The Effects of Embodiment in Virtual Reality on Implicit Gender Bias

Stephanie Schulze(✉), Toni Pence(✉), Ned Irvine(✉), and Curry Guinn(✉)

Univeristy of North Carolina at Wilmington, Wilmington, NC 28403, USA
{ss4134,pencet,irvinen,guinn}@uncw.edu

Abstract. Virtual reality allows users to have a virtual body that is different from their physical body, an idea known as embodiment. In previous research, embodiment in different types of avatars affected implicit attitudes. The purpose of this experiment was to discover how embodiment in different gendered avatars in virtual reality affects implicit gender bias. For embodiment, participants were placed in an office virtual environment with a male or female avatar. First, there was an orientation period where participants grew accustomed to their virtual body while looking at a mirror placed in front of them. Next, virtual humans of different genders walked in and out of the office with the mirror in view. Each participant completed a gender and leadership Implicit Association Test before and after the embodiment experience. The difference between post test scores and preliminary test scores indicates how implicit bias was affected.

Keywords: Virtual reality · Embodiment · Gender bias

1 Introduction and Related Work

Virtual reality has a multitude of applications. One use is to provide a virtual body in place of the physical one, known as embodiment. Whether the virtual body is similar or quite different than the physical body, it is possible for a user to transfer body ownership to the virtual body [1]. The idea of body ownership stems from the rubber hand illusion where a rubber hand is placed in front of a participant and is stroked at the same time as their real hand which remains hidden [2]. Participants reacted to the rubber hand as if it were their own [2]. In this case, ownership was achieved by synchronous visual and tactile stimulation as well as proprioception [2]. However, when a variation of the rubber hand illusion was implemented in virtual reality using a data glove, it was shown that visual and proprioception synchronicity along with motor activity created ownership [3]. Essentially, tactile information was not necessary. The illusion of ownership with a full-body avatar in virtual reality can also be created [1]. For example, adult males were able to take ownership of a female child avatar; when

Supported by University of North Carolina at Wilmington.

J. Y. C. Chen and G. Fragomeni (Eds.): HCII 2019, LNCS 11574, pp. 361–374, 2019.
https://doi.org/10.1007/978-3-030-21607-8_28

the avatar was threatened by another virtual human, heart rate accelerated [1]. In this experiment, for creating body ownership in virtual reality, having a first person perspective from the avatar was the most important condition while visuomotor synchronicity also was a key factor [1].

Being embodied and taking ownership of a virtual body can have an effect on implicit attitude and perceptions. Known as the Proteus Effect, users conform to the behavior that they believe others would expect them to have based on their avatar [4]. Exemplified in virtual reality, when embodied in a tall avatar as opposed to short, participants were more confident/aggressive in negotiation than those in the short avatar condition [4]. Because a common perception of tall people is that they are more confident, participants conformed to that behavior. Similarly, a participant approached a virtual human more closely when embodied in an attractive avatar over an unattractive avatar [4]. The perception of virtual objects within virtual reality was also shown to be affected based on the avatar [5]. In this experiment, participants were embodied in either a child avatar or an adult avatar scaled down to the same height as the avatar and asked to estimate the size of a series of virtual objects [5]. When there was strong body ownership, those embodied in the child avatar approximated the objects to be relatively twice their actual size [5]. This effect was not seen in those embodied in the adult avatar [5].

Adding to implicit attitude and perceptions, embodiment with body ownership in virtual reality can have an effect on bias. Participants who were embodied as either an elderly avatar or young avatar were asked to complete a word association task [6]. Those embodied in the old avatar were significantly more positive to the elderly than those embodied in the young avatar. There are several studies where implicit racial bias was affected by embodiment. In one in particular, the implicit racial bias against African Americans was actually higher for participants embodied in African American avatars than participants embodied in Caucasian avatars, regardless of participant race [7]. For this experiment participants engaged in an job interview in virtual reality with either an African American or Caucasian model [7]. An explanation as to why this bias occurred is because implicit racial bias typically exists in interviews [8]. Contrasting results occurred in a study also examining implicit racial body and embodiment, where the scenario was neutral. Participants were either embodied in a light skin avatar, a dark skin avatar, a purple 'alien' skin avatar, or not embodied at all [9]. While embodied, participants observed their body in a virtual mirror before a series of virtual humans, alternating in race, walked past. For those who were not embodied, or embodied in the purple skin avatar, implicit racial bias was affected less than those participants who were embodied in light skin and embodied in dark skin conditions. The implicit gender bias in participants embodied in the light skin avatar condition increased, while it decreased for those in embodied dark skin avatar condition. An extension of this experiment found that this decrease in implicit racial bias can be sustained for at least a week following embodiment experience [8].

