Effects of Transparency on Perceived Humanness: Implications for Rendering Skin Tones Using Optical See-Through Displays

Tabitha C. Peck, *Senior Member, IEEE*, Jessica J. Good, Austin Erickson, *Graduate Student Member, IEEE*, Isaac Bynum, and Gerd Bruder, *Member, IEEE*

Fig. 1. Illustration of two Adobe Mixamo virtual humans captured on the Microsoft HoloLens via its screen capture software. Due to the additive light model employed by this display and most current-state optical see-through displays, the Black person appears more transparent (less opaque) than the White person. This is apparent when observing how the background texture of the door is clearly visible through the Black person but is otherwise not as noticeable through the White person. In practice, this effect is more noticeable than it appears in this figure, as the screen capture software does not factor in effects of environment lighting when generating images.

Abstract—Current optical see-through displays in the field of augmented reality are limited in their ability to display colors with low *lightness* in the hue, saturation, lightness (HSL) color space, causing such colors to appear transparent. This hardware limitation may add unintended bias into scenarios with virtual humans. Humans have varying skin tones including HSL colors with low lightness. When virtual humans are displayed with optical see-through devices, people with low lightness skin tones may be displayed semi-transparently while those with high lightness skin tones will be displayed more opaquely. For example, a Black avatar may appear semi-transparent in the same scene as a White avatar who will appear more opaque. We present an exploratory user study (N = 160) investigating whether differing opacity levels result in dehumanizing avatar and human faces. Results support that dehumanization occurs as opacity decreases. This suggests that in similar lighting, low lightness skin tones (e.g., Black faces) will be viewed as less human than high lightness skin tones (e.g., White faces). Additionally, the perceived emotionality of virtual human faces also predicts perceived humanness. Angry faces were seen overall as less human, and at lower opacity levels happy faces were seen as more human. Our results suggest that additional research is needed to understand the effects and interactions of emotionality and opacity on dehumanization. Further, we provide evidence that unintentional racial bias may be added when developing for optical see-through devices using virtual humans. We highlight the potential bias and discuss implications and directions for future research.

Index Terms—Augmented reality, optical see-through displays, additive light model, transparency, race, skin tone, diversity.

1 Introduction

Virtual humans, such as human avatars or computer-controlled agents, are an integral component of a wide range of augmented reality (AR) applications, including remote collaboration [59, 61], assisting/collaborative tasks, entertainment/media, healthcare, and training [46]. Many of these tasks require social interaction between people and, unfortunately, biases already exist in interactive scenarios. For example, in healthcare a medical practitioner's racial bias can negatively influence their ability to accurately diagnose Black patients [19]. Racial bias in policing is well documented [21] and numerous high-profile shootings of Black individuals at the hands of police officers have sparked the Black Lives Matter movement. Bias, including racial bias, is so prevalent and harmful that many companies, from technology to banking, have put programs in place to combat bias and address diversity [47].

When developing technology it is imperative to understand how biases are added into systems in order to limit further propagation of harmful bias. In the present work we investigate whether AR environments that include virtual humans are adding unintentional bias into their scenarios due to current display technology. Imagery on optical see-through head-mounted displays (OST-HMDs), such as the Microsoft HoloLens or Magic Leap One, is typically generated us-

- *Tabitha C. Peck is with Davidson College. E-mail:tapeck@davidson.edu.*
- *Jessica J. Good is with Davidson College. E-mail:jegood@davidson.edu.*
- *Austin Erickson is with University of Central Florida. E-mail: ericksona@knights.ucf.edu.*
- *Isaac Bynum is with Davidson College. E-mail:isbynum@davidson.edu.*
- *Gerd Bruder is with University of Central Florida. E-mail:Gerd.Bruder@ucf.edu.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

ing an additive light model [18], where light emitted from the display blends together with the light from the physical environment. One consequence of this is that the color black, and some similar colors with low *lightness* in the hue, saturation, lightness (HSL) color space, cannot be rendered by adding light into the scene. These colors appear transparent on an OST-HMD, effectively not being rendered at all.

In order to represent the diversity of identities in the real world, virtual humans must possess varying skin tones and facial features. Based on current OST-HMD technology, a virtual human that is a visual representation of a Black person may appear semi-transparent in some lighting conditions, while in the same lighting condition a virtual human of a White person would appear opaque (see figure 1). This difference in opacity between virtual humans of varying skin-tones may impact a user's perception of, and response to, the virtual human by introducing additional unintentional bias into the scenario.

The systematic difference in opacity based on skin tone may lead users to perceive virtual humans of certain racial groups as less human. We investigated the perceived humanness of faces at varying opacity levels to highlight where AR developers may be adding unintentional bias into scenarios. In this exploratory work we investigate the effect of opacity-level on the perception of humanness for a small group of avatars and real human faces. Recognizing and highlighting where bias is added into development will enable researchers and developers to put into place bias mitigation strategies for better hardware and software development.

This paper makes several contributions to the field: We present an explorative user study that, to our knowledge, is the first investigation into how the opacity of virtual humans affects their perceived level of humanness. We show that the opacity level of a virtual human, as observed in optical see-through AR, has significant effects on user perception, causing more transparent virtual humans to be viewed as less human. We show that, due to the additive light model used in optical see-through displays, this effect is more prominent in virtual humans with darker skin tones, causing Black virtual humans to be perceived as less human compared to virtual avatars with lighter skin tones.

2 Background

To more fully present the technological issues with OST-HMDs as well as potential implications, we first discuss AR systems and the use of virtual humans within AR. We then describe how race is represented for virtual humans and the importance of studying the humanization (or dehumanization) of virtual humans.

2.1 Augmented Reality Systems

While a common goal of AR displays is to be able to portray imagery that is indistinguishable from reality, there are several current limitations that must be overcome [34]. Imagery on current-state OST-HMDs is typically displayed using an *additive light model* [18], where light emitted from the display blends together with the light from the physical environment. An effect of this light model is that the color black, and other colors with low lightness in the HSL color space, appear transparent to users or do not appear at all. For instance, this effect is observable when attempting to display black and white user interfaces (UIs) in either *light mode* or *dark mode* (see figure 2). Past research has shown that this phenomenon can have surprising effects on human perception of virtual content on OST-HMDs. For instance, it makes dark mode UIs more easily readable than light mode UIs at small font sizes, which goes against traditional conventions for other types of displays [12]. While research prototypes are being investigated to overcome this limitation, such as through *subtractive* light model displays [33] or displays with an *opacification* layer [5], they will not be available as consumer devices in the foreseeable future. Current consumer displays such as the Microsoft HoloLens 1 and 2 or the Magic Leap One are hence limited in the range of HSL colors that they can faithfully present.

To further complicate this, OST-HMDs also have to compete with the lighting that is present in the user's physical environment (see

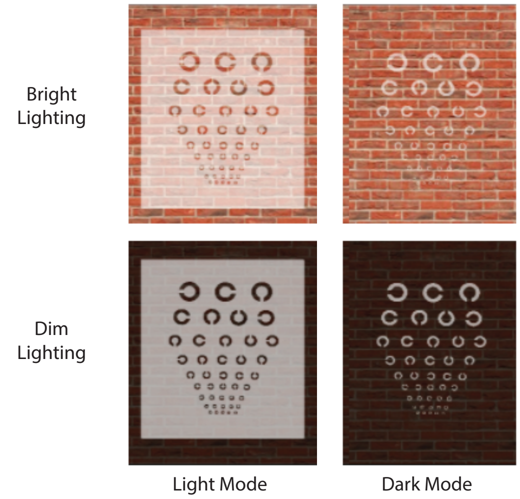

Fig. 2. This figure demonstrates how the color black will appear completely transparent when displayed on an optical see-through display. Black letters appear cut out of the background when using traditional light mode user interfaces (left), whereas the background appears completely transparent when using dark mode user interfaces (right). The lighting conditions in the physical environment further influence this, causing imagery to have less contrast in bright conditions (top), compared to dimly lit conditions (bottom).

figure 2). While this is typically not an issue for indoor or dimly-lit environments, if the user is in an area in which there is a high amount of ambient light, such as outdoors on a sunny day, then the virtual imagery may lose contrast to the point where it may appear washed out or even completely transparent [11]. This is the result of the display only being able to emit a limited amount of light compared to the light from the physical environment. The contrast between the display and environment decreases as the light from the environment intensifies. This reduction in contrast has the additional effect of further reducing the perceivable color space on the display, where colors with low lightness levels will gradually drop out of the perceivable color space as the environment lighting intensifies. This problem may be alleviated by limiting the intensity of the environment light that reaches the user, such as by using neutral density filters attached to the display. Conversely, the luminance capabilities of the display may be improved such that it is capable of emitting more intense light.

2.2 Virtual Humans

AR displays enable users to observe and interact with environments consisting of a real physical objects within the user's local environment and virtual imagery that appears to be collocated within the same environment. When virtual imagery is used to represent a real person or a virtual agent, it is referred to as a virtual human. Typically the appearance and behavior of the virtual human can be controlled by a person or a virtual agent in order to accomplish tasks such as navigating virtual environments, communicating with other users, and interacting with physical or virtual objects.

