https://doi.org/10.2352/EI.2026.38.8.IQSP-248

Perceptual Thresholds of Facial Lighting Chromaticity in Virtual Environments

Xinmiao Zhang, Sofie R. Herbeck, Christopher A. Thorstenson
Rochester Institute of Technology; Rochester, NY, USA

Abstract

Lighting chromaticity plays a critical role in the visual perception of rendered content embedded within virtual scenes. Consequently, much effort is made to reduce differences in lighting chromaticity between these, to avoid an unnatural combined appearance. However, it is not currently known "how" different these lighting conditions can be before people can notice them, or before it becomes detrimental to visual appearance. In this study, three psychophysical tasks are employed to assess participants' perception of simulated lighting conditions applied separately to rendered virtual backgrounds and virtual objects, with a focus on stimuli comprising human faces. The tasks assessed the influence of object characteristics (variable skin tone), and differences in lighting chromaticity (between scenes and objects) on visual assessments for perceived lighting matches, mismatches, and preferences. Results revealed that both object and lighting characteristics significantly influenced each perceptual judgment, in different ways. Chromaticity matches, mismatches, and preferences for facial stimuli depended on the scene chromaticity and skin tone but their patterns varied across tasks. The current work can provide guidance for virtual rendering based on visual perception of simulated lighting differences.

Introduction

Virtual representations of human faces play a critical role in remote communication, healthcare, education, professional gatherings, entertainment, and social interactions. In such contexts, facial appearance should closely resemble that of real individuals, enabling users to engage in experiences that are perceptually indistinguishable from interactions with actual people [1]. Facial appearance is rich with cues about a person's identity, health, and emotional state, which is rapidly and efficiently processed by others' visual and cognitive systems [2, 3, 4]. People form immediate trait impressions from facial appearance, often within just 100 milliseconds, and such judgments exhibit strong cross-cultural consistency and developmental stability [5]. Therefore, faces serve as an important stimulus class in vision science and social cognition research [6].

Past work has shown that people are especially sensitive to facial chromatic variations: for example, redness and yellowness shifts are more readily detected in faces than in non-face stimuli, and even subtle color changes can alter social judgments [7, 8, 4]. Further, memory for facial color has a significant impact on color perception [9], and preferred skin colors in digital imagery tend to be more saturated and slightly yellower than natural skin tones, with observers generally showing greater tolerance for changes in chroma than in hue [10]. Therefore, it is particularly important to understand how chromatic differences between virtual faces and scenes due to lighting conditions impact the appearance and evaluation of the combined environment.

Scene illumination estimation from an image is a challenging problem [12]. In virtual environments, consistent illumination is critical for merging background and virtual elements seamlessly [13, 11]. However, rendering realistic lighting often relies on inverse processes, where recovering scene illumination involves modeling soft shadows and interreflections—inherently difficult tasks[11]. Moreover, estimating lighting or other scene properties from a single image is highly under-constrained [15]. These limitations motivate the present study, which aims to determine the extent to which lighting conditions between virtual objects can differ before observers detect a discrepancy.

Understanding how real and virtual lighting interact to shape the perception of facial appearance across diverse skin tones is critical for optimizing visual realism and social communication in virtual environments [19]. When the lighting of virtual foreground objects is mismatched with that of the background environment, the resulting visual incongruity can reduce perceived naturalness for human faces [17]. The current work examines the extent to which lighting mismatches between virtual background scenes and faces become noticeable and how this sensitivity varies across different skin tones.

Methods

In this experiment, we simulated a background scene on an emissive display, and the virtual content was overlaid onto the same emissive display. This experiment simulated different lighting conditions on the virtual objects, which varied continuously along the Planckian locus, representing the chromatic variation of light sources as their correlated color temperature (CCT) changes. The experiment tested how people perceived the objects under these simulated lighting conditions, focusing on perceived matches, perceived mismatches, and preference, for the objects under different environmental lighting conditions. The experiment focused on perception of virtual faces having variable skin tones but included some non-face objects as comparison groups.

Apparatus

The experiment was conducted using a Dell 27-inch 4K PremierColor monitor with a resolution of 3840 x 2160 pixels. We used a CR-250 spectroradiometer to characterize this display; gamma-offset-gain (GOG) models were established to map color values between device RGB space and CIE 1931 XYZ space [20, 21]. Verification showed good performance between measured and model-predicted XYZ values (mean $\Delta E_{00} = 0.13$).