While there is extensive research on virtual embodiment, there is a lack of research on gender embodiment and the effects on implicit gender bias in virtual reality. The intent of this experiment is to discover if users can take ownership of a virtual body of a different gender and, with this body ownership, explore how embodiment affects implicit gender bias.

2 Materials and Methods

2.1 Experiment Design

The purpose of conducting this experiment is to discover if embodiment into each gender avatar reduces implicit gender bias. A 2×2 between groups design of gender of participant vs. gender of avatar was used for this experiment. Within each participant gender group, the gender of the avatar assigned to the participant alternated each time. The conditions were: male participant with male avatar, male participant with female avatar, female participant with male avatar, and female participant with female avatar. A gender and leadership Implicit Association Test (IAT) was completed before and after embodiment.

2.2 Technical Details

Physical Environment. The experiment was conducted in a private room within the computer and information science graduate lab on the University of North Carolina Wilmington's campus (UNCW). Participants used HTC Vive head mounted display (HMD) and controllers. The HMD offers a field of view of 110° and a resolution of 1080×1200 pixels per eye. Two HTC base stations set up in opposing corners track six degrees of freedom (x, y, z, yaw, pitch, roll) for the HMD and controllers.

Virtual Environment. The virtual environment participants were put in was a manager's office pictured in Fig. 1. The environment, "Manager Office Interior" by 3D Everything, was downloaded from the Unity Asset Store. What appears to be a white plane in Fig. 1 is a working mirror for those wearing the HMD. The mirror asset was abstracted from an HTC plugins demo scene which was downloaded from the Unity Asset Store.

Virtual Humans. The virtual humans and participant avatars were created using MakeHuman 1.1.1 software. This application allows you to fine tune details of a 3D model and apply a variety of hair and clothing. The models can also be exported with a skeleton to facilitate animation in Unity. Figure 2 features the basic customization elements with a model on the left and the model with the CMU-Compliant skeleton on the right. There were a total of eight models created: four female and four male. Out of these models, one female and one male were used for participant avatars. All models are 158 cm in height and Caucasian. The virtual humans that are used, exclusive of participant avatars, can be seen in Fig. 3.

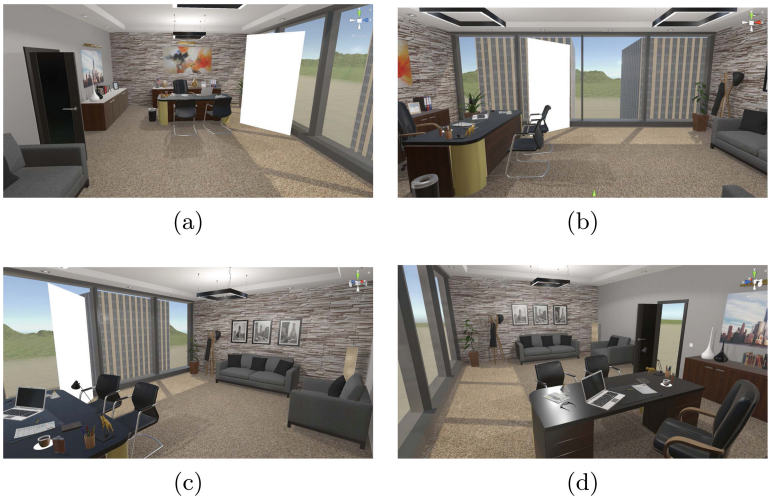

(a) (b)

(c) (d)

Fig. 1. The office environment

Animations. For the virtual humans in the scene, the idle neutral, walk forward, and turn short animation clips from Raw Mocap Data for Mecanim were used. These clips were used in conjunction with Unity's mecanim animation system to have each virtual human walk through the office door to the other side of the room, turn around, and walk out the office door. The first virtual human to walk through the door is a male. When he walks out, a female walks in. The animations continue like this, alternating genders, until all six virtual humans have walked out of the office.