One common use of virtual humans in AR is to allow remote communication and collaboration between people when it is not possible or is otherwise inconvenient to meet in person. While this domain has been more extensively explored with combinations of AR and virtual reality (VR) displays [6, 23, 37, 51], several recent works have explored this context solely using AR displays [57, 59, 61]. For instance, the AR telepresence project *Holoportation* reconstructed a user's appearance in real time and presented it to other remote AR users [48].

Virtual avatars are also sometimes used to change a user's self perception of their own body or identity through virtual embodiment illusions [22, 36]. A recent survey by Genay et al. investigated existing work in this domain, paying particular attention to how the user's sense of embodiment changes when applying various levels of avatarization to oneself in AR [20]. In their work, they describe a continuum of self-body avatarization with AR displays, ranging from using the person's real body (in collocated settings), to accessorization of the user's body with virtual imagery, to partial avatarization with additional virtual modifications, to full avatarization where the user appears as a purely virtual entity.

When not representing a real human, virtual humans are sometimes used to visually represent a virtual agent. A systematic review of this particular domain was conducted by Norouzi et al. [46], which indicated that the recent literature in AR focused on four main application areas of virtual agents: assistive/collaborative tasks, entertainment/media, healthcare, and training. A common theme in this domain is that virtual agents are often used as a stand in for a real person or a human-controlled avatar in contexts where it is expensive, dangerous, or otherwise difficult to have a real person.

2.3 Race

The concept of race generally includes a set of socially constructed categories designed to draw boundaries between advantaged and marginalized groups of people [45]. Racial categories are defined differently around the world and based on different features[1]. In the United States, one of the primary phenotypical indicators of race is skin tone [13, 54]. Although there is skin tone variation within racial groups, in broad terms Americans with darker skin are typically viewed as Black, while other Americans (of various ethnic backgrounds) are subsumed into the White category. Within racial groups, evidence suggests that there is often colorism, such that people with lighter skin tones are advantaged relative to those with darker skin tones [25, 32].

Although race is socially constructed, it carries high societal importance. Members of stigmatized racial groups bear the brunt of overt and covert discrimination as well as societal systems set up to disenfranchise and exclude them from equitable outcomes. The Black Lives Matter movement has highlighted ongoing systemic racism against Black Americans, emphasizing how race impacts the most basic of social outcomes such as educational access and medical care. Given its social importance, race is a highly accessible construct; research shows that people classify another individual's racial group within 200 milliseconds of encountering them [38]. Within AR, virtual humans may be created to represent differing racial groups based on several phenotypical features. Given the prominence of skin tone as a racial indicator however, it is likely that virtual human skin tone will be lightened or darkened to indicate racial group membership. Because of the current optical limitations of OST-HMDs, under some lighting conditions virtual humans representing Black individuals will appear to the user as semi-transparent, whereas lighter skinned virtual humans will be seen as more opaque. These appearances can increase the likelihood of dehumanization that could potentially affect how users perceive real humans of color.

2.4 Dehumanization

Dehumanization is a common psychological tool used to support the subjugation of certain groups of people. For example, historical justifications for slavery in the United States likened Black people to animals who needed to be tamed and controlled by White enslavers [35]. Viewing people as subhuman or animal-like has been shown to increase support for punitive criminal justice practices and justify state-perpetrated violence [21]. For example, dehumanizaton of Spanish protesters among police officers was associated with greater hostility and legitimization of violence against protesters [64].

Ongoing dehumanization of Black individuals in the United States, and elsewhere, is evidenced through research showing that Black medical patients are perceived to experience less physical pain compared to their White counterparts [56]. Indeed, Black pre-hospital patients are less likely to be administered pain medication by emergency personnel compared to all other racial groups [29], in part due to dehumanization [41], or super-humanization of Black people [60] such that the group is viewed as possessing supernatural qualities.

Although classic research on dehumanization has focused on viewing people as animalistic, technological advances suggest an additional method for dehumanization – perceiving people as robotic or mechanistic [27]. Consistent patterns in cross-cultural research have shown that animalistic dehumanization includes the perception that targets lack higher order cognitive capabilities and refined emotion, while robotic dehumanization primarily includes the perception that targets lack emotional capacity and desire [27].

We posit that the opacity of virtual humans may influence users' humanization of said virtual humans. Research on the *Uncanny Valley* hypothesis [44] shows that the degree of similarity to human appearance may influence users' positive or negative evaluation of an object. Specifically related to perceived humanness, Thompson et al. [55] found that the perceived humanness of avatars changed based on how the avatars moved. Within AR, avatars and virtual agents that are realistically humanoid, but more transparent, may fall into the Uncanny Valley such that they are perceived as eerie or supernatural. Thus, the Uncanny Valley literature parallels the dehumanization literature to suggest that virtual humans that are perceived as eerily robotic may elicit disgust and hostility [28]. If dark-skinned virtual humans are systematically more likely to be rendered semi-transparently due to technical limitations within AR, then engaging with Black virtual humans in AR may reinforce notions of robotic dehumanization of Black people. Identifying and highlighting this potential bias may guide researchers and developers into the creation and deployment of bias mitigation strategies and better system development.

3 Present Study

To our knowledge, no research to date has tested the impact of virtual human opacity on perceptions of humanness, in order to consider the implications of the current limitations of OST-HMD technology. VR and AR are used to investigate racial bias and to train professionals in various settings, including healthcare and criminal justice[2] [46, 49]. Given demonstrated inequities in pain perception across racial groups [41] and unequal use of force based on suspect race [53], it is particularly important for AR and VR training in these domains to work to ameliorate racial bias rather than reinforce it.

Therefore, in the present study we test how opacity and virtual human race impact perceptions of humanness. To do so, we utilized a mixed design comparing avatar and human images with varying races and opacity levels since avatars and humans may be perceived differently [63]. To determine whether any observed differences in humanness were based on racial group membership or based on skin tone, we included an East Asian target (light skin tone) in addition to a White (light skin tone) and Black (dark skin tone) target. Given that the present study utilized novel methodology, we included measures and manipulations to verify the validity of the work. For example, we included both human faces and avatar faces, as well as a Zombie target to verify the validity of the humanness measure; we expected zombie targets to be rated as less human than all other targets, and avatars to be rated as less human than human faces. We also included stereotyping and emotion measures to confirm our manipulation of race; if we have successfully manipulated target race, we expect differences in the stereotyping measures as a function of target race.

The present research was exploratory and we therefore did not make specific predictions. Instead, we sought to test whether opacity impacted perceptions of humanness and whether the effects varied as a function of target race or appearance.

4 Methods

4.1 Image Generation

In order to generate images that would be representative of what the user would observe on an OST-HMD we needed to understand how to model the transparency and color blending effects that are inherent in optical see-through AR imagery. While this effect has been formally

[1] https://understandingrace.org/GlobalCensus

[2] https://www.apexofficer.com/

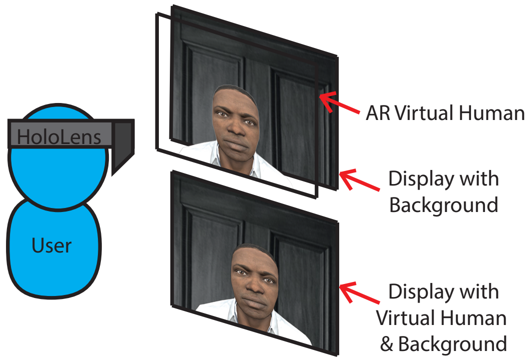

Fig. 3. Illustration depicting the setup of the perceptual matching task, in which a participant compares the appearance of an AR virtual human (top) to an identical virtual human displayed on a flat-panel display (bottom). Note that the AR imagery was displayed at the same depth as the flat panel display and appears in front of the background in this illustration only to emphasize that the top virtual human consists of a combination of AR imagery and a flat-panel display.

modeled by Gabbard et al. [18] for optical see-through displays in terms of the blending of light from the user's environment and light from the display, to our knowledge there does not exist a formulaic method of generating 2D images that capture such effects. However, in 2020 Gabbard et al. pointed out that color accuracy and robustness can be appropriately measured through the use of *perceptual matching* tasks where users are tasked with matching imagery displayed on an optical-see through display to imagery shown on a controlled display for which color parameters have already been established [17]. Thus to capture the effects of color blending and transparency for our experimental images we chose to perform a perceptual matching task.

The perceptual matching task was performed by one participant, who wore an AR HMD, the Microsoft HoloLens 1. They compared the appearance of an AR virtual human, displayed on the device in front of a flat panel display, to the appearance of the same virtual human displayed on a second flat panel display at the same distance as the AR virtual human (see figure 3.) The AR avatar was positioned to be in front of a display depicting the virtual background of a dark door image (see figure 4). A second display was directly beneath the flat panel display and was connected to a PC running a Photoshop 2021 project containing on separate layers 1) images of the same avatar segmented from the background and 2) the background image.

The participant adjusted the layer blending options and layer opacity levels of the virtual human in Photoshop to best match the appearance of the AR avatar on the HoloLens. Due to the potential impact of the tinted visor on the front of the HoloLens, the user would alternate between looking through the HMD at the AR avatar and lifting the HMD to see the avatar from the Photoshop project. The results of this task suggested that using the linear dodge (additive) layer blending option on the avatar's Photoshop layer, along with a layer opacity of 35 percent yielded images that closely matched each other.

