

Associative Memory and Trustworthiness of Artificial Faces in Young and Older Adults

Kylie O. Alberts and Alan D. Castel

University of California, Los Angeles, Los Angeles, CA, USA

IN PRESS

Experimental Aging Research

Author Note

Kylie O. Alberts  <https://orcid.org/0009-0007-6113-2903>

Alan D. Castel  <https://orcid.org/0000-0003-1965-8227>

This research was supported in part by the National Institutes of Health (National Institute on Aging; Award Number R01 AG044335 to Alan D. Castel). Please address correspondence to Kylie O. Alberts or Alan D. Castel, Department of Psychology, University of California, Los Angeles, CA 90095, albertsk@g.ucla.edu or castel@ucla.edu. The experiments reported in this article were not formally preregistered. The data and analysis code are available at <https://osf.io/eq2jf/> on the Open Science Framework. The authors certify that they have no affiliations with or involvement in any organization or entity with any financial or non-financial interest in the subject matter or materials discussed in this manuscript. The ideas and data appearing in the manuscript have not been disseminated before.

Abstract

Older adults generally show deficits in associative memory and increased trust in faces compared to young adults. However, little research has been conducted on older adults' associative memory and trust in artificial faces. The present study investigated young and older adults' perceived trustworthiness for real and artificial faces that were associated with either a scam or neutral condition. Participants viewed the faces before and after they were associated with either a scam or a neutral condition and subsequently rated each face on perceived trustworthiness. Participants were also tested on their memory for these associations. Both young and older adults rated faces associated with a scam as being less trustworthy. However, overall, older adults rated faces as more trustworthy than young adults. In addition, young adults were the only group to rate artificial faces as being less trustworthy than real faces, and older adults did not show this difference. Young and older adults also had similar accuracy for remembering the associations of real and artificial faces. However, only young adults had higher accuracy for real faces than artificial ones, while older adults showed no difference. These findings illustrate that older adults may perceive and remember artificial faces differently from young adults.

Associative Memory and Trustworthiness of Artificial Faces in Younger and Older Adults

The use of artificial faces has become increasingly prevalent in a variety of settings (e.g., Balas & Pacella, 2017; Billard et al., 2007; Mataric et al., 2009; Kanda et al., 2004). Artificial faces have been used positively in various contexts, such as robots with artificial faces supporting patients rehabilitating from strokes or robots helping children with autism (Billard et al., 2007; Mataric et al., 2009). However, artificial faces, also known as “deep fakes,” have also been used to impersonate real people in cybercrimes, such as imposter scams and identity theft (Raza et al., 2022; Ahmed, 2021). In 2022, older adults reported a median loss of \$1,000 for those 70 years or older, while young adults reported a median loss of \$550 (AARP, 2023). In more recent analyses, the Federal Trade Commission (FTC) has reported that the number of older adults who have lost \$10,000 or more has increased fourfold, and those who have lost more than \$100,000 have increased eightfold since 2020 (FTC, 2025). With the increasing use of generative artificial intelligence (AI), eighty-five percent of older adults in a recent poll have indicated that they are concerned about the increasing use of “deep fakes” in financial fraud (Williams, 2024). The Federal Bureau of Investigation (FBI) has also released a recent public service announcement warning the public of criminals exploiting generative AI to commit fraud on a large scale (FBI, 2024). The use of artificial faces may thus pose a threat to young and older adults who may be targeted by these kinds of scams and fraud. Therefore, the current study aims to investigate how young and older adults perceive the trustworthiness of artificial faces. In addition, the current study investigated people’s associative memory of artificial faces paired with pieces of information, particularly paired with a scam, in order to determine if there are age-related deficits in how people remember information paired with artificial faces.

Trustworthiness

Humans usually determine a face's trustworthiness based on various facial features and expressions (Rule et al., 2013; Oosterhof & Todorov, 2008). When encountering an artificial face, humans may use similar diagnostics to determine whether a face is trustworthy or not. However, this may be difficult because artificial face characteristics are inherently different than real faces (Balas et al., 2018). For instance, the "*uncanny valley phenomenon*" (MacDorman et al., 2009; MacDorman & Ishiguro, 2006) describes young adults' agreeability towards artificial faces, whereby humans find artificial faces to be more agreeable the more human-like they become, up until a point where the faces become too similar to real faces that they become unsettling for humans. This phenomenon shows that humans perceive real and artificial faces differently and may act differently when interacting with artificial faces.

Interestingly, older adults have not been shown to have the dip in agreeability when viewing artificial faces that is evident for young adults (Tu et al., 2020). Similar to young adults, older adults also have an increased agreeability towards artificial faces, the more human-like they become. However, unlike younger adults, older adults continue to like artificial faces the more human-like they become, without experiencing the dip in agreeability that young adults feel (Tu et al., 2020). Therefore, young and older adults seem to perceive artificial faces differently to some extent; however, few studies have tested this phenomenon in different contexts.

Studies have been conducted on the trustworthiness of these faces in young adults. For example, Balas and Pacella (2017) found that participants provided lower ratings of trustworthiness to artificial faces than real faces. The researchers hypothesized that people may view artificial faces as an "out-group" face category, meaning that they view such faces as

“different” from them and, therefore, less trustworthy. This finding aligns with prior research on the “other race” or the “other age” effect, whereby stimuli that appear different from oneself are perceived differently than those from an “in-group” (Balas & Pacella, 2017; Lee & Penrod, 2022; Meissner & Brigham, 2001; Rhodes & Anastasi, 2012). Other research has found that even labeling a face as being artificial can decrease the trustworthiness of faces relative to other faces labeled as real (Liefoghe et al., 2023). Therefore, at a baseline, artificial faces seem to be perceived as less trustworthy than real faces. The present study aims to investigate whether this effect also obtains in older adults, in order to provide some evidence regarding how older adults process, evaluate, and trust artificial faces.