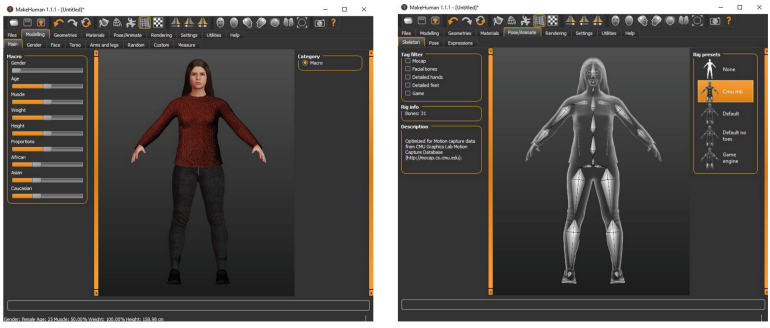

(a) (b)

Fig. 2. MakeHuman interface

The participants' avatar is animated using FinalIK by RootMotion. FinalIK takes the bones of a humanoid skeleton so when one bone is moved, related bones

are also moved creating animation that appears natural. This technique is called inverse kinematics. The position and rotation of the participants head and hands were tracked and reflected by their avatar in the virtual environment. While the exact position and rotation of the other parts of the participants body were not reflected, FinalIK calculated a relative position for corresponding bones of the virtual body.

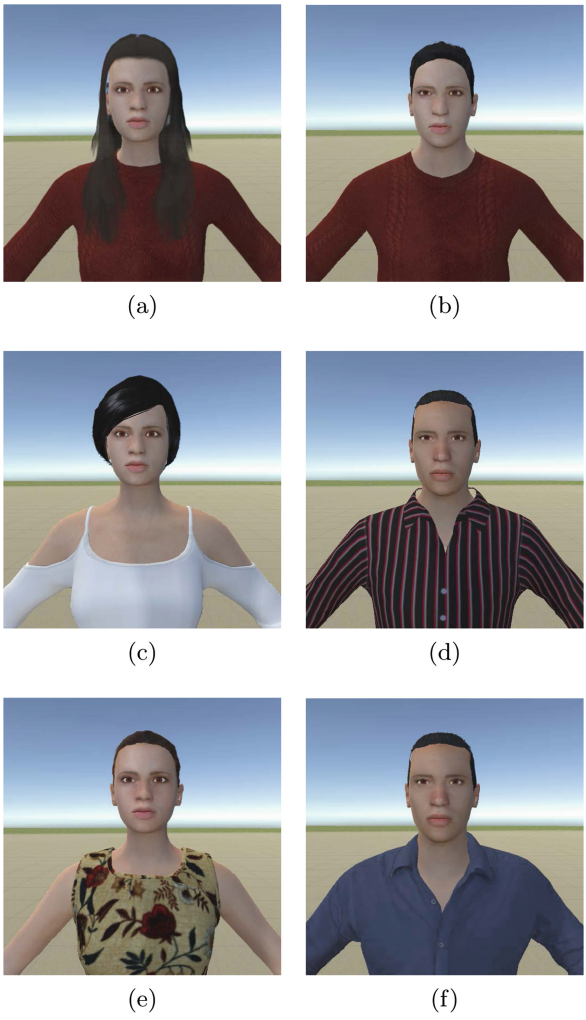

Fig. 3. The virtual humans used

2.3 Participants

A total of sixteen participants took part in this study. There were eleven males and five females. Broken down into each condition, male with male avatar had five participants, male with female avatar had six participants, female with male avatar had two participants, and female with female avatar had three participants. Fifth percent of the participants were aged eighteen to twenty-two. Thirty seven and a half percent of the participants were aged twenty-three to twenty-seven. Twelve and a half percent of the participants were over 47. Regarding race, fifteen participants were Caucasian and one participant was African American.

Participants were recruited from various classes in UNCW's computer science department and information technology department. Additional participants were acquired by asking people in the computer and information science building on UNCW's campus if they would like to participate.