It should be noted that there are several limitations with generating images in this manner. Most notable is that the resulting Photoshop parameters are dependent on both the physical lighting conditions present in the testing environment as well as the luminance capabilities of the OST-HMD. Decreasing the intensity of the physical lighting in the testing environment or increasing the luminance capability of the HMD would yield higher layer opacity parameters than what was obtained here, and vice versa. While this would be problematic if we wanted to compare avatars in different physical environments or different physical lighting conditions, these factors remain consistent across all avatar images used in this study. Thus this limitation should not impact our results. Images generated in this manner also do not take into account the effects of luminance non-uniformity on the OST- HMD [7, 39], which may or may not have effects on the user's perception of the avatars. However, this effect is difficult to quantify, even with the proper optical

(a) Avatars

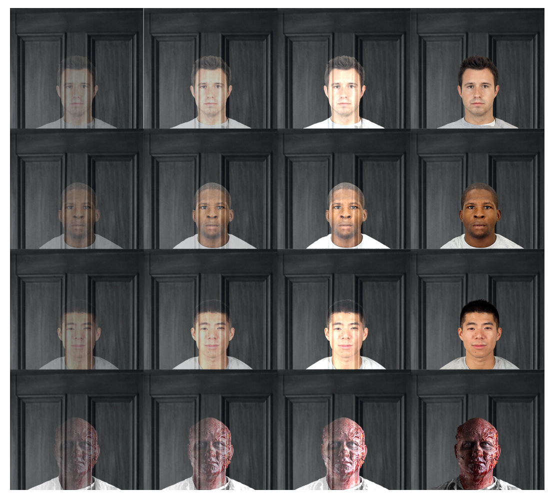

(b) Humans

Fig. 4. Visual stimuli used in the experiment: (a) avatars representing (from top to bottom) White, Black, East Asian, and Zombie categories, (b) humans representing (from top to bottom) White, Black, East Asian, and Zombie categories. For both (a) and (b), columns represent (from left to right) 35% opacity, 68% opacity, 100% opacity, and 100% opacity with color correction.

equipment, and so we leave investigation of this factor to future work.

We used the above-mentioned parameters to generate the first set of virtual human images that can be seen in the left-most columns of figure 4 a and b. These images provide a baseline for how virtual humans would appear on a current consumer optical see-through display, specifically the HoloLens 1, under indoor lighting conditions. To compare against these images, we generated three additional sets of images that consider how the appearance of the avatars would change as the display technology advances. The next two sets of images (second and third from the left column in figure 4 a and b) depicts how the avatars would appear should the luminance capability of the display increase to result in opacity parameters in a similar perceptual matching task of 68 percent and 100 percent, respectively. The value of 68 percent was chosen as a midway point between the current opacity values found for the HoloLens 1 and the maximum values possible at 100 percent opacity. For the final set of images (rightmost column in figure 4 a and b), we reverted the layer blending setting on the avatar from linear dodge (additive) to normal mode, which represents what users may see on a future OST display in which the issues of transparency and color blending are resolved. These image sets make up the four different opacity levels referenced in the remainder of the paper:

- **35% Opacity**: Current Display Opacity
- **68% Opacity**: Improved Display Opacity
- **100% Opacity**: Maximized Display Opacity (without color correction)
- **100% Opacity CC**: Maximized Display Opacity with Color Correction

Since these images would be viewed by participants on their own personal displays, it is possible that there were slight variations between how the images appeared to each participant. Such variations could potentially include: differences in brightness, contrast, and color tone of the image. However, since we are interested in comparing participants' collective perceptions between the different images rather than perceptual differences between participants, our results should not be negatively impacted by this variation. Further, such display variations should only reduce the likelihood of finding significant effects, as extreme display settings would make it more difficult to identify differences between images, effectively dampening significant effects with cases in which no effect was observed.

All 3D avatars used in our study were collected from the Microsoft Rocketbox Avatar Library [3] except for the zombie model, which was obtained from Adobe Mixamo [4]. The avatars were all retextured to appear as though they were wearing the same white collared shirt. For the real human conditions, images were gathered from the Chicago Face Database [5], except for the zombie image which was obtained from a stock image library. Images from the Chicago Face Database are standardized photographs of human faces that are intended for controlled scientific research. The faces selected for this experiment came from the main data set with neutral, non emotive faces. The door image used as a background in all conditions was similarly obtained from a stock image library. This exploratory study limited images to only male-appearing faces since racial stereotyping and dehumanization literature supports these effects on men. However, future work should investigate the effects on humanness caused by opacity and race for female and gender-neutral appearing faces as well.

4.2 Experimental Procedure

The procedure and all materials were approved by the Davidson College Institutional Review Board prior to beginning data collection. Participants initially accessed the study by clicking on the study URL from within the CloudResearch recruitment listing. The study was conducted online (hosted via the Qualtrics survey platform). Participants first read and completed informed consent and an eligibility checklist, which verified that they were at least 18 years old, had normal or corrected to normal vision, and were proficient in written and spoken English (since all study materials were presented in English).

Eligible, consenting participants were then randomly assigned to view and evaluate images from 2 of the 4 racial group conditions for a total of 16 images (4 opacity level $\times$ 2 appearance $\times$ 2 race). As recommended by Little and Jubin [42] we used a planned missing data design and chose to capture participants' evaluations of only 2 of the 4 racial groups in order to reduce participant fatigue and poor quality responding. Planned statistical imputation of missing data allowed for estimation of the other two racial groups for each participant. For each image, participants rated the figure using a series of 14 questions (see section 4.3). Participants were presented with the 16 images (and subsequent questions) in random order. Embedded within the image questions were two attention check questions which were used to screen out inattentive participants. After evaluating all images, participants completed demographic questions. Participants were compensated according to agreed upon rates with CloudResearch; participation lasted on average 19.35 minutes (min = 8.33, median = 15.30).

4.3 Measures

Participants evaluated the extent to which each figure seemed human, animal-like, robotic, competent, friendly, dangerous, angry, happy, creepy, and unearthly. Participants also evaluated the extent to which each figure had emotions, felt physical pain, had complex thoughts, and had control over their actions. All responses were given on a 100-point sliding scale ranging from "not at all" to "very much" with the initial selection snapped at the midpoint. Participants were required to move each slider before continuing to the next image.

Our primary interest was in assessing *humanness*, though we also included two items assessing emotion in order to control for possible baseline differences in facial expression across the images. Research has demonstrated that White participants tend to perceive faces of Black Americans as angrier than faces of White Americans [24,31] even when facial musculature/expression is held constant. We also included items assessing stereotypical perceptions of racial groups in order to verify that our manipulation of the racial groups was successful. Psychological research on the Stereotype Content Model indicates that social groups are stereotyped along two primary axes, warmth and competence [14]. Contemptuous prejudice is directed at groups that are stereotyped as low in both competence and warmth (e.g., Black Americans), while envious prejudice is directed at groups that are stereotyped as high in competence but low in warmth (e.g., Asian Americans). Groups that are stereotyped as high in both warmth and competence generally receive social admiration (e.g., White Americans).

Therefore, participants rated how *happy* and how *angry* each figure was, and we included these items in the overall experimental model in order to account for any variation in facial expression. We conducted separate analyses with the *friendly*, *competent*, and *dangerous* items in order to verify expected racial group differences in warmth and competence as well as the common stereotype that Black men are dangerous [8, 9].

For the remaining items assessing humanness, we conducted an exploratory principal components factor analysis with varimax rotation to determine how many factors best fit the data. Two factors with eigenvalues greater than 1 emerged. We retained items that loaded above .5 on only one of the two factors. The first factor (eigenvalue = 4.36, 48.41% variance explained) included 6 items (human, robot, unearthly, emotions, physical pain, and complex thoughts) and showed high scale reliability (Cronbach's alpha = .904). the second factor (eigenvalue = 2.01, 22.30% variance explained) included 3 items (animal, creepy, and control over actions) and demonstrated acceptable scale reliability (Cronbach's alpha = .762), but low conceptual meaning. The three items on the second factor seemed to be measures of humanness that did not fit the current study context well given that the targets in the present study were virtual (not animalistic) and static (did not have action). Therefore, we computed a single composite measure of *humanness* as the mean of the following items: human, 100-robot, 100-unearthly, emotions, physical pain, and complex thoughts.

4.4 Participants

160 participants, recruited and compensated through the CloudResearch platform, passed the attention checks and were included in the analysis. Participants self-verified that they met all inclusion criteria. A general population of participants were recruited to match the United States Census to support the generalizability of results. 49% of participants identified as women, 49% as men, and 2% as non-binary. The racial make-up of participants was 76% White, 12% Black of African American, 4% Hispanic/Latinx, 4% East Asian, 2% American Indian or Alaska Native, and the remaining 2% as other races. Education level ranged from 3% with some high school, 38% with a high school degree, 6% with a 2-year college degree, 33% with a 4-year college degree, 14% with a masters degree, and 6% with a doctoral degree. Household income demographics included 22% earning under \$25,000, 27% earning between \$25,000–\$50,000, 19% earning between \$50,000–\$75,000, 10% earning between \$75,000–\$100,000, 5% earning between \$100,000–\$125,000, 6% earning between \$125,000–\$150,000, 3% earning between \$150,000–\$175,000,and 8% over \$175,000. Finally, 36% were non-video-game players, 44% casual video-game players, 12% core-video game players, and 8% hard-core video game players.

