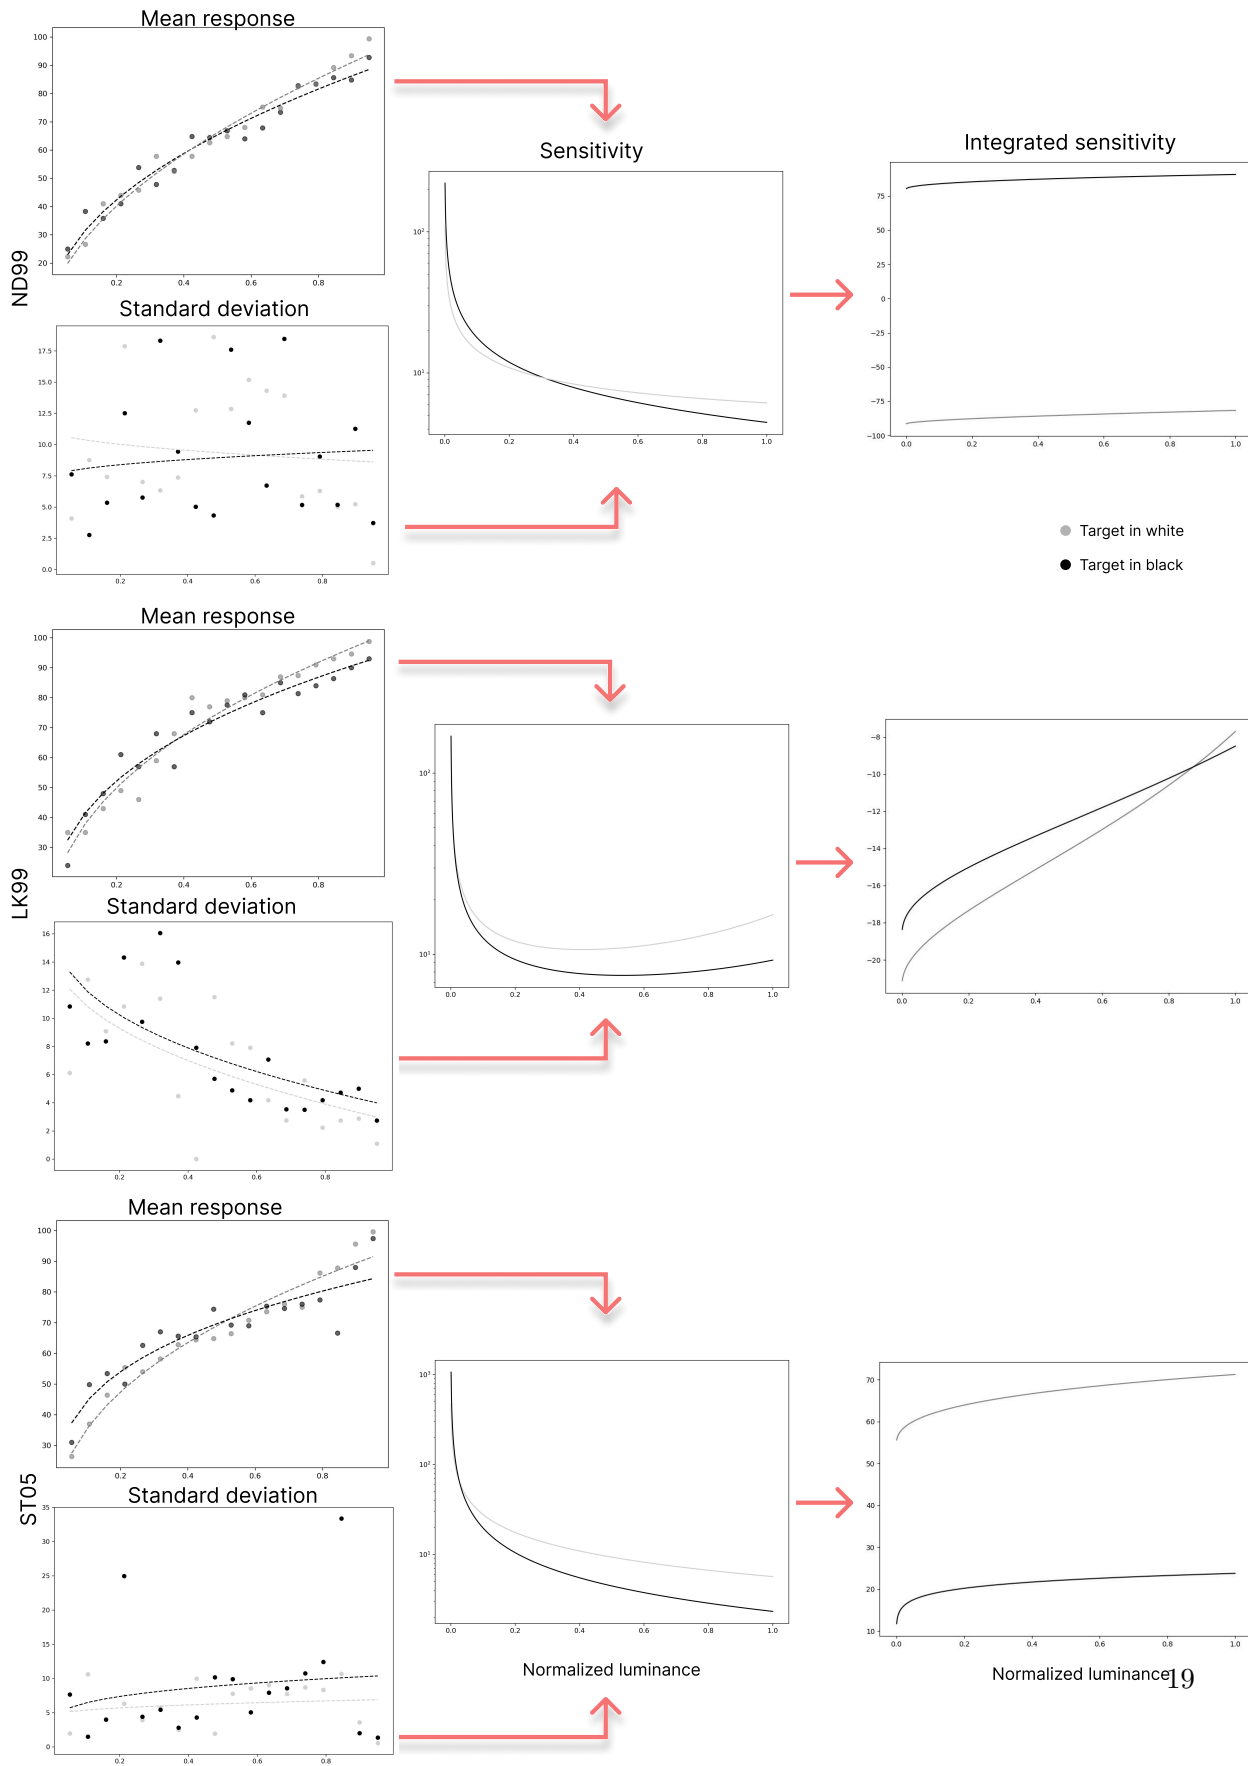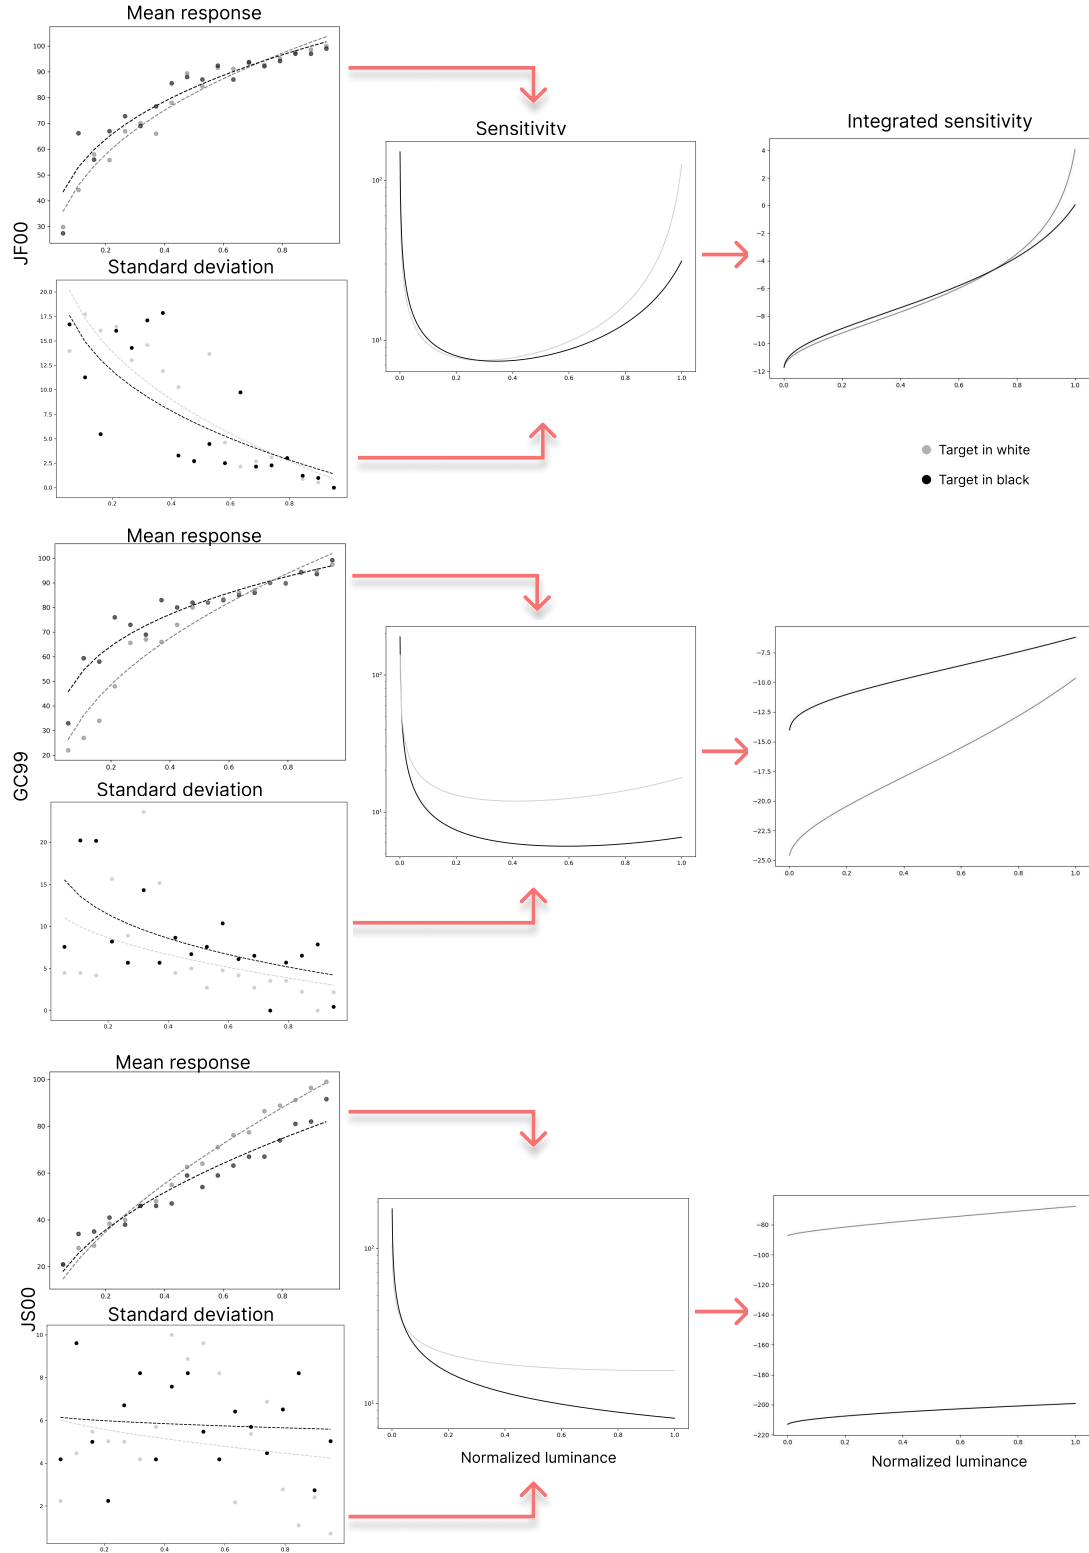
Figure 3.3: **The framework suggested by Zhou et al. (2024)**- obtaining mean response and standard deviation functions (top and bottom left) allows for the calculation of sensitivity function (center). Integrating the sensitivity function yields the intensity function (right).
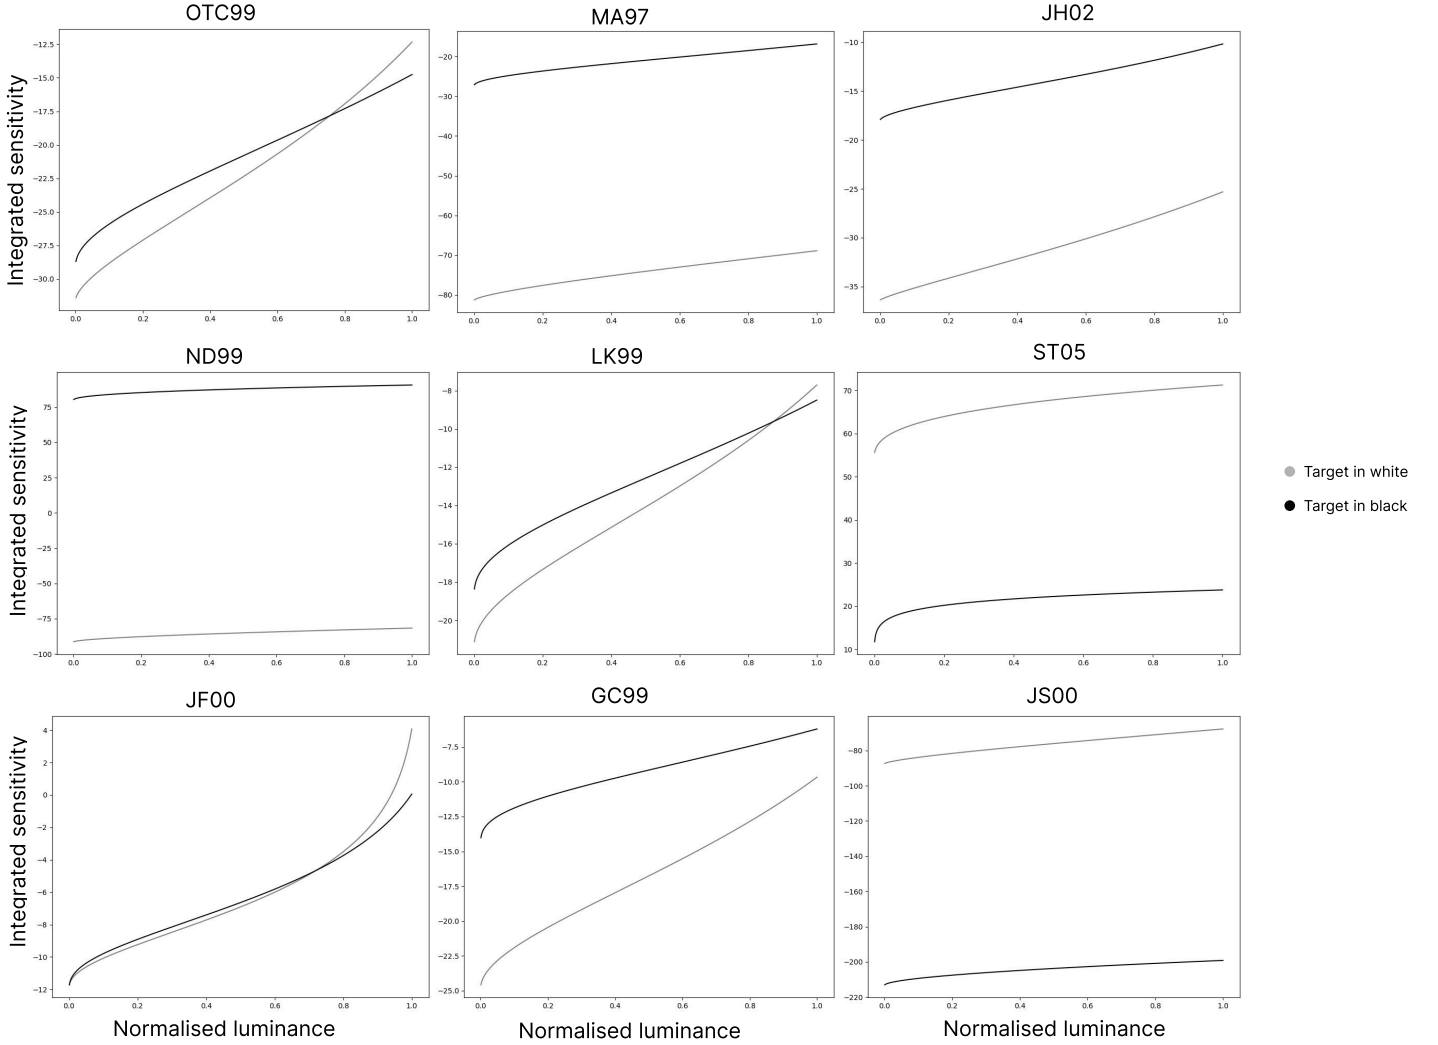
Figure 3.4: **Intensity (integrated sensitivity) functions for all participants.** Each plot shows intensity functions analytically derived by integrating Fisher sensitivity from individual magnitude estimation data, separately for targets in white (gray) and black (black) grating contexts. The x-axis shows normalized luminance, the y-axis shows relative intensity. Integration constants used for fitting to MLCM data are not applied here, so vertical offsets are arbitrary. Curves vary across participants, with some showing S-shapes and crossovers, others resembling logarithmic trends.