Stimuli

Background

Since the color-temperature scale does not align well with perceived color differences[22], we use the reciprocal color temperature unit, known as the 'mired,' measured in units of ($10^6\,\mathrm{K}^{-1}$), to describe changes in lighting chromaticity in a more perceptually-uniform way. Small intervals on this scale correspond to equally-perceptible changes across a broad range of color temperatures [23].

In this research, the D65 illuminant was selected as the reference point, representing a standard daylight condition with a CCT of approximately 6500 K (corresponding mired value 154 M, 'neutral'). Based on an investigation into our display gamut, the mired endpoints (89 M–269 M, 11236 K-3717 K) were chosen as the range where virtual object chromaticity shifted consistently before clipping became apparent. Within this range, we selected mired values of 109 M ('cool,' 9174 K) and 249 M ('warm,' 4016 K) to represent the two additional background lighting conditions.

The initial background scene was the same one used in previous work (see Figure 2)[24]. Using a reference illumination condition (D65), a chromatic adaptation transform (CAT) with full adaptation (D = 1) was applied to render all background lighting conditions. This approach has been used before to simulate the appearance of objects as if they were illuminated under lighting having different chromaticity [26]. The display model was subsequently applied to generate the final backgrounds.

Virtual objects

We employ the Monk Skin Tone Scale (MST) [31] to represent the wide range of human skin colors, which vary in melanin concentration [32]. To quantify skin pigmentation, we use the Individual Typology Angle (ITA)[32], a metric that correlates with measured skin melanin concentration. ITA is defined as $\mathrm{ITA} = \tan^{-1}\left(\frac{L^*-50}{b^*}\right)$ [32], where L^* indicates perceptual lightness. By calculating ITA for each MST value and analyzing its variation with L^*, we found that the MST over-represents both extremely light and dark tones. To reduce this over-representation, we averaged MST values at both ends of the scale, yielding six representative "target skin tones" for the present study.

The initial virtual facial objects also were the same ones used as in previous work (see Figure 1)[24]. We additionally selected two non-face objects—an achromatic cube, created in Blender[25], and a banana—to compare face observations to both achromatic and chromatic non-face objects. The cube was chosen for its basic, neutral shape, which lacks intrinsic meaning, and the fact that its lighting could be easily determined based on its own chromaticity. The banana, being chromatic and a familiar object with a memory color [9], was chosen to enable comparison with other chromatic objects (such as faces), which might also be expected to induce memory-color expectations.

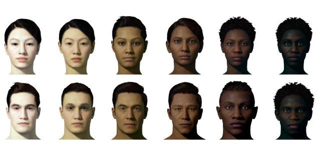

Figure 1. Twelve facial objects used in the current work.

We used the same CAT method on the virtual objects as for the background scene, except that they were simulated across a continuous mired range of 89–269 M in 1 M steps (180 images per virtual object). An example of the three backgrounds and a subset of the virtual objects over this range is shown in Fig. 2.

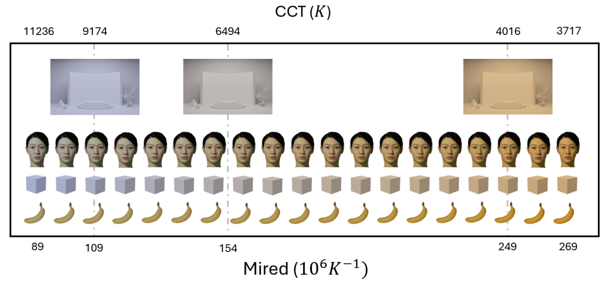

Figure 2. Example Background and Virtual Objects shown across the full mired range used in the experiment. In this example, virtual objects shown have 10-mired intervals (although the experimental stimuli had 1-mired intervals).

Procedure

Participants first completed an Ishihara color-vision test before receiving experimental instructions. Participants were told that they would be viewing virtual objects and completing tasks under the three background illuminant conditions. Following each background change, a 20s adaptation period (sufficient for 80% chromatic adaptation) was imposed [28]. During this adaptation period, participants were instructed to familiarize themselves with the new background lighting and do nothing else. The experiment comprised three main tasks: a perceived-match task, a perceived-difference task, and a preference task.

