Psychonomic Bulletin & Review
2008, 15 (6), 1093-1099
doi:10.3758/PBR.15.6.1093

Own- and other-race categorization of faces by race, gender, and age

Lun Zhao
Hebrew University of Jerusalem, Jerusalem, Israel
and Xuzhou Normal University, Xuzhou, China

and

Shlomo Bentin
Hebrew University of Jerusalem, Jerusalem, Israel

We investigated how visual experience with faces of a particular race affects subordinate group-level categorizations in Chinese and Israeli participants living in the respective countries. Categorization of faces by race, gender, and age was examined within subjects with participants who had only minimal experience with the other-race faces. As would be predicted by the previously documented other-race advantage effect, both Chinese and Israeli participants classified the race of the face more quickly and more accurately for other-race than for own-race faces. In contrast, the observers' race did not interact with the race of the rated face either for gender or for age categorization. The absence of these interactions suggests that the physiognomic characteristics that determine the gender and age of a face are universal, rather than race specific. Furthermore, these data suggest that determining the race of a face is not imposed as a first step in face processing, preempting the perception of other category-defining physiognomic characteristics.

Although face recognition is a well-documented domain of human perceptual expertise, this expertise is not equal across all faces. Particularly relevant to our present study are differences between processing own-race and other-race faces. Two major differences have been documented. The first is the own-race advantage in face identification (labeled *other-race effect*; ORE), expressed as easier identification of individual faces from one's own race, relative to other-race faces (for reviews, see Meissner & Brigham, 2001; Sporer, 2001). The second is the other-race advantage in race categorization (labeled *other-race advantage*; ORA), expressed as faster categorization of other-race than own-race faces by race (Levin, 1996; Valentine & Endo, 1992). Whereas the ORE is well established (e.g., Rhodes, Hayward, & Winkler, 2006; Valentine, 1991), the ORA has received less attention and has not been consistently found (e.g., Blascovich, Wyer, Swart, & Kibler, 1997; for a review, see Sporer, 2001). Yet, from a face perception perspective, the ORA, as well as the ORE, could be explained by the greater experience people usually have with own-race than with other-race faces.

According to an influential theoretical framework, faces are represented in memory as points in a multidimensional space where each dimension represents a perceptually relevant face feature (Valentine, 1991). These features are tuned by experience to capture subtle differences between individual faces. Hence, the distance between the representations of faces in the multidimensional space reflects their perceived similarity (Krumhansl, 1978). Since people have more experience with own-race faces, the perceptual dimensions underlying the space are not optimal for distinguishing among other-race faces and, consequently, such faces form a denser and more homogeneous *neighborhood* in the representational space than do own-race faces. This assumption has been empirically validated by testing the rated similarity of Caucasian and Chinese faces within and across races (Byatt & Rhodes, 2004).

The higher density should interfere with individual identification of other-race faces because each face can be more easily confused with its near neighbors (explaining the ORE). Higher density, however, could facilitate the categorization of other-race faces by race (explaining the ORA), because individual differences would not interfere with their grouping and because, in a homogeneous space, the probability that activating one exemplar would partially activate many other exemplars is higher, increasing the overall activation of other-race faces as a group.[1] Evidently the multidimensional model provides a parsimonious account for both other-race effects; however, this account is challenged, for example, by data showing reliable ORA for inverted faces (Levin, 1996). Since individual differences are less perceivable in inverted faces, the ORA should have been considerably attenuated. Yet it is possible that for upright faces, this model is correct, in which case sparse density might interfere with any type of grouping

S. Bentin, shlomo.bentin@huji.ac.il

of faces above the individual level. Could we extrapolate from race categorization to assume ORA for group categorization along other face-relevant dimensions?

Human faces may be grouped along different categorical dimensions. In addition to race, common dimensions are gender and age, which were the focus of this study. We focused on gender and age because these categories are natural and universal; hence, participants' experience with men and women, as well as with old and young faces, is probably comparable across races. Reduced experience with other-race faces might affect categorization by dimensions other than race in several ways. On the one hand, if the higher diversity perceived for own-race faces interferes with their grouping by shared physiognomic characteristics (such as those defining gender or age), we could expect finding an ORA for all kinds of group categorization of faces.[2] On the other hand, it is possible that the physiognomic characteristics of gender and/or age are not universal but vary across races. In that case, subordinate categorization by gender and/or age could be more efficient for own- than for other-race faces, showing ORE.