of absolute perceptual magnitude. The vertical offset in the plots has no inherent meaning and cannot be interpreted as such due to the ambiguity introduced by the integration constants. Because integration always yields a function defined up to an unknown additive term, the absolute positioning of each curve is arbitrary. Only the relative shape and curvature of each function carry valid information about perceptual structure. Any apparent differences in vertical alignment or crossover behavior are analytically irrelevant and should not be mistaken for evidence of stronger or weaker perceptual responses across contexts.

Looking at the sensitivity functions, it is noticeable that almost all participants follow the expected Weberian form, characterized by a monotonic decrease across the luminance range. The only exception is participant JF00- instead of a monotonic trend, JF00's sensitivity function follows a U-shape, reaching a minimum around normalized luminance of 0.4, before increasing again. This behavior resulted in a highly S-shaped intensity function, unmatched by any other participant.

Figure 3.5 shows the perceptual scales obtained from the MLCM experiment by Vincent et al. (2024). Due to the inherent ambiguity in the analytically derived intensity functions $I(s)$, specifically the unknown constants of integration, these MLCM scales were used as reference data for fitting. The fitting procedure adjusted a shared proportionality factor g and two independent vertical offsets (integration constants C_w and C_b) corresponding to the black and white context conditions. This alignment does not alter the overall shape of the functions, it only scales them uniformly and shifts them vertically, resolving the ambiguity introduced by the indefinite integral.

For each participant, the values of g and the integration constants were optimized to minimize the RMSE between the predicted intensity and the corresponding MLCM scale. This step is necessary because the MLCM-derived perceptual scales differ across observers- there is no single averaged scale-requiring separate fits for each pairing. In total, 72 fits were computed, covering all combinations of ME participants, contexts, and MLCM observers. RMSE values for all fits are reported in Table 5.3.

For improved visualization, the RMSE values are presented as a heatmap (Figure 3.6). The heatmap allows detailed comparison of fit quality across all participants and observers, highlighting patterns in the consistency of the fits. Inspection of the heatmap reveals systematic differences in fit quality across participants. Participants ND99 (average RMSE = 5.0%), MA97 (5.9%), LK99 (6.2%), and JS00 (6.3%) achieved the best fits, indicating strong correspondence between their intensity functions and the MLCM-derived perceptual scales. Participants OTC99 (6.9%) and ST05 (6.9%) showed moderate fit quality. In contrast, participants JF00 (7.9%), JH02 (8.0%), and GC99 (8.8%) exhibited the highest average RMSE values, suggesting greater deviations between the two measures for these individuals.