Perceived-match Task

The perceived-match task aimed to evaluate the rendered virtual objects' lighting that best matched the background lighting condition. In other words, the object that most closely resembled how it should appear if it were actually viewed within that environment.

Three different background lighting conditions were presented in separate blocks, with the order of blocks randomized across participants. Within each block, trials were also randomized. For each trial, participants were shown 9 image options, taken from a random object identity and presented in random position, which represented a slightly different simulated lighting condition relative to the background (see Figure 3). The 9 image options include 1 virtual object having identical mired as the background scene, and 8 other virtual objects having ±20 mired range in 5 mired increments.

Participants were asked to select the image that represented how it should appear under the same light as this room. Participants evaluated each virtual object under one background condition, before moving on to the next background condition. For this task, participants evaluated each virtual object (14) under each background scene (3), for a total of 42 trials.

Perceived-difference Task

The perceived-difference task aimed to identify the point at which the images appeared to be under categorically different

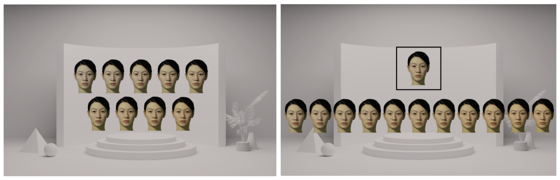

Figure 3. ***Left: Perceived-match task.*** *Observers selected the image that should appear under the same illumination as the room (i.e., best match to the scene lighting).* ***Right: Perceived-difference and preference tasks.*** *Observers advanced through the stimulus sequence and selected the first image that (i) appeared to be illuminated differently than the reference (perceived-difference), or (ii) appeared worse than the reference (preference).*

lighting conditions. In this task, the background and object block order and trial randomization were the same as in the previous task, but the presentation of the response options differed. Participants were presented with a reference image at the center of the screen, representing the "ideal" appearance of an object under the current background lighting condition (based on the CAT; see Figure 3). Below the reference image, a series of virtual objects was displayed, each depicting the same object identity but illuminated under a range of simulated lighting conditions.

The left-most image option was identical to the reference image, and each subsequent image increased its mired difference in increments of 3 mired in either the cooler or warmer direction (randomly selected). Participants were instructed: "Going from left-to-right, select the first image that appears as if it is under a different lighting condition than the reference image." Participants had the option to request the display of additional images for selection if needed (which expanded the possible mired range), and could also indicate if 'none of the presented images appeared to differ from the reference image' after they expanded the range in this way (which was recoded as the maximum possible value in that given direction). For this task, participants evaluated each virtual object (14) across each background scene (3), with mired changing in either warmer (positive) and cooler (negative) chromatic directions (2). This resulted in a total of 84 trials per participant.

Preference Task

The preference task aimed to assess the point at which differences in lighting between objects and backgrounds became detrimental to the preferred appearance of the objects.

In this task, the background and object block order, trial randomization, trial presentation, and response options were the same as in the previous task (see Figure 3). The only difference was in the instructions: participants were instructed, "Going from left-to-right, select the first image that looks worse than the reference image," where "worse" was defined as less preferable or less appealing. In total, participants completed 210 trials across all tasks, and the experiment lasted approximately 45 minutes.

Participants

Twenty-eight people ($M_{\text{age}} = 28, SD_{\text{age}} = 6.36$) participated in the study (some for extra course credit). Participants self-reported gender (10 male, 17 female, 1 non-binary) and ethnic identity (12 White or Caucasian, 14 Asian, 1 Multiracial, 1 preferred not to respond). One participant exhibited a color-vision deficiency (via both self-report and Ishihara screen), and was therefore excluded from subsequent analyses. Additionally, one participant's data were excluded from analysis due to repeated interruptions and irregular response patterns throughout the session, suggesting a lack of task engagement noted by the experimenter. Thus, the subsequent analyses were conducted with N = 26.