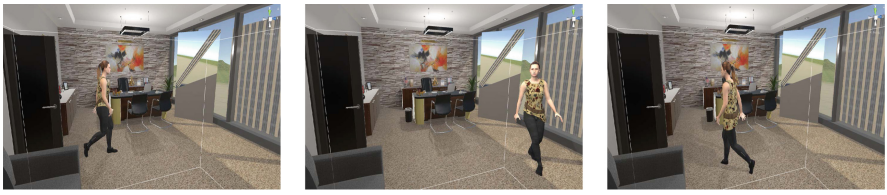

(a) (b) (c)

Fig. 4. Virtual humans walking in the office environment

2.4 Procedures

First, on a desktop computer, participants completed a preliminary questionnaire regarding demographics and previous experience with virtual reality. Then, also on the desktop computer, they completed a preliminary Implicit Association Test (IAT) on gender and leadership. Next, the participant put on the head mounted display (HMD) and took hold of the controllers to be placed in the immersive virtual office. In the environment, an avatar of the assigned gender is in place of the participant's physical body. The first time in this environment was the orientation period. A virtual mirror was placed in front of the participant so they could see their avatar. With controller functionality, they able were able to adjust avatar height and arm length so they corresponded with physical body proportions. The virtual mirror was placed in front of the participant as they performed a series of physical movements. The participant could look down at their virtual body to see their physical movements reflected in their avatar as well as in the mirror. Following the exercises, the participant was allowed 2 min to explore avatar movements. This process allowed the participants to become accustomed to their virtual avatar. After the orientation period, the participant was told a series of virtual humans will walk into the room. The environment remained identical to the environment in the orientation period, except the mirror was located further away to allow room for the virtual humans to walk in

front of the participant. The path the virtual human took through the virtual office is pictured in Fig. 4. The participant was then removed from the virtual environment and their avatar. Next, they took the gender and leadership IAT again on the desktop computer and results were recorded. Finally, the participant completed a post questionnaire regarding body ownership and what they thought the purpose of the experiment was.

2.5 Response Variables

Body Ownership. After being embodied in an avatar, participants answered questions relating to the level of body ownership they felt (Table 1). Each statement was presented with a one to five Likert scale with one being strongly disagree and five being strongly agree. While the variable Nervous was an inquiry, it was not used to determine body ownership.

Table 1. Post experience questionnaire statements

Variable	Statement
My body	I felt that the virtual body when looking down at myself was my own body
Two bodies	I felt as if I had two bodies
Mirror	I felt that the virtual body I saw when looking in the mirror was my own body
Features	I felt that my virtual body resembled my own (real) body in terms of shape, skin tone, or other visual features
Agency	I felt that the movements of the virtual body were caused by my own movements
Nervous	I became nervous when the other avatars approached me

Implicit Association Test. To measure implicit gender bias of participants, a gender and leadership implicit association test (IAT) was used. The IAT was created using FreeIAT 1.3.3. This test required users to rapidly categorize male and female faces of virtual humans used in the environment (Fig. 3) with words associated with a leader and a supporter (Table 2). The test consisted of five different stages each with five trials. Though the participants did not know, the first stage is a learning trial where users practice categorizing the virtual human images into female and male. Similarly, in the second stage users practice categorizing different words into leader and supporter. The third stage is no longer practice. Here, users are presented with both images and words to categorize into either Female-Leader group or Male-Supporter group. Stage four is another learning trial with the images, but the female and male category labels switch sides. This is to break up any familiarity with associations. The fifth and final stage is akin to the third stage. Both images and words are presented, but this time participants must categorize into either Male-Leader group or Female-Supporter group.

Table 2. Word categoriazation for IAT

Label	Words in group
Leader	leader, ambitious, determined, dedicated, assertive, manager
Supporter	supporter, understanding, sympathetic, compassionate, follower, assistant

The scores of the IAT are related to response time and accuracy in categorization [10]. The faster the reaction time, the more biased a user is to that category pairing. A positive score indicates that the user exhibits preference for Female-Leader and Male-Supporter pairings over Male-Leader and Female-Supporter pairings. Likewise, a negative score indicates that the user exhibits preference for Male-Leader and Female-Supporter pairings over Female-Leader and Male-Supporter pairings.