The average RMSE across all fits is 0.069, corresponding to an average deviation of 6.9%. The median RMSE is 0.065, which corresponds to deviation of 6.5%. The best fit achieved an RMSE of 0.036 (3.6% deviation), while the worst fit had an RMSE of 0.132 (13.2% deviation). It is noteworthy that vast majority of RMSE values (66 out of 72 $\approx$ 91%) fall below the 10% deviation mark.
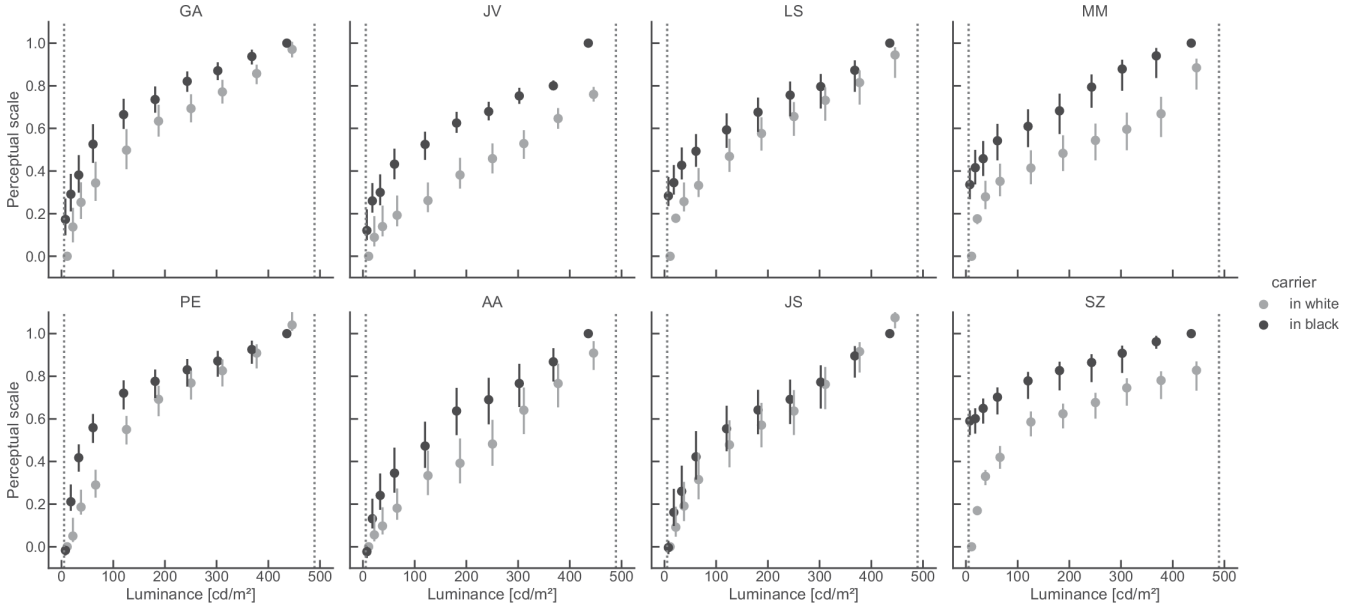
Figure 3.5: **Perceptual scales obtained via MLCM from Vincent et al. (2024).** Each plot shows the perceptual scale values derived from Maximum Likelihood Conjoint Measurement (MLCM) for one observer. The x-axis represents physical luminance in cd/m^2, while the y-axis shows the corresponding perceptual scale, normalized between 0 and 1. Black and gray points indicate targets embedded in black and white grating contexts, respectively. Error bars denote 95% confidence intervals. Most observers exhibit a consistent separation between contexts, with targets in black appearing perceptually brighter than those in white at equal luminance levels.
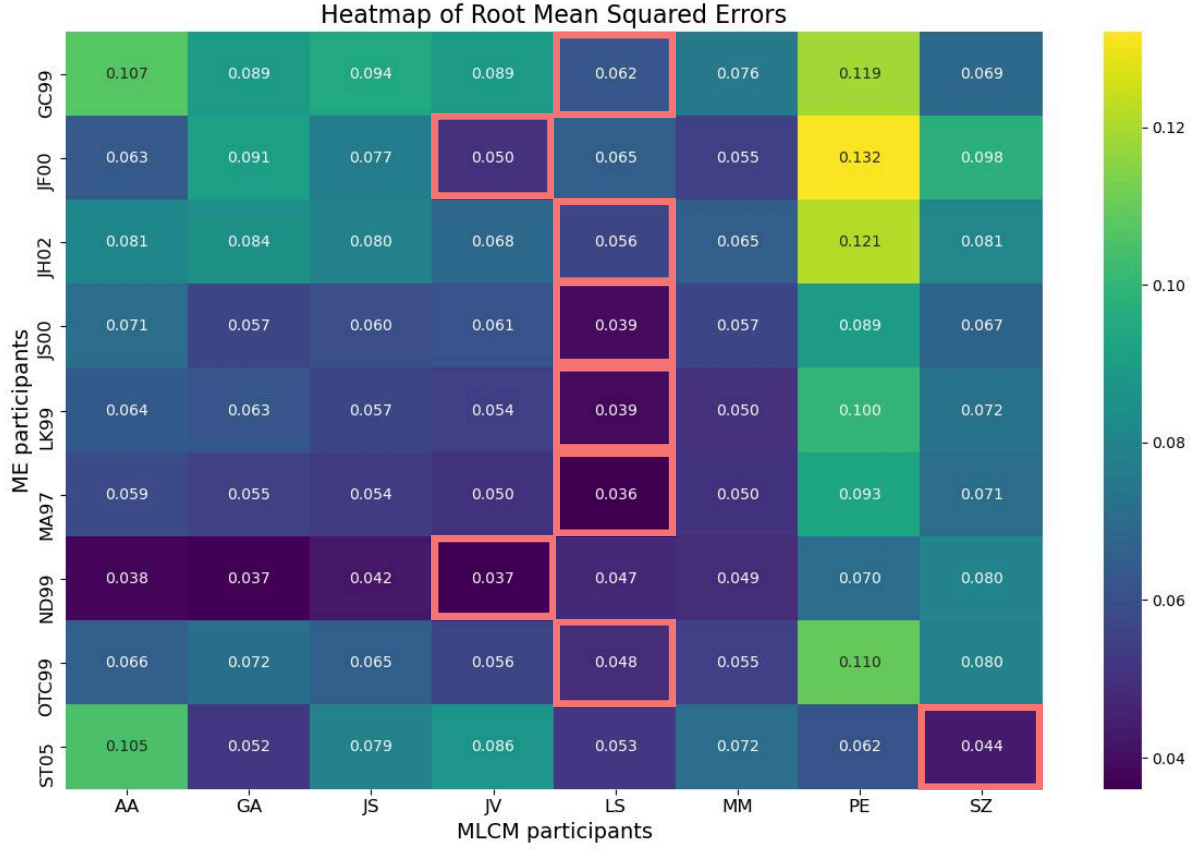
Figure 3.6: **Heatmap of root mean squared errors (RMSE) between ME-derived and MLCM-derived intensity functions.** Each cell represents the RMSE between the intensity function of an ME participant and the perceptual scale of an MLCM participant. Darker shade of a cell indicates lower error, reflecting stronger alignment and greater similarity in perceptual scaling across methods. Rows correspond to ME participants, columns to MLCM participants. The average RMSE across all comparisons is 0.069, with most values falling below 0.10. Red-outlined cells mark the lowest RMSE observed for each ME participant. Notably, many ME participants' intensity functions aligned best with the perceptual scale of observer LS from Vincent et al. (2024), suggesting that LS's scaling closely reflects the average structure observed across the ME data.
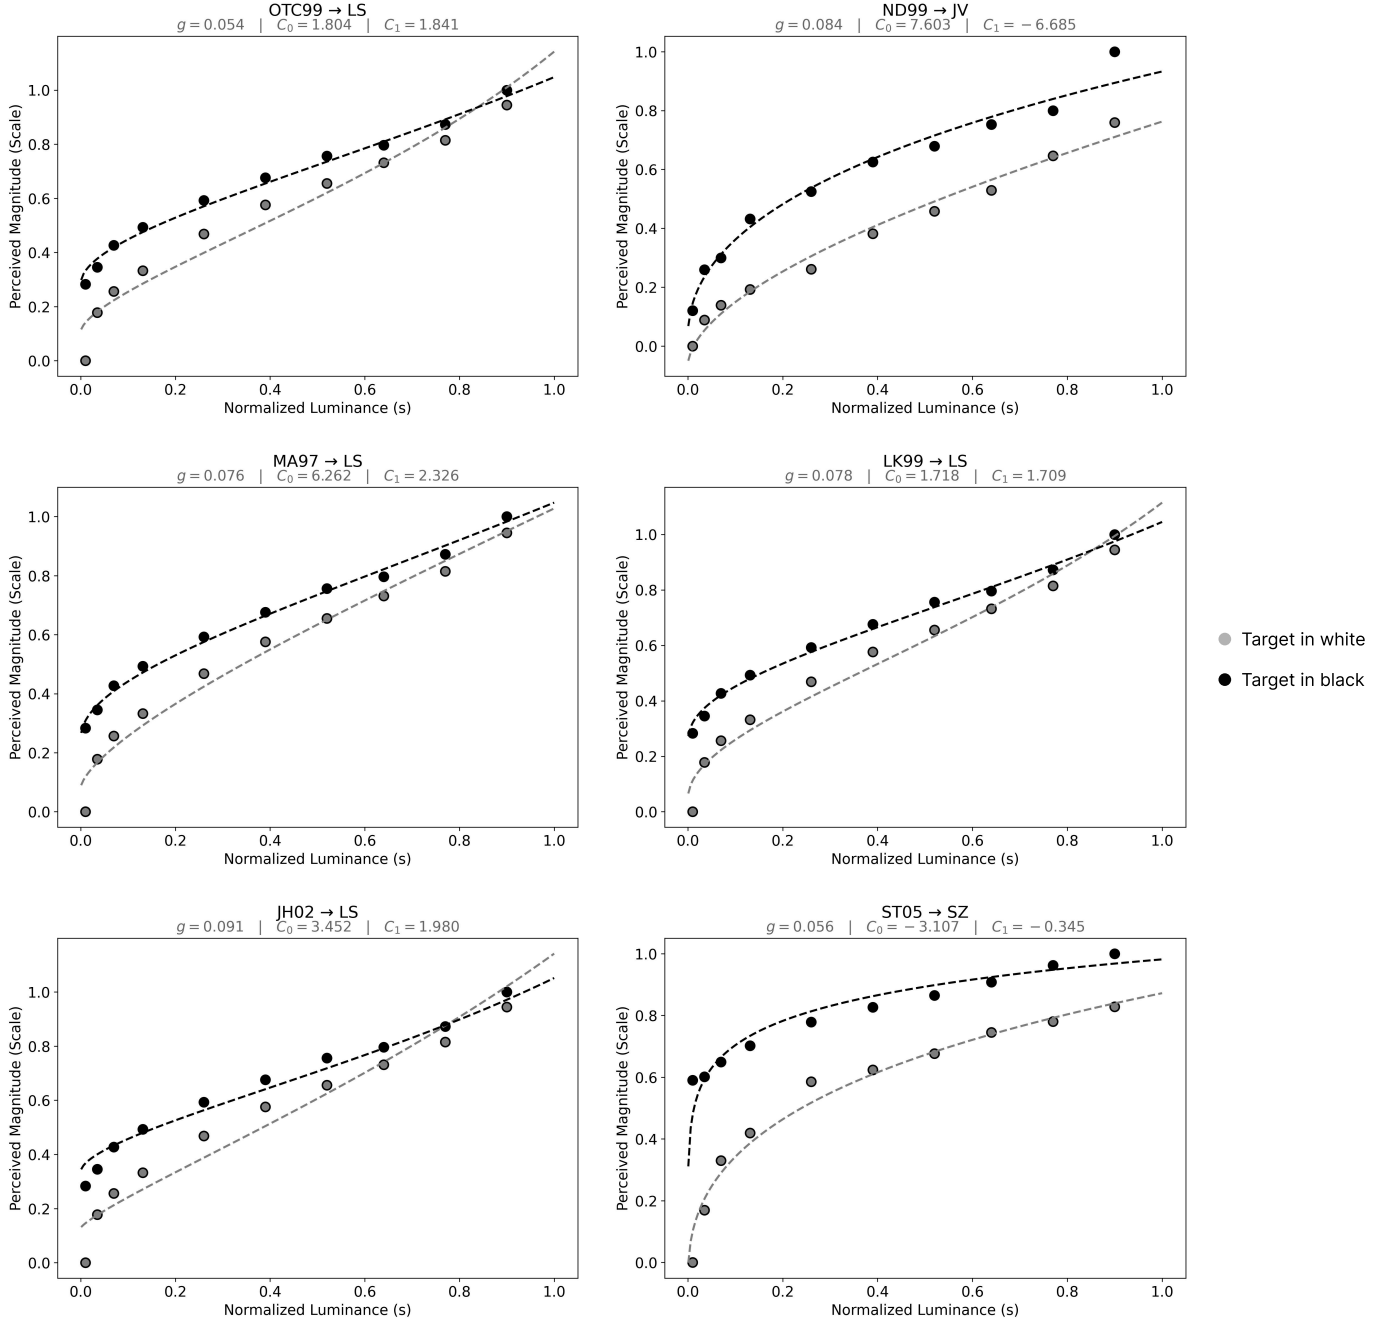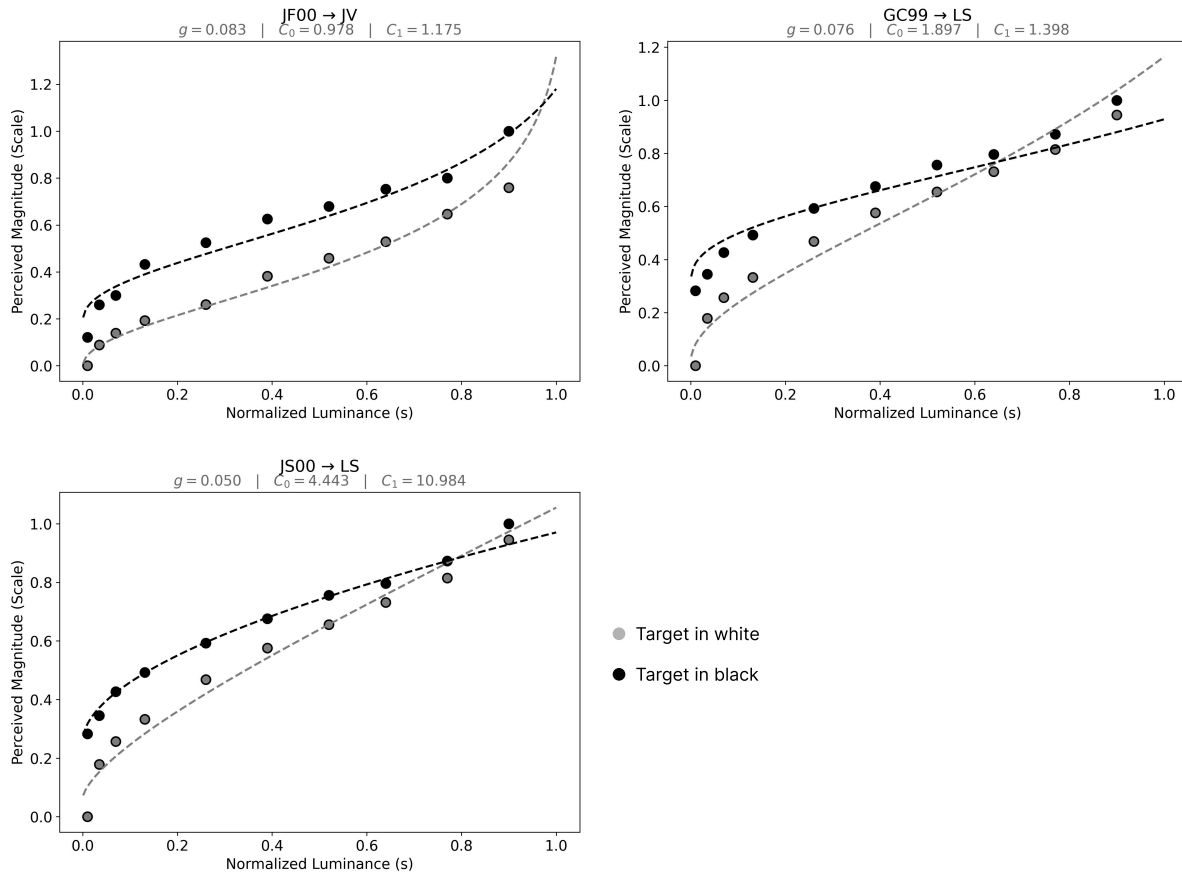
Figure 3.7: **Best fits per participant.** Each subplot shows the intensity function (dotted lines) fitted to a perceptual scale from Vincent et al. (2024). Points represent the MLCM data, while the annotated values indicate the fitted scaling factor g and integration constants C_w (target on white) and C_b (target on black).