Results

Perceived-Match Task

A repeated-measures ANOVA was conducted to assess how participants selected virtual objects that appeared to best match the background lighting conditions. The analysis included virtual object type (8; 6 facial skin tones, cube, banana), and background lighting conditions (3; cool, neutral, warm) as within-subject factors. The dependent variable was mired difference, calculated as the difference between the participant's selected mired value and the background mired value. A summary of these perceived-match results is shown in Figure 4.

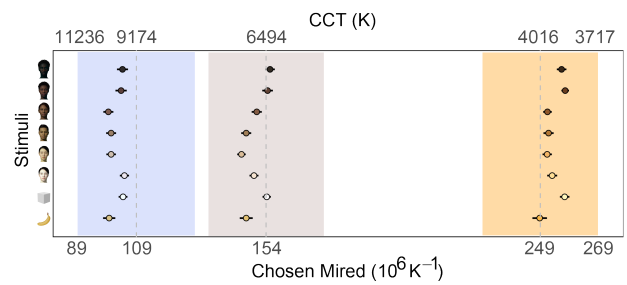

Figure 4. *Chosen mired values (mean and SE) for each virtual object across the cool, neutral, and warm background lighting conditions (represented by blue, gray, and orange background colors, respectively), in the perceived-match task. Dotted vertical lines represent the mired values of the background lighting conditions (109, 154, 249) under which the objects were viewed.*

There was a significant main effect of background lighting condition, $F(2,50) = 25.94$, $p < .001$. Under the cool background condition ($M = -6.94$, $SE = 1.37$), perceived matches for virtual objects shifted toward the cooler direction relative to the neutral condition ($M = -3.49$, $SE = 1.15$), p = .003. Conversely, under the warm background condition ($M = 4.53$, $SE = 1.30$), perceived matches shifted toward the warmer direction compared to the neutral condition, $p < .001$. The interaction between virtual object type and background lighting was not significant, $F(14,350) = 1.12$, $p = .34$, suggesting that this influence of background lighting was consistent across different virtual object types.

There was a significant main effect of virtual object type, $F(7,175) = 6.33$, $p < .001$, indicating that perceived lighting matches differed across object types. For facial stimuli, a significant quadratic trend, $F(1,125) = 11.32$, $p = .001$, revealed a nonlinear relationship in which mired differences initially tended to be more negative, and then more positive, as skin tone lightened. The cube ($M = 1.35$, $SE = .96$) produced a significantly closer perceived match to the background lighting than several other virtual objects ($ps < .0018$), but was not significantly different than skin tone 1 ($M = -1.41$, $SE = 1.50$, $p = 0.17$), skin tone 5 ($M = 1.28$, $SE = 1.64$, $p = 0.97$), or skin tone 6 ($M = 1.31$, $SE = 1.58$, $p = 0.98$).

Perceived-Difference Task

The results of this task revealed asymmetric lighting direction patterns most likely due to boundary limitations under the cool and warm background lighting conditions (see Figure 5; the cooler direction of the 'cool' background and warmer direction of the 'warm' background appear to be at the minimum and maximum boundary of the current study, respectively), suggesting that the responses in those directions were likely due to gamut constraints. Therefore, subsequent analyses focused exclusively on the neutral background lighting condition, where the full range was available for all virtual objects. A repeated-measures ANOVA was conducted to assess how participants selected virtual objects that appeared to be under categorically different lighting conditions than the background. The analysis included virtual object type (8; 6 facial skin tones, cube, banana) and direction of lighting shift (2; cooler, warmer) as within-subject factors. The dependent variable was the absolute mired difference, calculated as the unsigned difference between the participant's selected mired value and the background mired value (capturing the magnitude of mismatch regardless of direction). A summary of these perceived-difference results is shown in Figure 5.

