Psychological Science
http://pss.sagepub.com/

Unaware Person Recognition From the Body When Face Identification Fails

Allyson Rice, P. Jonathon Phillips, Vaidehi Natu, Xiaobo An and Alice J. O'Toole
Psychological Science published online 25 September 2013
DOI: 10.1177/0956797613492986

The online version of this article can be found at:
http://pss.sagepub.com/content/early/2013/09/25/0956797613492986

Published by:
SAGE
http://www.sagepublications.com

On behalf of:
Association for Psychological Science

Psychological Science OnlineFirst, published on September 25, 2013 as doi:10.1177/0956797613492986

Research Article

Psychological Science
XX(X) 1–9

DOI: 10.1177/0956797613492986
pss.sagepub.com

Unaware Person Recognition From the Body When Face Identification Fails

Allyson Rice[1], P. Jonathon Phillips[2], Vaidehi Natu[1], Xiaobo An[1], and Alice J. O'Toole[1]

[1]School of Behavioral and Brain Sciences, The University of Texas at Dallas, and
[2]National Institute of Standards and Technology, Gaithersburg, Maryland

Abstract

How does one recognize a person when face identification fails? Here, we show that people rely on the body but are unaware of doing so. State-of-the-art face-recognition algorithms were used to select images of people with almost no useful identity information in the face. Recognition of the face alone in these cases was near chance level, but recognition of the person was accurate. Accuracy in identifying the person without the face was identical to that in identifying the whole person. Paradoxically, people reported relying heavily on facial features over noninternal face and body features in making their identity decisions. Eye movements indicated otherwise, with gaze duration and fixations shifting adaptively toward the body and away from the face when the body was a better indicator of identity than the face. This shift occurred with no cost to accuracy or response time. Human identity processing may be partially inaccessible to conscious awareness.

Keywords

face perception, visual perception

Received 1/24/13; Revision accepted 4/29/13

The face is the most reliable, visually accessible cue to a person's identity. Humans are remarkably adept at face identification (Young & Bruce, 2011). Yet even from a distance, one can recognize a friend disembarking from a train or entering a shop across the street. The face is seen, but with less clarity than needed to resolve the features and facial configurations important for recognition (Burton, Bruce, & Hancock, 1999). In these cases, the body is often visible and may provide the necessary information for identification. From a visual perspective, body features are inherently lower resolution than facial features and consequently have the potential to provide a more robust visual signal than the face in challenging or distant viewing conditions.

Although the ability to identify others is crucial for survival, the scientific study of human recognition has focused almost exclusively on the face. Traditionally, psychological studies have concentrated on the body's ability to convey emotion (e.g., de Gelder et al., 2010), with limited emphasis on its use for identification. The few studies that have considered the role of the body in identifying people have focused on how combinations of faces and bodies support identification. Burton, Wilson, Cowan, and Bruce (1999), for example, asked participants to identify familiar and unfamiliar people from poor-quality videos by matching the person to an image. For familiar people, Burton, Wilson, et al. (1999) showed that obscuring the body or gait information in a video produced only a small decrement in identification accuracy, whereas obscuring the face was highly detrimental.

Employing a task commonly associated with eyewitness-identification scenarios, O'Toole et al. (2010) tested people's ability to match identity in pairs of images and videos of unfamiliar people, using stimuli that showed the entire person, the person with the face obscured, or the person with the body obscured. Although identification from the body was well above chance level, it was far less accurate than identification from the face. Notably, in static images, accuracy in identifying the entire person

Corresponding Author:
Alice J. O'Toole, School of Behavioral and Brain Sciences, GR4.1, The University of Texas at Dallas, Richardson, TX 75080
E-mail: otoole@utdallas.edu

was virtually identical to identification accuracy with the face alone. Combined, these results indicate that the body is useful for identification, but that people do not use it for this purpose if the face is visible and helpful. This fact was particularly surprising given that performance was not at ceiling level in either case.

Using a memory-based paradigm, Robbins and Coltheart (2012) replicated and extended these findings. Participants learned to name a small number of people from images that showed the whole person and were tested subsequently with brief (500-ms) presentations of "composite" images, as well as face-only and body-only images. Composite images were created by putting the head of one person onto the body of another person. On the composite trials, participants gave the name associated with the face used in the composite 90% of the time—the same level of identification accuracy that was achieved when participants made the identification using an image of the whole person. Consistent with the findings of O'Toole et al. (2010), results showed that when only the body was presented at test, accuracy was above chance level, again indicating the potential utility of the body for identification. Thus, it seems that when both the face and the body are useful for identification, people rely on the face (Burton, Wilson et al., 1999; O'Toole et al., 2010; Robbins & Coltheart, 2012).

From a behavioral perspective, the results of O'Toole et al. (2010) and Robbins and Coltheart (2012) suggest that people use a suboptimal strategy for identification, ignoring useful information in the body rather than integrating it into an identity decision. This leaves open the question of whether the body ever plays a meaningful role in identification when one can see the whole person. From a practical, adaptive point of view, in natural environments, the relative quality of identity information in faces versus bodies varies widely as a function of viewing conditions (e.g., illumination; Braje, Kersten, Tarr, & Troje, 1999) and person-specific factors (e.g., typicality of the face or body; e.g., Light, Kayra-Stuart, & Hollander, 1979). Consequently, identification may be more reliable from the face in some situations and from the body in others.