Tentative support for the latter alternative can be found in two studies suggesting that face categorization by gender (O'Toole, Peterson, & Deffenbacher, 1996) and estimation of a person's age from his/her face (Dehon & Brédart, 2001) are more accurate for one's own race than for other races. Although suggestive, these studies are not conclusive. In the gender categorization task, Caucasian and Japanese faces were presented in separate blocks, which might have increased the distinction between the two races. Moreover, both Caucasian and Asian observers performed better with Caucasian than with Japanese faces, raising questions about the equivalence of the gender cues provided by the two groups of faces. In the age estimation task, although the race-of-the-observer × race-of-the-face interaction was significant, a post hoc analysis revealed that whereas Caucasian observers estimated the age of Caucasians more accurately than the age of Africans, African observers were equally accurate with both types of faces. Regarding age, it is noteworthy that Dehon and Brédart examined age estimation, rather than categorization. Hence, whereas the other-race effect shown in their study suggests that precise estimation of age may involve diagnostic features learned by experience, the question still remains whether simple group-level categorization of faces by age is similarly efficient within and across races.

In addition to the caveats above, in those studies, reaction times (RTs) were not recorded, all the participants lived for many years in the same country, and the different types of face categorization were not compared within participants. Here, we explored whether the ORA is specific to categorization by race or is a general characteristic of group-level face categorization by comparing categorization of faces by race, gender, and age, for faces within and across the participants' race. To maximize the effect of experience, we tested Chinese participants living in China and Caucasian participants living in Israel, in an identical experimental design and using the same set of Chinese and Caucasian faces. Neither the Chinese nor the Israeli participants had much experience with the other-race faces.

METHOD

Participants

Thirty-six Chinese undergraduates (18 of them female) were recruited from Xuzhou Normal University (China), and 30 Israeli undergraduates (16 of them female) were recruited from the Hebrew University (Israel). The participants' ages were similar across countries, ranging from 20 to 30 years. They had normal or corrected-to-normal visual acuity and no history of psychiatric or neurological disorders. All the participants were right-handed, on the basis of self-report, and were paid for participation.

Stimuli

The stimuli were 200 face photographs without outer face features, 100 Chinese and 100 Caucasians. Within each race, 50 faces were men and 50 women, and within each gender, 25 were old and 25 were young (Figure 1). All the faces were unfamiliar to the participants. The young faces were photographs of Israeli and Chinese students (20–30 years old), and the old faces were downloaded from the Internet. All the stimuli were grayscale and were equated for luminance, brightness, and contrast. The stimuli were 360 × 360 pixels. Seen from a distance of 70 cm, they subtended 9.9º of visual angle (~12 cm).

Design and Procedure

The faces were categorized according to race, gender, and age in separate blocks. The participants were instructed to respond as quickly and as accurately as possible by pressing alternative buttons on a response key. The mapping of response hand to face category and the blocks' order were counterbalanced across participants. In each block, the stimuli were randomized and presented in sequences of 64, 64, and 72 stimuli, with a short break between sequences. Each face was presented at fixation for 300 msec, with an intertrial interval ranging randomly between 400 and 600 msec, measured after response. Each block started with 32 practice trials in which faces that were not used in the actual test were used. The experiment lasted approximately 20 min.

Accuracy rates and RTs (from the stimulus onset) were analyzed by a mixed-model ANOVA. For each participant, RTs that were more extreme than ±2 *SD*s from the mean in each condition were excluded (fewer than 2%). The between-groups factor was race of the observer (Chinese or Israeli) and the within-subjects factors were race of the face (Chinese or Caucasian) and task (race, gender, or age). When necessary, degrees of freedom were corrected for nonsphericity, using the Greenhouse–Geisser (G–GE) adjustment.[3]

RESULTS

Reaction Times

Mean RTs (for correct responses) for the different experimental conditions (two races of faces and three tasks) are presented in Figure 2. An ANOVA showed no main effect of race of the observer [$F(1,64) = 1.07$, $MS_e = 57{,}239$, $p = .30$, $\eta^2_p = .016$] and no task × race-of-the-observer interaction [$F(2,128) = 1.55$, $MS_e = 7{,}293$, $p = .22$; $\eta^2_p = .024$; G–GE = .68]. However, there was a significant main effect of task [$F(2,128) = 109.99$, $MS_e = 7{,}292$, $p < .001$, $\eta^2_p = .63$; G–GE = .94], which was further elaborated by univariate post hoc contrasts. These analyses showed that age decisions were faster (457.9 msec) than either gender decision [581.2 msec;