4 Discussion

This thesis tested whether the unified framework suggested by Zhou et al. (2024), linking perceived magnitude and discriminability via a common internal representation, generalizes to context-dependent visual illusions. Specifically, the framework was applied to White's illusion, where target patches with identical luminances perceived to be either brighter or darker, according to their spatial context. Perceived brightness ratings were collected through a magnitude estimation task across two conditions, targets embedded in black or white gratings, over a range of 18 different luminance levels. From these responses, mean response and standard deviation functions were derived, which yielded Fisher sensitivity, and upon integration, the intensity function. The resulted intensity functions were then fitted to the MLCM data from Vincent et al. (2024), resulting 72 different fits. Then the goodness of the fits was measured through calculating their root mean squared errors (RMSE).

The results showed that while the ME and MLCM experiment by Vincent et al. (2024) used different participants, the fitted intensity function manage to produce rather low RMSE scores for majority of the participants, with average of 6.9% deviation, which suggests the framework is robust enough to capture essential features of context-dependent visual phenomena.

4.1 Reversal of contextual modulation across luminance

A key empirical finding was the reversal of contextual modulation across the luminance scale, as seen in Figure 3.1. At low luminance levels, targets embedded in black grating appeared brighter than targets embedded in white grating, consistent with the expected assimilation effect of White's illusion. However, the assimilation effect reversed and turned into contrast effect around normalized luminance values between 0.6 and 0.8 for most participants, with some (such as LK99 and JS00) exhibiting an earlier reversal. This suggests that contextual modulation is not constant across luminance levels but depends on absolute stimulus intensity.

A possible explanation can be found in Betz, Shapley, Wichmann, and Maertens (2015). The paper demonstrated that White's illusion is primarily driven by the contrast at the orthogonal edges- the immediate boundaries between the target and its surrounding gratings. Strong contrast at these edges tends to reduce the perceived brightness of the target (contrast effect), while weak contrast favors assimilation- this is demonstrated in Figure 4.1. Targets on black grating experienced a strong luminance step at their boundary, promoting contrast-related suppression of their perceived brightness. Targets on white grating however, had much smaller luminance differences with
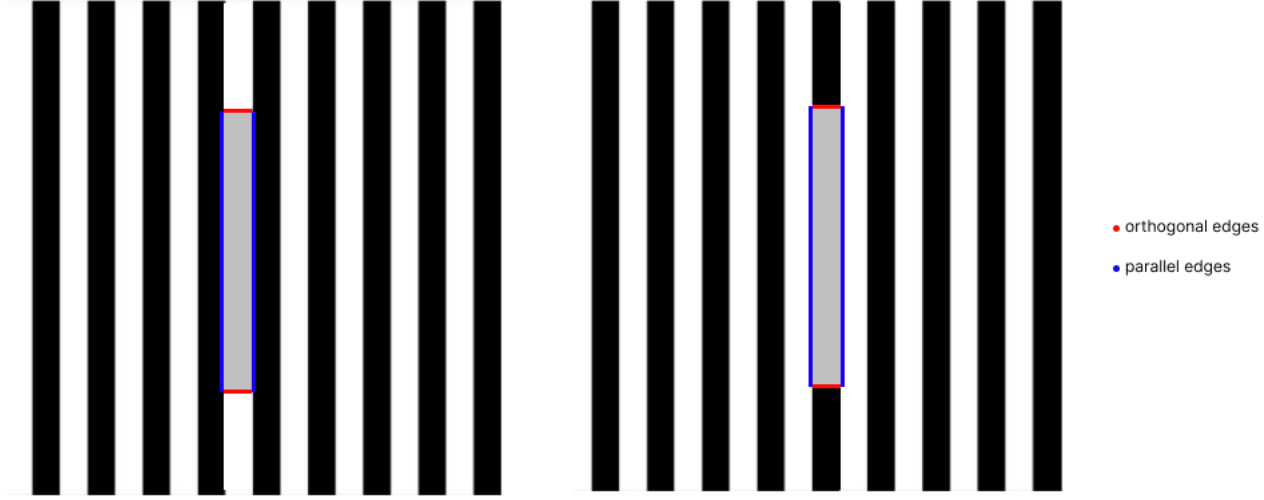
Figure 4.1: **Orthogonal edges drive assimilation effects in White's illusion.** Edges orthogonal to the grating (marked in red) have a strong impact on perceived brightness, as proposed by Betz et al. (2015), while edges parallel to the grating (marked in blue) exert minimal to no influence on the assimilation effect.

their surround. As a result, the usual direction of White's illusion reversed: targets on black grating perceived as darker than targets on white grating. These findings suggest that at high luminances, it is not absolute luminance alone, but the balance of local contrast across orthogonal edges, that determines whether contextual modulation favors assimilation or contrast effects. Figure 4.2, which was taken directly from the ME experiment, exemplifies this reversal.

4.2 Deviation from multiplicative noise assumption

The framework of Zhou et al. (2024) assumes that internal noise scales proportionally with the mean response- i.e., multiplicative noise. This implies a constant coefficient of variation across stimulus levels. However, this assumption was not supported by the collected data. For nearly all participants (except ST05 and ND99), standard deviation decreased with increasing luminance. Even for these two exceptions, the observed increase was minimal- both exhibited nearly flat noise profiles. This pattern contradicts the framework's core prediction that variability should grow with signal strength.

One explanation may lie in the structure of the magnitude estimation task. All luminance levels were repeated across five blocks in both contexts. High-luminance targets, being more visually salient, may have been more easily recalled or anchored to prior responses, thus reducing variability around extreme target luminances. In this view, standard deviation of mean responses reflects a combination of internal noise and
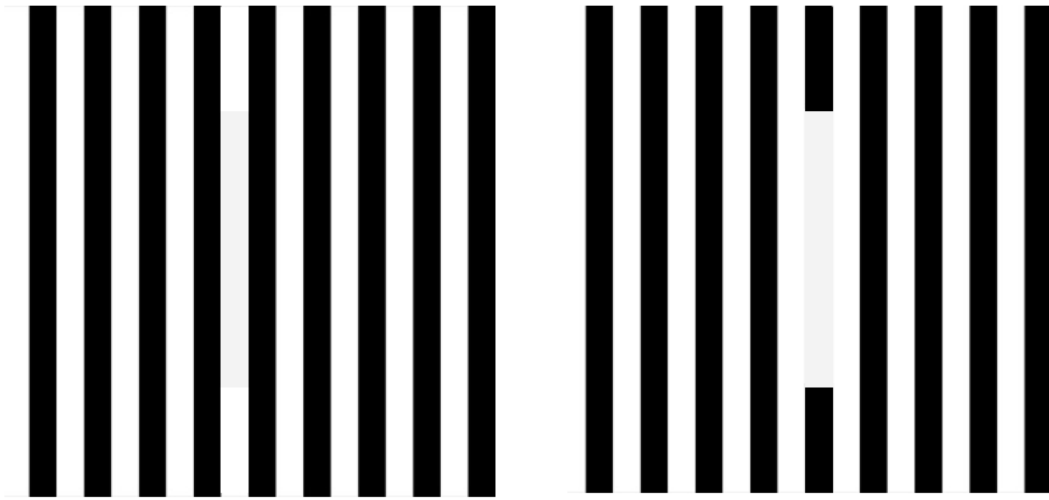
Figure 4.2: **Example of reversal of assimilation effect for high luminance targets.** The stimuli include target patches of identical, very high luminance, leading to a reversal of the typical assimilation pattern observed at lower luminances. Instead of the target appearing brighter when embedded in black grating (on the right), all the participants perceived the target embedded in a white grating (on the left) to be brighter, illustrating the luminance-dependent nature of contextual modulation.

task-specific artifacts such as anchoring, rather than noise in a strict perceptual sense. This interpretation assumes that standard deviation approximates internal noise, a key premise of the framework, but one that may not always hold under magnitude estimation tasks involving context effects.