[3] `https://github.com/microsoft/Microsoft-Rocketbox`
[4] `https://www.mixamo.com/`
[5] `https://www.chicagofaces.org/`

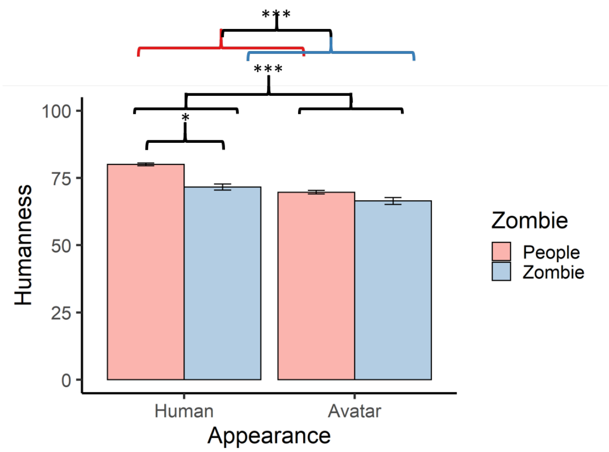

Fig. 5. A bar chart of average humanness with standard error bars. Zombie faces (blue) had significantly lower humanness compared to people (red). Further, avatar faces had lower humanness compared to human faces.

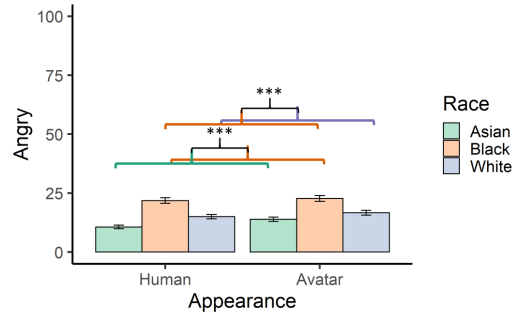

Fig. 6. A bar chart of the average angry score with standard error bars. Black faces (orange, middle) were rated angrier than both Asian (green, left) and White (purple, right) faces.

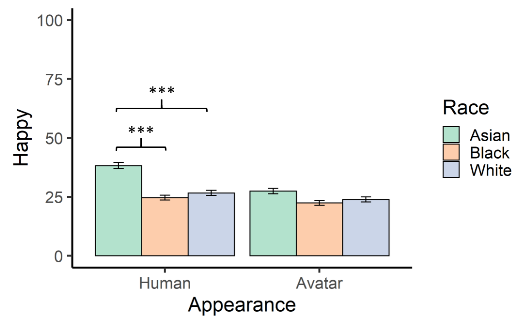

Fig. 7. A bar chart of the average happy score with standard error bars. Within the human faces (left), the Asian (green) face were rated happier than both the Black (orange) and White (purple) faces.

5 Results

Analysis was performed in R 4.0.0. Planned missing data were imputed using predictive-mean matching [43] and implemented with the mice 3.11.0 library. After imputation, submeasures were calculated as described in section 4.3.

5.1 Validation Analyses

To validate the *humanness* measure we performed a 2 (*zombie*: zombie, people) $\times$ 2 (*appearance*: avatar, human) within-participants analysis of variance (ANOVA). See figure 5. Zombies ($M = 69.02$, $SD = 30.21$) had significantly lower *humanness* compared to people ($M = 74.86$, $SD = 25.88$), $F(1,159) = 4.17$, $\boldsymbol{p} = \mathbf{.04}$, $\eta^2 = .02$. Further, humans ($M = 77.94$, $SD = 23.24$) had significantly higher *humanness* compared to avatars ($M = 68.96$, $SD = 29.87$), $F(1,159) = 180.76$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .03$. Finally, there was a significant interaction, $F(1,159) = 16.74$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .003$. When considering human faces, people had higher *humanness* compared to zombies, $t(175) = 2.88$, $\boldsymbol{p} = \mathbf{.004}$, $d = .37$. However, no significant difference was found between zombies and people in the avatar group, $t(175) = 1.10$, $p = .27$, $d = .11$.

These results support that the *humanness* metric did measure humanness since people had higher humanness compared to zombies and humans had higher humanness compared to avatars. The zombie faces were removed from the remainder of this analysis since they were included only to determine the construct validity of the humanness measure.

Because only one human and one avatar face was used to represent each racial category, we tested for expected differences in stereotyping based on racial category as a way of validating the manipulation of race. We tested for an effect of race on each of the *competent*, *friendly*, and *dangerous* measures using a 2 (*appearance*: avatar, human) $\times$ 3 (*race*: Asian, Black, White) ANOVA. Post-hoc analysis was performed using estimated marginal means with Tukey method adjustments for repeated tests by comparing race pairwise within each separate appearance face group. The highest-order significant effect is reported. Human faces ($M = 70.81$, $SD = .72$) were perceived to be significantly more *competent* compared to avatar faces ($M = 63.46$, $SD = .78$), $F(1,159) = 83.04$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .02$. A significant race $\times$ appearance interaction was identified for the *friendly* measure, $F(1,318) = 4.98$, $\boldsymbol{p} = \mathbf{.007}$, $\eta^2 = .002$. The Asian human face ($M = 34.59$, $SD = 1.49$) was rated significantly friendlier than the Black human face ($M = 24.77$, $SD = 1.21$), $t(362) = 2.76$, $\boldsymbol{p} = \mathbf{.02}$, $d = .86$. No differences were identified within the avatar faces. Finally, a significant race $\times$ appearance interaction was identified for the *dangerous* measure, $F(1,318) = 3.15$, $\boldsymbol{p} = \mathbf{.04}$, $\eta^2 = .001$. Consistent with expectations based on stereotypes [9], the Black human face ($M = 19.43$, $SD = 1.10$) was rated significantly more dangerous than the Asian human face ($M = 11.97$, $SD = .85$), $t(377) = 2.72$, $\boldsymbol{p} = \mathbf{.02}$, $d = .62$. No differences were identified within the avatar faces.

Having verified that we successfully manipulated racial category membership, we next sought to ensure that race was the only variable manipulated and that facial expression was controlled across the various targets. We therefore tested for an effect of race on the angry and happy measures using a 2 (appearance) $\times$ 3 (race) ANOVA. See figures 6 and 7. For the *angry* measure, significant main effects of both appearance, $F(1,159) = 10.98$, $\boldsymbol{p} = \mathbf{.001}$, $\eta^2 = .002$, and race, $F(1,318) = 7.86$, $\boldsymbol{p} = \mathbf{.0004}$, $\eta^2 = .03$, were found. Avatar faces ($M = 17.75$, $SD = .62$) were rated as angrier than human faces ($M = 15.81$, $SD = .58$), $t(159) = -3.31$, $\boldsymbol{p} = \mathbf{.001}$, $d = .07$. Further, Black faces ($M = 22.25$, $SD = .85$) were rated as angrier than both Asian faces ($M = 12.26$, $SD = .60$) and White faces ($M = 15.83$, $SD = .70$), $t(318) = -3.91$, $\boldsymbol{p} = \mathbf{.0003}$, $d = .08$ and $t(318) = 2.51$, $\boldsymbol{p} = \mathbf{.03}$, $d = .05$ respectively. This finding is consistent with past research showing that anger is more likely to be associated with Black faces compared to White faces [24, 31].

When considering the *happy* measure, significant main effects of both appearance, $F(1,159) = 49.95$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .01$, and race, $F(1,318) = 8.99$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .03$, were found. Additionally, there was a significant appearance $\times$ race interaction $F(1,318) = 15.01$, $\boldsymbol{p} < \mathbf{.0001}$, $\eta^2 = .007$. The Asian human face ($M = 38.27$, $SD = 1.32$) was perceived to be happier than both the Black human face ($M = 24.70$, $SD = 1.03$) and the White human face ($M = 26.63$, $SD = 1.08$), $t(405) = 5.43$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .63$ and $t(405) = 4.66$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .58$ respectively. No significant differences were identified between the avatar faces.

Because perceptions of emotion in facial expression differed across the various faces, we sought to use angry and happy as covariates in our primary analyses of the effect of opacity and race on humanness.

Table 1. Final linear regressions for the *humanness* measure considering opacity, race, angry, and happy. Human faces are in the left column and avatars are in the right column. Significance codes: *** < 0.001, ** < 0.01, *< 0.05, $\cdot < 0.1$.

