

The movement advantage in famous and unfamiliar faces: A comparison of point-light displays and shape-normalized avatar stimuli

Rachel J Bennetts; Jeesun Kim

University of Western Sydney

Darren Burke

University of Newcastle

Kevin R Brooks

Macquarie University

Simon Lucey and Jason Saragih

CSIRO

Rachel A Robbins

University of Western Sydney

Author note

Rachel Bennetts is now at the School of Design, Engineering, and Computing, Bournemouth University, UK.

Correspondence concerning this article should be addressed to Rachel Bennetts, School of Design, Engineering, and Computing, Bournemouth University, Poole House, Talbot Campus, Fern Barrow, Poole BH12 5BB. Email: rbennetts@bournemouth.ac.uk

Abstract

Facial movement may provide cues to identity, by supporting the extraction of face shape information via structure-from-motion, or via characteristic patterns of movement. Currently, it is unclear whether familiar and unfamiliar faces derive the same benefit from these mechanisms. This study examined the movement advantage by asking participants to match moving and static images of famous and unfamiliar faces to facial point-light-displays (PLDs) or shape-normalized avatars in a same/different task (Experiment 1). Experiment 2 also used a same/different task, but participants matched from PLD to PLD or from avatar to avatar. In both experiments, unfamiliar face matching was more accurate for PLDs than avatars, but there was no effect of stimulus type on famous faces. In Experiment 1, there was no movement advantage, but in Experiment 2, there was a significant movement advantage for famous and unfamiliar faces. There was no evidence that familiarity increased the movement advantage. For unfamiliar faces, results suggest that participants were relying on characteristic movement patterns to match the faces, and did not derive any extra benefit from the structure-from-motion cues in the PLDs. The results indicate that participants may use static and movement-based cues in a flexible manner when matching famous and unfamiliar faces.

A comparison of point-light-displays and avatars in famous and unfamiliar face recognition.

Facial movements provide a rich source of information about a person – we can use the way a face moves to help speech perception (Calvert & Campbell, 2003), or to infer emotions (Bassili, 1979), gender (Hill & Johnston, 2001), and even the identity of a person (Roark, Barrett, Spence, Abdi & O’Toole, 2003). There are several ways that facial movement can provide cues to identity. Firstly, movement may support the extraction of structural information (three-dimensional form) from the face, via structure-from-motion processes. Secondly, it may be possible to recognize someone based on the characteristic or idiosyncratic way they move their face and head (Roark et al., 2003). However, the relative importance of these different types of information for face recognition has yet to be established. Furthermore, it is unclear whether familiarity with a face helps us extract and use movement information in identity-based tasks such as naming or matching, and whether an effect of familiarity is present for both types of movement information. These questions were examined in the current study.

Many studies have found that familiar faces are recognized faster and more accurately when presented in motion (Knight & Johnston, 1997; Lander & Bruce, 2004; Lander, Christie & Bruce, 1999). However, the role of movement in unfamiliar face recognition is uncertain. Although some studies have found a benefit of movement when matching or learning unfamiliar faces (Hill & Johnston, 2001; Lander & Bruce, 2003; Pike, Kemp, Towell & Philips, 1997; Thornton & Kourtzi, 2002), other studies have found no effect of movement for unfamiliar faces (Bruce & Valentine, 1988; Bruce, Henderson, Newman & Burton, 2001, Experiment 2; Christie & Bruce, 1998; Shiff, Banka & de Bordes-Galdi, 1986). Where present, it is unclear whether the movement advantage for familiar and unfamiliar faces is due to structure-from-motion cues or the presence of characteristic movement patterns, or a combination of the two. Some studies have addressed this question by showing participants staged movements (e.g., Christie & Bruce, 1998; Pike et al., 1997) or changing the timing of the movements (e.g., Lander & Bruce, 2000; 2004), thus disturbing the

extraction of characteristic movement patterns; other studies have used uniformly shaped faces and heads (e.g., Hill & Johnston, 2001; Knappmeyer, Thornton, & Bulthoff, 2003), or only displayed expression and speech movements (e.g., Rosenblum, Niehus, & Smith, 2007; Thornton & Kourtzi, 2002), in order to limit the use of structure-from-motion cues to identity. However, no studies have directly compared the use of characteristic movement patterns and structure-from-motion, or investigated whether familiarity affects our use of each of these movement cues.

The robust effect of movement for the recognition of familiar, but not unfamiliar faces might reflect differences in the way that these types of stimuli are processed. That is, for familiar faces, people may have developed stored patterns of characteristic movements based on prior exposure to the face (Roark et al., 2003), in addition to well-developed structural representations (i.e., robust representations of the unchanging aspects of the face and head; Jenkins & Burton, 2011), and these may support the use of movement-based mechanisms during identity-based tasks such as naming or matching. Unfamiliar faces, on the other hand, do not have pre-existing representations of characteristic movements or stable face representations, which could make it difficult for people to extract and use movement-based cues to identity. This is particularly the case for characteristic movement patterns, which may take a longer time to build up than a structural representation of the face (O'Toole, Roark, & Abdi, 2002).

Alternatively, it is possible that the conflicting results from past studies were due to methodological differences, such as the task, or the type of stimulus used. For example, studies of the movement advantage in familiar faces have generally used degraded face images (e.g. Bruce & Valentine, 1988; Knight & Johnston, 1997; Lander & Bruce, 2000; Lander, Bruce, & Hill, 2001; Lander et al., 1999; Lander & Chuang, 2005; Rosenblum et al., 2007), whereas studies of the movement advantage in unfamiliar faces have used a mix of degraded and non-degraded stimuli (e.g., Pike et al., 1997; Shepherd, Ellis, & Davies, 1982; Shiff et al., 1986; Thornton & Kourtzi, 2002; Hill & Johnston, 2001). This is problematic because the movement advantage may be more apparent when static form-based cues, such as those that would be present in a photograph, are

harder to extract (for example, these cues might include the shape of someone's eyes or lips, skin tone, eye colour, etc.) (Knight & Johnston, 1997). Therefore, in order to resolve the question of whether the movement advantage differs across familiar and unfamiliar faces, it is necessary to equate the stimuli and methods across different levels of familiarity.