Alternatively, the failure of the multiplicative noise model may reflect a deeper limitation of the framework itself when applied to context-dependent phenomena such as assimilation. In illusions like White's, the perceived brightness is shaped by both the surrounding context of the target and the physical brightness of the target itself (seen in the reversal of the assimilation effect in higher target luminances), leading to modulations that cannot be attributed to stimulus strength alone. These contextual dependencies introduce additional complexity into the mean responses that may also influence response variability, which the framework by Zhou et al. (2024) interprets as noise. In this setting, internal noise no longer scales predictably with signal strength, as assumed by the framework. Instead, the interaction between target and surround introduces latent variables that potentially influences the relationship between mean response and standard deviation, violating the framework's core premise of proportional noise tied directly to the stimulus.

To explore this, a LOWESS analysis was applied to the standard deviation data (Figure 4.3). Unlike the global fit used in Equation 1.5, LOWESS is a non-parametric method that locally fits smooth trends without assuming any particular form. This revealed structures not captured by the parametric model. For several participants (e.g., MA97, JH02, LK99), standard deviation followed an inverted U-shape: low at luminance extremes and peaking around the mid-range. Others, such as ST05 and JF00, showed profiles more consistent with the original standard deviation fit, while the rest of the participants exhibited mixed patterns across contexts. These findings suggest that internal noise may follow non-monotonic structures in contextual illusions, possibly reflecting a mixture of perceptual and task-related factors. While the framework treats standard deviation as a proxy for internal noise, the current data highlight that this approximation is sensitive to experimental context and should be interpreted with caution.

In most participants, the LOWESS analysis revealed an inverted U-shaped noise profile, with elevated variability in the mid-range that was not captured by the fitted model. This implies that mid-range variability was systematically underestimated, resulting in inflated sensitivity estimates in that region. This causes the integrated intensity function to rise too steeply around mid-luminance, potentially exaggerating perceptual differences and shifting the apparent strength of contextual modulation. Yet, due to the flexibility of the integration step, these distortions were absorbed into the fit, leaving the direction and magnitude of bias difficult to recover from the final result.

4.3 Crispening-like patterns in noise- isolated effect or anchoring?

The vast majority of participants exhibited the expected sensitivity trend predicted by Weber's law. A single outlier was participant JF00, whose sensitivity function followed
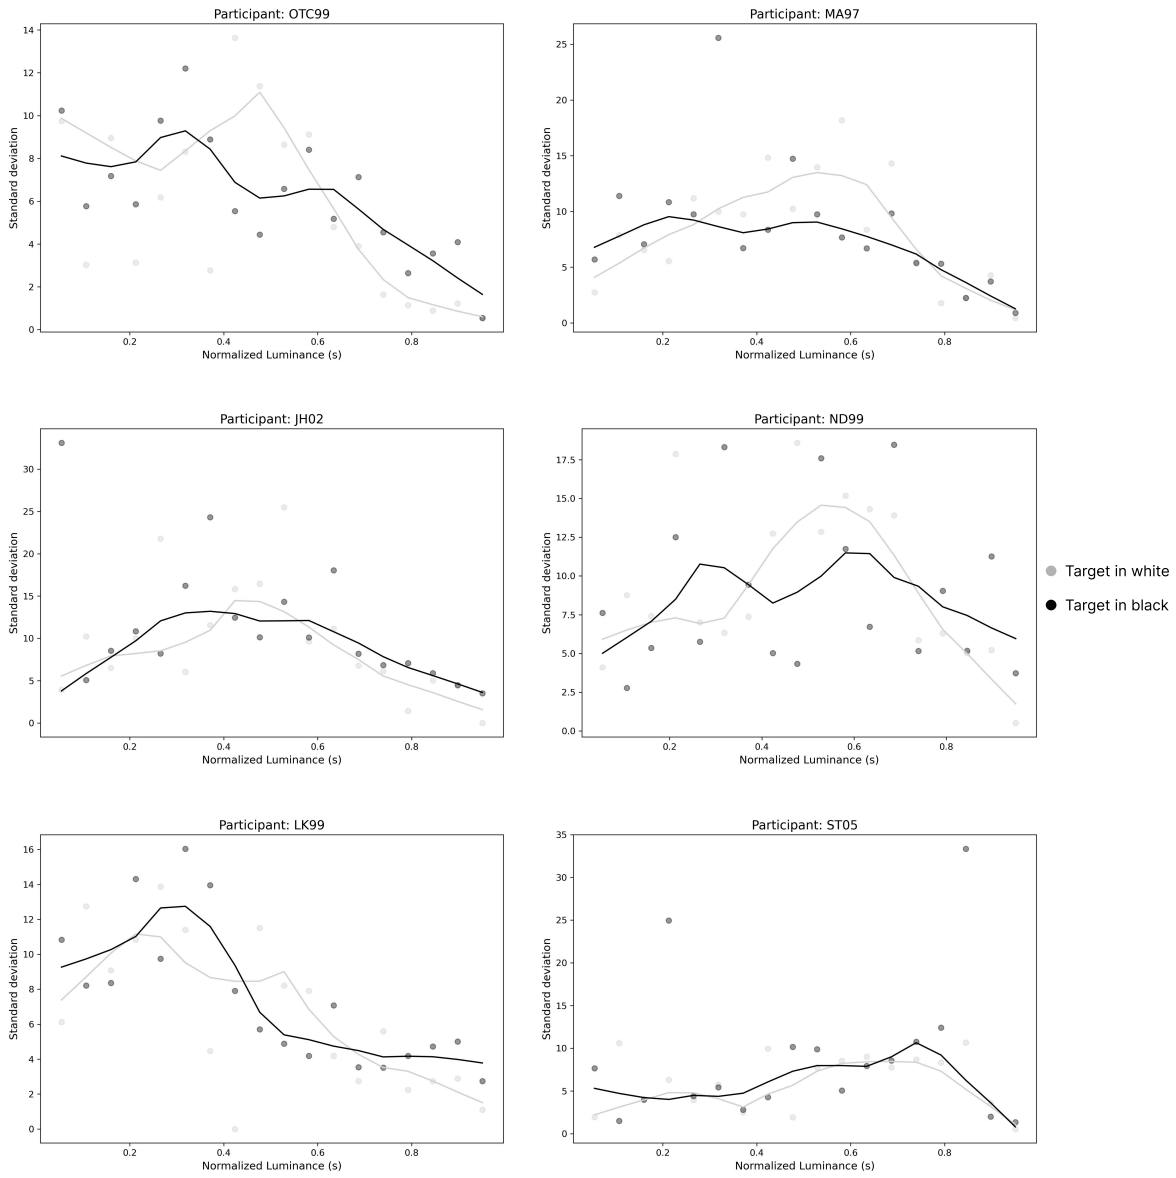
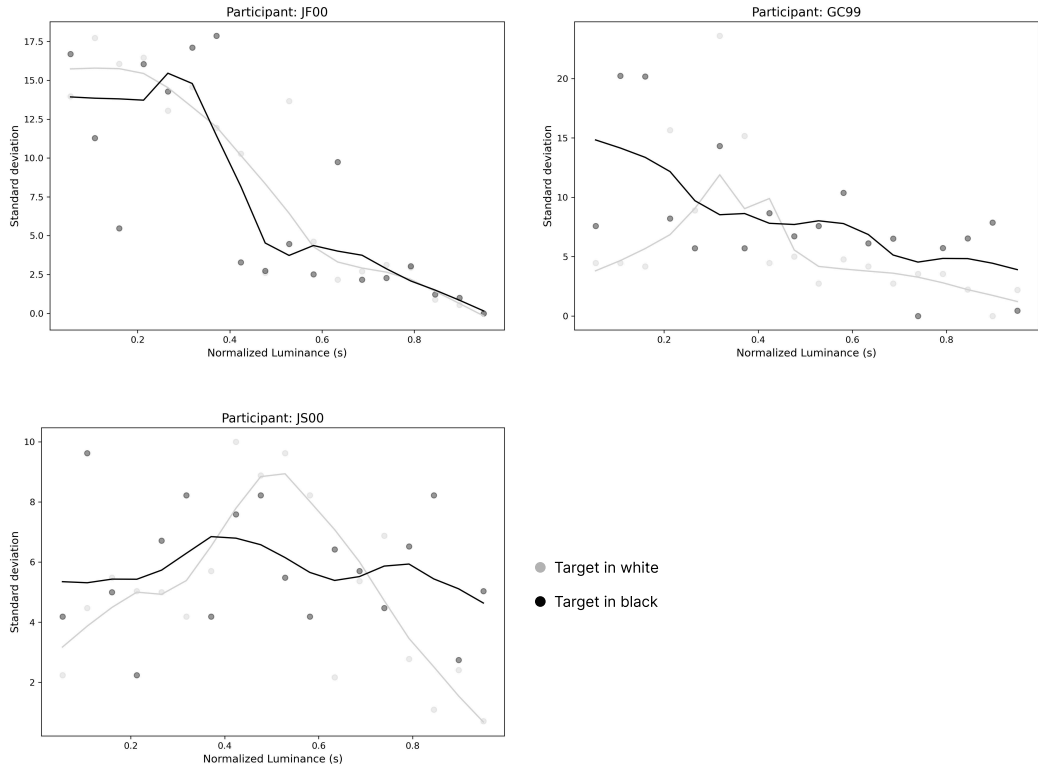Figure 4.3: **LOWESS fits to standard deviation functions across participants.** Each curve shows the smoothed trend of response variability as a function of luminance, separately for black and white context conditions. Many participants exhibit an inverted U-shape, with lower variability at luminance extremes and a peak in the mid-range. This trend suggests reduced internal noise when target luminance matches the background.
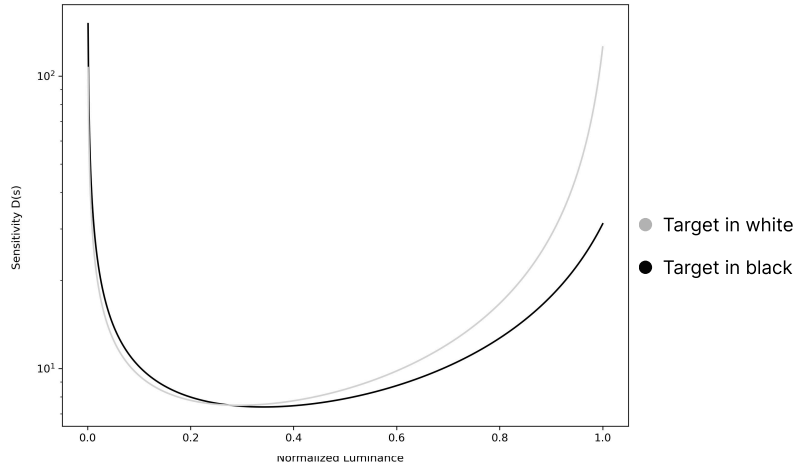
Figure 4.4: **Sensitivity functions for participant JF00 alignment with crispening effect.** The black and gray curves show the sensitivity for targets embedded in black and white grating contexts, respectively. Unlike the typical monotonic decrease observed in most participants, JF00 exhibits a U-shaped profile with a minimum around mid-luminance and elevated sensitivity at both low and high luminance levels. This unusual pattern aligns with the crispening effect, where discriminability peaks when the target's luminance closely matches the background, and falls off in intermediate regions.