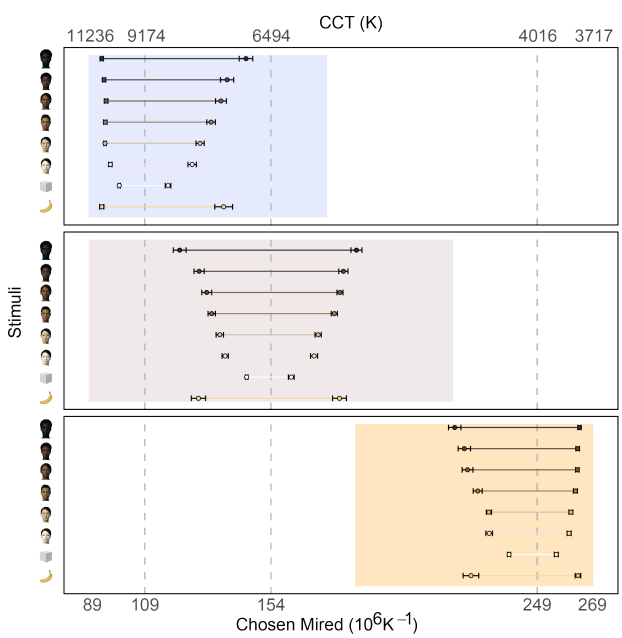

Figure 5. *Chosen mired values (mean and SE) for each virtual object across the cool, neutral, and warm background lighting conditions (represented by blue, gray, and orange background colors, respectively), in the perceived-difference task. Dotted vertical lines represent the mired values of the background lighting conditions (109, 154, 249) under which the objects were viewed.*

There was a significant main effect of virtual object type, $F(7,175) = 43.78$, $p < .001$. For facial stimuli, a significant linear trend was observed, $F(1,125) = 258.61$, $p < .001$, indicating that as skin tone lightened, smaller mired differences were needed before appearing as if the faces were under different lighting conditions. The cube ($M = 7.96$, $SE = 0.68$) exhibited the lowest difference range overall and differed significantly from all facial objects, $ps < .0001$. The banana ($M = 25.15$, $SE = 2.07$) yielded significantly larger difference values than most virtual objects, $ps < .007$, except skin 4 ($M = 23.77$, $SE = 1.61$), $t(25) = 0.74$, $p = 0.46$, and skin 5 ($M = 25.73$, $SE = 2.06$), $t(25) = -0.29$, $p = 0.77$.

There was no significant main effect of lighting shift direction, $F(1,25) = 0.56$, $p = 0.46$, nor a significant interaction between virtual object type and lighting shift direction, $F(7,175) = 0.67$, $p = 0.69$.

We additionally examined whether the observed differences in chosen mired ranges could be possibly explained by the perceptibility of color differences across skin tones as a function of mired change. Thus, we calculated ΔE_{00} color differences between each virtual object rendered under the chosen mired, and its identity rendered under the neutral background mired (see Figure 6). As shown, even though smaller mired differences were needed as faces became lighter, the corresponding color differences were roughly comparable. Therefore, it is likely that the observed patterns of results could be at least partly explained by differences in how the CATs impacted color appearance for different skin tones.

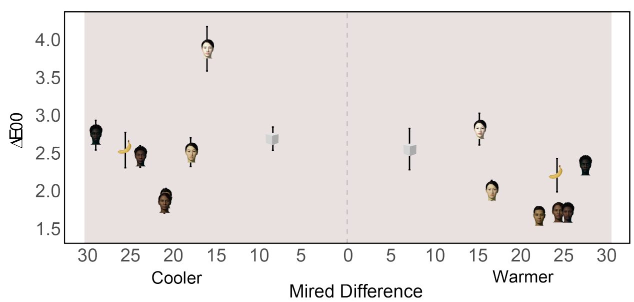

Figure 6. *ΔE_{00} (mean and SE) values are shown for each virtual object under neutral background lighting, across both cooler and warmer shift directions. As skin tones became darker, larger mired differences were needed to produce comparable ΔE_{00} color differences.*

Preference

For the same reasons as in the previous section, the analyses for the preference task focus exclusively on the neutral background lighting condition. The analyses for preference were identical to those conducted for the previous section, but with the dependent variable being the absolute mired difference calculated from preferred appearance choices (see Figure 7).

Results revealed a significant main effect of virtual object type, $F(7,175) = 38.31$, $p < .001$. Among facial stimuli, larger mired differences were needed to detrimentally impact preferred appearance as skin tones became darker. The cube exhibited the narrowest preference threshold range ($M = 14.2$, $SE = 1.73$), significantly lower than all facial objects ($ps < .0001$). The banana ($M = 36.4$, $SE = 2.88$) elicited significantly larger thresholds than skin 1 and skin 2 ($ps < .0001$), a smaller threshold than skin 6 ($p = .044$), and did not differ significantly from skin tones 3–5 ($ps > .057$).