To determine how the embodiment experience affects a participant, the IAT was adminstered both before and after the experience. The difference between the score after the experience and the score before the experience provides the change in IAT (changeIAT = postIAT − preIAT). In the changeIAT case, a positive score signifies a reduction in implicit gender bias. Similarly, a negative score indicates an increase in implicit gender bias.

3 Results and Discussion

3.1 Body Ownership

From the post experience questionnaire shown in Table 1, variables MyBody, Mirror, Features, and Agency are used to ascertain the degree of body ownership participants felt with their given avatar. The TwoBodies variable is not used because participants expressed confusion, which is reflected in the data. Figure 5 shows the mean of each body ownership variable per condition.

For the MyBody variable, ten participants indicated a four or five signifying they agree that the avatar body felt like their own. Three participants reported a three meaning they neither agree nor disagree the avatar body felt like their own. Finally, three participants marked a one or two, expressing that they disagree the avatar body felt like theirs. Overall, the mean response to MyBody was 3.625. Specified by condition, female with female avatar has the highest mean response of 4. Following is male with female avatar with a mean of 3.6666. Next is female with male avatar with a mean of 3.5. Lastly, those in male with male avatar condition reported a mean of 3.4.

When responding to the Mirror variable, ten participants replied with a four or a five. These participants agreed that they felt the virtual body when looking in the mirror was theirs. Neither agreeing nor disagreeing with the statement, two participants reported a 2. Four participants marked a two or a one. The mean

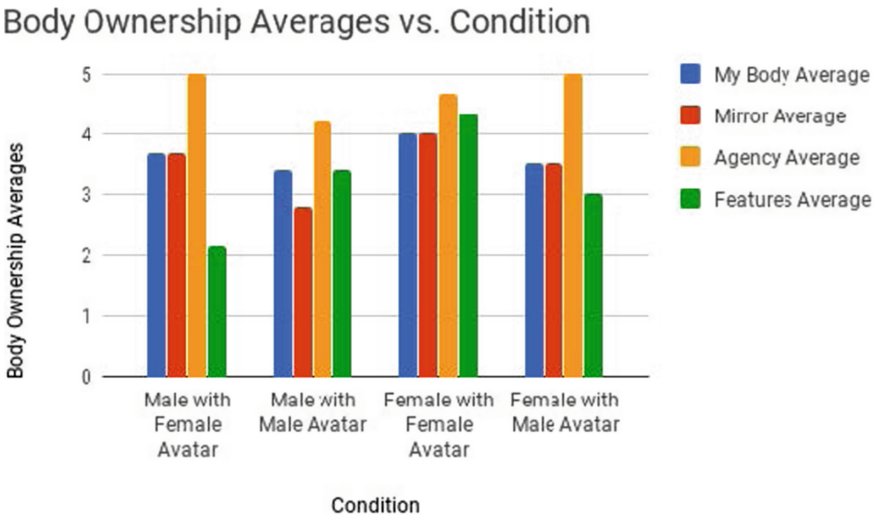

Fig. 5. Body ownership averages by condition

response for all participants for the Mirror variable was 3.4375. Having a mean over 3 are the conditions: male with female avatar, female with female avatar, and female with male avatar. Their means are 3.6666, 4, and 3.5 respectively. One thing to note is that for these three conditions, the Mirror mean is equivalent to the MyBody mean. Differing from the MyBody mean, the Mirror mean for the male with male avatar condition dips below 3 at a 2.8. This conveys that on average, male participants given a male avatar do not feel as if the virtual body they saw in the mirror was their own.

Participants responded to the Features variable to indicate to what degree they felt their virtual body resembled their physical body in terms of shape, skin tone, or other visual features. The variation of responses for Features is much more than the other variables. Six participants responded with a four or five, five participants responded with a three, and five participants responded with a two or one. However, the mean for all participants equates to 3.0625. While only slightly over three, this suggests that overall participants felt their virtual body resembled their physical one. Those in the conditions where participant gender matched their avatars, the Implicit Association Test resulted in higher means than those with mismatched gender. Female with female avatar amounted to 4.3333 and male with male avatar 3.4. The male with female avatar and female with male avatar had means 2.1667 and 3 respectively. The male with male avatar mean for Features is noticeably lower than the other conditions.