In the present study, we examined whether people can make use of bodies for identification when the face is present in an image but is uninformative as to identity. We conducted five experiments in which participants matched the identity of people in a special set of image pairs, chosen with the help of face-recognition algorithms. These algorithms were used to select pairs of *same-identity images*, images of the same person in which the faces appeared highly dissimilar, and pairs of *different-identity images*, images of different people in which the faces appeared highly similar. In Experiment 1, participants matched identity using the original images, which showed the face and body (including the neck, the shoulders, and part of the torso). In Experiment 2, they matched identity in pairs of images digitally edited to include only the face (face-only images). In Experiment 3, they matched identity in pairs of images edited to exclude the face (body-only images). In Experiment 4, participants again matched identity in the original images (as in Experiment 1), but they also rated their use of internal face and external face/body cues in making the identification decisions. In Experiment 5, we tracked the eye movements of participants while they matched identity in a subset of specially selected face-informative and body-informative images.

Identity Matching Experiments

The methods for the identity-matching task carried out in all experiments were similar, so we describe them together here, with later sections used to describe the rating procedure (Experiment 4) and the eye movement study (Experiment 5).

Method

Stimuli. Image pairs were selected from the database used in the 2006 Face Recognition Vendor Test (FRVT 2006), a U.S. Government–supported international competition of face-recognition systems (Phillips et al., 2010) conducted by the National Institute of Standards and Technology. The database was collected at the University of Notre Dame and contains images taken either outside or with ambient indoor illumination in a corridor. Images were acquired with a 6-megapixel Nikon D70 camera[1] (Phillips et al., 2012). Natural variations in facial expression and appearance (e.g., differences in clothing and hair style) were common in this data set.

Complete information about the FRVT 2006 competition can be found elsewhere (Phillips et al., 2012; Phillips et al., 2010). For present purposes, our goal was to use data from algorithms participating in the FRVT 2006 to select pairs of images that contained *no computationally useful information in the face*. To find image pairs with no computationally useful information in the face, we proceeded as follows. The task of algorithms in the FRVT 2006 was to assign similarity scores to pairs of images. The similarity score for each image pair was the algorithm's estimate of the likelihood that the two images showed the same person (i.e., a higher similarity score indicated a higher likelihood that the two images showed the same person). Although algorithms had access to the entire image (including the neck, the shoulders, and part of the torso), similarity scores were computed using only the face.

More concretely, algorithms rated the similarity in all possible pairs of two sets of images. Set 1 consisted of 1,085 images of 457 individuals. Set 2 contained 1,085 different images of the same individuals. Thus, each algorithm produced a 1,085 × 1,085 similarity matrix of 1,177,225 similarity scores, where element $s_{i,j}$ of the matrix contained the similarity between the *i*th image in Set 1 and the *j*th image in Set 2. The matrix contained similarity scores for 3,297 same-identity pairs and 1,173,928 different-identity pairs. We used a statistical fusion of identification estimates (i.e., the computed similarity between the two faces) from three top-performing algorithms from the FRVT 2006 competition (see Algorithm Fusion in the Supplemental Material available online for a description of the fusion procedure).

A schematic of the similarity-score distributions for same-identity and different-identity pairs appears in Figure 1a. To find image pairs in which the face failed as an identity cue, we sampled from the extreme tails of the two distributions. For same-identity pairs, we sampled 50 image pairs from extreme left tail of the same-identity distribution (highly dissimilar pairs of images showing the same person; see example pairs in Fig. 1b). For same-identity pairs, each image was taken on a different day to ensure that the person was not wearing the same clothing in both images (so that clothing was not a cue to identity). For different-identity pairs, we sampled 50 image pairs from the extreme right tail of the different-identity distribution (pairs of highly similar images showing different people, see example pairs in Fig. 1c). These served as the stimuli for Experiments 1 through 4. A subset of these pairs was selected for use with the eye tracker in Experiment 5.