a U-shaped profile, reaching its minimum around a normalized luminance of 0.4 and rising again thereafter (see Figure 4.4). This atypical pattern recalls the crispening effect described by Whittle (1992), which refers to enhanced sensitivity for brightness differences when a target's luminance closely matches its background. Whittle showed that small deviations from the background luminance are perceived more acutely than larger deviations, which become less discriminable. In JF00's case, sensitivity was highest at luminance levels where the target closely resembled the background grating- a very dark patch embedded in black grating, and a very bright patch embedded in white grating, dropping in the mid-range where contrast with the surround was less defined.

Interestingly, the inverted U-shaped noise profile found in the LOWESS analysis aligns with this interpretation. Across majority of participants, standard deviation was lower at the luminance extremes and peaked at intermediate levels, suggesting that internal noise might be minimized when target and surround luminance are similar.

It is important to note that the LOWESS curves were not used solely as a non-parametric smoothing tool to reveal structure that may be too subtle to detect by eye in the raw data. LOWESS makes no assumptions about functional form and should not be interpreted as an analytical fit. Instead, it offers a qualitative view of trends in the data- such as the inverted U-shape.

However, since only one participant showed a sensitivity function that aligns with crispening effect(JF00), it is rather unlikely to assume that crispening is the underlying explanation for the group-wide noise pattern. A more plausible account is anchoring: participants were repeatedly exposed to the same luminance values, and extreme values

are easier to recognize and recall, leading to more consistent responses. In contrast, mid-range luminances are harder to identify precisely, resulting in greater response variability. Therefore, while the LOWESS noise trend matches the structure predicted by the crispening effect, it more likely reflects task-induced anchoring than genuine perceptual enhancement.

4.4 Alignment between ME-derived and MLCM-derived intensity

The comparison between intensity functions derived from the ME experiment and perceptual scales obtained through MLCM showed strong overall consistency. The average RMSE across all fits was 6.9%, with a median of 6.5%. In 91% of cases (66 out of 72), the deviation remained below 10%. Given that the ME and MLCM data were collected from different observers, this level of correspondence suggests a high degree of alignment.

This apparent success may be partly explained by the smoothing effect of the integration step: it reduces the influence of local deviations in sensitivity, making the resulting intensity functions more robust to mismatches in the noise model. In other words, even if the noise functions deviate from theoretical assumptions, a well-fitted mean response can still yield integrated intensity curves that resemble perceptual data.

However, these results must be interpreted with caution. Both the integration constants (C_w, C_b) and the scaling factor (g), were treated as free parameters and adjusted to minimize RMSE for each fit. These parameters cannot be derived directly from the ME data alone due to inherent mathematical ambiguities in the integration step. The 6.9% average error therefore reflects a best-case fit under these constraints. While this does not show that the framework independently reproduces the MLCM scales, it does indicate that, once appropriately scaled and shifted, it can approximate them closely. This suggests a potential for high alignment, even under contextual modulation.

Nonetheless, the alignment remains noteworthy. Despite known biases in magnitude estimation such as anchoring, and despite the deviations observed in the noise fits, the resulting intensity functions still captured key features of the perceptual structure. This includes the overall shape of the curves and the direction of contextual modulation. The vertical separation between conditions, however, was directly influenced by the fitted integration constants and cannot be considered an independent result. Still, the findings offer partial support for the idea that perceived magnitude and sensitivity reflect a shared internal representation, even in context-dependent cases like White's illusion.

RMSE values also varied across participants, suggesting individual differences in how well the framework generalizes. ND99, MA97, LK99, and JS00 had the lowest RMSEs on average (5.0%–6.3%), while JF00, JH02, and GC99 showed the highest average deviations (7.9%–8.8%).

4.5 Limitations

The framework suggested by Zhou et al. (2024) relies on a few mathematical assumptions. It claims that noise grows proportionally with the mean response- a logical assumption, considering that no stimulus should produce near-zero noise in the internal representation of an observer. However, this assumption was contradicted by the collected noise data, where noise decreased on average while perceived intensity grew.

In addition, although the integration process introduces an unknown constant C (as shown in Equation 2.6), this constant cannot be determined directly from the ME data. Instead, C was estimated by fitting the integrated sensitivity functions to the perceptual scales obtained from MLCM. This fitting provides a value of C that optimizes the correspondence between the analytically derived intensity functions and observed perceptual scales, but it may not reflect the true internal mapping. The ability to flexibly fit C in the current analysis allowed better alignment with such context-dependent differences, but it does not resolve the fundamental ambiguity introduced by the integration step.

A further limitation is that the ME and MLCM measurements were obtained from different participant groups. Since perceptual encoding can vary across individuals, comparisons between the two methods inevitably introduce ambiguity. A more conclusive test would require collecting ME and MLCM data from the same observers, allowing a direct within-subject comparison of the derived intensity functions.

Perceptual calibration likely occurred during the initial trials, as participants familiarized themselves with the task and internalized the stimulus scale. This may have led to imprecise or inconsistent responses at the beginning of the experiment, introducing noise unrelated to the perceptual process of interest. Discarding the first trial or incorporating a practice phase could help reduce this type of variability.

A known issue of the ME procedure is anchoring. Participants may consciously or unconsciously base their ratings on previous responses, rather than evaluating each stimulus independently, also known as anchoring effect. One strategy to minimize anchoring effects is to spread trials over multiple days, allowing participants to forget prior anchors and rely more on their immediate perceptual judgments. Another approach is to vary the stimulus set across blocks: instead of repeating the same 18 luminance levels across all blocks, multiple versions of the experiment could be created with differently spaced luminance values. This would reduce the ability to memorize specific stimulus-response mappings and encourage genuine estimation on each trial. Additionally, the participant sample was not balanced with respect to gender. Notably, the two female participants were the only ones to exhibit positive slopes in their standard deviation functions, suggesting a possible demographic influence on internal noise properties. A more balanced participant pool would be necessary to systematically investigate such effects.

4.6 Conclusion

This thesis tested whether the unified framework proposed by Zhou et al. (2024), which links perceived magnitude and discriminability through a shared internal representation,

extends to context-dependent phenomena like White's illusion. In this framework, perceived magnitude corresponds to the mean internal response, while sensitivity is defined as its derivative divided by internal noise.

Magnitude estimation and standard deviation data were used to compute sensitivity functions, which then were integrated in order to achieve intensity functions. The intensity functions were then fitted to perceptual scales obtained via MLCM by Vincent et al. (2024), using the unknown integration constant and a scaling constant as free parameters. Despite relying on different participant groups, the alignment between the two was strong: 91% of the 72 fits had an RMSE below 10%, with a mean deviation of 6.9%. This level of consistency suggests that the framework can approximate perceptual structure even in context effects like White's illusion.

However, a key assumption of the model- that internal noise scales proportionally with the mean response, was not supported by the data. Parametric fits often showed decreasing noise with increasing luminance, and LOWESS smoothing revealed inverted U-shaped noise profiles in many participants. While this resembles the crispening effect, it is more likely explained by anchoring: extreme luminance values are easier to recognize and recall, leading to more stable responses at the luminance extremes.

This conflict can be reconciled by the smoothing effect of the integration step (Equation 2.6). Although the fitted noise functions systematically deviated from the framework's multiplicative assumption, showing a mixed pattern of multiplicative and additive components, the resulting intensity functions still closely matched the MLCM perceptual scales. This suggests that accurate modeling of the mean response is sufficient to recover the global shape of perceived intensity, even if the local noise structure is inaccurate. In this sense, the framework demonstrates robustness: it can capture large-scale perceptual trends even when its assumptions about noise do not fully hold.

Future work should test the framework using ME and MLCM data from the same participant pool. This would eliminate between-subject variability and allow a direct comparison of intensity functions and perceptual scales. Future studies should also minimize anchoring effects to determine whether the inverted U-shape in noise reflects a systematic property of how brightness is encoded under contextual modulation or whether it is a byproduct of the task design.

5 Appendix

5.1 Driving the integrated sensitivity functions

Take the integral:

$$\int \frac{akx^{a-1}}{ckx^a + b}\, dx \tag{5.1}$$

This integral can be solved by substitution. Let:

$$u = ckx^a + b \tag{5.2}$$

Then the derivative of u with respect to x is

$$\frac{du}{dx} = ackx^{a-1} \tag{5.3}$$

Solving for dx:

$$dx = \frac{1}{ackx^{a-1}}\, du \tag{5.4}$$

Substituting into the original integral:

$$\int \frac{akx^{a-1}}{u} \cdot \frac{1}{ackx^{a-1}}\, du = \int \frac{1}{cu}\, du \tag{5.5}$$

$$= \frac{1}{c}\log(u) + C \tag{5.6}$$

Substituting back for u, the final solution:

$$\int \frac{akx^{a-1}}{ckx^a + b}\, dx = \frac{1}{c}\log(\alpha(ckx^a + b)) + C \tag{5.7}$$

5.2 Mean response RMSE and bias

Table 5.1: **RMSE and Bias values for all participants across context conditions.** The table records root mean squared error (RMSE) and perceptual bias for targets embedded in white and black grating phases, enabling assessment of model fit and contextual modulation for each participant.

Participant	RMSE		Bias	
	in white	in black	in white	in black
OTC99	4.309	3.291	0.164	0.055
MA97	3.921	6.056	-0.233	-0.184
JH02	5.285	5.495	0.268	0.098
ND99	3.253	3.856	0.102	0.091
LK99	4.297	4.424	0.026	-0.084
ST05	4.070	5.906	0.052	-0.018
JF00	3.959	5.762	-0.081	-0.084
GC99	5.211	4.856	-0.222	-0.057
JS00	2.696	4.588	0.233	0.222

5.3 Standard deviation RMSE and bias

Table 5.2: **RMSE and Bias values for all participants across context conditions.** The table records root mean squared error (RMSE) and perceptual bias for targets embedded in white and black grating phases, enabling assessment of model fit and contextual modulation for each participant.