The Agency variable refers to how much a participant felt that the movements of the virtual body were caused by their own movements. Fifteen participants indicated a four or five while one participant indicated a two. Equating to 4.6875, the overall mean for Agency was higher than the overall mean for any other variable. In all conditions, the mean was above four. The mean for both the

male with female avatar condition and female with male avatar condition was 5. The mean for male with male avatar condition was 4.2 and the mean for female with female avatar was 4.6667.

3.2 Implicit Association Test

To gauge the effect embodiment has on implicit gender bias, participants took a gender and leadership implicit association test prior to embodiment (preIAT) and after embodiment (postIAT). Recall that a positive score indicates an implicit bias against males while a negative score indicates an implicit bias against females. The further the score is away from zero, the stronger the bias. Subtracting preIAT scores from postIAT scores results in the change of IAT scores: either reduced or increased.

Looking at the left figure of Fig. 6, the mean postIAT score increases from the mean preIAT score for the male with female avatar, male with male avatar, and female with female avatar conditions. This means on average participants in these conditions became more biased against females. In all of these conditions, the average participant initially revealed implicit gender bias against females in the preIAT with mean scores of $-.6218$, $-.82954$, and $-.15393$ respectively. However, for the remaining condition, female with male avatar, a reduction in bias against males is seen. The mean postIAT score is reduced from the mean preIAT score. In this condition, the average participant initially revealed implicit gender bias against males in the preIAT with a mean score of .3544.

When analyzing the left figure of Fig. 6, the degree of bias in both conditions with male participants is noticeably greater than the degree of bias in both conditions with female participants. The right figure of Fig. 6 organizes the data by gender to more closely examine this insight. On average, female participants initially express implicit gender bias against males with a mean of .0494 for preIAT scores. After the embodiment, there is a switch and implicit gender bias against males is expressed with a mean of $-.077$. For the female average scores, both the preIAT and post IAT scores never exceed $\pm.1$. This is not the case for the male participant average scores. The mean preIAT score is $-.7162$ and the mean postIAT score is $-.8606$. Both of IAT mean scores for males surpasses $-.1$, implying that the male participants have stronger implicit gender biases than female participants.

The results of individual participants can be evaluated for more specification. While thirteen participants scored negative on the preIAT, indicating implicit gender bias against females, three participants scored positive on the preIAT, indicating implicit gender bias against males. For the purpose of evaluating the change in IAT scores equivalently across data, the preIAT and postIAT scores of these participants are negated. Therefore, a negative change in IAT implies an increase in implicit bias against males while a positive change in IAT implies a decrease in implicit bias against males. The other thirteen participants' change in IAT remains interpreted as increase in implicit bias against females for negative values and decrease in implicit bias against females for positive values. As the data is represented together, a negative change indicates an increase in implicit

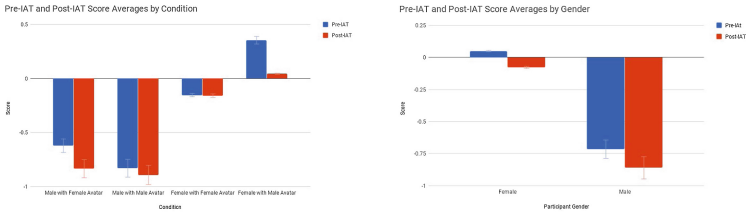

(a) (b)

Fig. 6. Pre-IAT and Post-IAT score averages by condition and by gender

gender bias while a positive change indicates a reduction of implicit gender bias, regardless of which bias.

In Fig. 7, each participant's change in IAT score is provided. On average, the mean change in IAT was a .0976, meaning participants experienced a slight improvement in gender bias. The participants appear to be well-balanced in whether their implicit bias increased or decreased. Seven of the participants improved their IAT score after being embodied, implying a reduction in implicit gender bias. The other nine participants exemplified a worse IAT score after being embodied signifying an increase in implicit gender bias. The mean of the degree of improvement for those participants whose bias decreased was .83717. This is a larger degree of change than for those whose bias increased. Here, the mean degree of decrement was −.4776. Grouping by gender, IAT change was relatively equitable with females on average experiencing a .18536 change and males on average experiencing a .0577 change.