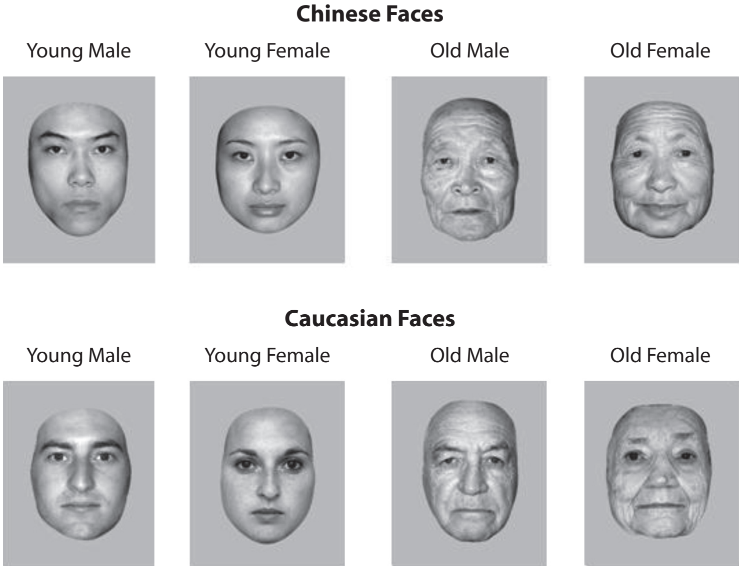

Figure 1. Examples of Chinese and Caucasian faces of different genders and ages.

$t(65) = 11.33$, $p < .001$, $d = 0.97$] or race decisions (596.7 msec; $t = 12.94$, $p < .001$, $d = 1.25$), with no difference between the latter two tasks ($t = 1.81$, $p = .23$, $d = 0.14$). The main effect of race of the face was also significant [$F(1,64) = 7.54$, $MS_e = 746$, $p < .01$, $\eta_p^2 = .11$] but was qualified by significant race-of-the-face × race-of-the-observer [$F(1,64) = 23.07$, $MS_e = 746$, $p < .001$, $\eta_p^2 = .26$] and race-of-the-face × task [$F(2,128) = 3.83$, $MS_e = 1{,}243$, $p < .05$, $\eta_p^2 = .056$] interactions. Importantly, the three-way interaction was also significant [$F(2,128) = 25.31$, $MS_e = 1{,}243$, $p < .001$, $\eta_p^2 = .28$; G–GE = .68].

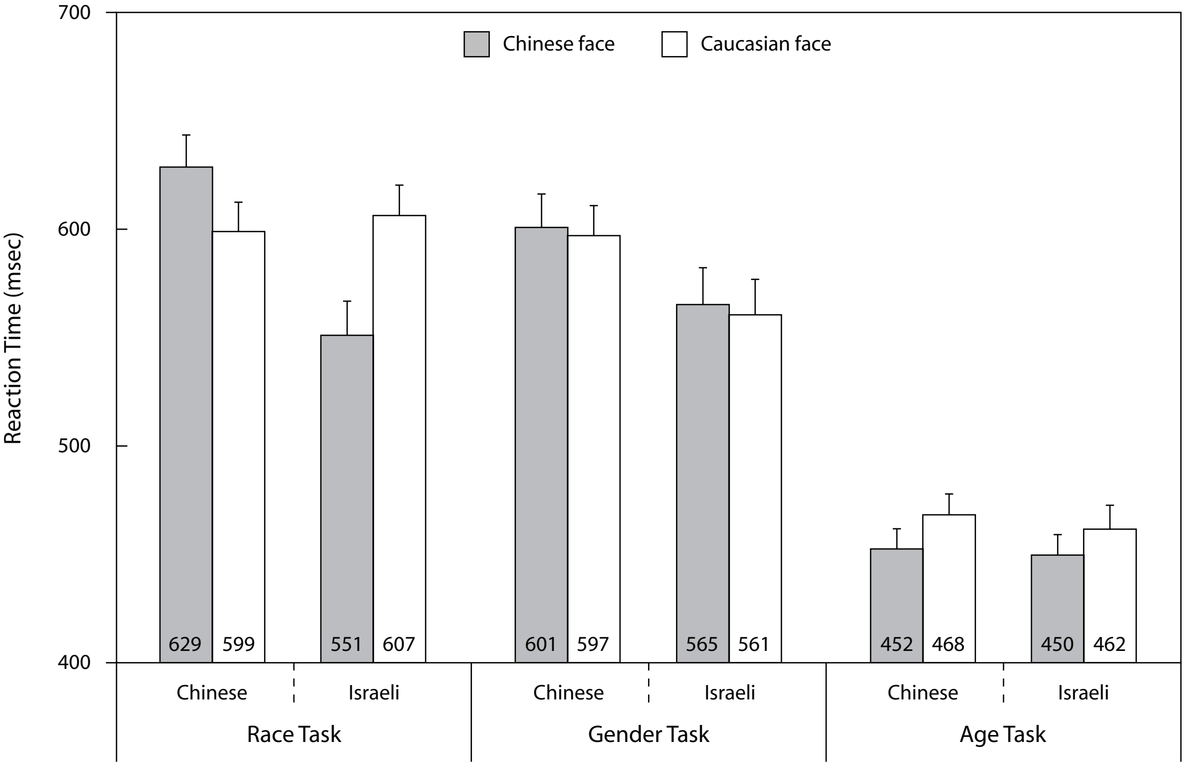

Figure 2. Reaction times of Chinese and Israeli participants for subordinate categorization of Chinese and Caucasian faces by race, gender, and age. The error bars are standard errors of the means.