Only a few movement-based face recognition studies have directly compared performance with familiar and unfamiliar faces using equivalent methods and stimuli (e.g., Bruce et al., 2001; Lander & Davies, 2007; Roark, O'Toole, Abdi, & Barrett, 2006). However, the results of these studies are inconsistent, and the majority focused on experimentally familiar faces (e.g., Lander & Davies, 2007; Roark et al., 2006), rather than faces that are familiar to participants because of prolonged real-world exposure. Only one study has compared the movement advantage in personally familiar and unfamiliar faces: Bruce et al. (2001, Experiment 1) asked participants to match familiar and unfamiliar people from low quality CCTV stimuli (moving and static) to high quality photographs. They found that familiarity helped participants overall (see also Burton, Wilson, Cowan, & Bruce, 1999), but there was no movement advantage for either familiar or unfamiliar faces. However, as noted above, other work has found that the movement advantage may only be apparent when the stimuli are highly degraded. Knight and Johnston (1997) found that a movement advantage was present for famous faces that were presented upright and negative (reversed contrast), but not when the faces were presented upside down or as normal black and white images. It is possible that the use of less degraded stimuli in the Bruce et al. study resulted in participants relying primarily on static form-based cues, thus minimizing the effect of movement.

The present study addressed the same question as Bruce et al. (2001) – that is, does our level of real-world familiarity with a face influence the presence and magnitude of the movement advantage? To that end, we compared matching performance for static and moving images of famous and unfamiliar faces. However, to minimize the chance that participants would rely on static form-based cues, we presented participants with degraded images of the faces, rather than photographs or CCTV footage. This study also examined the relative use of characteristic

movement patterns and structure-from-motion information in famous and unfamiliar faces, by comparing the movement advantage across stimuli that preserved individual motion patterns, but eliminated the variability in facial structure across the different faces (shape-normalized avatars); and stimuli that preserved individual structure and motion patterns (point-light-displays, or PLDs).

Shape-normalized avatars are created by tracking movement from multiple face areas and projecting it onto one standard facial form – for example, an “average” face (e.g., Hill & Johnston, 2001). Avatars preserve movement information, but provide minimal identifying structural information (since the shape of the face is normalized) – in other words, participants must rely primarily on characteristic movement patterns to match the faces. Consequently, any movement advantage with avatars is likely to arise from recognition or matching of characteristic movement patterns, rather than structure-from-motion. Previous studies have shown that the characteristic face and head movements depicted by avatars are sufficient to perform identity- and sex-matching tasks, even when structural cues are limited (Hill & Johnston, 2001; Knappmeyer et al., 2003).

PLDs are created by tracking multiple points on the face, and then editing the video so only the movement of the points is visible (Bassili, 1979; Johansson, 1973). Because PLDs only show the motion of a limited number of points on the face, they contain less movement information than avatars (which convey movement information across the whole face). However, PLDs preserve some cues to the shape of the head and face, some of which may be available from static frames, some of which may be extracted by structure-from-motion processes. Therefore, any movement advantage from PLDs could be mediated by characteristic movement patterns, structure-from-motion, or both. Like avatars, PLDs have been used in establishing the presence of a movement advantage in face recognition (Bruce & Valentine, 1988; Rosenblum et al., 2007).

Although both avatars and PLDs have been used widely in face processing research, only one study has directly compared performance with avatars and PLDs using natural movements, with sex classification as the dependent variable. Hill, Jinno, and Johnston (2003) used motion capture to record the face and head movements of 12 actors, and then presented the resulting

recordings as PLDs, shape-normalized PLDs, or shape-normalized avatars. Participants were equally good at discriminating the sex of the face whether it was presented as a PLD, a shape-normalized PLD, or shape-normalized avatar. However, sex classification is quite different from matching based on perceived identity, and the two tasks may rely on different types of movement (Hill & Johnston, 2001). Consequently, it is unclear whether participants should perform better when viewing PLDs or avatars in an identification-based task. Logically, PLDs provide more valid cues to identity (structure and motion) than avatars (motion alone), which should result in a greater movement advantage for PLDs than avatars. This is particularly true for familiar faces, which should benefit both from enhanced structural information and characteristic movement patterns. On the other hand, PLDs have been criticized for presenting a sparse array of movement information (Knight & Johnston, 1997), which might only be sufficient to support individual recognition in highly familiar faces (e.g., Rosenblum et al., 2007). Therefore, participants may need the extra movement information carried by avatars in order to extract characteristic movement patterns from a face. If so, any movement advantage would be greater with avatars. By comparing these stimulus types, we aimed to identify the origins of the movement advantage, and determine whether it differs for familiar and unfamiliar faces.

In this paper we present two experiments comparing identity matching performance for famous and unfamiliar faces presented as PLDs or shape-normalized avatars. We conducted two experiments with a same/different matching task: In Experiment 1, participants were presented non-degraded images to be matched to avatars/PLDs. As the non-degraded images provide clear identity information, the processing of three-dimensional form and characteristic movement information has the potential to be influenced by the stored face representations, and we expected strong effects of familiarity. However, non-degraded images may also direct participants' attention to static form-based cues, which may reduce the movement advantage. In Experiment 2, participants were asked to match pairs of avatars or PLDs. This procedure should minimize participants' attention to static form-based cues, leading to a stronger movement advantage than in Experiment 1. However,

participants do not have access to clear identity cues, which could inhibit matching based on stored structural or characteristic movement patterns. Consequently, we expected a weaker effect of familiarity in Experiment 2 than in Experiment 1.

Experiment 1: Matching Between Formats – Video to PLD and Video to Avatars

Participants were presented with a still image or moving clip of a famous or unfamiliar person, followed by a still image or moving clip of a PLD or an avatar. In line with the findings of Bruce et al. (2001), we predicted that familiar faces (in this case, famous faces) would be matched more accurately than unfamiliar ones. Overall, it was expected that participants would perform worst when presented with two static images, and best when presented with two moving clips. For famous faces, any non-degraded image (i.e., both moving and static) should allow people to access stored movement patterns. Therefore, we expected that famous faces would be matched relatively well regardless whether the initial, non-degraded clip was presented in motion or not. On the other hand, we expected that unfamiliar faces (that do not have any stored movement information) would be better matched from moving clips than static images.