	Humans $R^2 = .34$					**Avatars** $R^2 = .39$				
	β	Std. Error	t-value	p		β	Std. Error	t-value	p	
(Intercept)	82.97	2.72	30.51	**<0.001**	***	75.35	3.76	20.03	**<0.001**	***
gender	–	–	–	–		-5.30	2.04	-2.60	**0.01**	*
happy	-0.02	0.06	-0.38	0.70		0.21	0.08	2.56	**0.01**	*
angry	-0.67	0.16	-4.29	**<0.001**	***	-0.74	0.15	-4.92	**<0.001**	***
Opacity	9.54	3.31	2.88	**<0.01**	**	12.13	4.35	2.79	**<0.01**	**
RaceBlack	-2.15	3.74	-0.58	0.56		8.70	4.98	1.74	*0.08*	.
RaceWhite	-5.58	3.69	-1.51	0.13		4.91	4.86	1.01	0.31	
happy:angry	0.00	0.00	0.64	0.53		0.00	0.00	0.10	0.92	
happy:Opacity	-0.04	0.07	-0.65	0.52		-0.43	0.11	-3.95	**<0.001**	***
angry:Opacity	-0.17	0.23	-0.76	0.45		-0.18	0.20	-0.86	0.39	
happy:RaceBlack	0.14	0.09	1.60	0.11		-0.05	0.13	-0.43	0.67	
happy:RaceWhite	0.10	0.08	1.17	0.24		-0.12	0.12	-1.04	0.30	
angry:RaceBlack	0.12	0.17	0.68	0.50		0.27	0.18	1.55	0.12	
angry:RaceWhite	-0.03	0.19	-0.18	0.86		-0.12	0.19	-0.65	0.52	
Opacity:RaceBlack	1.19	4.69	0.25	0.80		-5.23	6.23	-0.84	0.40	
Opacity:RaceWhite	3.60	4.66	0.77	0.44		-6.63	6.10	-1.09	0.28	
happy:angry:Opacity	0.00	0.00	0.98	0.33		0.00	0.00	1.79	*0.07*	.
happy:angry:RaceBlack	0.00	0.00	-0.55	0.59		-0.01	0.00	-1.51	0.13	
happy:angry:RaceWhite	0.00	0.00	0.37	0.71		0.00	0.00	0.54	0.59	
happy:Opacity:RaceBlack	-0.18	0.12	-1.58	0.12		0.17	0.18	0.96	0.34	
happy:Opacity:RaceWhite	-0.07	0.11	-0.60	0.55		0.11	0.17	0.66	0.51	
angry:Opacity:RaceBlack	0.44	0.24	1.80	*0.07*	.	0.03	0.24	0.13	0.90	
angry:Opacity:RaceWhite	0.55	0.26	2.11	**0.04**	*	0.21	0.25	0.82	0.41	
happy:angry:Opacity:RaceBlack	0.00	0.00	-0.91	0.36		0.00	0.01	-0.28	0.78	
happy:angry:Opacity:RaceWhite	-0.01	0.01	-1.81	*0.07*	.	0.00	0.01	-0.76	0.45	

However, the angry and happy measures both failed the assumption of homogeneity of regression slopes, a necessary assumption for analysis of covariance. Therefore, the two items were entered as factors rather than covariates in our subsequent analyses. Finally, avatar and human faces were analyzed separately since *humanness* was significantly different between the two groups.

5.2 Human Faces

Multiple regression was used to test experimental effects on *humanness* for human faces. Hierarchical models were built, sequentially adding opacity, angry, race, and happy measures, in that order. Starting with opacity, adding angry significantly improved the model fit, $\chi^2(2) = 758.53$, $\boldsymbol{p} < \mathbf{.0001}$. Adding race, again significantly improved model fit, $\chi^2(8) = 34.42$, $\boldsymbol{p} < \mathbf{.0001}$, and finally adding happy improved model fit, $\chi^2(12) = 40.91$, $\boldsymbol{p} < \mathbf{.0001}$. Adding participant gender as a covariate did not significantly improve the model and it was therefore removed. The final regression can be seen in table 1.

Significant main effects of angry, $t(1775) = -4.29$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .20$, and opacity, $t(1753) = 2.88$, $\boldsymbol{p} = \mathbf{.004}$, $d = .14$, were found, such that lower perceptions of anger and greater opacity were associated with higher ratings of humanness. These main effects were modified by a higher-order angry×opacity×race 3-way interaction, $t(1775) = 2.11$, $\boldsymbol{p} = \mathbf{.03}$, $d = .10$. The significant 3-way interaction was evaluated for each race by comparing slopes, pairwise, with Tukey adjustments. See figure 8. No significant differences in slope were identified for the Asian human face. However, for both the Black and White human faces significant differences were found between 100% opacity with color correction and all other levels. For the Black face, the slope at 100% opacity with color correction ($\beta = -.29$) was less steep compared to 100% opacity ($\beta = -.35$), $t(481) = -3.50$, $\boldsymbol{p} = \mathbf{.003}$, 68% opacity ($\beta = -.41$), $t(481) = -4.39$, $\boldsymbol{p} = \mathbf{.0001}$, and 35% opacity ($\beta = -.48$), $t(484) = -5.46$, $\boldsymbol{p} < \mathbf{.0001}$. A similar pattern was seen for the White face, the slope at 100% opacity with color correction ($\beta = -.31$) was less steep compared to 100% opacity ($\beta = -.57$), $t(484) = -4.31$, $\boldsymbol{p} = \mathbf{.0001}$, 68% opacity ($\beta = -.45$), $t(484) = -3.73$, $\boldsymbol{p} = \mathbf{.001}$, and 35% opacity ($\beta = -.56$), $t(481) = -4.62$, $\boldsymbol{p} < \mathbf{.0001}$. In other words, anger was less predictive of humanness for the White and Black fully opaque faces with color correction compared to the White and Black faces at all other opacity levels. Whereas the effect of anger on humanness did not differ based on opacity for Asian faces. For all faces, lowering opacity reduced humanness and faces that were perceived as angrier were perceived as less human. Further, the Black face was perceived to be significantly angrier than both the Asian and White faces (see figure 6) thus having a greater effect on perceived humanness at lower opacity levels (see figure 8, left).

5.3 Avatar Faces

The same hierarchical regression analysis was performed to test experimental effects on humanness for avatars. Starting with opacity, adding race significantly improved the model fit, $\chi^2(4) = 11.31$, $\boldsymbol{p} = \mathbf{.02}$. Adding angry, again significantly improved model fit, $\chi^2(6) = 947.68$, $\boldsymbol{p} < \mathbf{.0001}$, adding happy again improved the model fit $\chi^2(12) = 67.40$, $\boldsymbol{p} < \mathbf{.0001}$, and finally adding participant gender as a covariate improved the model $\chi^2(2) = 7.12$, $\boldsymbol{p} = \mathbf{.03}$. The final regression can be seen in table 1.

A significant main effect of gender (male, female, non-binary) was found, $F(2,157) = 3.80$, $\boldsymbol{p} = \mathbf{.03}$, $\eta^2 = .05$. Women ($M = 66.66$, $SD = 1.00$) gave significantly lower humanness ratings to avatars compared to men ($M = 72.62$, $SD = .89$), $t(157) = 2.74$, $\boldsymbol{p} = \mathbf{.02}$, $d = .15$. There was also a significant main effect of angry, $t(1794) = -4.92$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .23$. As found with human faces, avatars that were perceived to be angrier were perceived to be less human, ($\beta = -.62$), $\boldsymbol{p} < \mathbf{.0001}$.

Significant main effects of happy, $t(1793) = 2.57$, $\boldsymbol{p} = \mathbf{.01}$, $d = .12$, and opacity, $t(1759) = 2.80$, $\boldsymbol{p} = \mathbf{.005}$, $d = .13$, were qualified by the higher-order significant happy × opacity interaction, $t(1770) = -3.94$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .19$. The significant 2-way interaction was evaluated by comparing slopes, pairwise, with Tukey adjustments. Significant differences in slope were identified between 35% opacity ($\beta = .16$) and 100% opacity with color correction ($\beta = -.27$), $t(1752) = 4.47$, $\boldsymbol{p} < \mathbf{.0001}$, $d = .21$. Trends were found between 35% opacity and both the 100% opacity without color correction ($\beta = -.19$) and 68% opacity ($\beta = -.07$), $t(1746) = -2.43$, $p = .07$, $d = .12$ and $t(1740) = 2.44$, $p = .07$, $d = .12$ respectively. See figure 9. In essence, for the