Participant	RMSE		Bias	
	in white	in black	in white	in black
OTC99	3.473	2.162	0.00	0.00
MA97	4.844	4.890	0.00	0.00
JH02	6.258	6.323	0.00	0.00
ND99	4.949	4.918	0.00	0.00
LK99	3.259	2.916	0.00	0.00
ST05	3.134	7.942	0.00	0.00
JF00	2.996	4.328	0.00	0.00
GC99	5.341	4.135	0.00	0.00
JS00	2.719	1.966	0.00	0.00

5.4 Mean response and standard deviation residuals plots
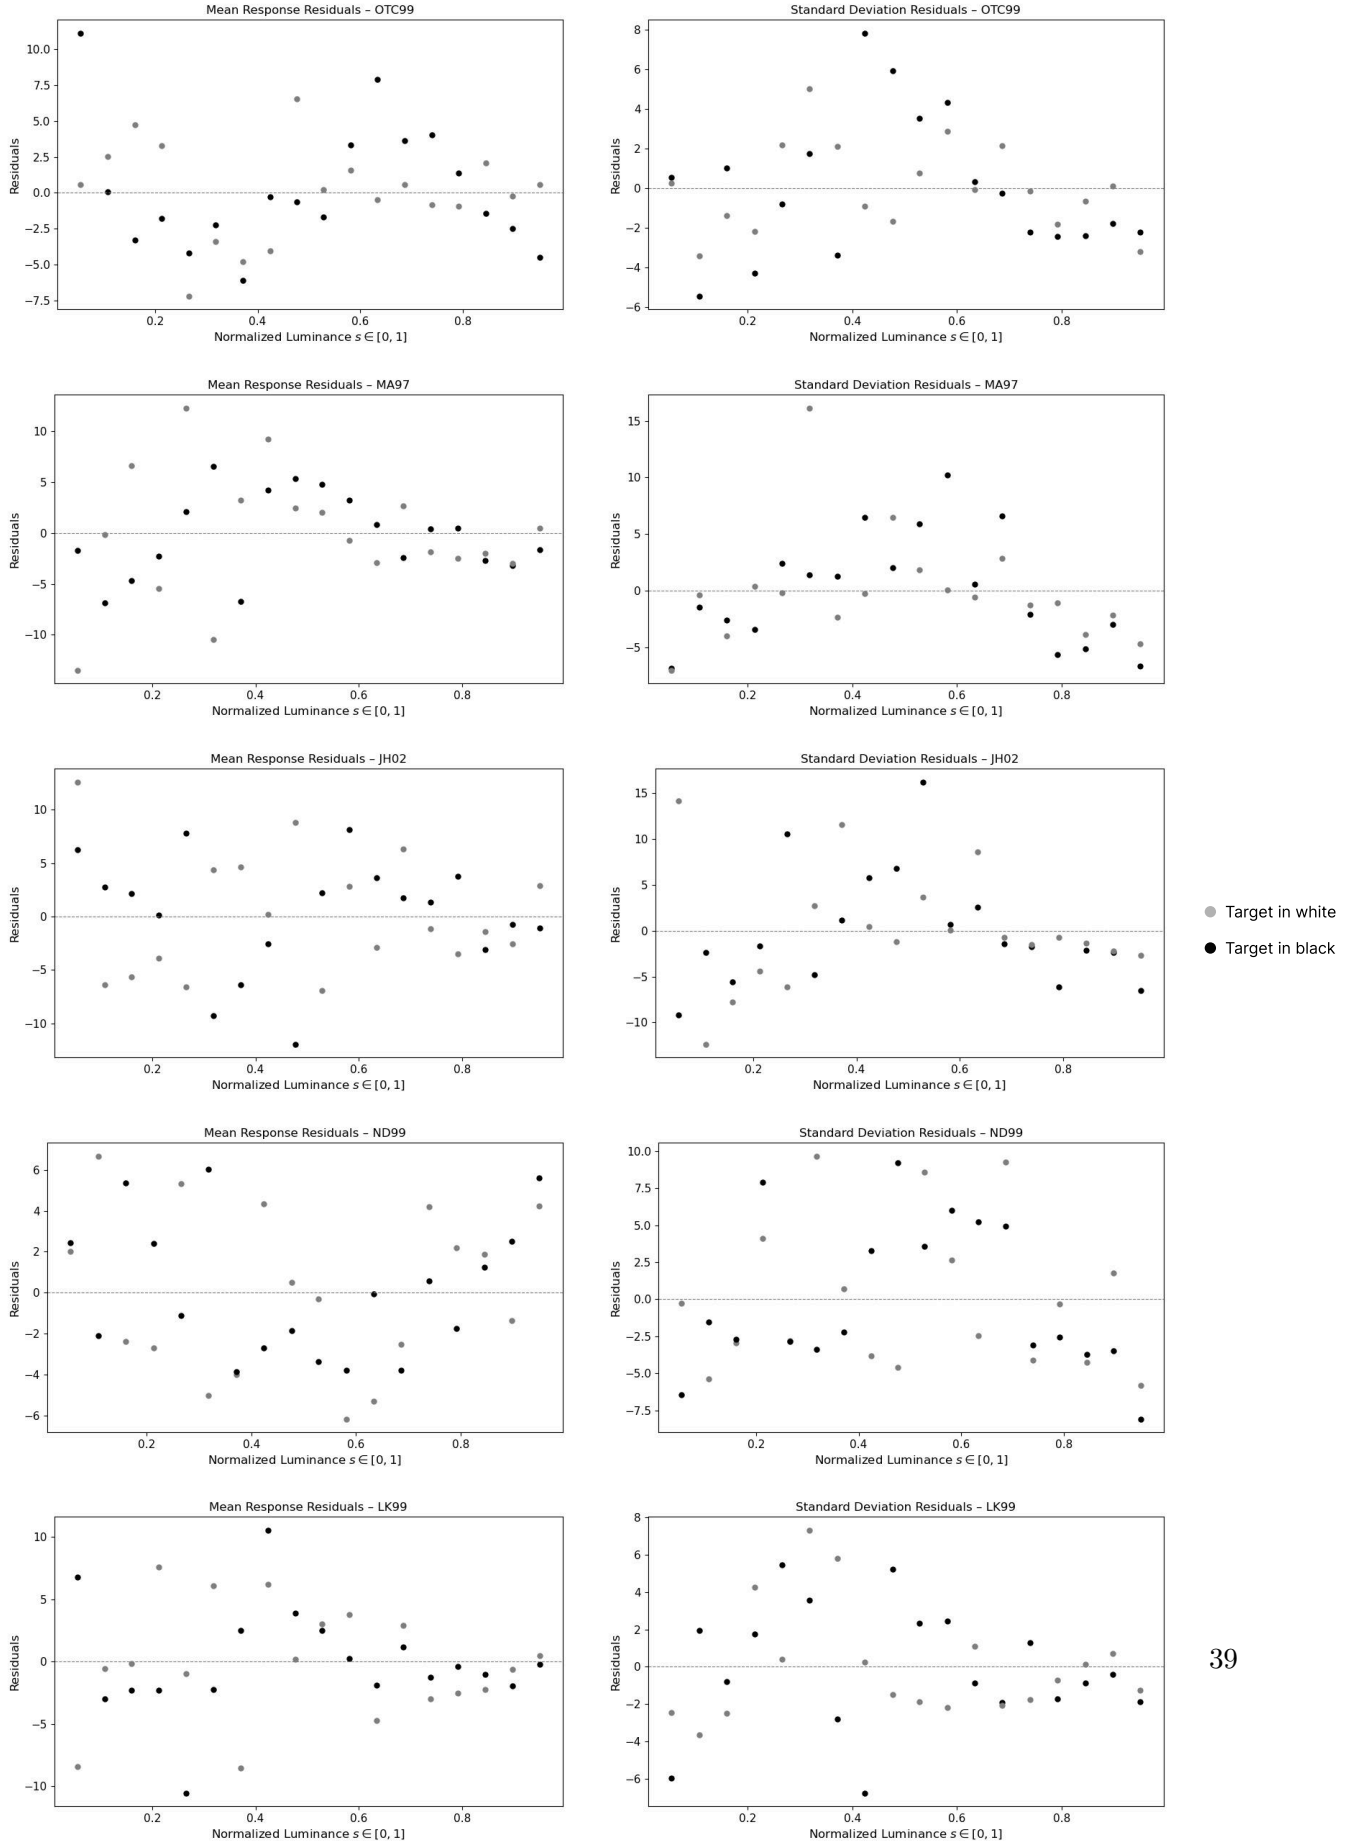
5 Appendix
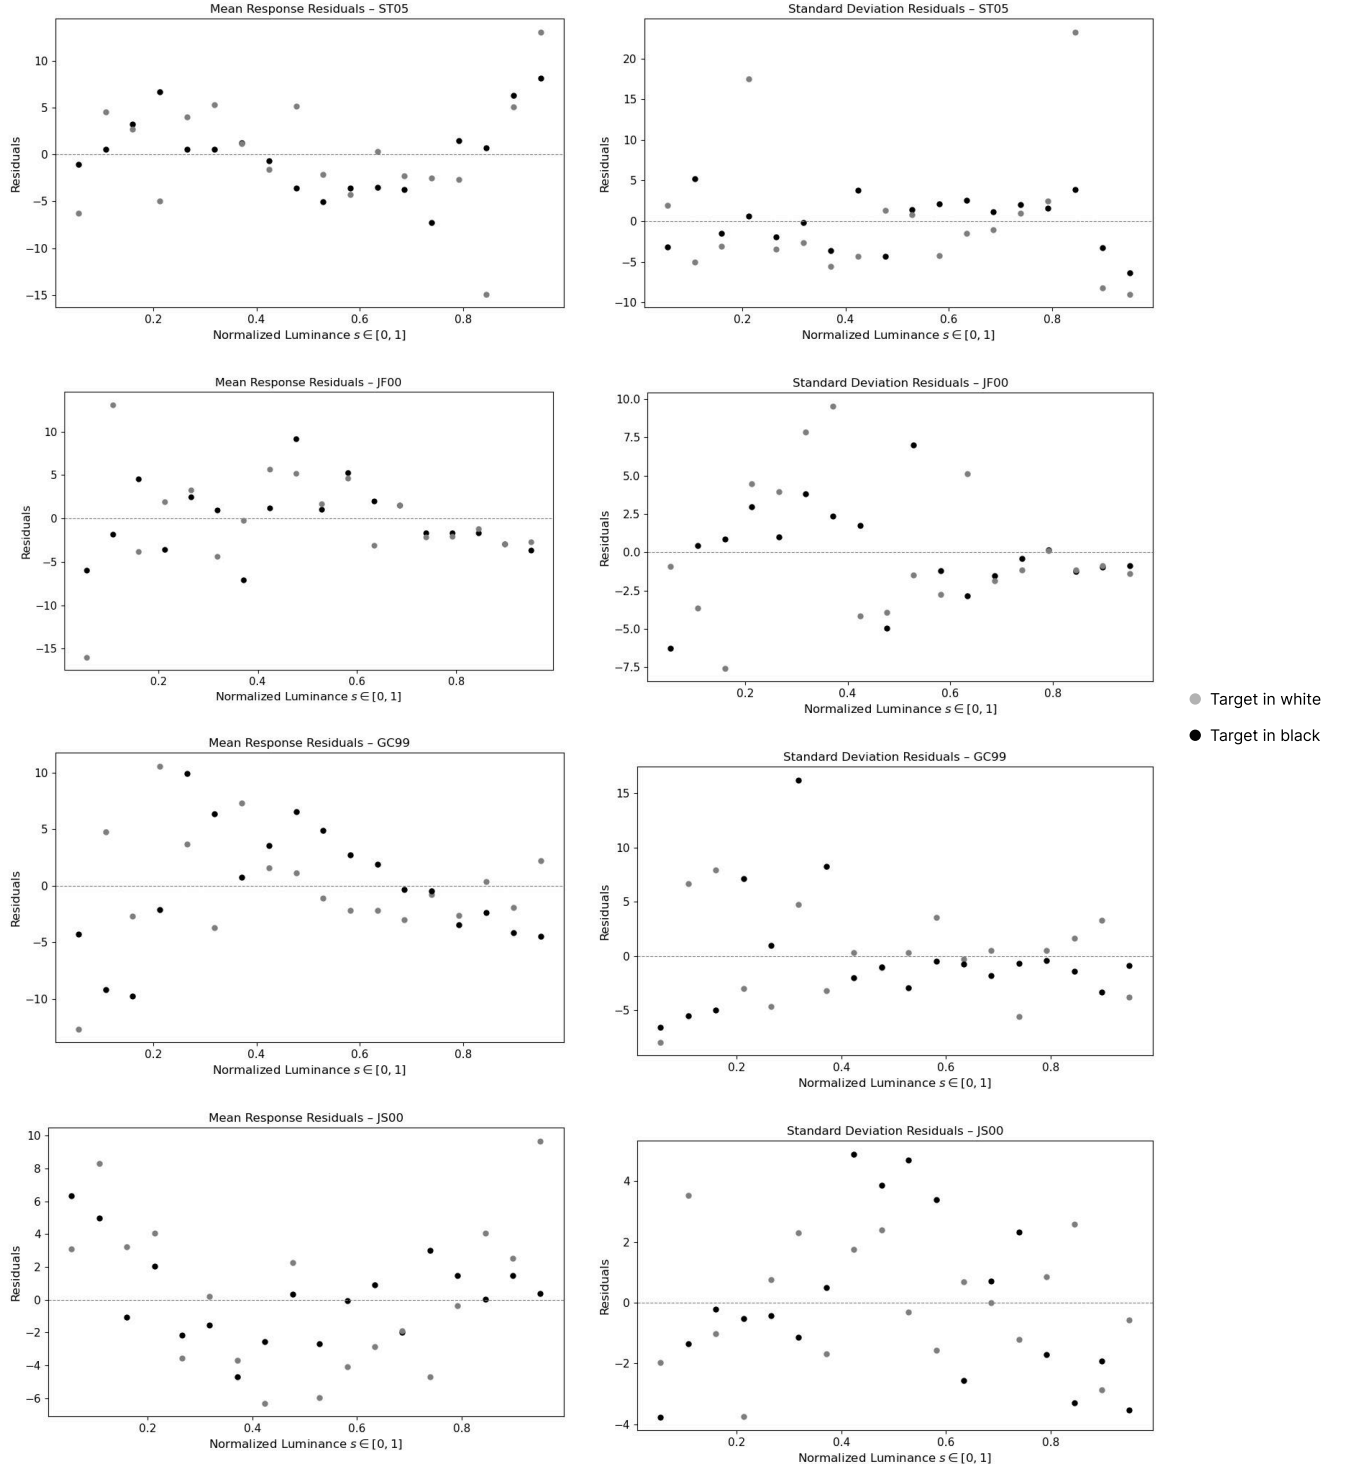
Figure 5.1: **Residuals of model fits for mean response and standard deviation functions across luminance levels and context conditions.** RMSE values varied between participants, with some showing larger errors in one context, but bias remained low throughout. For standard deviation fits, bias was zero in all cases, indicating symmetric residuals. Model fit quality differed across participants, with some showing good alignment and others higher residuals.