As mentioned above, it is unclear whether people should be able to match faces better to PLDs or avatars. If people are better at matching to PLDs, it suggests that a match between the form cues present in the non-degraded and degraded images (e.g. face shape and structure-from-motion information) can facilitate performance in an identity-matching task, above and beyond the use of characteristic movement patterns. On the other hand, if there is no advantage for PLDs, or if avatars are matched better, it suggests that people are concentrating primarily on characteristic movement patterns, rather than form cues, to match the faces.

In addition to structure-from-motion and characteristic movement patterns, it is also possible that a movement advantage could arise from both PLDs and avatars for another reason: moving stimuli, which consist of a series of static frames, carry more visual information than static stimuli. Previous studies have addressed this issue by comparing moving stimuli to displays with multiple static images or movement distortions (Lander et al., 1999; Lander & Bruce, 2000; Rosenblum et

al., 2002; Pike et al., 1997), and all have found an advantage for veridical movement, over and above the contribution of extra static information. As such, we chose to present the simplest test of the movement advantage, by comparing performance for single static frames and 2s moving clips.

Methods

Participants. Thirty-eight undergraduate students (27 female) from the University of Western Sydney, aged between 18 and 59 years (mean age 23.5 years) participated in this experiment in return for course credit. All reported normal or corrected-to-normal vision. Four participants' data were excluded from analysis due to their failure to follow the instructions. One participant's data was excluded due to failure to recognize any of the famous faces. Thus, the final sample consisted of 33 participants.

Stimuli and Materials. A set of video images of six highly familiar (famous) and six unfamiliar adult males was obtained from the online content of two talk shows. The clips were chosen based on a pilot study, in which a separate group of 34 undergraduate students viewed three 8s clips each of 21 male adult faces (63 clips in total). All clips showed the person speaking, facing towards the camera in an interview situation. The clips were selected to show each face from approximately the same viewpoint and distance from the camera, and to exclude extreme facial movements. Participants used a 7-point scale to rate each face on three dimensions: familiarity, distinctiveness of movement, and amount of movement (always in that order, using different clips of the face for each dimension). Participants were also asked to name or otherwise identify any face they indicated was familiar. Distinctiveness was defined as how much a movement would stand out in a crowd, or how individual or idiosyncratic it was to the person. Participants were instructed to try and ignore how familiar the person was, or how distinctive their face shape or features may be, and give a rating based purely on the movement of the face and head. When rating amount of movement, participants were again instructed to try and ignore how familiar the person in the video was to them, and how distinctive their movements were, and concentrate purely on how much the face and head moved.

Based on the ratings of familiarity, the six highest and six lowest rated faces were selected as the “famous” and “unfamiliar” faces for the present study. Famous faces had a mean familiarity rating of 6.46, unfamiliar faces had a mean familiarity rating of 1.43. For each face, a single clip was selected for conversion into PLDs and avatars. Famous and unfamiliar faces had mean distinctiveness ratings of 3.85 and 3.32; and mean amount of movement ratings of 4.42 and 4.47, respectively. Separate two-tailed t-tests were conducted on the distinctiveness and amount of movement ratings for the chosen famous and unfamiliar clips; neither was significant (distinctiveness: $p = .39$; amount: $p = .93$). To create the different identity trials, each clip was paired with another clip from the same familiarity group (i.e., all famous faces were paired with other famous faces). The pairings were chosen based on similar ratings of familiarity, distinctiveness, and amount of movement, and remained consistent across both experiments, to negate any response bias a participant might have displayed towards one particular clip.

Each clip was cut into four separate 2 s videos, which were used as the non-degraded sample clip, and also as the basis for the creation of PLDs and avatars in these experiments. Point-light displays (PLDs) were created by tracking the movement of 27 facial regions using Motion (Apple), then superimposing small grey dots onto each point that mimicked the movement of the underlying region. Finally, the background was set to black. The location of the points was based on the PLDs used by Hill et al. (2003), with additional points added to the cheeks, orbits and temples to make the image more “face-like” (Figure 1). Pupils were not included in the PLDs, as the computer tracking was not accurate enough to follow eye-movements.

Figure 1 about here

Avatars were created using a custom face-tracking program (Saragih, Lucey & Cohn, 2010; Saragih, Lucey & Cohn, 2011). The program tracks 66 points on the face, (including eyes and pupils) and creates an “avatar” which mimics the movements of the original video sequence, but displays a uniform shape and texture for all actors. In these experiments, a white “mask” was chosen for the avatar shape (Figure 2).

Figure 2 about here

The mask was chosen to be a neutral image that did not resemble any of the faces, and was presented on a black background. In total, 144 videos were created: four 2 s videos per stimulus type (non-degraded videos, PLDs and avatars), for each of the 12 identities. There was also a corresponding static image, created by taking a single frame at random from each 2 s PLD, avatar and original video clip. Sample stimuli can be viewed in the supplementary materials for this article. As the frame was taken at random, the static image did not always show a neutral expression (i.e., it preserved any idiosyncratic face and head poses or expressions). All video and static images measured 960 x 540 pixels, and were presented on a black surround. All videos were presented at 25 frames per second. The experiment was run on a MacBook Pro using Superlab 4.0.3, and images were presented on a BENQ E2200 HD 22-inch monitor (1920 x 1080 pixels).

Design and Procedure. Participants completed a same/different identity matching task. The experiment was a fully repeated measures design, 2 (familiarity: famous/unfamiliar) x 2 (stimulus type: PLD/avatar) x 4 (movement of clips: dynamic/dynamic (D/D); dynamic/static (D/S); static/dynamic (S/D); static/static (S/S)). To prevent simple matching of identical movement sequences and to encourage a focus on identity matching, the movement sequences and/or static pictures shown in the two images were extracted from non-overlapping sections of the original

video clip. Each condition contained twenty-four trials (two same and two different trials for each identity), resulting in 384 trials per participant. Trials were presented in two blocks, and different clip pairings were presented in each block, to prevent learning of particular movement sequences. Within each block, the order of presentation of trials was randomized. Block presentation order and stimulus pairings were counterbalanced between participants.