In general, older adults are more trusting than young adults (Alberts et al., 2025; Bailey & Leon, 2019). Older adults have also been found to base their analysis of trustworthiness on familiarity rather than specifically on a memory of whether something is actually trustworthy (Skurnik et al., 2005). Memory for trustworthiness may partially be linked to an emotional response, as young adults tend to have better memory for faces perceived as untrustworthy than faces perceived as trustworthy (Rule et al., 2012). Prior research has also shown that, despite older adults being more trusting of faces than young adults, young and older adults both equally adapted their perceived trustworthiness to be less trusting of faces after they were associated with scams (Alberts et al., 2025).

Memory

Artificial faces are remembered less well than real faces (Balas & Pacella, 2015; Balas & Pacella, 2017). For example, Balas and Pacella (2015) tested two groups of younger adults (one that viewed real faces and the other that viewed artificial faces) who were asked to study these faces for a later recognition test. Participants were better at remembering the real faces than the

artificial faces, therefore adding to the theory that participants may view artificial faces as the “out-group” and be less likely to remember the faces (Balas & Pacella, 2017). However, no prior research has investigated associative memory for such faces, despite artificial faces often being paired with information, nor has prior research investigated older adults’ memory for artificial faces.

In prior associative memory research, older adults have shown a deficit in associative memory for name-to-face pairings (e.g., Naveh-Benjamin et al., 2004) and face-to-face pairings (e.g., Rhodes et al., 2008). In addition, older adults have shown a deficit in remembering negative information, such as negative monetary information (Castel et al., 2016). Older adults have also shown a deficit in associative memory for real faces associated with scams and donations when compared to young adults (Alberts et al., 2025) and tend to rely more on gist-familiarity than more detailed specific recollection (Castel & Craik, 2003; Jacoby, 1999), which could influence how people remember associative information about products, prices, advertisements and potentially-fraudulent offers (e.g., Alberts & Castel, 2025; Castel, 2005; McGillivray & Castel, 2010; Murphy et al., 2023; Whatley et al., 2014). The present study investigated whether older adults’ deficits in associative memory were also present when viewing artificial faces, as this may be a situation where older adults struggle with binding information, given the more limited familiarity of the artificial faces (Kamp et al., 2018).

Current Study

Due to the increased use of artificial faces, it is important to investigate how young and older adults perceive the trustworthiness of these faces when paired with pieces of information. Surprisingly, very little research has explored memory of artificial faces beyond the study by Balas and Pacella (2015). In addition, to our knowledge, no research has tested the associative

memory of artificial faces and confidence in remembering these faces, or older adults' perceptions of trustworthiness and/or memory of artificial faces. Given the increasing prevalence of artificial faces, this represents an important population to test. The present study used a similar design as Alberts et al. (2025), in which young and older adults rated the trustworthiness of real faces before and after those faces were associated with a scam and extended this work by examining both real and artificial faces. More specifically, it examined associative memory and trustworthiness by pairing real and artificial faces with a scam or neutral label.

We hypothesized that our results would generally replicate some of the main findings from Alberts et al. (2025) and extend this to artificial faces. We predicted that older adults would, on average, be more trusting than young adults. We also predicted that, despite young adults being more accurate at remembering the paired associations than older adults, both young and older adults would still adjust their ratings of faces paired with scams as being less trustworthy once the association is made. In terms of artificial faces, we hypothesized that, due to older adults generally having higher levels of trust towards others (Castle et al., 2012), older adults would also be more trusting of artificial faces than young adults. Consistent with prior research, we predicted that real faces would be perceived as being more trustworthy than artificial faces (Balas & Pacella, 2017). We also predicted that participants would rate the faces associated with scams as being less trustworthy after viewing the face paired with a scam condition. We also predicted that participants would be more accurate in remembering the associations for real than artificial faces (Balas & Pacella, 2015), and that this would be the case for older adults because both age groups can exhibit an "other group" bias (Wylie et al., 2015). In addition, we hypothesized that artificial faces associated with scams would be the least trustworthy of all the associations.

Method

Transparency and Openness

Informed consent was acquired, and the study was completed in accordance with the UCLA Institutional Review Board. This study was not formally preregistered. Data were analyzed, and all figures were made using R Studio (RStudio Team, 2020), specifically using the *tidyverse* (Wickham et al., 2019), *emmeans* (Lenth, 2023), *lmer4* (Bates et al., 2015), *ggplot2* (v3.3.3; Wickham, 2016), *purrr* (Wickham & Henry, 2025), and *cocor* packages (Diedenhofen & Musch, 2015). All information needed to reproduce the analyses is available on OSF, including stimuli, data, and analysis code (Alberts et al., 2025).

Participants

Ninety-six young adult participants (age range 18-29 years old; $M_{age} = 20.21$, $SD = 1.45$) were recruited from the UCLA undergraduate subject pool, and 96 (age range 54-81; $M_{age} = 69.31$, $SD = 4.72$) older adult participants were recruited through Prolific to participate in the study. Young adults received course credit for participating in the study, with one hour of participation equaling one course credit granted. Older adults were compensated US\$10/hr. An a priori power analysis, using G*Power 3.1.7 (Faul et al., 2007), indicated that for a linear regression with four predictors (age, face type, condition, and time), assuming $\alpha = .05$ and $\text{power} = .80$, 192 participants would be needed to reliably detect a medium effect size ($n_p^2 = .06$).

Materials

Faces in this study were taken from Balas and Pacella (2015) and used with permission. Balas and Pacella (2015) created the artificial faces from faces included in the Chicago Face Database (Ma et al., 2015). The artificial faces were created using the PhotoFit tool in the software FaceGen. All of the faces had neutral expressions and were cropped to only include the

face and not include any other information, such as beards or mustaches (Figure 1). At the end of the experiment, participants completed the Scam Susceptibility Questionnaire (James et al., 2014), which is a five-item scale whereby participants rate their agreement to different prompts about scams on a Likert scale. An example question from this scale would be a question such as “I answer the phone whenever it rings, even if I do not know who is calling.” Participants would respond with a rating from 1 (*strongly disagree*) to 7 (*strongly agree*).