Lighting shift direction also had a significant effect on participants' preferences, $F(1,25) = 21.62$, $p < .001$. Preference

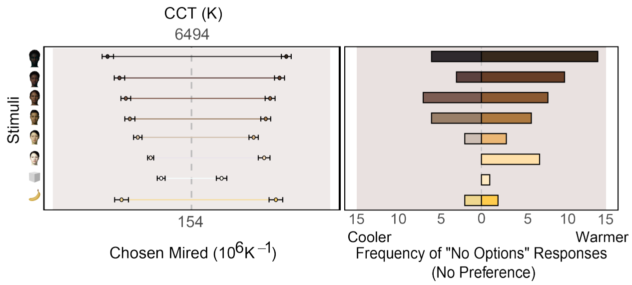

Figure 7. ***Left:*** *Chosen mired values (mean and SE) for each virtual object under neutral background lighting conditions in the perceived-preference task.* ***Right:*** *Frequency of responses where no stimulus options were judged as being less preferable ("No Worse") than the reference stimulus, as a function of lighting direction from the background (warmer vs. cooler). Virtual objects tended to be more preferred in the warmer direction regardless of the magnitude of the change, particularly for facial stimuli.*

thresholds were significantly lower in the cooler direction ($M = 28.2$, $SE = 2.41$) than in the warmer direction ($M = 34.5$, $SE = 2.66$), $t(25) = -4.65$, $p = .0001$, indicating that participants generally evaluated objects that became relatively warmer as more preferable to those that changed in the cooler direction. A significant interaction was also observed between object type and lighting shift direction, $F(7,175) = 3.63$, $p = .001$, showing that this pattern held true for all virtual objects, except for the cube, which did not differ in preference as a function of lighting direction $p = 0.97$.

To further understand participants' preference criteria, we separately examined the frequency of responses whereby no stimulus options were judged as being less preferable than the reference stimulus (see Figure 7). Under neutral lighting conditions, such responses tended to be more frequent in the warmer direction. This supports the notion that participants tended to prefer warmer appearances regardless of the magnitude of the lighting difference, particularly for facial objects.

Discussion

The current study assessed how similar (or different) lighting conditions between virtual objects and scenes needed to be before they appeared to match, mismatch, and impact preference. By systematically varying virtual object characteristics (facial vs. non-facial, skin tone) and lighting conditions (cool, neutral, warm), we quantified the thresholds at which lighting mismatches became perceptually noticeable or detrimental to appearance.

The perceived-match task results highlight how background lighting conditions influence participants' ability to match virtual object appearances accurately. Based on previous findings, we expected that as skin tone became lighter, perceived match accuracy would improve, reflected by facial objects having mired values more closely aligned with the background lighting condition. But our results more closely align with the findings of Mbatha et al. [29], showing that lighter skin tones yielded the most accurate lighting matches, medium skin tones the least accurate, and darker skin tones falling in between. A possible explanation is that participants may have prioritized chroma-matching, particularly because the background scenes were achromatic. The cube, with its achromatic appearance, served as a consistent baseline for evaluating lighting, while the banana, due to its saturated yellow hue, was more often accurately matched under warmer lighting. These findings demonstrate that participants' perceived matches are influenced not only by background lighting conditions but also by the characteristics of the object.

The perceived-difference task provides threshold ranges at which lighting differences between the virtual object and the background become perceptually distinguishable, across various object types and background lighting conditions. To better understand why perceived-difference thresholds varied across skin tones, we computed the ΔE_{00} color differences between each facial object and its identity rendered under the neutral background lighting condition. This analysis revealed that facial stimuli having darker skin tones exhibited smaller ΔE_{00} values for the same mired step sizes compared to lighter skin tones. This likely contributed to their broader threshold ranges, as participants required larger lighting changes to detect a noticeable difference.

The preference task revealed a clear influence of lighting direction on participants' preferred appearances. Across all conditions, participants consistently showed a bias toward warmer appearances, particularly for facial stimuli. This is consistent with some past research showing that people tend to prefer more saturated, redder, and yellower faces [10]. Therefore, in the case of preferences, it may be the case that the magnitude of lighting change was not as impactful as the direction of the color change.