Every trial of the main experiment began with a fixation cross, presented in the centre of the screen for 200 ms. The fixation cross was extinguished and replaced by two facial images, both 2 s long, presented sequentially and separated by a 500 ms grayscale noise mask (see supplementary materials for an example of a trial). To prevent matching based purely on the location of the face onscreen, the initial face image was randomly offset from the centre of the screen by 40 pixels to the left or right, and the subsequent face image was offset 40 pixels in the opposite direction. Participants were asked to indicate via key press whether the two facial images showed the same person, or two different people. Participants were instructed to respond as quickly and as accurately as possible after the second video had finished. Once the key press was recorded, the next trial began immediately. In Experiment 1 the first stimulus shown in each trial was always non-degraded (unmanipulated). The second stimulus was either a PLD or an avatar that was static or dynamic.

Participants were tested individually in a darkened room. They completed 12 practice trials (without feedback) prior to beginning the main experiment, and received several breaks throughout testing. Following the main experiment, each participant was shown one non-degraded video of each of the 12 identities, and asked to rate the video for familiarity on a seven point scale, (1 = unfamiliar; 7 = highly familiar), and to name the person (or provide other unambiguous identity information) if they were famous. Participants' data were excluded from analysis if they did not rate all famous faces 6 or higher and unfamiliar faces 2 or lower, or if they could not name at least three of the six famous faces. This study was approved by the University of Western Sydney Human Ethics Committee and was carried out in accordance with Australian and international ethical guidelines.

Results and Discussion

Signal Detection Theory analysis: A criterion-independent measure of sensitivity (d') was calculated for each participant in each condition (Macmillan & Creelman, 2005). Hit and/or false positive scores of zero or one were replaced with values of 0.042 and 0.948 respectively (these correspond to $1/2N$ and $1-1/2N$ respectively, in accordance with Macmillan & Kaplan, 1985). A d' of 0 represents chance performance, while a d' of 3.46 represents perfect performance in a condition after the above adjustment. One sample t-tests were carried out to compare the resulting d' scores to chance performance levels. Participants performed significantly above chance in 13 out of 16 conditions (all $ps < .05$, see Table 1). All three conditions that were not significantly above chance involved matching unfamiliar faces to avatars – the D/S, S/D and S/S conditions ($ps > .08$).

 Table 1 about here

A 2 x 2 x 4 ANOVA was carried out on the d' scores.¹ The main effect of familiarity was significant: Overall, famous faces were matched significantly better than unfamiliar faces, $F(1,32) = 6.33, p = .017, \eta_p^2 = .16$. There was also a border-line trend for PLDs to be matched better than avatars, $F(1,32) = 3.99, p = .054, \eta_p^2 = .11$. However, both of these main effects were qualified by a significant interaction between familiarity and stimulus type, $F(1,32) = 4.73, p = .037, \eta_p^2 = .13$. Pairwise comparisons (Bonferroni corrected) for each level of familiarity revealed that, for famous faces, there was no significant difference between PLDs and avatars ($p = .845$); but unfamiliar faces were recognized significantly better from PLDs than avatars ($p = .003$). Pairwise comparisons within each stimulus type showed that, when presented with avatars, participants were significantly

¹ Analyses were also carried out to examine the effects of block, but no main effects or interactions were significant ($ps > .2$), therefore block was excluded from any further analysis.

better at matching famous faces than unfamiliar faces ($p = .002$); when presented with PLDs, no significant difference emerged between famous and unfamiliar faces $p = .689$).

The ANOVA revealed no main effect of movement, $F(3,93) = .992$, $p = .40$, $\eta_p^2 = .031$, and no interactions between movement and any other variables ($ps > .15$). However, in the interest of investigating the familiarity effect, Bonferroni-corrected pairwise comparisons for famous vs. unfamiliar faces were carried out in each movement condition and for each stimulus type. There was no significant difference between famous and unfamiliar faces across the majority of conditions ($ps > .14$), with the only significant case being S/S images presented as avatars, where famous faces were matched significantly better than unfamiliar faces, $p = .006$.

The interaction between familiarity and stimulus type suggests that an exact match between form cues is less important for famous faces than unfamiliar faces. Famous faces were matched relatively well from both PLDs and avatars, even when both stimuli were static. One possible explanation for this result is that our representations of famous faces include characteristic pose and expression cues, which participants may have been using to match the faces when the form cues were redundant (i.e., in the avatar condition). This is an extension of an idea proposed by Burton, Jenkins, and Schweinberger (2011), that our mental representation of a familiar person incorporates not only the unchangeable, structural aspects of their face, but the way a face varies across our encounters of it (see also Haxby, Hoffman & Gobbini, 2000). In other words, seeing a non-degraded image of a familiar face allowed our participants to easily access stored representations of the person, leading to relatively good performance in famous face trials, regardless of whether the faces were moving or static, or whether they contained useful structural information or not.

On the other hand, unfamiliar faces were matched relatively poorly from avatars, but quite well from PLDs. Matching of unfamiliar faces appears to have been dominated by shape or structural information, such as the cues present in PLDs. When those cues were present, participants could use them to match unfamiliar faces almost as well as familiar faces. However, when the

structural information was uninformative, as in the avatar stimuli, participants failed to perform above chance level in most conditions.

Movement advantage: To further investigate whether moving stimuli were matched better than static stimuli, a secondary analysis was conducted. The d' value for the S/S movement condition was subtracted from the values of each other condition. The resulting “movement advantage” scores are shown in Figure 3, where zero represents no advantage, and larger values indicate a more substantial performance improvement for moving vs. static stimuli. Movement advantage scores for each condition were subjected to one sample t-tests to compare them to 0 (no movement advantage), followed by a 2 (familiarity) x 2 (stimulus type) x 3 (movement of clips) ANOVA. No conditions were significantly greater than 0 ($ps > .17$), and there were no main effects or interactions ($ps > .05$).