Table 1
Mean Reaction Times (in Milliseconds; With Standard Deviations) of Chinese and Israeli Participants in the Gender and Age Categorization Tasks Split by Age (in the Gender Task) and by Gender (in the Age Task)

		Gender Task				Age Task			
		Young		Old		Male		Female	
Participants	Faces	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Chinese	Caucasian	566	127	628	130	481	91	455	18
(n = 36)	Chinese	581	124	618	140	454	81	451	78
Israeli	Caucasian	545	127	575	130	471	90	453	80
(n = 30)	Chinese	555	124	576	140	453	81	446	78

The three-way interaction was explored with separate race-of-the-face × race-of-the-observer ANOVAs for each task. For race decisions, the race-of-the-face × race-of-the-observer interaction was significant [$F(1,64) = 30.37$, $MS_e = 1{,}962$, $p < .001$, $\eta_p^2 = .32$], reflecting the fact that for both Chinese and Israelis, other-race faces were classified more quickly than own-race faces (Figure 2). Although the ORA was slightly stronger for Israelis (55.3 msec) than for Chinese (30.0 msec), this difference was not significant [$t(64) = 1.63$, $p = .11$, $d = 0.42$]. In contrast to race categorization, there was no race-of-the-face × race-of-the-observer interaction either for gender, [$F(1,64) < 1.00$] or for age categorization [$F(1,64) < 1.00$]. More detailed analyses showed that gender categorization was faster for young (562 msec) than for old (600 msec) faces [$F(1,64) = 98.63$, $MS_e = 932$, $p < .001$, $\eta_p^2 = .61$], and although the age effect on gender categorization was larger for Caucasian faces (47.4 msec) than for Chinese faces [29.7 msec; $F(1,64) = 5.37$, $MS_e = 868$, $p < .025$, $\eta_p^2 = .08$], it was significant both for race of the faces and for race of the observers [$t(65) = 9.30$, $p < .001$, $d = 0.37$, and $t(65) = 5.21$, $p < .001$, $d = 0.24$, respectively; Table 1]. An ANOVA including the gender of the observer as an additional between-subjects factor showed that the male and female participants categorized the face's gender equally quickly [$F(1,62) < 1$], and there was no gender-of-the-observer × gender-of-the-face interaction [$F(1,62) = 1.02$]. Age categorization was faster for Chinese faces (451 msec) than for Caucasian faces (465 msec) [$F(1,64) = 38.70$, $MS_e = 165$, $p < .001$, $\eta_p^2 = .38$] and was faster for female faces (452 msec) than for male faces (465 msec) [$F(1,64) = 38.56$, $MS_e = 304$, $p < .001$, $\eta_p^2 = .38$]. A significant race-of-the-face × gender-of-the-face interaction [$F(1,64) = 18.60$, $MS_e = 264$, $p < .001$, $\eta_p^2 = .23$] reflected the fact that the gender-of-the-face effect on age categorization was significant for Caucasian faces [$t(65) = 6.78$, $p < .001$, $d = 0.28$] but only approached significance for Chinese faces [$t(65) = 1.807$, $p = .075$, $d = 0.06$; see Table 1].

Accuracy

The percentages of correct responses in the six experimental conditions (three tasks and two races of faces) are presented in Figure 3. These data were analyzed using the same statistical model as that for RTs.

The ANOVA showed that, across tasks, Chinese (90.5%) processed faces more accurately than did Israelis (87.7%) [$F(1,64) = 5.9$, $MS_e = 139$, $p < .025$, $\eta_p^2 = .08$], but the race-of-the-observer × task interaction was not significant [$F(2,128) = 2.27$, $MS_e = 26$, $p = .117$, $\eta_p^2 = .034$]. As for RTs, there was a main effect of task on accuracy [$F(2,128) = 129.03$, $MS_e = 26$, $p < .001$, $\eta_p^2 = .67$, G–GE = .84]. Post hoc comparisons showed that age was more accurately classified (94.2%) than was race (88.1%) [$F(1,65) = 82.13$, $MS_e = 14$, $p < .001$, $\eta_p^2 = .56$], which was more accurately classified than was gender (85.1%) [$F(1,65) = 24.55$, $MS_e = 12$, $p < .001$, $\eta_p^2 = .27$]. The main effect of the race of the face was also significant [$F(1,64) = 4.10$, $MS_e = 29$, $p < .05$, $\eta_p^2 = .06$], but was qualified by a significant interaction with task [$F(2,128) = 19.21$, $MS_e = 43$, $p < .001$, $\eta_p^2 = .23$, G–GE = .79] and by a three-way interaction [$F(2,128) = 5.12$, $MS_e = 43$, $p < .01$, $\eta_p^2 = .074$].