Procedure

The procedure followed the design of Alberts et al. (2025) and was built using PsychoPy (Pierce et al., 2019). In the pre-study phase, participants viewed 36 faces for 6 seconds each and then rated them on how trustworthy they appeared on a scale of 0 (*not trustworthy at all*) to 10 (*very trustworthy*). Participants rated the faces for trustworthiness immediately after viewing the face. Half of the faces were real and half were artificial faces (see Pre-Study Phase Honesty Rating in Figure 2).

In the study phase, participants were told that they would view faces paired with a label of “scam” (indicating that the person committed a scam/fraud) or paired with a neutral label (indicating that the person did not commit a scam/fraud). Participants were told that they needed to remember the faces and their labels for a later memory test. In this phase, participants viewed 24 faces randomly selected from the previously rated 36 faces one-at-a-time for 6 seconds. Each face was randomly paired with a scam or neutral label for each new participant. (see Study Phase, in Figure 2).

In the post-study phase, participants’ memory for the associations was tested. In the test phase, participants viewed the same 24 faces from the study phase and were asked which label the face was paired with (scam or neutral; see Test Phase, in Figure 2). Participants also

indicated their confidence in the accuracy of their response on a scale of 0 (*not confident at all*) to 10 (*very confident*). Responses were self-paced. After the test phase, participants were shown all 36 faces from the Pre-Study phase and were once again asked how honest the faces seemed (see Post-Study Phase Honesty Rating, Figure 2). At the end of the experiment, participants completed the Scam Susceptibility Questionnaire (James et al., 2014).

Results

Accuracy

Participants' accuracy for remembering the association paired with each face (i.e., the probability of selecting the correct association) is presented in Figure 3. These data were analyzed by fitting a logistic mixed effects model using the *glmer* function using R Version 4.3.1 (R Core Team, 2021). The model included a three-way interaction for Age (young vs. old) X Face Type (real vs. artificial) X Condition (scam vs. neutral). A likelihood ratio test indicated that including participants' scam susceptibility scores in the model, did not significantly improve fit of the model so we did not include it in our final model, $\chi^2(1) = .14, p = .71$. All variables were dummy coded with "older adults" as the comparison for age, "artificial" as the comparison for face type, and "neutral" as the comparison for condition. We also included a random intercept for participants and face images. To test the simple effects of the model, we compared estimated marginal means of the full model using *emmeans* and *pairs* functions from Version 1.8.4 of the *emmeans* R package (Lenth, 2023). There was no significant effect of age, $b = .03, SE = .14, z = .22, p = .83$, face type, $b = .18, SE = .17, z = 1.02, p = .31$, or condition, $b = .14, SE = .13, z = 1.09, p = .28$. However, there was a significant face type by age interaction, $b = .42, SE = .18, z = 2.33, p = .02$. Young adults were less accurate at remembering associations for artificial faces ($M = .64, SD = .48$) than real faces ($M = .72, SD = .45$), $OR = .65, 95\% CI [0.44,$

0.96], $z = 2.87$, $p = .02$. There was no significant difference in accuracy for remembering artificial ($M = .62$, $SD = .49$) and real faces ($M = .68$, $SD = .47$) for older adults, $OR = .79$, 95% CI [0.54, 1.15], $z = 1.63$, $p = .36$.

Honesty Ratings

Participants' honesty ratings are presented in Figure 4. We also used a linear mixed-effect model to analyze the honesty ratings for faces in the pre-study phase and the post-study phase. Our model included a four-way interaction for Age (young vs. old) X Condition (scam vs. neutral vs. no association) X Time (pre vs. post) X Face Type (real vs. artificial). Similar to our analysis for accuracy, we compared estimated marginal means to test the simple effects of the model, and all predictors were dummy coded with "post" as the comparison condition for the time variable. The no association condition includes faces that were only included in the pre- and post-honesty ratings and were not included in the study or test phase. The model revealed a significant effect of condition, whereby neutral faces ($M = 4.91$, $SD = 2.44$), $b = .48$, $SE = .04$, $z = 11.16$, $p < .001$, and faces paired with no association ($M = 5.16$, $SD = 2.43$), $b = .73$, $SE = .18$, $z = 4.04$, $p < .001$, were rated as more trustworthy than faces associated with a scam ($M = 4.43$, $SD = 2.47$). There was also a significant effect of face type, where real faces ($M = 5.18$, $SD = 2.41$) were found to be rated as more trustworthy than artificial faces ($M = 4.48$, $SD = 2.47$), $b = .47$, $SE = .23$, $t = 2.01$, $p = .046$. Older adults ($M = 5.27$, $SD = 2.45$) also had higher honesty ratings than young adults ($M = 4.40$, $SD = 2.40$), $b = 1.21$, $SE = .19$, $t = 6.55$, $p < .001$.

Honesty Ratings: Two-Way Interaction

There was also a significant interaction between face type and age. Young adults rated real faces ($M = 4.97$, $SD = 2.34$) as more honest than artificial faces ($M = 3.83$, $SD = 2.32$), $b = 1.15$, $SE = .17$, $z = 6.67$, $p < .001$. There was no difference for older adults ($M_{\text{real}} = 5.39$, $SD_{\text{real}} =$

2.46; $M_{\text{artificial}} = 5.14$, $SD_{\text{artificial}} = 2.44$), $b = .25$, $SE = .17$, $z = 1.46$, $p = .46$. In addition, older adults rated artificial, $b = 1.32$, $SE = .15$, $z = 8.82$, $p < .001$, and real faces, $b = .42$, $SE = .15$, $z = 2.82$, $p = .02$, as more honest than young adults.