Figure 3 about here

Conclusions from Experiment 1: In sum, Experiment 1 showed a significant familiarity effect, driven by differences in performance for famous and unfamiliar avatars. However, there was no evidence of a movement advantage, regardless of familiarity or the type of stimulus used.

The finding of a familiarity effect, but no movement advantage when matching from a high quality to a degraded image replicates the results of Bruce et al., (2001), and extends on their findings in two ways. First, the fact that the familiarity effect was driven by avatars, not PLDs, suggests that the benefit of familiarity does not arise solely because participants are better at matching shape or structure information for familiar than unfamiliar faces. Second, Bruce et al.’s study only used static high quality images, whereas the current study used both moving and static non-degraded images. Our results show that the lack of movement advantage in Bruce et al.’s study

did not arise because participants could not match movements – even presenting high quality moving images of famous faces is not sufficient to elicit a movement advantage when matching to PLDs or avatars.

There are several possible explanations for the lack of movement advantage. Firstly, it is possible that participants relied on form cues – static face-shape information – throughout the experiment. Participants could have been using primarily shape-based cues for the famous and unfamiliar PLDs, which contain some structural information even in the absence of movement. When presented with avatars, which had no valid distinguishing shape cues, participants may have relied on characteristic static face/head poses to match videos of famous faces with their avatar counterparts (a strategy that is not available for the unfamiliar faces). Note that this does not suggest that participants could not extract movement information, but merely that they paid attention to and based their decisions on the static cues in the stimuli.

Secondly, it is possible that participants were unable to extract enough useful movement information from PLDs and avatars to support a movement advantage. This would be surprising, as movement advantages have been demonstrated with similar avatar and PLD stimuli before (Hill & Johnston, 2001; Rosenblum et al., 2007). Finally, it is possible that the static form-based information in the non-degraded images interfered with the processing of movement information by attracting participants' attention away from the movement information (particularly for the famous faces). That is, people may have been distracted by the fact that some of the faces were famous, or paid attention to the expressions, eye gaze, or myriad other social cues that are present in a non-degraded face image, rather than focusing on characteristic movements or extracting structural cues from the faces. Further, matching movements across different types of stimuli might have been too demanding (see also Liu & Chaudhuri, 1997; 2000). In other words, having different form and texture cues in the non-degraded image and the PLDs/avatars might have confused participants and resulted in poor performance.

To test whether the presence of a familiarity effect without a movement advantage in Experiment 1 was driven by the presence of a non-degraded image, Experiment 2 repeated the same/different matching task, but replaced the non-degraded image with a degraded one of the same type (i.e., either two PLDs or two avatars).

Experiment 2: Matching within formats: PLD to PLD and avatar to avatar.

In Experiment 2, we investigated whether participants could match two degraded images, (i.e., PLD to PLD and avatar to avatar). If the nature of the movement information in the PLD and avatar stimuli caused the lack of movement advantage in Experiment 1, similar results should arise in Experiment 2. However, if the lack of movement advantage was either due to distraction by the non-degraded images or difficulties in matching across different stimulus types, the results for both types of stimulus should significantly improve.

As each trial only contained degraded stimuli (either PLDs or avatars), we expected the effect of familiarity to diminish considerably, because unless the movements were so distinctive and idiosyncratic as to provide a cue to the famous identity, participants should not be able to draw on memory to match characteristic head poses or typical movement sequences. Furthermore, if the superior matching performance for famous faces in Experiment 1 was due to different levels of interest or attention for famous and unfamiliar face trials – for example, participants paying more attention to trials in which a famous face appeared – presenting two degraded images should eliminate the matching advantage for famous faces altogether.

Finally, given the results of Experiment 1 (that PLDs lead to better performance than avatars for unfamiliar faces), we expected similar results in the current study. However, since we also expected the effects of familiarity to diminish significantly, we predicted that PLDs would be matched better than avatars for both famous and unfamiliar faces.

Methods

Participants. Seventeen undergraduate students (12 female) from the University of Western Sydney, aged between 18 and 45 years (mean age 22 years) participated in this experiment in return

for course credit. All reported normal or corrected-to-normal vision. One participant's data was excluded from analysis due to their failure to follow the instructions, leaving a total of 16.

Stimuli, Design and Procedure. The experiment was identical to Experiment 1, except that each trial showed either two PLDs or two avatars, and that the image manipulation types were presented blocked (participants viewed all PLDs in one block and all avatars in another). Block presentation order and stimulus pairings were counterbalanced between participants.

Results and Discussion

Signal Detection Theory analysis: Once again, a d' score was calculated for each participant in each condition, and one sample t-tests were carried out to compare each condition to chance. The d' results for Experiment 2 are shown in Table 2. As in Experiment 1, performance was above chance in the majority of conditions ($p < .05$ in 12 out of 16 conditions, see Table 2). Participants failed to perform significantly above chance level for the S/S condition for PLDs (famous and unfamiliar) and unfamiliar avatars, and for the S/D condition for unfamiliar avatars. This pattern of results is markedly different from Experiment 1, where participants only failed to perform above chance levels for unfamiliar faces with avatar test images. The results from Experiment 2 suggest that removing the non-degraded image encouraged participants to rely more heavily on movement-based cues, rather than matching known famous face images to similar-shaped/posed avatars and PLDs.

Table 2 about here

A 2 x 2 x 4 ANOVA was carried out on the d' scores. In accord with the results of Experiment 1, there were significant effects for familiarity and stimulus type. Overall, famous faces were matched better than unfamiliar faces, $F(1,15) = 17.46, p = .001, \eta_p^2 = .54$, and PLDs were

matched better than avatars, $F(1,15) = 8.56, p = .010, \eta_p^2 = .36$. The interaction between familiarity and stimulus type was not significant, $F(1,15) = .738, p = .40, \eta_p^2 = .05$, but planned comparisons revealed no significant difference between matching performance for famous faces presented as PLDs and avatars, $p = .092$, whereas unfamiliar faces were matched significantly better from PLDs than from avatars, $p = .013$.