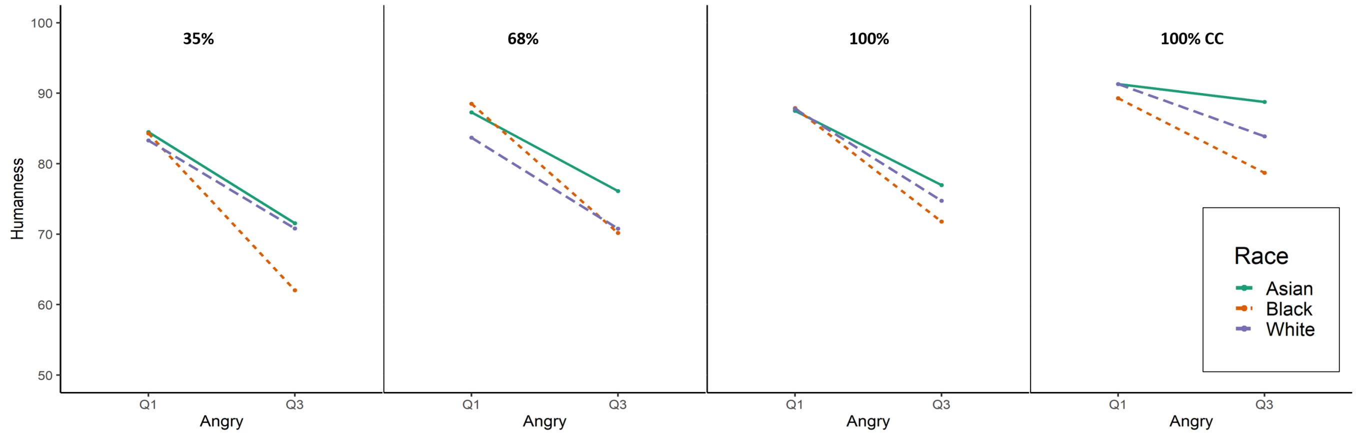

Fig. 8. A visual representation of the 3-way *humanness* × *opacity* × *race* interaction for human faces. Each graph displays data for one opacity level increasing from left (35% opacity) to right (100% with color correction) for each race (Asian: green, solid, Black: orange, small-dash, White: purple, long-dash). The humanness level (y-axis) based on each race's perceived angry level (x-axis) at quartile 1 (left) and quartile 3 (right). As opacity increases, humanness increases. As perceived angryness increases, humanness decreases.

most transparent avatars (35%), perceived happiness predicted greater humanness, whereas for the most opaque avatars (100% with color correction), happiness predicted lower humanness.

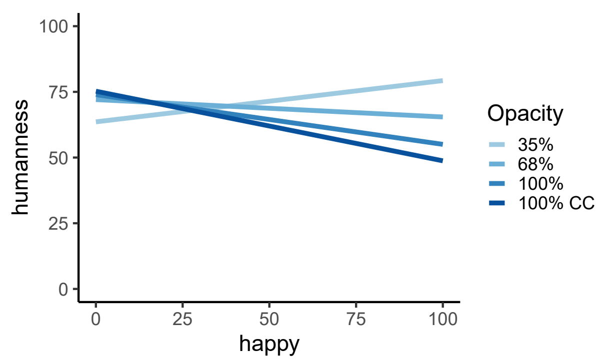

Fig. 9. The relationship between *happy* and *humanness* for avatar faces at each opacity level.

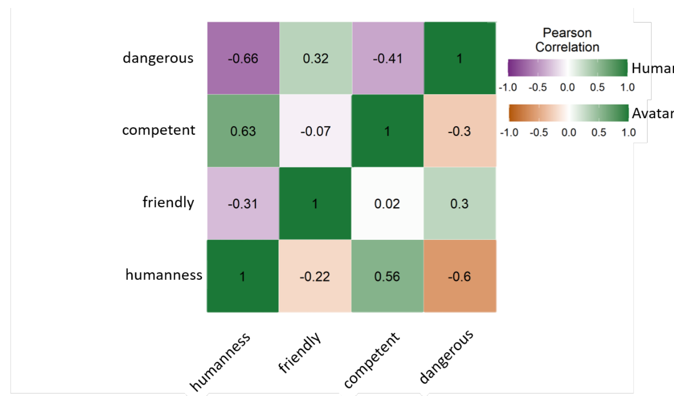

Fig. 10. Pearson correlation coefficient matrix between humanness, competent, friendly, and dangerous. The avatar face correlations are in the top left triangle and the human faces are in the bottom right.

5.4 Humanness Implications

Although the present study was exploratory, with the main goal being to test the effect of opacity on *humanness*, our final analyses explored possible implications of humanness. In other words, given that reduced opacity predicts lower ratings of humanness, what are the possible downstream effects of reduced humanness? To explore this, we computed bivariate correlation coefficients between humanness and the measures of stereotyping (competent, friendly, dangerous). The results are shown in figure 10. Humanness significantly correlated with friendly, competent, and dangerous for both avatars and humans with all $\boldsymbol{p} < \mathbf{.0001}$. Of note, greater humanness was associated with less danger (Human: $r = -.6$, Avatar: $r = -.66$), greater competence (Human: $r = .56$, Avatar: $r = .63$), and less friendliness (Human: $r = -.22$, Avatar: $r = -.31$).

6 Discussion

The primary goal of the present study was to explore whether a known inconsistency in the rendering of certain colors with OST-HMDs injected unintentional racial bias into AR applications. Because the color black appears transparent, this means that dark-skinned (i.e., Black) virtual humans are perceived by AR users as more transparent in the same environments in which White virtual humans are seen as less transparent. We tested whether opacity impacted perceptions of *humanness* broadly, and as a function of virtual human race. We chose humanness as our outcome measure because of the strong existing literature on the importance of humanness in virtual technology (e.g., Uncanny Valley research [44]) and the historic and ongoing role of dehumanization in perpetuating violence and exploitation of racial groups [21, 35].

Overall, across both static human and avatar faces, we found that opacity affected humanness. As opacity increased, perceptions of humanness increased. Women perceived avatar faces to be less human than men. Notably, we did not find that that race moderated this effect. In other words, regardless of human and avatar race, more transparent figures were perceived as less human. However, current OST-HMDs render lighter versus darker skin tones at different levels of transparency. While certain ambient lighting conditions can make all virtual avatars more or less transparent, transparency varies as a function of color (HSL) and therefore skin tone. Regardless of ambient lighting, virtual humans depicting White individuals will be rendered as more opaque than virtual humans depicting Black individuals. Assuming that the same dehumanization effects transfer from static images to OST-HMDS Black avatars will effectively be perceived as less human compared to their White counterparts.

Dehumanization has been shown to contribute to antipathy and violence across many different domains. For example, dehumanization of women is associated with men's willingness to rape and sexually harass women [52], dehumanization of Japanese and Haitian victims of natural disasters predicts less willingness to provide aid [3], and Christian participants' dehumanization of Muslims predicts increased willingness to torture Muslim prisoners of war [58]. On the positive side, perception of humanness is thought to be a powerful social cognition that can contribute to ameliorating racial prejudice [1, 2]. Even considering across species, research shows that highlighting the similarity of animals to humans (i.e., humanizing animals) increases moral concern for the welfare of animals, as well as moral concern for the welfare of marginalized human groups [4]. The importance of humanness for intergroup relations is clear, suggesting that the selective alteration of humanness through differential opacity rendering of certain racial groups by OST-HMDs is an issue worthy of continued study.

Although we initially included measures of emotion to control for potential variation in facial expression across the avatar and human images, we observed interesting interactive patterns of emotion and opacity in our data. Among human faces, anger tended to be associated with lower perceptions of humanness, particularly at lower levels of opacity. As opacity decreased, the slope between anger and dehumanization increased for both the White and Black faces. Further exacerbating the issue, Black Americans are viewed as angrier than White Americans [24, 31] a pattern that was replicated in our data. The Black human was perceived as angrier than the other humans which has a greater dehumanizing affect at lower opacity levels (see figure 8). Current display technology is limited to displaying virtual humans at the lowest opacity level we tested (35 Opacity); assuming that the results transfer to OST-HMDs, this suggests that Black virtual humans may be more dehumanized than virtual humans with lighter skin-tones.

We also found that the relationship between ratings of happiness and ratings of humanness varied as a function of opacity. For most levels, we observed a negative relationship such that happier-looking avatars appeared less human; however, this pattern was reversed for the most transparent avatars. Altogether, these exploratory findings regarding emotionality indicate that researchers and developers of AR technology need to consider how emotion is created within the virtual world. In the present study we utilized static images of faces. Cross-cultural research demonstrates remarkable consistency in the ability to detect and portray basic emotions through facial expressions [10]. However, emotion can be conveyed through a variety of methods, including voice, postural shifts, and gestures; an understanding of these factors leads to better development of embodied conversational agents or virtual avatars [40]. The findings of the present study suggest that AR researchers and developers should consider opacity when designing virtual agents meant to convey particular emotions.

6.1 Implications

AR is already in use in a variety of real-world applications, from medical training, to vehicle repair, to educational settings. Our results demonstrated that perceived humanness of static images was affected by opacity. Future work is needed to demonstrate that these results transfer to OST-HMDs. Assuming they transfer, the implications of opacity differences become important. Consider that a physician may use AR technology to assist in visualizing a patient's anatomy during surgery [15]. If the patient is Black and therefore their skin tone renders as partially transparent, results of the present study suggest that they may be viewed as less human. This perception could potentially exacerbate pre-existing tendencies to view Black individuals as having higher pain tolerances than other racial groups [56], resulting in less pain medicine being administered.

In the present study, we observed that perceived humanness was negatively correlated with perceived dangerousness and positively associated with perceived competence. Though exploratory, these relationships could suggest clear implications for real-world AR use. AR and VR technology is already being used by, and marketed toward, military and police for use in both training and live patrol. In the United States, Black men are associated with danger [9], and Black suspects are disproportionately shot and killed by police [53]. AR training applications that render Black virtual humans as more transparent than White virtual humans may reinforce military and police trainees' perceptions of danger through lowered humanness. In a lower stakes, yet still important, use of AR, colleagues may engage in AR remote collaborations [62]. If the virtual human representing a Black collaborator is more transparent, that may reduce perceptions of competence due to lowered humanness.