In addition, there was a significant interaction of condition and time. Faces that were associated with a scam were rated as less trustworthy after the association was made ($M_{\text{pre}} = 4.78$, $SD_{\text{pre}} = 2.58$; $M_{\text{post}} = 4.09$, $SD_{\text{post}} = 2.58$), $b = .69$, $SE = .06$, $z = 11.38$, $p < .001$. In addition, faces that did not have any association were rated as less trustworthy in the post-honesty ratings ($M_{\text{pre}} = 5.30$, $SD_{\text{pre}} = 2.30$; $M_{\text{post}} = 5.02$, $SD_{\text{post}} = 2.54$), $b = .28$, $SE = .06$, $z = 4.68$, $p < .001$. There was no difference between pre ($M = 4.85$, $SD = 2.27$) and post-honesty ratings ($M = 4.97$, $SD = 2.60$) for faces with a neutral association, $b = .12$, $SE = .06$, $z = 2.04$, $p = .32$. For pre-honesty ratings, faces with no association was rated higher than faces later associated with a scam and this neared significance, $b = .52$, $SE = .19$, $z = 2.83$, $p = .052$. There was no difference between pre-honesty rating scores for faces later associated with a neutral label and those later associated with a scam, $b = .07$, $SE = .06$, $z = 1.18$, $p = .85$, or those that had no association, $b = .45$, $SE = .19$, $z = 2.45$, $p = .14$. For post-honesty ratings, faces associated with neutral, $b = .88$, $SE = .06$, $z = 14.60$, $p < .001$, and ones that had no association, $b = .93$, $SE = .19$, $z = 5.02$, $p < .001$, were rated as more trustworthy than faces associated with scams. There was no difference between post-honesty ratings for faces associated with neutral information and those that had no association, $b = .05$, $SE = .19$, $z = .26$, $p = .99$.

There was also a significant interaction between condition and face type. Artificial faces that were associated with a scam ($M = 4.17$, $SD = 2.44$) were rated as less trustworthy than neutral ($M = 4.51$, $SD = 2.45$), $b = .33$, $SE = .06$, $z = 5.50$, $p < .001$. There was no difference between artificial faces associated with neutral and no association ($M = 4.78$, $SD = 2.48$), $b = .28$,

$SE = .25$, $z = 1.09$, $p = .88$, nor a difference between artificial faces associated with a scam and those that had no association, $b = .61$, $SE = .25$, $z = 2.40$, $p = .16$. Real faces that were associated with neutral labels ($M = 5.32$, $SD = 2.36$), $b = .62$, $SE = .06$, $z = 10.27$, $p < .001$, and ones that had no association ($M = 5.54$, $SD = 2.32$), $b = .84$, $SE = .25$, $z = 3.31$, $p = .01$, were rated as more trustworthy than faces associated with scam labels ($M = 4.70$, $SD = 2.47$). Real faces that were associated with neutral information were rated as more trustworthy than artificial faces, $b = .81$, $SE = .21$, $z = 3.86$, $p = .002$. There were no differences between real and artificial faces associated with scam, $b = .53$, $SE = .21$, $z = 2.50$, $p = .12$, nor those with no association, $b = .76$, $SE = .29$, $z = 2.60$, $p = .10$.

Honesty Ratings: Three-Way Interaction

Lastly, there was also a significant three-way interaction between condition, time, and face type (see Figure 4). Pre-honesty ratings for real faces ($M = 5.16$, $SD = 2.25$) later associated with scams, were rated as more trustworthy than artificial faces ($M = 4.39$, $SD = 2.30$), $b = .77$, $SE = .22$, $z = 3.51$, $p = .02$. There was no difference between pre-honesty ratings for artificial and real faces that were later associated with neutral ($M_{real} = 5.18$, $SD_{real} = 2.14$; $M_{artificial} = 4.52$, $SD_{artificial} = 2.35$), $b = .67$, $SE = .22$, $z = 3.06$, $p = .09$, or no association ($M_{real} = 5.64$, $SD_{real} = 2.19$; $M_{artificial} = 4.95$, $SD_{artificial} = 2.37$), $b = .69$, $SE = .30$, $z = 2.32$, $p = .46$. For post-honesty ratings, there was no difference between real ($M = 4.23$, $SD = 2.59$) and artificial faces ($M = 3.95$, $SD = 2.55$) associated with scams, $b = .28$, $SE = .22$, $z = 1.30$, $p = .98$, or no association ($M_{real} = 5.43$, $SD_{real} = 2.43$; $M_{artificial} = 4.61$, $SD_{artificial} = 2.57$), $b = .82$, $SE = .30$, $z = 2.77$, $p = .19$. However, there was a significant difference for faces associated with neutral information, whereby real faces ($M = 5.45$, $SD = 2.56$) were rated as more trustworthy than artificial faces ($M = 4.50$, $SD = 2.56$), $b = .96$, $SE = .22$, $z = 4.37$, $p < .001$. For faces associated with scams, there

was a significant difference between artificial faces before and after the association was made, whereby artificial faces had higher honesty ratings before the association was made than after, $b = .44$, $SE = .09$, $z = 5.21$, $p < .001$. Real faces that were associated with a scam were also rated as less trustworthy after the association, $b = .92$, $SE = .09$, $z = 10.89$, $p < .001$. For faces associated with neutral, there was no difference for pre and post honesty ratings for artificial faces that were associated with neutral, $b = .02$, $SE = .09$, $z = .24$, $p = 1.00$, nor for real faces that were associated with neutral, $b = .27$, $SE = .09$, $z = 3.13$, $p = .08$. For faces that were not associated with anything, artificial faces were rated as less trustworthy in the post-test honesty ratings than pre-test honesty ratings, $b = .35$, $SE = .09$, $z = 4.08$, $p = .002$. However, there was no difference between pre- and post-test honesty ratings for real faces that had no association, $b = .22$, $SE = .09$, $z = 2.54$, $p = .32$.

For post-honesty ratings, neutral real faces, $b = 1.21$, $SE = .09$, $z = 14.23$, $p < .001$, and real faces with no association, $b = 1.20$, $SE = .26$, $z = 4.58$, $p < .001$, were rated as more trustworthy than real faces that were associated with scams. There was no difference between real neutral faces and real faces with no association, $b = .018$, $SE = .26$, $z = .07$, $p = 1.00$. In addition, neutral artificial faces were rated as more trustworthy than scam artificial faces, $b = .54$, $SE = .09$, $z = 6.38$, $p < .001$. There was no difference between neutral artificial faces and artificial faces with no association, $b = .11$, $SE = .26$, $z = .44$, $p = 1.00$, nor between scam artificial faces and faces with no association, $b = .68$, $SE = .26$, $z = 2.52$, $p = .33$.