As in the RT analysis, separate race-of-the-observer × race-of-the-face ANOVAs showed that the two factors interacted only for race categorization [$F(1,64) = 5.46$, $MS_e = 68$, $p < .025$, $\eta_p^2 = .079$]. The same interactions for gender and age categorizations were $F(1,64) < 1$ and $F(1,64) = 1.95$, $MS_e = 8$, $p = .167$, $\eta_p^2 = .03$, respectively. The post hoc analysis of the simple effects in the race categorization task showed the expected trend toward an ORA. Chinese categorized the race of Caucasian faces more accurately (91.4%) than the race of Chinese faces (88.4%), and Israelis categorized the race of Chinese faces more accurately (88.1%) than the race of Caucasian faces (84.4%), but despite the significant interaction, neither of these two simple effects was significant.

In the gender categorization task, a main effect of race of the face [$F(1,64) = 57.13$, $MS_e = 21.5$, $p < .001$, $\eta_p^2 = .47$] and the absence of an interaction between race of the face and race of the observer [$F(1,64) < 1$] indicated that both groups categorized the gender of Caucasian faces (88.1%) more accurately than the gender of Chinese faces (82.0%). A more detailed analysis showed that young faces were categorized by gender more accurately (91.0%) than were old faces (79.1%) [$F(1,64) = 362.26$, $MS_e = 25$, $p < .001$, $\eta_p^2 = .85$], and notwithstanding the interaction of this effect with race of the face [$F(1,64) = 18.30$, $MS_e = 37$, $p < .001$, $\eta_p^2 = .22$], it was significant for both Chinese faces [$t(35) = 14.16$, $p < .001$, $d = 2.2$] and Caucasian faces [$t(35) = 12.78$, $p < .001$, $d = 1.5$]. Importantly, none of these effects interacted with race of the observer (all Fs < 1; see Table 2). ANOVA including the gender of the observer as an additional factor showed