This mirrors the effects found in Experiment 1, and indicates that participants still relied more heavily on structural cues to match unfamiliar than famous faces, even when there were no obvious cues that some of the faces were famous (i.e., no non-degraded images for comparison). Pairwise comparisons within each stimulus type also closely followed the pattern of results found in Experiment 1: for avatar stimuli, famous faces were matched better than unfamiliar faces, $p = .009$, but there was no significant effect of familiarity for PLD stimuli, $p = .057$.

Unlike in Experiment 1, the ANOVA showed a significant main effect of movement of clips, $F(1.85, 27.72) = 9.94, p = .001, \eta_p^2 = .40$ (Greenhouse-Geisser correction applied for sphericity violations), and a significant stimulus type by movement of clips interaction, $F(2.48, 37.17) = 4.04, p = .019, \eta_p^2 = .212$, reflecting the fact that, for avatars, there were no significant differences between any of the movement conditions ($ps > .34$), whereas PLDs in the S/S condition were matched worse than in any other movement condition ($ps < .012$). As in Experiment 1, Bonferroni-corrected pairwise comparisons were carried out to assess the difference between famous and unfamiliar faces in all conditions. Once again, famous faces presented as avatars were matched significantly better than unfamiliar faces in the S/S condition ($p = .001$), but the familiarity difference in all other conditions was not significant ($ps > .07$).

One puzzling element of the d' statistics was that participants were very good at discriminating between famous faces presented as avatars when both images were static. This is puzzling because the majority of identifying static cues (texture, structure, features) had been rendered uninformative in the avatar condition. Furthermore, the effect was isolated to famous face avatars: the same effect did not arise for famous faces presented as PLDs, or unfamiliar faces

presented as avatars. It is possible the effect arose because the static famous face avatars were more similar than the static unfamiliar face avatars (e.g., perhaps the famous faces were consistently pictured in a similar pose, whereas the unfamiliar faces were more varied), or because the famous face stimuli in the “different” trials were more dissimilar than the equivalent unfamiliar face stimuli (e.g., perhaps Ben Stiller and Kyle Sandilands consistently turned their faces in different directions, whereas the unfamiliar faces were more homogenous).

To test this effect, an analysis of image dissimilarity was carried out on all the static avatar pairs used in Experiment 2. Image dissimilarity was calculated as the root-mean-square difference in greyscale values across each pair of images, after the greyscale values of the pixels had been normalised (Henson, Mouchlianitis, Matthews, & Kouider, 2008; Vuilleumier, Henson, Driver, & Dolan, 2002). In this analysis, identical images would have a score of 0 and completely dissimilar images would have a score of 1. Since the images were offset, and participants may have considered this when judging the similarity of the image pairs, the image analysis was run at every possible pixel offset, and the minimum score was retained (this represents the closest possible match between the two images).

A 2 (familiarity) x 2 (same or different trial) ANOVA was conducted on the resulting minimum scores. Both main effects were significant, familiarity: $F(1,92) = 18.70, p < .0005, \eta_p^2 = .17$; same/different: $F(1,92) = 41.82, p < .0005, \eta_p^2 = .31$; but the interaction failed to reach significance, familiarity x same/different: $F(1,92) = 2.98, p = .087, \eta_p^2 = .03$. Planned pairwise comparisons on the interaction between familiarity and same/different revealed that dissimilarity statistics for famous and unfamiliar face pairs in “same” trials were not significantly different, $p = .07$, but famous face pairs in the “different” trials were more dissimilar than unfamiliar face pairs, $p = .002$. This may suggest that the matching advantage for static famous face avatars could have arisen in the “different” trials – the mismatched famous faces were more dissimilar than the mismatched unfamiliar faces. However, that pattern of results is not borne out by analyses of the hit and correct rejection scores. If image dissimilarity explained the results for famous faces presented

as static avatars, famous faces should have had a similar hit rate, and a significantly higher CR rate, than unfamiliar faces. Pairwise comparisons (with Bonferroni-correction) show that famous face avatars had significantly more hits ($p = .022$) and CRs ($p = .013$) than unfamiliar face avatars in the S/S condition. Perceptual similarity between two static images may explain why famous faces were correctly rejected more than unfamiliar faces, but it is still unclear why famous faces were correctly matched more often than unfamiliar faces in the static-to-static avatar condition.

In sum, the d' results for Experiment 2 replicated the effect of familiarity found in Experiment 1 – famous faces were matched better than unfamiliar faces overall, and particularly in the avatar condition. The difference between famous and unfamiliar faces was most evident in the S/S avatar condition, where famous faces were matched surprisingly well. However, in Experiment 2, unlike Experiment 1, participants were not aware that some of the faces presented were famous until after the experiment. As they had no way of knowing who the famous people were, the results from Experiment 2 eliminate the possibility that participants were *consciously* looking for particular facial characteristics or poses, or that they were paying more attention to the trials with famous faces in them.

The d' results also confirm that PLDs were discriminated better than avatars overall. This effect arose primarily because unfamiliar faces were matched better from PLDs than from avatars, but it is unclear why. As mentioned in the introduction, PLDs and avatars differ in two ways – PLDs contain unique structural information, whereas avatars do not; and avatars contain a face form (i.e., eyes, skin texture etc), whereas PLDs do not. The PLD advantage may have arisen because participants were using the structural information in PLDs to improve their matching performance, or it may have arisen because participants were distracted by the facial form cues in avatars (i.e., they were trying to match the faces on the basis of the eyes and skin texture, rather than the movement information). To distinguish between these possibilities, we conducted separate $2 \times 2 \times 4$ ANOVAs on the hit and CR rates for Experiment 2. If participants were distracted by the form cues in the avatars, we would expect more false alarms (and hence less CRs) to avatars than to PLDs –

that is, participants should be more likely to incorrectly match two faces simply because they look superficially similar. On the other hand, hit rates should either be unaffected, or be higher for avatars than PLDs. This was not the case: CR rates for PLDs and avatars were not significantly different overall, $F(1,15) = 2.27, p = .153, \eta_p^2 = .13$. When broken down by familiarity, CR rates for PLDs and avatars were not significantly different for unfamiliar faces, $p = .096$, or for famous faces, $p = .357$. On the other hand, PLDs resulted in significantly more hits than avatars overall, $F(1,15) = 5.16, p = .038, \eta_p^2 = .26$, although comparisons for famous and unfamiliar faces individually failed to reach significance ($ps > .05$). These results suggest that participants were using the extra structural information in the PLDs to improve their matching performance, rather than being impeded or misled by the presence of uniform facial form cues.