These implications are in need of direct testing. However, the results of the present study suggest multiple downstream consequences of dehumanization based on virtual human transparency in AR.

6.2 Limitations & Future Directions

The present study represents the first test of the effects of differential transparency on the perception of virtual humans' humanness. As such, all measures and materials were created for this research. The faces that were utilized were all young men, and only one specific face was used to represent each human/avatar $\times$ race category. Additionally, the Zombie human face was likely identifiable as a White man with zombie makeup. Future research should test the effects of opacity on multiple virtual human faces from each racial category as well as varying identities and varying representations (e.g., full body, interactive). Future research should also consider whether user identity characteristics impact perceptions based on transparency including cross-gender and cross-race interactions; VR researchers have called for increased research on diverse representation [50] that should apply to both users and avatars.

It is also important to consider the measurement of humanness. In the present study we combined multiple facets of humanness such as the experience of pain or the possession of emotions into one overarching measure of humanness. Measuring humanness in relation to the uncanny valley [30] especially concerning avatars should be further investigated. There may be important distinctions between these factors that have differential implications based on opacity. Future research should operationalize humanness in varied ways. A critical next step is to go beyond self-report measures of humanness and assess behavior directly, possibly including social presence [26]. Future research could investigate if opacity affects use of force among police officers, physician behavior toward virtual humans, or workplace team outcomes such as work product quality or efficiency. Directly assessing the behavioral implications of opacity and race within virtual humans in AR will determine the importance and urgency of developers creating designs to counteract these effects.

Since a single testing environment was used in the study, we cannot produce a model of how the results of the study would change with respect to factors such as background color and environment illuminance. However, previous research has already examined the effects that such factors have on user perception in optical see-through AR displays [11, 18]. Based on the existing research, increasing environment illuminance or lightening the background color would have the effect of reducing the contrast between the background and virtual human in the image, thus appearing more transparent, which would likely result in the images overall being rated as less human. Decreasing environment luminance or darkening the background color would have the opposite effect, increasing contrast and thus reducing transparency, likely producing results that were overall rated as being more human. Changing the hue of the background color would in turn shift the hues that comprise the virtual human towards that of the background color. This would result in virtual humans that appear more red, green, or blue than intended. Since it is not usual to see virtual humans of such hues, it is likely that overall the results would shift and be rated as less human overall. Future work could verify that this is the case by incorporating environment illuminance and background color as independent variables in a similar user study involving virtual humans.

The varying levels of opacity used in the study conditions show interesting implications for future work in AR displays, in that even if the contrast between the virtual imagery and the user's physical environment is increased, the imagery will still appear transparent due to the issue of color blending between virtual imagery and the user's physical environment. While progress is being made to understand and resolve the limitation of color blending [17, 18], we are likely years away from being able to adjust the coloration of virtual imagery in real time in order to correct for its effects. Further, even when color blending is resolved and solutions are integrated into consumer AR displays, we will still be faced with the issue that optical see-through displays cannot present imagery that is darker than the user's physical environment without blocking incoming light from the environment. This means that in certain conditions, such as bright settings or settings with high lightness colors, the appearance of a dark virtual image may still appear transparent. Research into techniques to correct this issue for optical see-through displays is still at an early stage. This is the case for both optical see-through *head-mounted* displays as well as *heads-up* displays, such as those integrated into a driver's windshield in the automotive industry [16]. As a result, it seems likely that the implications of the study presented here will continue to remain a relevant problem for the foreseeable future.

7 Conclusion

In this paper, we have presented an exploratory study on how user perceptions of virtual humans are affected by opacity. The virtual humans were designed to simulate the current limitations of optical see-through head-mounted display associated with the display technology, notably opacity and color blending. Our results suggest that virtual humans that appear more transparent are perceived as less human than more opaque virtual humans. Given that virtual humans with darker skin tones are more likely to be rendered as more transparent, the implication is that dark skinned virtual humans may be perceived as less human than those with lighter skin tones.

While these display limitations will eventually be overcome with time, it is likely that these issues will remain prominent in consumer optical see-through displays for the next several iterations of the technology. Because of this, current and future applications for such displays should carefully consider how these factors potentially introduce unintended bias into applications involving virtual humans.

Acknowledgments

Funding for this work was supported by Davidson College's Faculty Study and Research program. This material includes work supported in part by the National Science Foundation under Collaborative Award Numbers 1800961, 1800947, and 1800922 (Dr. Ephraim P. Glinert, IIS) to the University of Central Florida, University of Florida, and Stanford University, respectively, and the Office of Naval Research under Award Numbers N00014-21-1-2578 and N00014-17-1-2927 (Dr. Peter Squire, Code 34).

References

[1] F. Albarello, R. Crisp, and M. Rubini. Promoting beliefs in the inalienability of human rights by attributing uniquely human emotions through multiple categorization. *The Journal of social psychology*, 158(3):309–321, 2018.

[2] F. Albarello and M. Rubini. Reducing dehumanisation outcomes towards blacks: The role of multiple categorisation and of human identity. *European Journal of Social Psychology*, 42(7):875–882, 2012.

[3] L. Andrighetto, C. Baldissarri, S. Lattanzio, S. Loughnan, and C. Volpato. Human-itarian aid? two forms of dehumanization and willingness to help after natural disasters. *British journal of social psychology*, 53(3):573–584, 2014.

[4] B. Bastian, K. Costello, S. Loughnan, and G. Hodson. When closing the human–animal divide expands moral concern: The importance of framing. *Social Psychological and Personality Science*, 3(4):421–429, 2012.

[5] O. Cakmakci, Y. Ha, and J. Rolland. A compact optical see-through head-worn display with occlusion support. In *Third IEEE and ACM International Symposium on Mixed and Augmented Reality*, pp. 16–25, 2004.

[6] M. L. Chenechal, T. Duval, V. Gouranton, J. Royan, and B. Arnaldi. Vishnu: virtual immersive support for helping users an interaction paradigm for collaborative remote guiding in mixed reality. In *2016 IEEE Third VR International Workshop on Collaborative Virtual Environments (3DCVE)*, pp. 9–12, 2016.

[7] S. de Cunsel. Evaluation of augmented reality (AR) displays performance based on human visual perception. In B. C. Kress and C. Peroz, eds., *Digital Optics for Immersive Displays II*, vol. 11350, pp. 47 – 53. International Society for Optics and Photonics, SPIE, 2020.

[8] P. G. Devine. Stereotypes and prejudice: Their automatic and controlled components. *Journal of personality and social psychology*, 56(1):5–18, 1989.

[9] P. G. Devine and A. J. Elliot. Are racial stereotypes really fading? the princeton trilogy revisited. *Personality and social psychology bulletin*, 21(11):1139–1150, 1995.

[10] P. Ekman, W. V. Friesen, M. O'sullivan, A. Chan, I. Diacoyanni-Tarlatzis, K. Heider, R. Krause, W. A. LeCompte, T. Pitcairn, P. E. Ricci-Bitti, et al. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of personality and social psychology*, 53(4):712–717, 1987.

[11] A. Erickson, K. Kim, G. Bruder, and G. F. Welch. Exploring the limitations of environment lighting on optical see-through head-mounted displays. In *Symposium on Spatial User Interaction*, SUI '20, pp. 1–8. Association for Computing Machinery, New York, NY, USA, 2020.

[12] A. Erickson, K. Kim, A. Lambert, G. Bruder, M. P. Browne, and G. F. Welch. An extended analysis on the benefits of dark mode user interfaces in optical see-through head-mounted displays. *ACM Transactions on Applied Perception*, 18(3):1–22, May 2021.

[13] C. Feliciano. Shades of race: How phenotype and observer characteristics shape racial classification. *American Behavioral Scientist*, 60(4):390–419, 2016.

[14] S. T. Fiske, A. J. Cuddy, P. Glick, and J. Xu. A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of personality and social psychology*, 82(6):878–902, 2002.

[15] H. Fuchs, H. Yang, T. Peck, A. Bulysheva, et al. Methods, systems, and computer readable media for image guided ablation, Feb. 23 2016. US Patent 9,265,572.

[16] J. L. Gabbard, G. M. Fitch, and H. Kim. Behind the Glass: Driver Challenges and Opportunities for AR Automotive Applications. *Proceedings of the IEEE*, 102(2):124–136, 2014.

[17] J. L. Gabbard, M. Smith, C. Merenda, G. Burnett, and D. R. Large. A perceptual color-matching method for examining color blending in augmented reality head-up display graphics. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–18, 2020.

[18] J. L. Gabbard, J. E. Swan, J. Zedlitz, and W. W. Winchester. More than meets the eye: An engineering study to empirically examine the blending of real and virtual color spaces. In *2010 IEEE Virtual Reality Conference (VR)*, pp. 79–86, 2010.

[19] H. J. Geiger. Racial and ethnic disparities in diagnosis and treatment: a review of the evidence and a consideration of causes. *Unequal treatment: Confronting racial and ethnic disparities in health care*, 417:1–38, 2003.

[20] A. C. S. Genay, A. Lecuyer, and M. Hachet. Being an avatar "for real": a survey on virtual embodiment in augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–21, 2021.