Discussion

This study investigated young and older adults' perceived trustworthiness and associative memory for real and artificial faces. Each face was rated for trustworthiness before and after being paired with a scam or neutral condition to assess how perceived trustworthiness could

change after a face is paired with potentially untrustworthy information, following a similar design we have used in prior research (Alberts et al., 2025). To our knowledge, no prior studies have investigated associative memory for information paired with artificial faces or studied older adults' perceived trustworthiness of these faces, both of which have important implications and applications for older adults.

Based on prior findings of older adults being more trusting than young adults (Alberts et al., 2025; Bailey & Leon, 2019; Castle et al., 2012), we predicted that older adults would show more trust than young adults for the real and artificial faces. Our findings were consistent with our hypothesis, with older adults rating both real and artificial faces as being more trustworthy than young adults. Interestingly, only young adults rated real faces as being more trustworthy than artificial faces (Balas & Pacella, 2017), and older adults did not show this difference and instead rated the real and artificial faces similarly. This may reflect how older adults perceive artificial faces. For the uncanny valley phenomenon, older adults have been shown not to have a dip in agreeability when viewing increasingly human-like faces (Tu et al., 2020). Therefore, this effect may also translate to trustworthiness, whereby older adults do not find particularly human-like artificial faces to be untrustworthy. If this interpretation were to be true in the real world, older adults would likely not find artificial faces used in scams, or “deep fakes,” to be untrustworthy. Therefore, older adults may be more likely to fall for these particular scams. Another possibility could be that older adults may not be as attuned to determining whether a face is artificial versus real. Since we did not tell the participants that some of the faces were artificial, it could be that older adults could not distinguish between the real and artificial faces and thus rated them as similarly trustworthy. Therefore, young adults may be viewing artificial faces as the “out-group” while older adults cannot or do not distinguish between them (Balas &

Pacella, 2017; Meissner & Brigham, 2001; Rhodes & Anastasi, 2012). This interpretation may be credible as studies have found that young and older adults have similar negative attitudes and lack of trust towards faces that are labeled as computer-generated (Heezen, 2023; Liefoghe et al., 2023). Therefore, it is likely that older adults would be less trusting of the artificial faces in the present study if they were informed that the faces were artificial. So, the lack of difference in trust shown in this study could indicate that older adults are not able to identify some of the faces as being artificial. If this interpretation is the case, interventions should be created to help older adults better identify common features of artificial content, such as mechanical-sounding voices and irregular blinking “deep fake” videos, in hopes of reducing potential scam susceptibility (Han et al., 2024; Hsu et al., 2020; Zhai et al., 2025). These theories should be investigated more thoroughly to determine what is driving this effect, as warning older adults may be useful for practical reasons.

Young and older adults were similarly accurate in remembering the associations between faces and their condition labels, inconsistent with our prior work (Alberts et al., 2025). However, young adults were more accurate at remembering information that was paired with real faces than information paired with artificial faces, and there was no difference for older adults in this condition. Older adults were not more accurate at remembering the information paired with real faces than artificial faces. This may suggest that older adults either perceive faces differently than young adults, as shown in the uncanny valley phenomenon (Tu et al., 2020), or older adults have trouble distinguishing between real and artificial faces.

When the real and artificial faces were associated with a scam, however, both young and older adults decreased their honesty ratings for those faces and deemed those faces as being more untrustworthy. There was no difference between faces associated with a neutral condition for

their pre- and post-honesty ratings. These findings are consistent with prior research showing that both young and older adults deem faces associated with a scam as being less trustworthy after the association is made (Alberts et al., 2025). Interestingly, faces that were not associated with any stimuli were rated less trustworthy in the post-honesty ratings than in the pre-honesty ratings, in contrast to Alberts et al. (2025). In the present study, participants may adopt a general mistrust of all faces once some faces were paired with a scam; therefore, if someone did not remember what a face was paired with (even if there was no pairing) they may have rated it as less trustworthy. This would create a difference in responses between faces paired with the neutral label and those not paired with anything, because participants may remember the faces that were paired with the neutral label and infer that those faces were not untrustworthy. In contrast, faces that were not paired with anything would not cue the memory of an associated label, making them riskier to trust. Despite this inconsistency, the present study highlights important findings about young and older adults' perception and memory of real and artificial faces. Notably, it is the first study to demonstrate that young adults both trust and remember real faces more so than artificial faces in contrast to older adults, who do not show these differences. Thus, older adults may perceive artificial faces differently from younger adults, and these differences may impact trust and memory, as well as a variety of other factors that play a role in face judgements.

We do note that one potential limitation of this study could be that young adults were given course credit for their participation in the study, while older adults were given money for their participation in the study. Therefore, different motivations in these studies could have led to a difference in results, and perhaps the older adults represent a more active/healthy sample than what may be representative of older adults in general (cf., Greene & Naveh-Benjamin, 2022),

which could explain the similar performance in memory in the present study. Another potential limitation could be that the degree to which the faces are artificial could lead to different results, especially if older adults have difficulty distinguishing between real and artificial faces. Therefore, these findings may not reflect older adults' perception of all types of artificial faces, but the present work shows that older adults may not spontaneously be able to detect these differences. Finally, eyesight and other health measures were not collected for this experiment, which could lead to a confound where those with poorer eyesight may have a harder time distinguishing between the faces (Creighton et al., 2019). Therefore, future studies may benefit by including such measures. Lastly, the faces used in the experiment were used in prior work and were counterbalanced for real and artificial faces, meaning that every face was shown as either its real or artificial version to different participants. However, the faces were not controlled for age or race (as that was not a measure of interest in our study), and this could lead to potential biases in the perceived trustworthiness of the faces. Therefore, future research could examine how these factors influence trust and memory in younger and older adults.

The present study showed that both young and older adults can remember associations with real and artificial faces; however, young adults are more accurate in remembering associations with real faces than artificial ones, while older adults show no difference. In addition, despite older adults rating all of the faces as being more trustworthy than young adults did, only young adults showed a difference in perceived trustworthiness of real and artificial faces. Young adults rated real faces as being more trustworthy than artificial faces, and older adults showed no difference. However, both young and older adults rated the faces as being less trustworthy when associated with a scam. This suggests that young and older adults may have different perceptions of artificial faces. These findings could have implications for how older

adults interact with artificial intelligence as it becomes more advanced. Older adults could potentially be more trusting of artificial faces, such as “deep fakes” that are being used, and may be less able to distinguish between real and artificial faces, leading to detrimental outcomes, especially when “deep fakes” are used in scams.