Movement Advantage: To investigate the movement advantage, and in particular the interaction between stimulus type and movement from the d' analysis, the d' value for the S/S movement condition was subtracted from the values of each other condition and subjected to secondary analyses. The movement advantage score for each condition is shown in Figure 4. The movement advantage failed to reach significance for the three famous avatar conditions ($ps > .4$) and the unfamiliar avatars in the S/D movement condition ($p = .194$); all other conditions showed a significant movement advantage, $ps < .05$. The lack of movement advantage for famous avatars appears to have arisen because of the extremely high performance in the S/S avatar condition for famous faces. As discussed above, it is unclear why participants performed so well in this condition – perhaps they paid attention to characteristic facial expressions or head poses, and this strategy was simply more effective for famous faces than unfamiliar faces. This possibility is discussed further below.

Figure 4 about here

A 2 x 2 x 3 ANOVA showed all main effects were significant. There was a larger movement advantage for unfamiliar faces than for famous faces, $F(1,15) = 5.64$, $p = .031$, $\eta_p^2 = .27$, although both famous and unfamiliar faces showed an overall movement advantage score greater than 0 when collapsed across movement condition and stimulus type, $ps < .05$. PLDs showed a greater movement advantage than avatars, $F(1,15) = 7.80$, $p = .014$, $\eta_p^2 = .34$. When averaged across famous and unfamiliar faces and movement condition, avatars did not show any significant movement advantage ($p = .147$), but PLDs did ($p = .001$), which suggests that the superior matching performance for PLDs in the d' analysis was not simply a consequence of better static cues in the PLDs. Finally, the effect of movement of clip was also significant in this analysis, $F(2,30) = 3.94$, $p = .030$, $\eta_p^2 = .21$. Pairwise comparisons (Bonferroni-corrected) reveal that the D/D condition resulted in a larger movement advantage than the S/D condition ($p = .033$), although all three conditions (D/D, D/S, S/D) led to a significant movement advantage when collapsed across all other variables ($ps < .025$). No interactions were significant, $ps > .1$, but planned pairwise comparisons between stimulus type and familiarity revealed that unfamiliar faces did not show a significant difference in the movement advantage for PLDs and avatars, $p = .324$, whereas famous faces showed a larger movement advantage for PLDs than avatars, $p = .001$. The movement advantage for PLDs was not significantly different for famous and unfamiliar faces, $p = .641$, whereas the movement advantage for avatars was greater for unfamiliar than for famous faces, $p = .008$.

Conclusions from Experiment 2: Overall, Experiment 2 showed a strong, significant movement advantage for famous and unfamiliar faces presented as PLDs, and for unfamiliar faces presented as avatars. It is interesting that unfamiliar faces had a larger overall movement advantage than famous faces, but this was largely driven by the fact that famous face avatars were matched well in the S/S condition, minimizing the movement advantage for famous avatars. When static

matching performance was equivalent, as in the PLD condition, familiarity had no significant effect on the movement advantage.

It is still unclear why participants performed so well when matching two static famous face avatars, but comparatively poorly when the same faces were shown as PLDs, or when the avatars showed unfamiliar faces. It is possible that participants based their decisions in static trials on cues such as the position of the head or idiosyncratic expressions (e.g., a lopsided smile). These cues may be clearer with avatars than PLDs, because of the presence of a facial form. This could also explain why the avatar advantage appeared with famous, but not unfamiliar faces: for the famous faces, participants may have been able to match the face and head poses to stored representations based on past encounters with the face. Characteristic head poses/expressions may be particularly relevant for famous faces – people used to being on camera may be more conscious of their appearance, and strike similar characteristic poses/expressions deliberately (“iconic poses”, Carbon, 2008). Further research with personally familiar faces may establish whether this effect arises purely from familiarity, or as a result of some training or representation style unique to famous faces.

Regardless of why participants performed so well when matching static famous avatars, the finding that unfamiliar faces showed a larger movement advantage than famous faces in the avatar condition is still noteworthy, as it indicates that people may use movement to “top-up” their representation of an unfamiliar face to the same level as their representation of familiar faces. In other words, movement may act as a compensatory mechanism, which can help recognition when it is difficult to access form information (as suggested by Knight & Johnston, 1997; Roark et al., 2003), and also when our mental representation of the face is underdeveloped, as in the case of unfamiliar faces compared to famous faces.

General Discussion

In this study, we compared matching performance for famous and unfamiliar faces presented as point-light-displays and shape-normalized avatars. Overall, famous faces were

matched better than unfamiliar faces, and PLDs were matched better than avatars. When matching from a non-degraded image to a degraded face, movement did not improve matching performance (Experiment 1). However, moving stimuli were recognized significantly better than static stimuli when two degraded face images were presented in each trial (Experiment 2). When it occurred, the movement advantage for unfamiliar faces was a similar size for PLDs and avatars. The movement advantage for famous faces was larger for PLDs than avatars, but interpretation of this result was hampered by unusually high performance for static-to-static comparisons of famous faces in the avatar condition. Interestingly, there was no significant difference between the movement advantage for famous and unfamiliar PLDs. This indicates that, when static matching performance is similar for famous and unfamiliar faces, and both characteristic movement patterns and structure-from-motion cues are available, familiarity has no significant effect on the use of movement information for face matching. Below we will discuss the familiarity results, examining the role of structure and movement cues in famous and unfamiliar faces and the implications of these findings for future research.

Familiarity - overall

We found an overall effect of familiarity in both experiments: In line with previous studies (Bruce et al., 2001; Burton et al., 1999), familiar faces (specifically, famous faces presented as avatars) were matched better than unfamiliar faces in both of the present experiments. Notably, this effect arose even when participants were unaware that the trials contained famous faces (Experiment 2), which suggests two things: first, that the results were not simply a result of people paying more attention to famous faces; and second, in some circumstances (such as a matching task used in this study), a familiarity advantage may not require overt recognition.