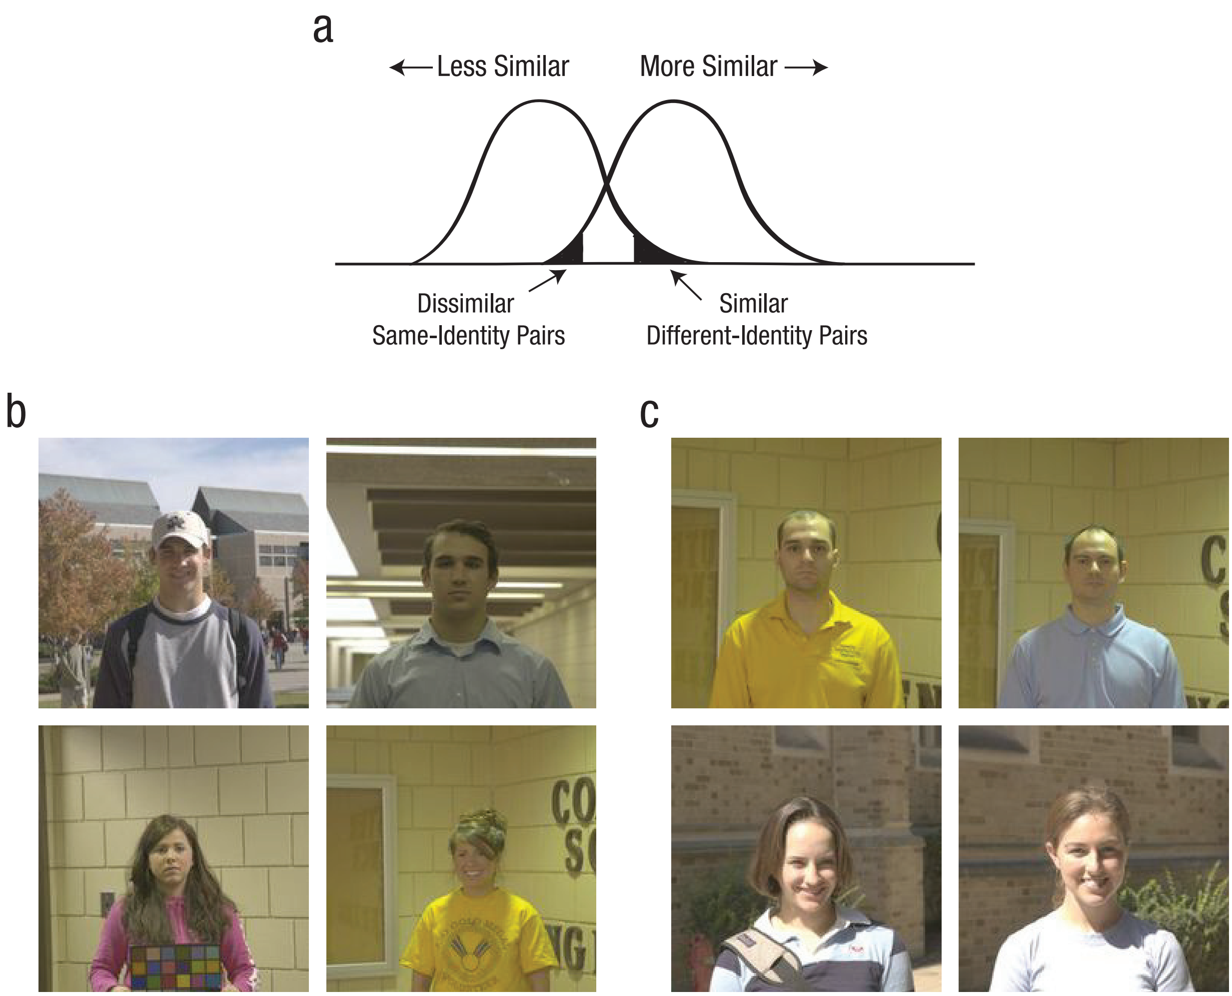

Fig. 1. Stimulus-sampling paradigm (a) and example stimuli showing dissimilar images of the same people (b) and similar images of different people (c). Panel (a) shows a schematic of similarity-score distributions for same-identity and different-identity pairs. Higher similarity scores indicate a higher likelihood that the images show the same person. Sampling from the extreme left of the same-identity distribution yields highly dissimilar images of the same people, as shown in (b), in which the images in the top row and the images in the bottom row are of the same person. Sampling from the extreme right of the different-identity distribution yields highly similar images of different people, as shown in (c), in which the images in the top row and the images in the bottom row are of different people.

Participants. Undergraduate students from the School of Behavioral and Brain Sciences at the University of Texas at Dallas participated in the experiments in exchange for a research credit in a psychology course. The sample sizes of the five experiments were as follows—Experiment 1: $N = 17$ (12 females, 5 males); Experiment 2: $N = 21$ (15 females, 6 males); Experiment 3: $N = 24$ (17 females, 7 males); Experiment 4: $N = 23$ (16 females, 7 males); Experiment 5: $N = 20$ (12 females, 8 males).[2] In Experiment 5, one male participant was excluded because of eye-tracking artifacts that resulted in signal loss on more than 20% of trials.

Stimulus processing. The original (unedited) images were used in Experiments 1, 4, and 5. For Experiment 2, we created face-only versions of the images by digitally covering everything in the image but the face with a uniform gray background. For Experiment 3, we created body-only images by covering the face with a beige oval. Examples of the edited stimuli are shown in Figure 2a.

Procedure. On each trial, participants viewed a pair of images presented side by side on a computer screen. Participants reported on their perception of the identity of the people shown in the images by choosing one of five responses: 1, *Sure they are the same person*; 2, *Think they are the same person*; 3, *Don't know*; 4, *Think they are different people*; and 5, *Sure they are different people*. This rating served as the *human-generated similarity score*, similar to that produced by the algorithms. The images and prompt ("Are these two people the same?") remained visible until a response was entered. Presentation order for the 100 image pairs was randomized for each participant.

Results

Performance in the identity-matching experiments was measured in two ways. First, we computed receiver-operator-characteristic (ROC) curves from the distribution of ratings for same-identity and different-identity image pairs. The ROC gives a more complete picture of performance than does a standard yes/no (same/different) paradigm. Second, to test for statistical differences across experiments, we used the summary measure of d'. The hit rate was defined as the proportion of same-identity pairs correctly judged to be images of the same person, and the false-alarm rate was defined as the proportion of different-identity pairs incorrectly judged to be images of the same person. Response ratings of 1 or 2 were considered same-identity judgments, and responses of 3, 4, and 5 were considered different-identity judgments.[3]