This study has limitations that should be addressed in future research. First, the range of lighting conditions used in the perceived-difference and preference tasks was unexpectedly smaller than likely needed, for the warm and cool conditions, which prevented reliable inferences from those conditions. Future work should consider refining the facial objects set—extremely light and dark skin tones posed challenges related to display gamut limitations and are also less commonly represented in measured data [16]. Finally, background lighting conditions were simulated digitally via CATs rather than implemented through actual changes in the physical room lighting. As a result, this may have constrained participants' ability to make naturalistic lighting judgments.

Conclusion

Virtual environments rely on lighting to support perceptual coherence between rendered objects and their background scenes, particularly when rendering human faces with varying skin tones. In this study, we simulated various background lighting conditions and assessed participants' perceptual responses across three tasks: perceived-match, perceived-difference, and preference. By systematically varying virtual object properties (faces with variable skin tones vs. non-faces) and lighting conditions (cool, neutral, warm), we found that both virtual object characteristics and background lighting conditions significantly influenced appearance judgments, which varied depending on the particular task.

Together, these findings provide a perceptual foundation for informing face rendering in virtual environments and highlight the importance of accounting for skin tone diversity and perceptual biases in different lighting conditions. Future work should expand these findings to dynamic and interactive environments, where lighting mismatches may further affect user engagement and social presence.

References

[1] Liang Cui, Jia Liu, Virtual Human: A Comprehensive Survey on Academic and Applications, IEEE Access, 11, 123830–123845 (2023).

[2] Carlos F. Benitez-Quiroz, Radhakrishna Srinivasan, Alejandro M. Martinez, Facial color is an efficient mechanism to visually transmit emotion, Proc. Natl. Acad. Sci., 115, 3581–3586 (2018).

[3] Rachael E. Jack, Philippe G. Schyns, The human face as a dynamic tool for social communication, Curr. Biol., 25, R621–R634 (2015).

[4] Christopher A. Thorstenson, The Social Psychophysics of Human Face Color: Review and Recommendations, Social Cognition, 36, 247–273 (2018).

[5] Leslie A. Zebrowitz, First Impressions From Faces, Curr. Dir. Psychol. Sci., 26, 237–242 (2017).

[6] Kurt Hugenberg, John P. Wilson, Faces are central to social cognition, in The Oxford Handbook of Social Cognition, Oxford University Press, Oxford, 2013, pg. 167–193.

[7] Kok Wei Tan, Ian D. Stephen, Colour Detection Thresholds in Faces and Colour Patches, Perception, 42, 733–741 (2013).

[8] Christopher A. Thorstenson, Adam D. Pazda, Andrew J. Elliot, Subjective Perception of Color Differences Is Greater for Faces than Non-Faces, Social Cognition, 35, 299–312 (2017).

[9] Maryam Hasantash, Rosa Lafer-Sousa, Arash Afraz, Bevil R. Conway, Paradoxical impact of memory on color appearance of faces, Nat. Commun., 10, 3010 (2019).

[10] Haiyin Zeng, Ming Ronnier Luo, Colour and tolerance of preferred skin colours on digital photographic images, Color Res. Appl., 38, 30–45 (2013).

[11] Celong Liu, Lingyu Wang, Zhong Li, Shuxue Quan, Yi Xu, Real-time lighting estimation for augmented reality via differentiable screen-space rendering, IEEE Trans. Vis. Comput. Graph., 29(4), 2132–2145 (2022).

[12] Steven D. Hordley, Scene illuminant estimation: past, present, and future, Color Res. Appl., 31(4), 303–314 (2006).

[13] Katrien Jacobs, Céline Loscos, Classification of illumination methods for mixed reality, Comput. Graph. Forum, 25(1), 29–51 (2006).

[14] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, Xiaowei Zhou, Modeling indirect illumination for inverse rendering, Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 18643–18652 (2022).

[15] Yiqin Zhao, Chongyang Ma, Haibin Huang, Tian Guo, LITAR: Visually coherent lighting for mobile augmented reality, Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., 6(3), 1–29 (2022).

[16] Yuheng Lu, Kinjiro Xiao, Michael Pointer, Rui He, Shu Zhou, Ahmed Nasseraldin, Supatcha Sueeprasan, Cheng Gao, Chuanliang Li, Aamir Sohaib, et al., The International Skin Spectra Archive (ISSA): a multicultural human skin phenotype and colour spectra collection, Sci. Data, 12, 487 (2025).