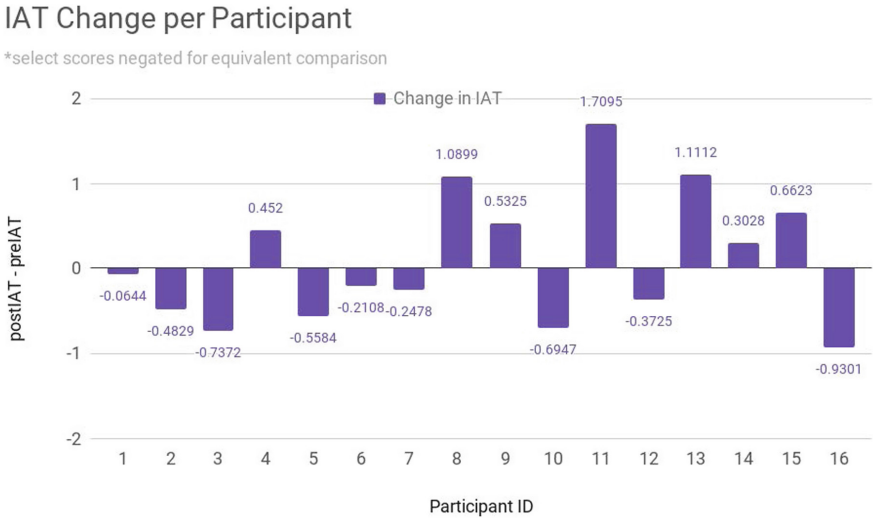

Fig. 7. IAT change per participant

4 Discussion

When discussing the results for body ownership and implicit association test, it is important to note that there were relatively few participants that took part in this study. Moreover, the genders of participants was disproportionate. There were only five female participants and eleven male participants. Therefore the results should be interpreted with caution.

4.1 Body Ownership

The effects of embodiment on implicit attitudes and perceptions are known to be stronger the more a user takes ownership of the virtual body. For example, when looking at the effects of embodiment on implicit racial bias, those who were in the not embodied condition experienced significantly lower body ownership; these participants also experienced less of a change in IAT scores compared to those who were embodied in either a light skin or dark skin avatar [9]. In this experiment, body ownership was achieved with first person perspective from the avatar and visuomotor synchronicity [9]. When comparing perspectives (first person vs. third person) as well as visual movement synchronicity (synchronous vs. asynchronous) it was discovered that first person perspective and synchronous movements were the ideal conditions for creating the body ownership illusion in virtual reality [1].

To achieve body ownership in this experiment, the participants were all placed in first person perspective and physical movements were synchronously reflected by their avatar in the virtual environment. The variable Agency, referring to how much control of the avatars movements the participant felt, was most positively responded to with an overall mean of 4.6875. This signifies that visuomotor synchronicity can be achieved using just a head mounted display and tracked controllers.

Looking at the results of body ownership pictured in Fig. 5, the mean response for the My Body variable across all conditions was slightly above 3, indicating that an average participant felt as if the virtual body were their own. However this is a relatively low mean when comparing with a study using embodiment into different races [9]. Here, the mean for embodied conditions was around 4. Given that first person perspective was used and visuomotor synchronicity was achieved for both experiments, there may be another factor affecting body ownership. One reason for why the My Body mean is lower may be the amount of time participants spent embodied. In the experiment, participants spent around four minutes in the orientation period, depending on how quickly the participant completed the exercises, and two minutes in the period where virtual humans walked by. This equates to a total embodiment time of around six minutes. In the racial embodiment experiment, participants spent five minutes in the orientation period and six and a half minutes in the approach period for a total embodiment time of eleven and a half minutes [9]. For The racial embodiment experiment with sustained effects, participants spent even longer being embodied with a five minute orientation period and a ten minute scenario period for a total

embodiment time of fifteen minutes [9]. The total embodiment time for both the implicit racial bias experiments is around double of the total embodiment time for this experiment.

Overall, using first person perspective and visuomotor synchronicity, participants were able to take ownership of a virtual body even if the body was the opposite gender. Future research may want to explore how the total time spent embodied affects the degree of body ownership.