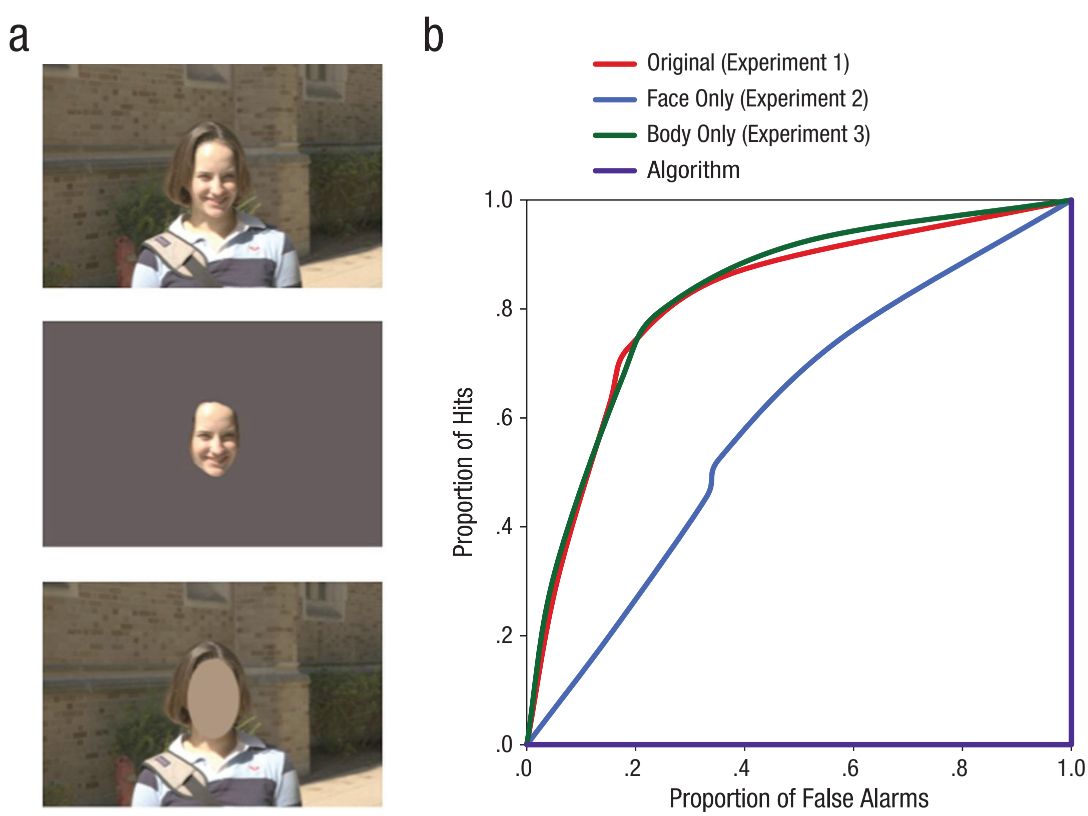

Fig. 2. Example stimuli and results. Panel (a) shows example original (Experiments 1, 4, and 5), face-only (Experiment 2), and body-only (Experiment 3) images of a person. The graph in (b) shows performance as indicated by receiver-operating-characteristic (ROC) curves as a function of condition. By design, the image pairs we selected using data from the algorithm yielded a "reverse" ROC curve, with identification incorrect for all image pairs.

Identification-accuracy results for the original images are shown in Figure 2b. Identification accuracy for the original images (Experiment 1) was well above chance level, $d' = 1.504$, $SE = 0.15$, $t(16) = 9.97$, $p < .001$ (two-tailed). Identification accuracy for the body alone (Experiment 3) was also well above chance level, $d' = 1.500$, $SE = 0.09$, $t(23) = 15.91$, $p < .0001$ (two-tailed), and was indistinguishable from accuracy for the original image (Experiment 1), $F(1, 61) < 1$, n.s. (two-tailed). Identification accuracy for the face alone (Experiment 2) was poor, $d' = 0.3637$, $SE = .15$, but greater than chance level, $t(20) = 2.39$, $p < .02$ (two-tailed). Combined, these results indicate that body information, rather than face information, accounted for participants' accuracy in identifying people in the original unedited images, which were chosen by the algorithm to have poor-quality information for face identification.

Given the pattern of results, which suggested nearly exclusive reliance on the body for identification, in Experiment 4, we asked participants to rate a set of features on the basis of how often they used each feature in making their identification decisions. This question was aimed at evaluating participants' metacognitive knowledge of their reliance on face versus body features. With a new set of participants, we repeated Experiment 1 with identity matching of the original images, but we followed this task with a questionnaire asking participants to report their reliance on a set of internal-face and external-face and -body features (see Feature-Use Questionnaire in the Supplemental Material to view the questionnaire). Specifically, participants rated how often they used features in their identification decisions, using a 5-point scale from 1, *never used*, to 5, *used all of the time*. Participants overwhelmingly reported greater reliance on the internal face than on external features and body cues (Fig. 3).