[17] Anku Anku, Susan P. Farnand, Preferred White Balance for Applications Using Virtual Backgrounds, Color Res. Appl., 46, 1227–1244 (2021).

[18] Ana Elisa Gomes, Sérgio M. C. Nascimento, João M. M. Linhares, Cross-cultural preferences for optimal daylight illumination in viewing human faces, J. Opt. Soc. Am. A, 42, B525–B533 (2025).

[19] Mike Seymour, Kai Riemer, Judy Kay, Actors, Avatars and Agents: Potentials and Implications of Natural Face Technology for the Creation of Realistic Visual Presence, J. Assoc. Inf. Syst., 19(10) (2018).

[20] Roy S. Berns, Methods for Characterizing CRT Displays, Displays, 16, 173–182 (1996).

[21] Mark Fairchild, David Wyble, Colorimetric Characterization of the Apple Studio Display (Flat Panel LCD), Rochester Institute of Technology, Rochester, NY, 1998.

[22] Irwin G. Priest, A Proposed Scale for Use in Specifying the Chromaticity of Incandescent Illuminants and Various Phases of Daylight, J. Opt. Soc. Am., 23, 41–45 (1933).

[23] Shoji Tominaga, Brian A. Wandell, Natural Scene-Illuminant Estimation Using the Sensor Correlation, Proc. IEEE, 90, 42–56 (2002).

[24] Yanmei He, Christopher A. Thorstenson, Facial color matching in optical see-through augmented reality, *Journal of Vision*, 25, 16–16 (2025).

[25] Blender Online Community, Blender – a 3D Modelling and Rendering Package, Blender Foundation, Version 4.0 (2024).

[26] Mark D. Fairchild, Color Appearance Models, Wiley, Hoboken, NJ, 2013, pg. 1–441.

[27] Kaida Xiao, Joseph M. Yates, Fadi Zardawi, Supatcha Sueeprasan, Nan Liao, Leila Gill, Sophie Wuerger, Characterising the variations in ethnic skin colours: A new calibrated data base for human skin, Skin Res. Technol., 23, 21–29 (2017). https://doi.org/10.1111/srt.12295

[28] Mark D. Fairchild, Lisa Reniff, Time Course of Chromatic Adaptation for Color-Appearance Judgments, J. Opt. Soc. Am. A, 12, 824–833 (1995).

[29] Success K. Mbatha, Marthinus J. Booysen, Rensu P. Theart, Skin Tone Estimation under Diverse Lighting Conditions, J. Imaging, 10, 109 (2024).

[30] Jonathan P. Sacksner, Tushar Kaundinya, Krutarth Daftary, Kelsey L. Jackson, Roopal V. Kundu, Skin tone preferences and their influence on skin care behaviors, Arch. Dermatol. Res., 314, 961–966 (2022).

[31] Ellis Monk, The Monk Skin Tone Scale, 2023.

[32] Yue Wu, Toshiyuki Tanaka, Makio Akimoto, Utilization of Individual Typology Angle (ITA) and Hue Angle in the Measurement of Skin Color on Images, BioImages, 28, 1–8 (2020).

Author Biography

Xinmiao Zhang received her B.S. in Software Engineering from Harbin Normal University. She is currently a Ph.D. student in Color Science at the Munsell Color Science Laboratory, Rochester Institute of Technology. Her research focuses on face perception in real and extended reality environments.

Sofie Herbeck received their B.A. in Computer Science & Theater and Performance Studies from the University of California, Berkeley, where they also worked as a research assistant in a vision science lab. They are currently a Color Science Ph.D. student in the Munsell Color Science Laboratory at the Rochester Institute of Technology. Sofie's dissertation work focuses on individual color-matching functions and the methods used to characterize them.

Christopher Thorstenson received his MS in color science from the Rochester Institute of Technology (2017) and his PhD in social psychology from the University of Rochester (2019). He is currently an Assistant Professor in the Program of Color Science and the Munsell Color Science Laboratory at the Rochester Institute of Technology. His research focuses on the visual and social perception of social agents in real and extended reality environments.