References

- AARP. (2023, February 28). *Americans Lost Record-Breaking \$8.8 Billion to Scams in 2022*.
AARP. <https://www.aarp.org/money/scams-fraud/info-2023/ftc-consumer-losses.html>.
- Ahmed, S. (2021). Fooled by the fakes: Cognitive differences in perceived claim accuracy and sharing intention of non-political deepfakes. *Personality and Individual Differences*, 182, 111074.
- Alberts, K. O., & Castel, A. D. (2025). Memory for free: Gist-based false recall of an advertisement in young and older adults. *Applied Cognitive Psychology*, 39, e70041.
- Alberts, K. O., Whatley, M. C., & Castel, A. D. (2025). Associative memory for honest and dishonest faces in younger and older adults. *Psychology and Aging*, 40, 610-627.
- Bailey, P. E., & Leon, T. (2019). A systematic review and meta-analysis of age-related differences in trust. *Psychology and Aging*, 34(5), 674.
- Balas, B., & Pacella, J. (2015). Artificial faces are harder to remember. *Computers in human behavior*, 52, 331-337.
- Balas, B., & Pacella, J. (2017). Trustworthiness perception is disrupted in artificial faces. *Computers in Human Behavior*, 77, 240-248.
- Balas, B., Tupa, L., & Pacella, J. (2018). Measuring social variables in real and artificial faces. *Computers in Human Behavior*, 88, 236-243.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Billard, A. (2017). On the mechanical, cognitive and sociable facets of human compliance and their robotic counterparts. *Robotics and Autonomous Systems*, 88, 157-164.

- Castle, E., Eisenberger, N. I., Seeman, T. E., Moons, W. G., Boggero, I. A., Grinblatt, M. S., & Taylor, S. E. (2012). Neural and behavioral bases of age differences in perceptions of trust. *Proceedings of the National Academy of Sciences*, 109(51), 20848-20852.
- Castel, A. D. (2005). Memory for grocery prices in younger and older adults: The role of schematic support. *Psychology and Aging*, 20, 718–721.
- Castel, A. D., & Craik, F. I. M. (2003). The effects of aging and divided Attention on memory for item and associative information. *Psychology and Aging*, 18(4), 873–885.
- Castel, A. D., Friedman, M. C., McGillivray, S., Flores, C. C., Murayama, K., Kerr, T., & Drolet, A. (2016). I owe you: Age-related similarities and differences in associative memory for gains and losses. *Aging, Neuropsychology, and Cognition*, 23(5), 549-565.
- Han, C., Mitra, P., & Billah, S. M. (2024). Uncovering human traits in determining real and spoofed audio: Insights from blind and sighted individuals. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1-14).
- Heezen, S. L. C. (2023). Generational differences in general attitudes towards AI and the level of trust in artificial faces [Master's thesis, Utrecht University]. UU Student Theses Repository. <https://studenttheses.uu.nl/handle/20.500.12932/45178>
- Hsu, C. C., Zhuang, Y. X., & Lee, C. Y. (2020). Deep fake image detection based on pairwise learning. *Applied Sciences*, 10(1), 370.
- Creighton, S. E., Bennett, P. J., & Sekuler, A. B. (2019). Classification images characterize age-related deficits in face discrimination. *Vision Research*, 157, 97-104.
<https://www.sciencedirect.com/science/article/pii/S0042698918301317>

- Diedenhofen B, & Musch, J. (2015). “cocor: A Comprehensive Solution for the Statistical Comparison of Correlations.” PLoS ONE, 10(4), e0121945.
doi:10.1371/journal.pone.0121945, <http://dx.doi.org/10.1371/journal.pone.0121945>.
- Faul, F., Erdfelder, E., Lang, A.G. *et al.* G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39, 175–191 (2007).
- Federal Bureau of Investigation. (2024, December 3). *Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud* [Press release].
<https://www.ic3.gov/PSA/2024/PSA241203>
- Federal Trade Commission. (2025, August). *FTC Data Show a More Than Four-Fold Increase in Reports of Impersonation Scammers Stealing Tens and Even Hundreds of Thousands from Older Adults*. Federal Trade Commission. <https://www.ftc.gov/news-events/news/press-releases/2025/08/ftc-data-show-more-four-fold-increase-reports-impersonation-scammers-stealing-tens-even-hundreds>
- Greene, N. R., & Naveh-Benjamin, M. (2022). Online experimentation and sampling in cognitive aging research. *Psychology and Aging*, 37(1), 72-83.
- James, B. D., Boyle, P. A., & Bennett, D. A. (2014). Correlates of susceptibility to scams in older adults without dementia. *Journal of Elder Abuse & Neglect*, 26(2), 107–122.
- Kamp, S.-M., Bader, R., & Mecklinger, A. (2018). Unitization of word pairs in young and older adults: Encoding mechanisms and retrieval outcomes. *Psychology and Aging*, 33(3), 497–511.
- Kanda, T., Hirano, T., Eaton, D., & Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: A field trial. *Human–Computer Interaction*, 19(1-2), 61-84.

- Lee, J., & Penrod, S. D. (2022). Three-level meta-analysis of the other-race bias in facial identification. *Applied Cognitive Psychology*, 36(5), 1106-1130.
- Lenth, R. (2023). emmeans: Estimated Marginal Means, aka Least-Squares Means. R package version 1.8.5, <https://CRAN.R-project.org/package=emmeans>.
- Liefooghe, B., Oliveira, M., Leisten, L. M., Hoogers, E., Aarts, H., & Hortensius, R. (2023). Are natural faces merely labelled as artificial trusted less?. *Collabra: Psychology*, 9(1), 73066.
- Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago Face Database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47(4), 1122–1135.
- MacDorman, K. F., Green, R. D., Ho, C. C., & Koch, C. T. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25(3), 695-710.
- MacDorman, K. F., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies*, 7(3), 297-337.
- Matarić, M., Tapus, A., Winstein, C., & Eriksson, J. (2009). Socially assistive robotics for stroke and mild TBI rehabilitation. *Advanced technologies in rehabilitation*, 145, 249-262.
- Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law*, 7(1), 3-35.
- McGillivray, S., & Castel, A. D. (2010). Memory for age-face associations: The role of generation and schematic support. *Psychology and Aging*, 25, 822-832.