This left us with a paradox. The recognition data clearly indicated the use of body information for identification. However, the subjective ratings suggested that participants were unaware of how important the body was in their decision. To measure how participants allocated attention to a person as they made an identity decision, we used eye tracking. Our focal question was,

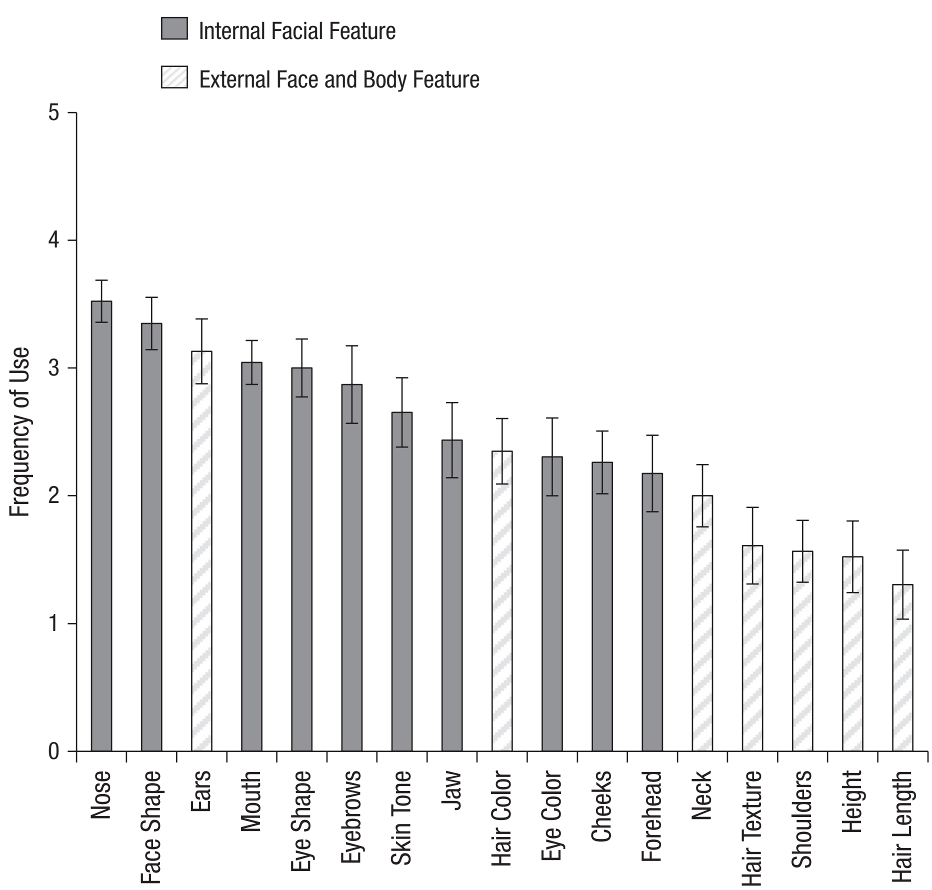

Fig. 3. Results from Experiment 4: self-reported reliance on features for identification, as a function of type of feature. Participants ranked their reliance on features using 5-point scales, with higher ratings indicating greater usage. Errors bars represent standard errors.

can attention be driven automatically on the basis of the location (or locations) of the most useful identity information?

Eye-Tracking Experiment

Method

In our final experiment, we tracked eye movements as participants matched identity using a subset of the original, unedited image pairs, which were selected so that the face or the body was the better cue to identity. Selection of these face-informative and body-informative stimuli was done using an item analysis from the data collected in Experiments 1 through 3 (see the Method section below).

We tested three hypotheses about how attention guides identification. The first was that people begin by fixating the face but redirect attention to the body if the face fails. This hypothesis predicts different patterns of gaze depending on whether the face or the body is the better cue to identity. It also predicts that more time should be spent looking at body-informative image pairs than face-informative image pairs, because people should begin with the face and then move to the body if the face does not support an identification decision. The second hypothesis was that people attend naturally to the internal face and use peripheral vision to process identity from the body. This hypothesis predicts no difference in the pattern of gaze or in the time needed to perform the task for face-informative versus body-informative pairs. The third, data-driven hypothesis was that the human visual system efficiently directs attention to the identity-informative parts of the person. This hypothesis predicts different gaze patterns depending on whether the face or body has better identity information. Critically, it also predicts no difference in the time spent looking at face-informative versus body-informative pairs.

Stimuli. Because the identity pairs were the same as those used in Experiments 1 through 3, we were able to use an item analysis to find pairs of images in which the body or the face was the better cue to identity. Body-informative image pairs were defined as those for which identification accuracy with only the body (Experiment 3) was greater than with only the face (Experiment 2). Face-informative pairs were those for which identification accuracy with only the face (Experiment 2) was greater than with only the body (Experiment 3). As expected from the results of the first three experiments, a large majority of the pairs (77 of 100) in this special data set were body-informative pairs. The remaining 23 pairs qualified as face-informative pairs. To make equal numbers of same-identity and different-identity face-informative stimulus pairs, we reduced the number of face-informative pairs to 20 (10 same-identity pairs and 10 different-identity pairs), eliminating the pairs with the smallest face advantage. We used the same procedure to select body-informative same-identity and different-identity pairs, again ranking pairs according to the difference in *d*′s for body-only (Experiment 3) minus face-only (Experiment 2) identification accuracy. Starting with the largest differences, we selected 20 body-informative pairs (10 same-identity pairs and 10 different-identity pairs). This yielded 40 pairs in total (20 face-informative and 20 body-informative pairs) for the eye-tracking experiment.

Apparatus. Eye movements were recorded with a Tobii T60 XL eye tracker (Tobii Technology, Stockholm, Sweden) attached to a 24-in. widescreen monitor. While the participant viewed stimuli on the monitor, the eye tracker calculated the location of the eyes' gaze at a rate of 60 Hz. We defined areas of interest (AOI) in each image. The AOI for the internal face was defined using the internal hairline and jaw, excluding the ears. The AOI for the external face/body region was defined as the rest of the person (see Fig. 4).