4.2 Implicit Association Test

Looking at the results of preIAT scores and postIAT scores in Fig. 6, one interesting observation is that the overall degree of implicit gender bias was noticeably different between male and female participants, regardless of what gender avatar was assigned to them. For all of the conditions, the postIAT score was one that is more bias towards males/against females than preIAT score. For those in the female with male avatar condition, this means there was actually a reduction in implicit gender bias because the preIAT scores were bias against males. For all other conditions this means there was an increase in implicit gender bias. But again, all postIAT scores in comparison to preIAT scores represent a change to favor males. Although the scenario itself remained neutral to biases, the virtual environment was a managers office which could have existing implicit gender biases associated with it. This is comparable to the experiment in which those embodied in an African American model became more biased against African Americans because the scenario was an interview which involves pre-existing bias [7]. Likewise, this could explain why those in the male with female avatar and female with female avatar conditions became more favorable to males.

Something to consider is that in this experiment, each IAT stage only had five trials. In a similar experiment, where an IAT was used to measure effects of embodiment, 20 or 40 trials were used for each stage depending on the stage. This is a considerably large amount of trials compared to this experiment. Having more trials could provide a more accurate representation of implicit gender bias and therefore better exemplify the effect of embodiment on said bias.

5 Conclusion

Virtual reality is a powerful tool that can be used to change the way people think or feel, even implicitly. Using embodiment, users can take ownership of a completely virtual body, even if it is a different size, shape, race, or gender than the physical body. Many studies have shown how being embodied in different virtual bodies can affect implicit attitudes. More specifically, there are studies addressing how embodiment can reduce, or in some cases increase, racial implicit bias or age- related implicit bias. However, there is no research examining how implicit gender bias is affected by embodiment in virtual reality. This experiment shows that users are able to take ownership of a virtual body, even one of a different gender, using visuomotor synchronization and first person perspective. With this

body ownership, the average implicit gender bias of participants became more favorable to males than before the embodiment experience. Future research in this area should focus on how the total time embodied affects body ownership of different gendered avatars. In addition, replicating this study with more participants could provide more evidence that implicit gender bias is affected by gender embodiment.

References

1. Slater, M., Spanlang, B., Sanchez-Vives, M.V., Blanke, O.: First person experience of body transfer in virtual reality. PloS one **5**(5), e10564 (2010)
2. Botvinick, M., Cohen, J.: Rubber hands feel touch that eyes see. Nature **391**(6669), 756 (1998)
3. Sanchez-Vives, M.V., Spanlang, B., Frisoli, A., Bergamasco, M., Slater, M.: Virtual hand illusion induced by visuomotor correlations. PloS one **5**(4), e10381 (2010)
4. Yee, N., Bailenson, J.: The proteus effect: the effect of transformed self-representation on behavior. Hum. commun. Res. **33**(3), 271–290 (2007)
5. Banakou, D., Groten, R., Slater, M.: Illusory ownership of a virtual childbody causes overestimation of object sizes and implicit attitude changes. Proc. Nat. Acad. Sci. **110**(31), 12846–12851 (2013)
6. Yee, N., Bailenson, J.N.: Walk a mile in digital shoes: the impact of embodied perspective-taking on the reduction of negative stereotyping in immersive virtual environments. In: Proceedings of PRESENCE, vol. 24, p. 26 (2006)
7. Groom, V., Bailenson, J.N., Nass, C.: The Influence of racial embodiment on racial bias in immersive virtual environments. Soc. Influ. **4**(3), 231–248 (2009)
8. Banakou, D., Hanumanthu, P.D., Slater, M.: Virtual embodiment of whitepeople in a black virtual body leads to a sustained reduction in their implicit racial bias. Front. Hum. Neurosci. **10**, 601 (2016)
9. Peck, T.C., Seinfeld, S., Aglioti, S.M., Slater, M.: Putting yourself in the skin of a black avatar reduces implicit racial bias. Conscious. Cogn. **22**(3), 779–787 (2013)
10. Greenwald, A.G., Nosek, B.A., Banaji, M.R.: Understanding and using the implicit association test: I. an Improved scoring algorithm. J. Pers. Soc. Psychol. **85**(2), 197 (2003)