- Murphy, D. H., Schwartz, S. T., Alberts, K., Siegel, A. L. M., Carone, B. J., Castel, A. D., & Drolet, A. (2023). Clinically studied or clinically proven? Memory for claims in print advertisements. *Applied Cognitive Psychology*, 37, 1085-1093.
- Naveh-Benjamin, M., Guez, J., Kilb, A., & Reedy, S. (2004). The associative memory deficit of older adults: further support using face-name associations. *Psychology and Aging*, 19(3), 541-546.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105(32), 11087-11092.
- Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods*. 10.3758/s13428-018-01193-y
- R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- R Core Team (2021) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna. <https://www.R-project.org>
- Raza, A., Munir, K., & Almutairi, M. (2022). A novel deep learning approach for deepfake image detection. *Applied Sciences*, 12(19), 9820.
- Rhodes, M. G., & Anastasi, J. S. (2012). The own-age bias in face recognition: a meta-analytic and theoretical review. *Psychological Bulletin*, 138(1), 146-174.
- Rhodes, M. G., Castel, A. D., & Jacoby, L. L. (2008). Associative recognition of face pairs by younger and older adults: The role of familiarity-based processing. *Psychology and Aging*, 23, 239-249.

- Rule, N. O., Krendl, A. C., Ivcevic, Z., & Ambady, N. (2013). Accuracy and consensus in judgments of trustworthiness from faces: behavioral and neural correlates. *Journal of Personality and Social Psychology*, 104(3), 409-426.
- Rule, N. O., Slepian, M. L., & Ambady, N. (2012). A memory advantage for untrustworthy faces. *Cognition*, 125(2), 207-218.
- Skurnik, I., Yoon, C., Park, D. C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research*, 31(4), 713-724.
- Whatley, M. C., & Castel, A. D. (2022). The role of metacognition and schematic support in younger and older adults' episodic memory. *Memory & Cognition*, 50, 601-616.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4, <https://ggplot2.tidyverse.org>.
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L.D., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T.L., Miller, E., Bache, S.M., Müller, K., Ooms, J., Robinson, D., Seidel, D.P., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K., Yutani, H. (2019). "Welcome to the tidyverse." *Journal of Open Source Software*, 4(43), 1686.
- Wickham, H., & Henry, L. (2025). *purrr: Functional Programming Tools*. R package version 1.1.0.9000, <https://github.com/tidyverse/purrr>.
- Williams, Alicia R. (2024). Older Adults Express High Concern and Limited Knowledge About Scams and Fraud Involving Artificial Intelligence. *AARP Research*. <https://doi.org/10.26419/res.00873.001>
- Wylie, L. E., Bergt, S., Haby, J., Brank, E. M., & Bornstein, B. H. (2015). Age and lineup type differences in the own-race bias. *Psychology, Crime & Law*, 21(5), 490-506.

Zhai, Y., Xue, X., Guo, Z., Jin, T., Diao, Y., & Jeung, J. (2025). Hear Us, then Protect Us: Navigating Deepfake Scams and Safeguard Interventions with Older Adults through Participatory Design. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1-19).



Figure 1 Example of a real face included in the experiment (on the left) and that face when turned into an artificial version (on the right).

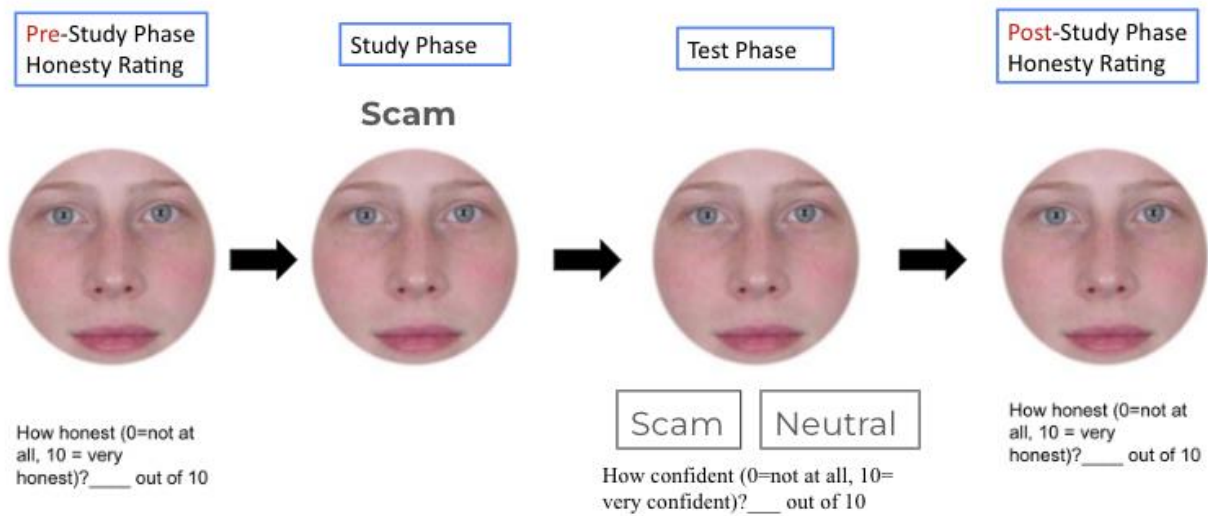


Figure 2 The general procedure used in this experiment.