Procedure. The procedure of Experiment 5 was identical to that used in Experiment 1, except that the task was performed on a subset of stimuli while eye movements were tracked.

Results

Eye movement patterns were scored in the following manner. The first two measures were magnitude-normalized measures of the tendency to look at the face relative to the entire person. *Internal-face fixations* were defined as the proportion of fixations to the internal-face AOI relative to the entire person. *Internal gaze duration* was defined as the duration of looking at the internal-face AOI relative to the entire person; values over .50 indicated a bias to fixate the internal face versus the external face and body. As we discuss below, all conditions yielded a face bias. The second two measures were the absolute number of fixations and the overall duration of looking across the internal and external AOIs. These latter measures gauged the overall processing time allocated to these AOIs. We note that other studies have normalized eye movements to faces and bodies on the basis of the relative size of the two areas (e.g., Birmingham, Bischof, & Kingstone, 2008a, 2008b; Fletcher-Watson, Findlay, Leekam, & Benson, 2008; cf. Bindeman, Sheepers, Ferguson, & Burton, 2010). In the present case, we focused on a shift in eye movement patterns as a function

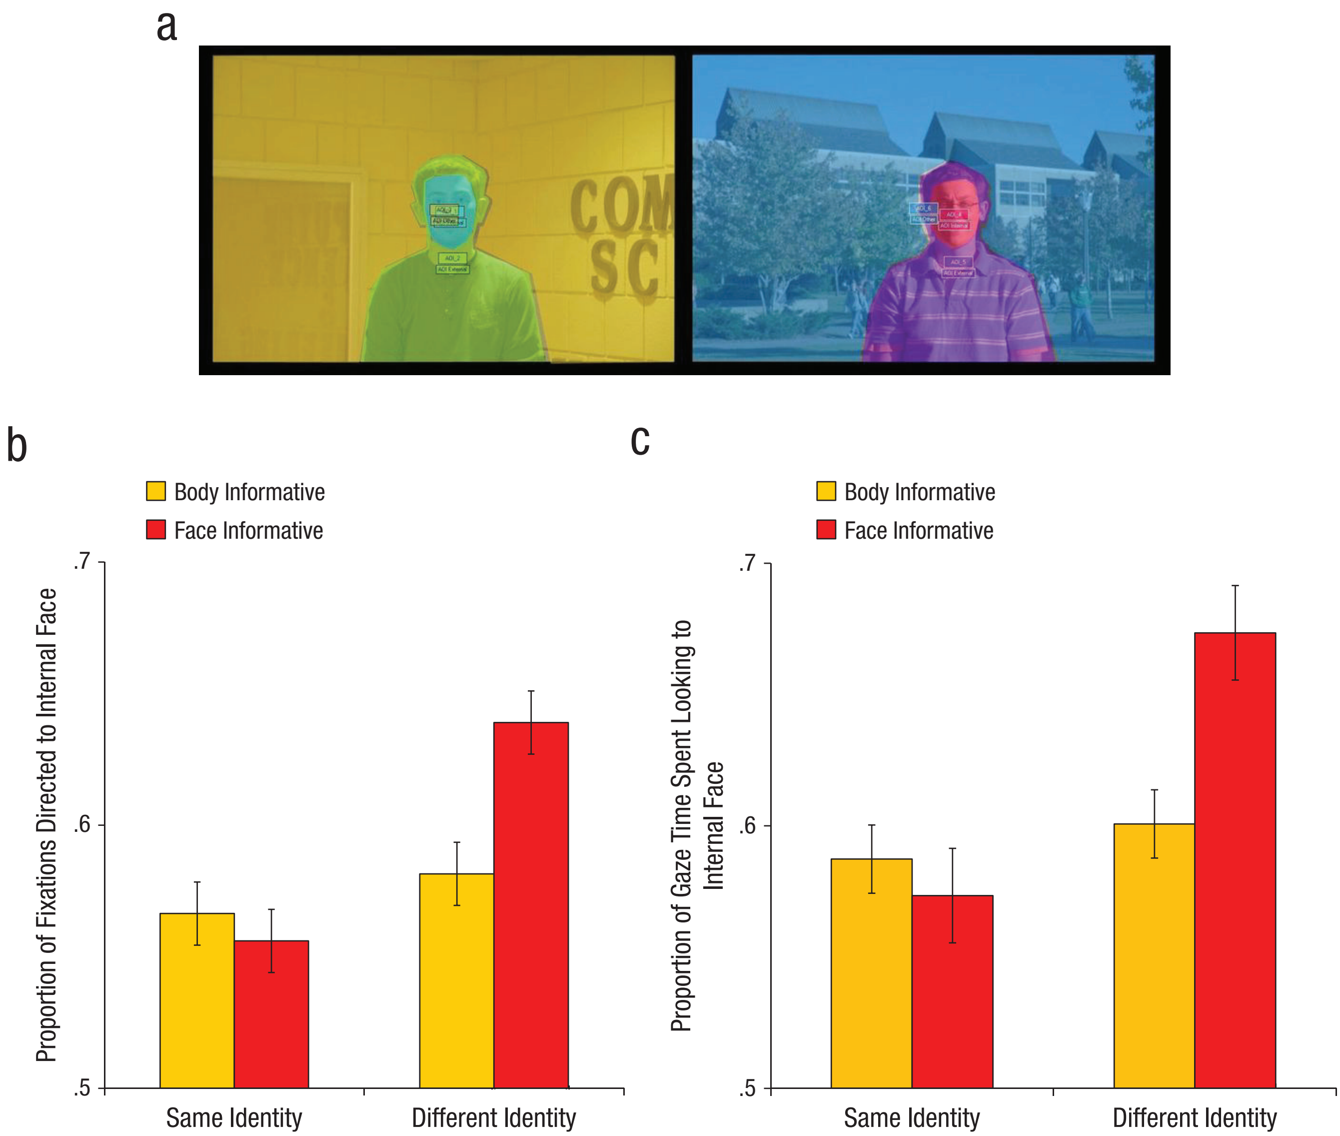

Fig. 4. Screenshot of the image-editing process used to demarcate areas of interest and results from Experiment 5. Panel (a) shows the areas of interest (AOIs) defined in the images (internal face, body without face, and other). The graphs in (b) and (c) show the proportion of fixations to the internal face relative to the person and the proportion of gaze time spent looking to the internal face relative to the person as a function of type of image pair. Error bars represent standard errors.

of the value of the identity information in the face versus the body, using absolute time as a measure of the efficiency of the strategy.

For each eye-movement measure, we computed a within-subjects analysis of variance on information type (face informative vs. body informative) and match status (same-identity pairs vs. different-identity pairs). For the relative measure of internal-face fixations, there was a main effect of information-type, $F(1, 18) = 4.97$, $p < .05$, $\eta_p^2 = .21$, and match status, $F(1, 18) = 10.43$, $p < .01$, $\eta_p^2 = .37$. These main effects were qualified by a significant interaction between match status and information type, $F(1, 18) = 7.06$, $p < .02$, $\eta_p^2 = .29$. Analogous results were found for the relative measure of internal-face gaze time, which showed a marginal main effect of information type, $F(1, 18) = 3.90$, $p < .06$, $\eta_p^2 = .18$, and a main effect of match status, $F(1, 18) = 10.31$, $p < .01$, $\eta_p^2 = .36$. Again, the main effects were qualified by a significant interaction between match status and information type, $F(1, 18) = 5.51$, $p < .03$, $\eta_p^2 = .23$. In both cases, the results indicated a weaker face bias for body-informative different-identity image pairs compared with face-informative different-identity image pairs. No difference in face bias was found for the same-identity pairs, possibly because of the longer gaze durations found for same-identity pairs relative to different-identity pairs.