[21] P. A. Goff, J. L. Eberhardt, M. J. Williams, and M. C. Jackson. Not yet human: implicit knowledge, historical dehumanization, and contemporary consequences. *Journal of personality and social psychology*, 94(2):292–306, 2008.

[22] M. Gonzalez-Franco and T. C. Peck. Avatar embodiment. towards a standardized questionnaire. *Frontiers in Robotics and AI*, 5:74, 2018.

[23] R. Grasset, P. Lamb, and M. Billinghurst. Evaluation of mixed-space collaboration. In *Fourth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'05)*, pp. 90–99, 2005.

[24] A. G. Halberstadt, V. L. Castro, Q. Chu, F. T. Lozada, and C. M. Sims. Preservice teachers' racialized emotion recognition, anger bias, and hostility attributions. *Contemporary Educational Psychology*, 54:125–138, 2018.

[25] L. Hannon. White colorism. *Social Currents*, 2(1):13–21, 2015.

[26] C. Harms and F. Biocca. Internal consistency and reliability of the networked minds measure of social presence. 2004.

[27] N. Haslam, Y. Kashima, S. Loughnan, J. Shi, and C. Suitner. Subhuman, inhuman, and superhuman: Contrasting humans with nonhumans in three cultures. *Social Cognition*, 26(2):248–258, 2008.

[28] N. Haslam and S. Loughnan. Dehumanization and infrahumanization. *Annual review of psychology*, 65:399–423, 2014.

[29] H. A. Hewes, M. Dai, N. C. Mann, T. Baca, and P. Taillac. Prehospital pain management: disparity by age and race. *Prehospital Emergency Care*, 22(2):189–197, 2018.

[30] C.-C. Ho and K. F. MacDorman. Measuring the uncanny valley effect. *International Journal of Social Robotics*, 9(1):129–139, 2017.

[31] K. Hugenberg and G. V. Bodenhausen. Facing prejudice: Implicit preju-

dice and the perception of facial threat. *Psychological Science*, 14(6):640–643, 2003.

[32] M. Hunter. The consequences of colorism. In *The melanin millennium*, pp. 247–256. Springer, 2013.

[33] Y. Itoh, T. Langlotz, D. Iwai, K. Kiyokawa, and T. Amano. Light attenuation display: Subtractive see-through near-eye display via spatial color filtering. *IEEE Transactions on Visualization and Computer Graphics*, 25(5):1951–1960, 2019.

[34] Y. Itoh, T. Langlotz, J. Sutton, and A. Plopski. Towards indistinguishable augmented reality: A survey on optical see-through head-mounted displays. *ACM Comput. Surv.*, 54(6), July 2021.

[35] K. Jacoby. Slaves by nature? domestic animals and human slaves. *Slavery and Abolition*, 15(1):89–99, 1994.

[36] K. Kilteni, R. Groten, and M. Slater. The sense of embodiment in virtual reality. *Presence: Teleoperators and Virtual Environments*, 21(4):373–387, 2012.

[37] K. Kiyokawa, H. Iwasa, H. Takemura, and N. Yokoya. Collaborative immersive workspace through a shared augmented environment. In B. Gopalakrishnan and S. Murugesan, eds., *Intelligent Systems in Design and Manufacturing*, vol. 3517, pp. 2 – 13. International Society for Optics and Photonics, SPIE, 1998.

[38] J. T. Kubota and T. Ito. Rapid race perception despite individuation and accuracy goals. *Social neuroscience*, 12(4):468–478, 2017.

[39] Y.-H. Lee, T. Zhan, and S.-T. Wu. Prospects and challenges in augmented reality displays. *Virtual Real. Intell. Hardw.*, 1(1):10–20, 2019.

[40] M. Lhommet and S. C. Marsella. Expressing emotion through posture. *The Oxford handbook of affective computing*, 273:273–285, 2014.

[41] J. Lin, A. Drain, A. Goharzad, and P. Mende-Siedlecki. What factors fuel racial bias in pain perception and treatment?: A meta-analysis across 40 experimental studies. pp. 1–40, 2020.

[42] R. J. Little and D. B. Rubin. *Statistical analysis with missing data*, vol. 793. John Wiley & Sons, 2019.

[43] T. D. Little, T. D. Jorgensen, K. M. Lang, and E. W. G. Moore. On the joys of missing data. *Journal of pediatric psychology*, 39(2):151–162, 2014.

[44] M. Mori. Bukimi no tani [the uncanny valley]. *Energy*, 7:33–35, 1970.

[45] A. Morning. *The nature of race*. University of California Press, 2011.

[46] N. Norouzi, K. Kim, G. Bruder, A. Erickson, Z. Choudhary, Y. Li, and G. Welch. A Systematic Literature Review of Embodied Augmented Reality Agents in Head-Mounted Display Environments. In F. Argelaguet, R. McMahan, and M. Sugimoto, eds., *ICAT-EGVE 2020 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*, pp. 101–111. The Eurographics Association, 2020.

[47] I. N. Onyeador, S.-k. T. Hudson, and N. A. Lewis Jr. Moving beyond implicit bias training: Policy insights for increasing organizational diversity. *Policy Insights from the Behavioral and Brain Sciences*, 8(1):19–26, 2021.

[48] S. Orts-Escolano, C. Rhemann, S. Fanello, W. Chang, A. Kowdle, Y. Degtyarev, D. Kim, P. L. Davidson, S. Khamis, M. Dou, V. Tankovich, C. Loop, Q. Cai, P. A. Chou, S. Mennicken, J. Valentin, V. Pradeep, S. Wang, S. B. Kang, P. Kohli, Y. Lutchyn, C. Keskin, and S. Izadi. Holoportation: Virtual 3D Teleportation in Real-Time. In *Proceedings of the Symposium on User Interface Software and Technology*, pp. 741–754, 2016.

[49] T. C. Peck, J. J. Good, and K. Seitz. Evidence of racial bias using immersive virtual reality: Analysis of head and hand motions during shooting decisions. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2502–2512, 2021.

[50] T. C. Peck, L. E. Sockol, and S. M. Hancock. Mind the gap: The underrepresentation of female participants and authors in virtual reality research. *IEEE transactions on visualization and computer graphics*, 26(5):1945–1954, 2020.

[51] T. Piumsomboon, G. A. Lee, J. D. Hart, B. Ens, R. W. Lindeman, B. H. Thomas, and M. Billinghurst. *Mini-Me: An Adaptive Avatar for Mixed Reality Remote Collaboration*, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2018.

[52] L. A. Rudman and K. Mescher. Of animals and objects: Men's implicit dehumanization of women and likelihood of sexual aggression. *Personality and Social Psychology Bulletin*, 38(6):734–746, 2012.

[53] K. Scott, D. S. Ma, M. S. Sadler, and J. Correll. A social scientific approach toward understanding racial disparities in police shooting: Data from the department of justice (1980–2000). *Journal of Social Issues*, 73(4):701–722, 2017.

[54] A. Smedley and B. Smedley. *Race in North America: Origin and evolution of a worldview*. Routledge, 2018.

[55] J. C. Thompson, J. G. Trafton, and P. McKnight. The perception of humanness from the movements of synthetic agents. *Perception*, 40(6):695–704, 2011.

[56] S. Trawalter, K. M. Hoffman, and A. Waytz. Racial bias in perceptions of others' pain. *PloS one*, 7(11):e48546, 2012.

[57] S. Utzig, R. Kaps, S. M. Azeem, and A. Gerndt. Augmented reality for remote collaboration in aircraft maintenance tasks. In *2019 IEEE Aerospace Conference*, pp. 1–10, 2019.

[58] G. T. Viki, D. Osgood, and S. Phillips. Dehumanization and self-reported proclivity to torture prisoners of war. *Journal of Experimental Social Psychology*, 49(3):325–328, 2013.

[59] K. Waldow, A. Fuhrmann, and S. M. Grünvogel. Investigating the effect of embodied visualization in remote collaborative augmented reality. In P. Bourdot, V. Interrante, L. Nedel, N. Magnenat-Thalmann, and G. Zachmann, eds., *Virtual Reality and Augmented Reality*, pp. 246–262. Springer International Publishing, Cham, 2019.

[60] A. Waytz, K. M. Hoffman, and S. Trawalter. A superhumanization bias in whites' perceptions of blacks. *Social Psychological and Personality Science*, 6(3):352–359, 2015.

[61] B. Yoon, H.-i. Kim, G. A. Lee, M. Billinghurst, and W. Woo. The effect of avatar appearance on social presence in an augmented reality remote collaboration. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 547–556, 2019.

[62] B. Yoon, H.-i. Kim, G. A. Lee, M. Billinghurst, and W. Woo. The effect of avatar appearance on social presence in an augmented reality remote collaboration. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 547–556, 2019.

[63] K. Yu, G. Gorbachev, U. Eck, F. Pankratz, N. Navab, and D. Roth. Avatars for teleconsultation: Effects of avatar embodiment techniques on user perception in 3d asymmetric telepresence. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4129–4139, 2021.

[64] A. Zlobina and R. Andujar. Police violence during citizens' protests: Dehumanization processes, legitimation of violence, and hostility. *Peace and Conflict: Journal of Peace Psychology*, 2021.