5.5 Detailed fit results- best to worst

ME participant	MLCM participant	g	C0	C1	RMSE
MA97	LS	0.076	6.262	2.326	0.036
ND99	JV	0.084	7.603	-6.685	0.037
ND99	GA	0.098	8.971	-7.788	0.037
ND99	AA	0.104	9.365	-8.438	0.038
LK99	LS	0.078	1.718	1.709	0.039
JS00	LS	0.050	4.443	10.984	0.039
ND99	JS	0.109	9.962	-8.889	0.042
ST05	SZ	0.056	-3.107	-0.345	0.044
ND99	LS	0.086	7.949	-6.806	0.047
OTC99	LS	0.054	1.804	1.841	0.048
ND99	MM	0.077	7.087	-5.961	0.049
MA97	JV	0.073	5.865	2.178	0.050
MA97	MM	0.067	5.537	2.170	0.050
JF00	JV	0.083	0.978	1.175	0.050
LK99	MM	0.069	1.522	1.626	0.050
ST05	GA	0.077	-4.482	-0.924	0.052
ST05	LS	0.068	-3.927	-0.760	0.053
MA97	JS	0.095	7.728	2.671	0.054
LK99	JV	0.075	1.514	1.586	0.054
JF00	MM	0.076	1.021	1.240	0.055
OTC99	MM	0.048	1.599	1.744	0.055
MA97	GA	0.085	6.941	2.536	0.055
OTC99	JV	0.052	1.601	1.717	0.056
JH02	LS	0.091	3.452	1.980	0.056
JS00	MM	0.044	3.911	9.778	0.057
JS00	GA	0.056	4.919	12.195	0.057
LK99	JS	0.098	2.035	1.898	0.057
MA97	AA	0.090	7.207	2.526	0.059
JS00	JS	0.063	5.443	13.503	0.060
JS00	JV	0.048	4.092	10.381	0.061
ST05	PE	0.089	-5.265	-1.184	0.062
GC99	LS	0.076	1.897	1.398	0.062
LK99	GA	0.087	1.873	1.845	0.063
JF00	AA	0.102	1.166	1.283	0.063
LK99	AA	0.092	1.829	1.794	0.064
JH02	MM	0.081	3.048	1.862	0.065
OTC99	JS	0.067	2.145	2.066	0.065
JF00	LS	0.084	1.142	1.263	0.065
OTC99	AA	0.064	1.938	1.957	0.066
JS00	SZ	0.039	3.620	8.934	0.067

JH02	JV	0.087	3.166	1.843	0.068
GC99	SZ	0.060	1.633	1.405	0.069
ND99	PE	0.112	10.206	-9.014	0.070
JS00	AA	0.059	5.024	12.686	0.071
MA97	SZ	0.059	5.007	2.114	0.071
ST05	MM	0.059	-3.414	-0.540	0.072
OTC99	GA	0.060	1.966	1.988	0.072
LK99	SZ	0.061	1.467	1.633	0.072
GC99	MM	0.066	1.662	1.341	0.076
JF00	JS	0.107	1.320	1.346	0.077
ST05	JS	0.085	-5.009	-1.186	0.079
OTC99	SZ	0.042	1.526	1.727	0.080
JH02	JS	0.114	4.197	2.234	0.080
ND99	SZ	0.066	6.265	-4.942	0.080
JH02	AA	0.108	3.873	2.113	0.081
JH02	SZ	0.071	2.806	1.839	0.081
JH02	GA	0.101	3.774	2.135	0.084
ST05	JV	0.063	-3.755	-0.734	0.086
GC99	GA	0.084	2.060	1.493	0.089
GC99	JV	0.071	1.650	1.270	0.089
JS00	PE	0.064	5.555	13.815	0.089
JF00	GA	0.093	1.229	1.344	0.091
MA97	PE	0.096	7.831	2.784	0.093
GC99	JS	0.094	2.241	1.499	0.094
JF00	SZ	0.063	1.001	1.267	0.098
LK99	PE	0.099	2.072	1.996	0.100
ST05	AA	0.078	-4.683	-1.073	0.105
GC99	AA	0.087	2.001	1.405	0.107
OTC99	PE	0.068	2.174	2.157	0.110
GC99	PE	0.095	2.293	1.603	0.119
JH02	PE	0.114	4.218	2.321	0.121
JF00	PE	0.105	1.332	1.421	0.132

References

Betz, T., Shapley, R., Wichmann, F. A., & Maertens, M. (2015). Testing the role of luminance edges in white's illusion with contour adaptation. *Journal of vision*, *15*(11), 14–14.

Fechner, G. T. (1860). *Elemente der psychophysik* (Vol. 2). Breitkopf u. Härtel.

Garcia-Marques, T., & Fernandes, A. (2023). Perceptual anchoring effects: Evidence of response bias and a change in estimates sensitivity. *Brain and Behavior*, *13*(11), e3254.

Gescheider, G. (1997). Psychophysics: the fundamentals 3rd ed.-psycnet. *Psychophysics: the fundamentals*.

Kingdom, F., & Prins, N. (2016). *Psychophysics: A practical introduction*. Academic Press. Retrieved from `https://books.google.de/books?id=3sHQBAAAQBAJ`

Knoblauch, K., Maloney, L. T., Knoblauch, K., & Maloney, L. T. (2012). Maximum likelihood conjoint measurement. *Modeling psychophysical data in R*, 229–256.

Krueger, L. E. (1989). Reconciling fechner and stevens: Toward a unified psychophysical law. *Behavioral and Brain Sciences*, *12*(2), 251–267.

Mertens, A., Mertens, U. K., & Lerche, V. (2021). On the difficulty to think in ratios: a methodological bias in stevens' magnitude estimation procedure. *Attention, Perception, & Psychophysics*, *83*, 2347–2365.

Sowden, P. T. (2012). Psychophysics.

Stevens, S. S. (1957). On the psychophysical law. *Psychological review*, *64*(3), 153.

Stevens, S. S., & Galanter, E. H. (1957). Ratio scales and category scales for a dozen perceptual continua. *Journal of experimental psychology*, *54*(6), 377.

Vincent, J., Maertens, M., & Aguilar, G. (2024). What fechner could not do: Separating perceptual encoding and decoding with difference scaling. *Journal of Vision*, *24*(5), 5–5.

White, M. (1979). A new effect of pattern on perceived lightness. *Perception*, *8*(4), 413–416.

Whittle, P. (1992). Brightness, discriminability and the "crispening effect". *Vision Research*, *32*(8), 1493–1507.

Zhou, J., Duong, L. R., & Simoncelli, E. P. (2024). A unified framework for perceived magnitude and discriminability of sensory stimuli. *Proceedings of the National Academy of Sciences*, *121*(25), e2312293121.