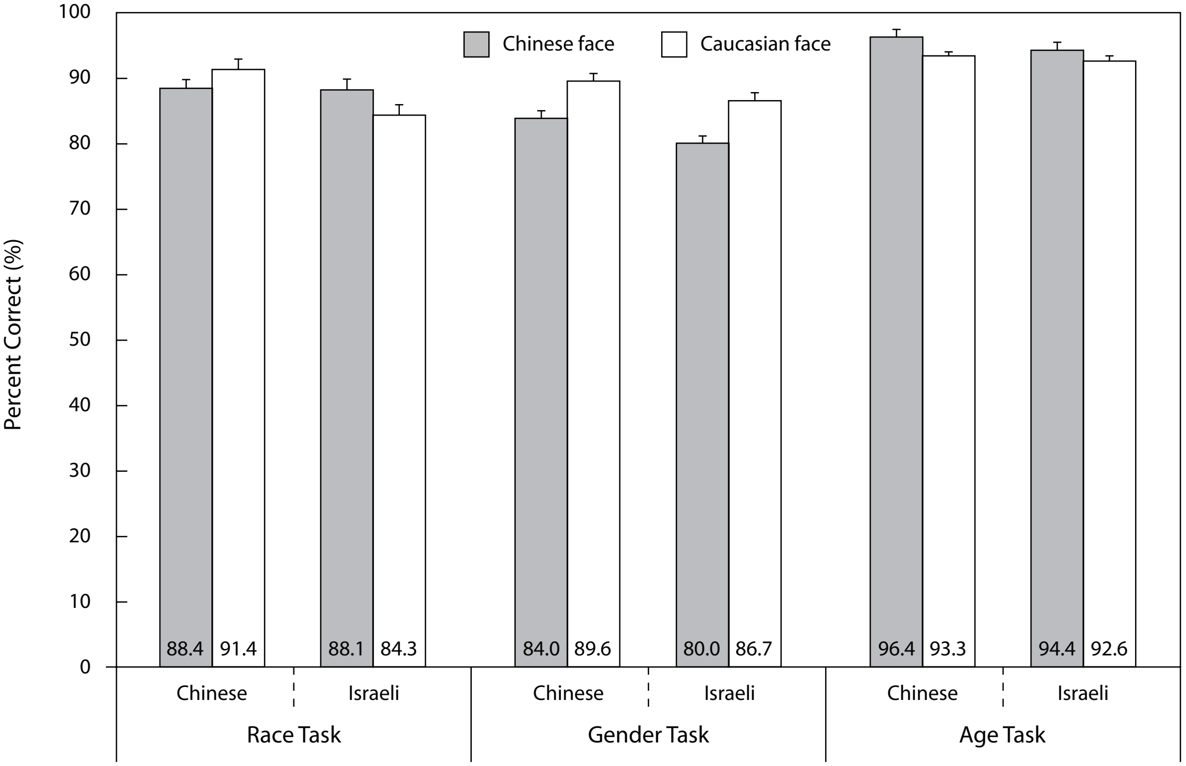

Figure 3. Percentages of correct responses of Chinese and Israeli participants for subordinate categorization by race, gender, and age. The error bars are standard errors of the means.

that male and female observers categorized the gender of faces equally accurately [$F(1,62) < 1$], and there was no interaction between gender of the observer and gender of the face [$F(1,62) < 1$].

Chinese faces were categorized by age more accurately than were Caucasian faces [95.4% vs. 93.0%, respectively; $F(1,64) = 25.66$, $MS_e = 8$, $p < .001$, $\eta_p^2 = .29$]. In addition, there was a main effect of gender of the face [$F(1,64) = 32.26$, $MS_e = 32$, $p = .001$, $\eta_p^2 = .33$], suggesting that the age of female faces was categorized more accurately (95.6%) than that of male faces (92.8%). This effect was qualified by a significant interaction with race of the face [$F(1,64) = 56.43$, $MS_e = 18$, $p < .001$, $\eta_p^2 = .47$], as well as by a second-order interaction of race of the face × gender of the face × race of the observer [$F(1,64) = 5.14$, $MS_e = 18$, $p < .05$, $\eta_p^2 = .07$]. Separate ANOVAs showed that the interaction between gender of the face and the race of the face was significant for both Chinese [$F(1,35) = 48.12$, $MS_e = 20$, $p < .001$, $\eta_p^2 = .58$] and Israeli [$F(1,29) = 14.20$, $MS = 28$, $p < .001$, $\eta_p^2 = .33$] participants. Across groups, the age of Caucasian females (95.8%) was categorized more accurately than was that of Caucasian males (90.1%) [$t(65) = 8.42$, $p < .001$, $d = 1.3$], whereas for Chinese faces, the ages of males and females were categorized equally accurately [95.5% and 95.4%, respectively; $t(6) = 0.16$, $d = 0.02$].

Finally, in order to explore possible response bias effects (Meissner, Brigham, & Butz, 2005; Slone, Brigham, & Meissner, 2000; Sporer, 2001), we used a signal detection analysis, calculating the d' and the response bias criterion (c), for each task with a distinction between own-race and other-race faces in the gender and age tasks. The pattern of d's was identical to that described above, and there were no significant response biases.[4]

Table 2
Mean Percentages Correct (With Standard Deviations) of Chinese and Israeli Participants in the Gender and Age Categorization Tasks Split by Age (in the Gender Task) and by Gender (in the Age Task)

		Gender Task				Age Task			
		Young		Old		Male		Female	
Participants	Faces	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Chinese	Caucasian	94.4	5.4	84.7	8.4	90.1	7.8	96.6	4.2
($n = 36$)	Chinese	90.7	6.6	77.3	8.4	96.8	4.8	96.1	4.2
Israeli	Caucasian	90.6	5.5	82.9	8.2	90.3	7.7	94.9	4.4
($n = 30$)	Chinese	88.5	6.6	71.6	8.2	94.0	4.9	94.7	3.8

DISCUSSION

In the present study, we investigated how visual experience with faces of a particular race affects categorization of faces by race, gender, and age in Chinese and Israelis, using a within-subjects design with participants who had only minimal experience with the other-race faces. As would be predicted by previously documented ORA, both groups of participants classified the race of the face more quickly and more accurately for other-race than for own-race faces. In contrast, race of the observer did not interact with race of the face either for gender or for age categorization. The absence of these interactions suggests that the physiognomic characteristics that determine the gender of a face and allow its categorization by age are universal, rather than race specific. Furthermore, the absence of an interaction between race of the face and race of the observer in the gender and age categorization tasks suggests that determining the race of a face is not imposed as a first step in the processing of faces from other races, preempting the perception of other category-defining physiognomic characteristics in other-race faces, as is suggested by the *race feature* theory (Levin, 1996). Yet these data do not speak against the possibility that classification by race could be based on race-distinctive features (Levin, 2000).