The absolute number of fixations and absolute gaze-duration measures yielded a main effect of match type, indicating more fixations, $F(1, 18) = 8.06$, $p < .01$, $\eta_p^2 = .31$, and longer gaze duration, $F(1, 18) = 8.85$, $p < .01$, $\eta_p^2 = .33$, for same-identity images compared with different-identity images. Mean gaze duration for same-identity stimuli was 1,228.4 ms longer than for different-identity images, a difference that was due to participants' being more conservative in affirming a same-identity match than in rejecting a different-identity mismatch. Critically, there was no effect of information type, which means that participants did not spend more time processing face-informative pairs than body-informative pairs, making the results inconsistent with Hypothesis 1 (entailing a time-consuming reallocation of attention to the body if the face fails). Also, there was no interaction between information type and match status for either measure. Thus, the quality of identity information in the face compared with the body did not affect overall processing time. Finally, we speculate that the higher absolute processing times for same- relative to different-identity pairs may have been related to the lack of a difference in the strength of the face bias for face-informative and body-informative same-identity pairs. Specifically, it seems possible that automatic and subtle changes in eye movement patterns, such as those found with the different-identity pairs, may be easiest to detect when processing is fast and efficient.

Discussion

These experiments demonstrated that the body can be the deciding factor in person identification, even when people are unaware of using it for this purpose. Feature-use ratings revealed strikingly limited conscious access to the critical role that the body played in identification decisions. Instead, eye movements provided evidence for the importance of the body, indicating an efficient and adaptive gaze strategy tailored to optimizing identification on a stimulus-by-stimulus basis. This strategy is reminiscent of the data-driven eye movement effects found for synthetic images created with natural-scene statistics (Najemnik & Geisler, 2005, 2008). Given the complexity of human faces and bodies as visual stimuli, our results suggest the contribution of complex high-level visual-scene statistics to eye movement controls. When individuals look at people in natural scenes, eye movements are directed to the body up to 40% of the time (Bindeman et al., 2010). The present study establishes identification as one goal of these fixations. We speculate also that if the stimuli in the present study had consistently included more than just the torso, there would have been a potential for the lower body area to include additional identity information that captured attention and improved identification. These results also highlight the importance of considering multiple measures of person processing, including performance data, gaze patterns, and subjective feature-use reports.

Participants' limited awareness of body processing is analogous to other reports concerning the processing of faces, with both suggesting that basic components of the person-processing system are inaccessible to conscious awareness. Facial-expression perception has been found using electromyography in patients with cortical blindness (Tamietto et al., 2009). Likewise, familiar-face "recognition" has been detected using electrodermal skin conductance in prosopagnosics (Bauer, 1984; Tranel & Damasio, 1985).