Figure 3 *The proportion of correct selections based on the type of information each faces (real and artificial) that was paired with for young and older adults. Error bars reflect the standard error of the mean.*

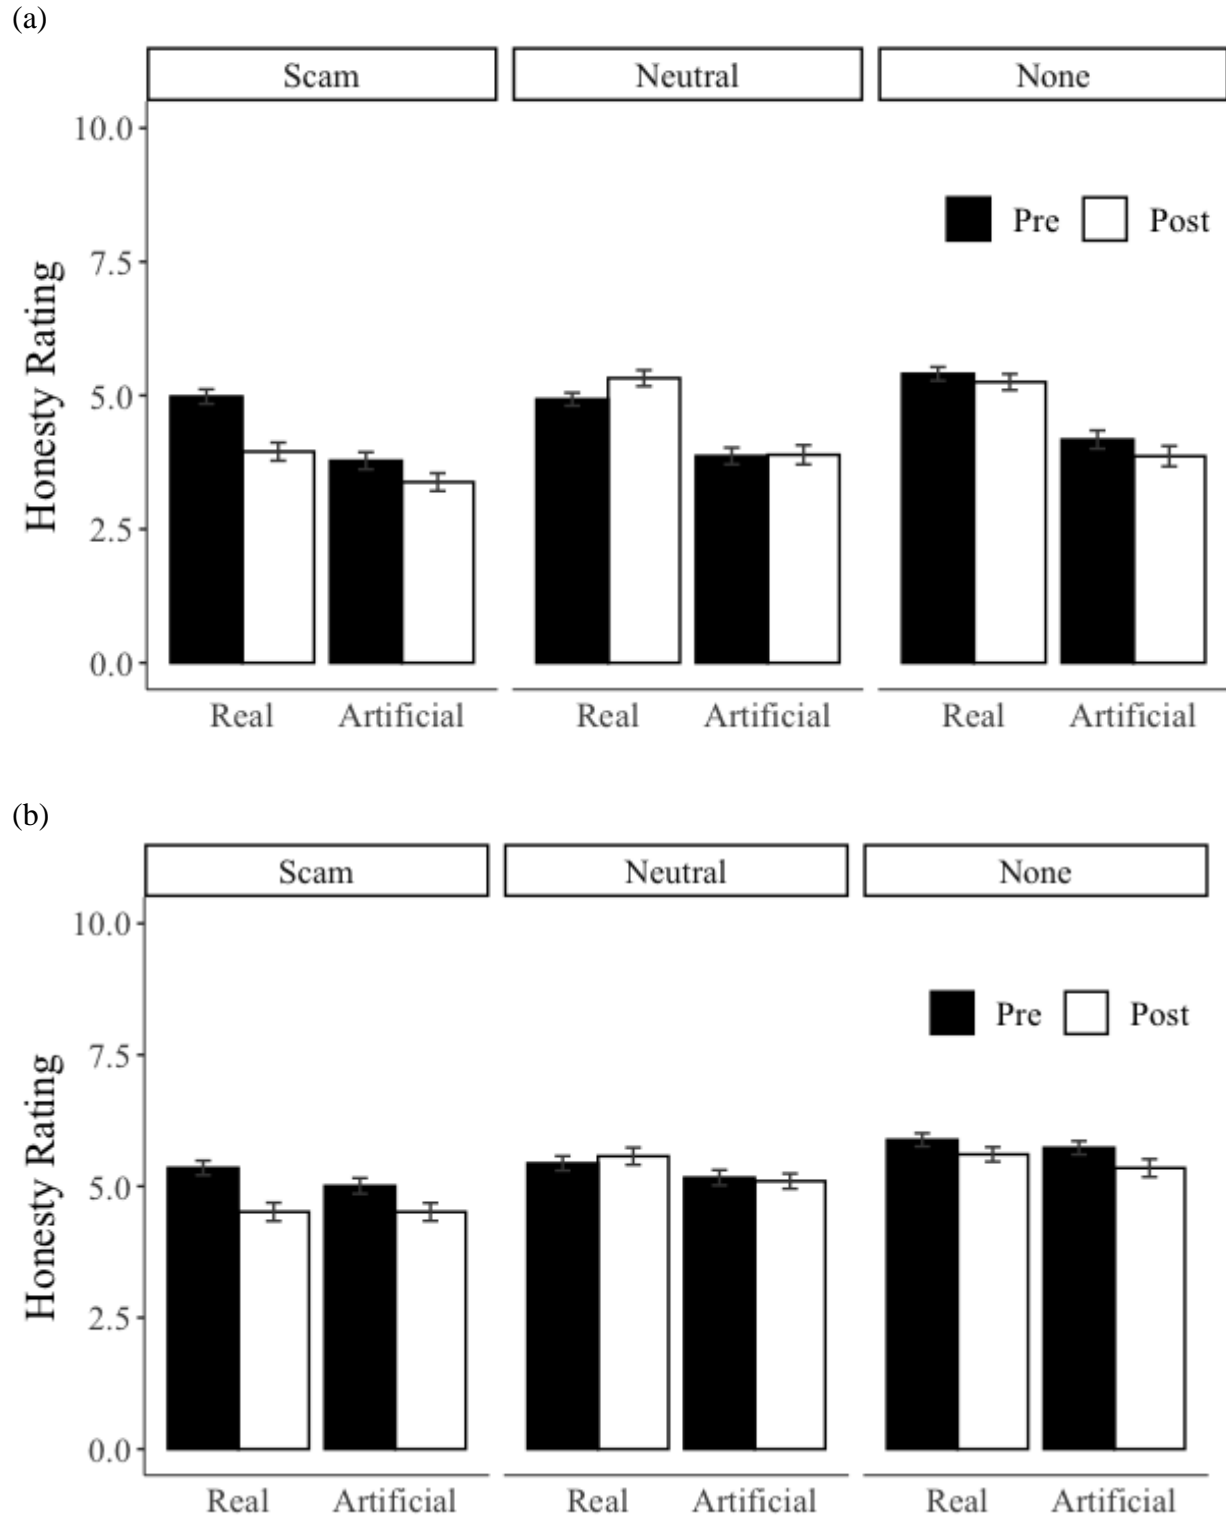


Figure 4 Young (a) and older (b) adults' pre- and post-honesty ratings for real and artificial faces. Error bars reflect the standard error of the mean.