Interestingly, both groups classified the gender of Caucasian faces more accurately than the gender of Chinese faces but classified the age of Caucasian faces less accurately and more slowly than the age of Chinese faces. The better distinction between young and old Chinese faces could be explained by the particular stimulus selection (constrained by availability on the Internet)[5] and, therefore, should be considered with caution. In contrast, the main effect of race of the face on categorization of faces by gender could hardly be explained by stimulus selection artifacts, because all the young faces were portraits of males and females chosen at random among students at the Hebrew University and the Xuzhou Normal University. In addition, all the faces were presented without any head or facial hair. Since hair is usually more abundant in Caucasian than in Asian males, it could have cued gender discrimination better in the former group. Indeed, it has been shown that hair is one of the most important cues in memory for a face's gender (Wright & Sladden, 2003), as well as for race (Maclin & Malpass, 2003) and face recognition (e.g., O'Donnell & Bruce, 2001). Finally, note that, as in the present pattern, O'Toole et al. (1996) found that regardless of the participants' races, gender categorization was less accurate for Japanese than for Caucasian faces. On the basis of the reasons above, we suggest that physiognomic gender characteristics are probably more distinct in Caucasian faces than in Chinese faces.

The ORA for race categorization in Israeli and Chinese participants replicates and extends similar patterns found with participants living in the same country (Bruyer, Leclere, & Quinet, 2004; Levin, 1996; Valentine & Endo, 1992). This outcome demonstrates that faster categorization of other-race faces cannot be attributed solely to other races' forming minority groups (cf. Bartsch & Judd, 1993; Mullen, 1991). Rather, experience in the processing of face features associated with a particular race accounts, at least in part, for this effect. However, the present results do not replicate the previously reported ORA for gender (O'Toole et al., 1996) and age estimation (Dehon & Brédart, 2001). Both these discrepancies could be explained by important methodological differences between the present and previous studies, as well as by a difference in the stimuli that have been presented. In addition to the differences mentioned in the introduction, in the O'Toole et al. study of gender categorization, the faces were exposed for a brief duration and were backward masked. This procedure may have increased task difficulty, imposing decisions based only on first-glance (putatively global) impressions, which could have been more accurate for own- than for other-race faces. Note, however, that despite that difference, gender categorization in both experiments was similarly accurate ($d' = 2.57$ and 1.99 for Caucasian and Chinese faces, respectively, in our experiment, as compared with $d' = 2.52$ and 2.08 for Caucasian and Japanese faces, respectively, in O'Toole et al., 1996). Finally, although the Asian faces in their study were Japanese, none of the participants were Japanese. The 15 Oriental participants were actually Koreans, Chinese, Vietnamese, and Filipino; hence, for all the participants, the Japanese faces, as well as the Caucasian faces, represented, indeed, people from a different ethnic group. Consequently, it is not clear what the origin of the ORE in that study was.

The difference between the present results and the ORA for age estimation reported by Dehon and Brédart (2001) shows that although age might be an ill-defined criterion for categorization, it is not influenced by race, at least when the differences between the age groups across both races are conspicuous (as is revealed by the significantly shorter RTs in the age categorization task than in the race and gender categorization tasks).[6] In contrast, age estimation is a more difficult perceptual task than is categorization by age; hence, it is more sensitive to subtle age distinctions that could benefit from experience.

Notwithstanding the differences between the results of the present study and those in previous reports, and the possibility that under certain circumstances, gender and age information can be read better from own-race than from other-race faces, most studies have converged in showing ORA for categorization of faces by race, but not by gender or age. Although not all possible dimensions along which faces can be grouped have been investigated, at the very least the present results suggest that ORA is task specific, rather than reflecting different priorities of processing race, gender, or age (i.e., different entry points to the conceptual system) for own- and other-race faces. Along with Valentine's (1991; Valentine & Endo, 1992) multidimensional space model of face encoding and classification, we propose that the reduced sensitivity to individual defining features renders other-race faces a more homogeneous group and, hence, easier to define. This hypothesis is also supported by a study showing that the ORA is reduced if the other-race faces have been previously familiarized by training (Bruyer et al., 2004), as well as by new data showing that the down-shift categorization for faces disappears when the homogeneity among them is increased (D'Lauro, Tanaka, & Curran, 2008). This pattern goes along with Murphy and Brownell's (1985) differentiation hypothesis suggesting that the entry point of recognition is jointly determined by the observer's perceptual expertise with stimuli in the relevant domain and the homogeneity of the stimulus environment.