At the outset, we noted that psychological studies have focused primarily on the body's role in social communication. How we use bodies for identification has remained a missing piece of our understanding of how we identify and socially communicate with whole people in the real world. The present study demonstrates that the strong reliance on the face for identification is not a definitive limitation of human identity processing (cf. Burton et al., 1999; O'Toole et al., 2010; Robbins & Coltheart, 2012). Instead, these findings indicate that human identity processing involves a complex but flexible system that takes into account the social tendency to look at faces, the critical need to identify other people, and the quality of information needed for these competing tasks.

Author Contributions

All authors designed the experiments. A. Rice and V. Natu implemented the identity-matching experiments, and A. Rice implemented the eye-tracking experiment. A. Rice and V. Natu tested the participants. P. J. Phillips, X. An, and A. O'Toole coordinated face-recognition-algorithm data for stimulus selection. A. Rice analyzed the data, and V. Natu and X. An conducted statistical analyses. A. Rice prepared the figures, and A. O'Toole supervised the project and composed the manuscript with the assistance of A. Rice.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by funding from the Technical Support Working Group of the Department of Defense. P. J. Phillips was supported in part by the Federal Bureau of Investigation.

Supplemental Material

Additional supporting information may be found at http://pss.sagepub.com/content/by/supplemental-data

Notes

1. The identification of any commercial product or trade name does not imply endorsement of or recommendation of it by the National Institute of Standards and Technology.
2. Although other researchers (Alaerts, Nackaerts, Meyns, Swinnen, & Wenderoth, 2011) have found an advantage for females in recognizing emotions from point-light displays of face and body movement, we found no difference between male and female participants in matching accuracy in any of our experiments.
3. To compute a *d*′ score from the rating-scale data, it was necessary to assign ratings to indicate same-identity versus different-identity judgments. We decided arbitrarily to break the scale between 2 and 3 to compute *d*′. However, for thoroughness, all analyses were repeated with the scale divided between 3 and 4. Both analyses yielded the same pattern of results.

References

Alaerts, K., Nackaerts, E., Meyns, P., Swinnen, S. P., & Wenderoth, N. (2011). Action and emotion recognition from point light displays: An investigation of gender differences. *PLoS ONE*, *6*(6), e20989. Retrieved from http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0020989

Bauer, R. M. (1984). Autonomic recognition of names and faces in prosopagnosia: A neuropsychological application of the Guilty Knowledge Test. *Neuropsychologia*, *22*, 457–469.

Bindeman, M., Sheepers, C., Ferguson, H. J., & Burton, A. M. (2010). Face, body, and center of gravity mediate person detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, *36*, 1477–1485.

Birmingham, E., Bischof, W. F., & Kingstone, A. (2008a). Gaze selection in complex social scenes. *Visual Cognition*, *16*, 341–355.

Birmingham, E., Bischof, W. F., & Kingstone, A. (2008b). Social attention and real-world scenes: The roles of action, competition and social content. *Quarterly Journal of Experimental Psychology*, *61*, 986–998.

Braje, W., Kersten, D., Tarr, M. J., & Troje, N. (1999). Illumination effects in face recognition. *Psychobiology*, *26*, 371–380.

Burton, A. M., Bruce, V., & Hancock, P. J. B. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, *23*, 1–31.

Burton, A. M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-quality video. *Psychological Science*, *10*, 243–248.

de Gelder, B., Stock, V. D., Meeren, H. K., Sinke, C. B., Kret, M. E., & Tamietto, M. (2010). Standing up for the body. Recent progress in uncovering the networks involved in the perception of bodies and bodily expressions. *Neuroscience Biobehavioral Reviews*, *34*, 513–527.

Fletcher-Watson, S., Findlay, J. M., Leekam, S. R., & Benson, V. (2008). Rapid detection of person information in a naturalistic scene. *Perception*, *37*, 571–583.

Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory*, *5*, 212–228.

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*, 387–391.

Najemnik, J., & Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, *8*(3), Article 4. doi:10.1167/8.3.4

O'Toole, A. J., Weimer, S., Dunlop, J., Barwick, R., Ayyad, J., & Phillips, P. J. (2010). Recognizing people from dynamic video: Dissecting identity with a fusion approach. *Vision Research*, *51*, 74–83.

Phillips, P. J., Beveridge, J. R., Draper, B. A., Givens, G., O'Toole, A. J., Bolme, D., . . . Wiemer, S. (2012). The Good, the Bad, and the Ugly Face Challenge Problem. *Image and Vision Computing*, *30*, 177–185.

Phillips, P. J., Scruggs, W. T., O'Toole, A. J., Flynn, P. J., Bowyer, K. W., Schott, C. L., & Sharpe, M. (2010). FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE: Transactions on Pattern Analysis and Machine Intelligence*, *32*, 831–846.

Robbins, R., & Coltheart, M. (2012). The effects of inversion and familiarity on face versus body cues to person recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 1098–1104.

Tamietto, M., Catelli, L., Vighetto, S., Perozzo, P., Geminiani, G., Weisenkrantz, L., & de Gelder, B. (2009). Unseen facial and body expressions trigger emotional reactions. *Proceedings of the National Academy of Sciences, USA*, *106*, 17661–17666.

Tranel, D., & Damasio, A. R. (1985). Knowledge without awareness: An autonomic index of facial recognition by prosopagnosics. *Science*, *228*, 1453–1454.

Young, A. W., & Bruce, V. (2011). Understanding person perception. *British Journal of Psychology*, *102*, 959–974.