In conclusion, the present study demonstrates that higher expertise with own-race faces is not consequential for all kinds of subordinate group-level categorizations of faces. Moreover, it does not support a view suggesting that even for other-race faces, race is not always the first feature to be distinguished, prior to other facial characteristics, such as gender or age. Rather, the present data suggest that subordinate group-level categorizations of faces are task specific and flexible, being modulated by universal, as well as by race-specific, physiognomic characteristics.

AUTHOR NOTE

The study was supported by NIMH Grant R01 MH 64458 to S.B. We are indebted to Assaf Harel for important comments on an earlier version of the manuscript. Correspondence concerning this article should be addressed to S. Bentin, Department of Psychology, Hebrew University of Jerusalem, Jerusalem 91905, Israel (e-mail: shlomo.bentin@huji.ac.il).

REFERENCES

Bartsch, R. A., & Judd, C. M. (1993). Majority–minority status and perceived ingroup variability revisited. *European Journal of Social Psychology*, **23**, 471-483.

Blascovich, J., Wyer, N. A., Swart, L. A., & Kibler, J. L. (1997). Racism and racial categorization. *Journal of Personality & Social Psychology*, **72**, 1364-1372.

Bruyer, R., Leclere, S., & Quinet, P. (2004). Ethnic categorisation of faces is not independent of face identity. *Perception*, **33**, 169-179.

Byatt, G., & Rhodes, G. (2004). Identification of own-race and other-race faces: Implications for the representation of race in face space. *Psychonomic Bulletin & Review*, **11**, 735-741.

Dehon, H., & Brédart, S. (2001). An "other-race effect" in age estimation from faces. *Perception*, **30**, 1107-1113.

D'Lauro, C., Tanaka, J. W., & Curran, T. (2008). The preferred level of face categorization depends on discriminability. *Psychonomic Bulletin & Review*, **15**, 623-629.

Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, **85**, 445-463.

Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 1364-1382.

Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross race recognition deficit. *Journal of Experimental Psychology: General*, **129**, 559-574.

Maclin, O. H., & Malpass, R. S. (2003). Last but not least. *Perception*, **32**, 249-252.

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, & Law*, **7**, 3-35.

Meissner, C. A., Brigham, J. C., & Butz, D. A. (2005). Memory for own- and other-race faces: A dual-process approach. *Applied Cognitive Psychology*, **19**, 545-567.

Mullen, B. (1991). Group composition, salience, and cognitive representations: The phenomenology of being in a group. *Journal of Experimental Social Psychology*, **27**, 297-323.

Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **11**, 70-84.

O'Donnell, C., & Bruce, V. (2001). Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception*, **30**, 755-764.

O'Toole, A. J., Peterson, J., & Deffenbacher, K. A. (1996). An "other-race effect" for categorizing faces by sex. *Perception*, **25**, 669-676.

Rhodes, G., Hayward, W. G., & Winkler, C. (2006). Expert face coding: Configural and component coding of own-race and other-race faces. *Psychonomic Bulletin & Review*, **13**, 499-505.

Slone, A. E., Brigham, J. C., & Meissner, C. A. (2000). Social and cognitive factors affecting the own-race bias in whites. *Basic & Applied Social Psychology*, **22**, 71-84.

Sporer, S. L. (2001). Recognizing faces of other ethnic groups: An integration of theories. *Psychology, Public Policy, & Law*, **7**, 36-97.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology*, **43A**, 161-204.

Valentine, T., & Endo, M. (1992). Towards an exemplar model of face processing: The effects of race and distinctiveness. *Quarterly Journal of Experimental Psychology*, **44A**, 671-703.

Wright, D. B., & Sladden, B. (2003). An own gender bias and the importance of hair in face recognition. *Acta Psychologica*, **114**, 101-114.

NOTES

1. Similar predictions result from a second version of the multidimensional space theory (the *norm-based model*) that assumes the existence of a face *prototype* at the origin of the multidimensional space and slightly different computational principles. Since the face prototype is based on experience, it represents own-race faces better than it does other-race faces. However, as was elegantly demonstrated by Levin (1996), distinctiveness of exemplars from a prototype is probably not the most important factor determining the ORE (but see Sporer, 2001, for an opposite view).

2. For simplicity, we will use the acronym ORA also when referring to an other-race advantage in categorization of face groups on dimensions other than race.

3. For simplicity, the uncorrected degrees of freedom are presented along with the G–GE epsilon.

4. The data are available upon request from the second author.

5. For example, it is possible that the group of older Chinese people was actually older than the older Caucasians. Note, however, that the categorization by age was very accurate for both groups.

6. It is possible, however, that if the distinction between ages had been more difficult, age categorization also could have revealed a race-of-the-participant × race-of-the-face interaction.

(Manuscript received December 27, 2007;
revision accepted for publication June 13, 2008.)