According to classic models of face recognition (Bruce & Young, 1986; Burton, Bruce, & Hancock, 1999), familiar faces are recognized more easily than unfamiliar faces because we have a better view-invariant representation of familiar faces – in other words, a better representation of their shape and structure. However, it is unlikely that the advantage for familiar face matching arose

purely because participants were more familiar with the shape or three-dimensional structure of the face. If this were the case, we would have expected famous PLDs, which preserved face shape, to be matched significantly more accurately than avatars, which used a standardized shape. We would also expect famous PLDs to be matched significantly better than unfamiliar PLDs. However, neither experiment found this pattern of results: for famous faces, there were no significant differences between PLDs and avatars; for PLDs, there were no significant differences between famous and unfamiliar faces. This leads to the question: what information were participants using to match the famous faces?

As discussed in Experiment 1, it is possible that familiarity with a face also leads to familiarity with the stereotypical way a person holds their head, smiles, or grimaces. These cues may be sufficient to perform simple matching tasks quite well, even when shape and static form-based cues are rendered uninformative. For example, characteristic poses may activate stored face representations, but at a sub-threshold level – that is, not enough to link the face to personal semantic information, but sufficient that when another face with the same characteristic pose cues appears, participants can match them effectively. This sub-threshold activation may explain how a familiarity effect occurred in the absence of overt naming or recognition of the famous faces. However, as these pose/expression cues are insufficient for overt recognition – probably because they are less reliable than static shape, three-dimensional form, or even characteristic movement information – we would expect more cues would be needed to give rise to a familiarity effect in a naming task (e.g., Knight & Johnston, 1997; Lander et al., 1999; Lander & Bruce, 2000), or a matching task with multiple options (e.g., Bruce & Valentine, 1988; Rosenblum et al., 2007).

Unlike famous faces, unfamiliar faces were matched significantly better from PLDs than avatars in both experiments. This indicates that participants were most likely using shape-based information to match the unfamiliar faces – when shape cues could be matched (as in the PLD conditions for both experiments), participants were relatively accurate; when shape cues could not be used to match the faces because they were conflicting (as in the avatar condition for Experiment

1) or uninformative (as in the avatar condition of Experiment 2), performance decreased dramatically.

Interestingly, when structural information was available (PLDs), accuracy for unfamiliar faces increased to the same level as famous faces. It is unclear why prior exposure to a face did not confer an advantage for PLDs, whereas it did for avatars. One possible explanation is that cues such as structure and characteristic pose information are not strictly additive, but are used in a flexible, ad hoc manner – participants may have reached their highest level of performance when provided with structural cues, and not searched for or based their decision on any other cues (i.e., they may not have had to access their stored representations of familiar faces when they could simply match face shape or structure). When structural cues were rendered uninformative, however, participants may have been forced to fall back on other cues such as characteristic pose information (which would only be available for familiar faces), leading to an advantage for famous faces that was isolated to the avatar condition.

Familiarity - movement

Despite the overall familiarity advantage, Experiment 2 found a larger movement advantage for unfamiliar faces than for famous ones. This is quite unusual: typically, previous studies on famous faces have often found a significant movement advantage (Knight & Johnston, 1997; Lander & Bruce, 2004; Lander et al., 2001), whereas several studies using unfamiliar faces have failed to find any effect of movement (Bruce & Valentine, 1988; Bruce et al., 2001; Christie & Bruce, 1998; Shiff et al., 1986). Based on these results, and the results of previous studies comparing the movement advantage in familiar and unfamiliar faces, we expected a larger movement advantage for famous faces (Roark et al., 2006); or no difference in the movement advantage regardless of familiarity (Lander & Davies, 2007). This raises two questions: why didn't participants show a reliable movement advantage for famous faces; and why did Experiment 2 find a movement advantage for unfamiliar faces where many previous studies have not?

Flexible use of different facial cues, as proposed above, may account for the pattern of familiarity results for the movement advantage. Participants may have only used movement cues when they found the static information insufficient, such as for PLDs, which may not have presented the facial expressions and head poses clearly enough for familiarity-based matches; and for unfamiliar faces, when participants could not rely on characteristic facial expressions. It is unclear why previous studies using famous faces have not found similar results – possibly because some studies used neutral static images of famous faces (e.g., Knight & Johnston, 1997), which may have reduced the amount of distinctive face and head poses in their static conditions, thereby increasing the movement advantage. Furthermore, as mentioned above, characteristic head pose cues may be useful in simple same/different tasks, but insufficient in more difficult naming tasks (e.g., those used by Knight & Johnston, 1997; Lander et al., 1999; Lander & Bruce, 2000; Lander & Chuang, 2005).

The use of a matching task could also explain why this study found a movement advantage for unfamiliar faces. Previous studies of unfamiliar face have generally used old/new recognition tests, which require participants to recall faces after a long delay (in some cases, up to 24 hours; Schiff et al.). It may be relatively easy to extract and compare dynamic cues but quite difficult to store or retrieve characteristic movement information across longer time periods. This is possibly because without a person associated with the movement (i.e., without a “person identity node”, as suggested in the Bruce & Young, 1986, model of face recognition), memories of the unfamiliar face’s characteristic movement patterns interfere with each other. If so, this would explain why using a same/different task, which allows participants to compare movements more directly (as in the current study), is more conducive to finding a movement advantage in unfamiliar faces (see also Thornton & Kourtzi, 2002).

It is interesting to note that the movement advantage for unfamiliar faces was a similar size for PLDs and avatars, even though unfamiliar static PLDs were matched significantly better than avatars. This indicates that participants may have used static shape information as a cue, but the

movement advantage was primarily driven by matching characteristic motion patterns, rather than structure-from-motion. Note that this does not mean that structure-from-motion is not used in unfamiliar face recognition – Pike et al. (1997) showed that participants were more accurate at recognizing unfamiliar faces that were learnt from rigid rotational motion than from multiple static images, and Farivar, Blanke, and Chaudhuri (2009) showed that participants could match unfamiliar faces based on structure-from-motion cues alone. However, in both studies, structure-from-motion was the only available movement cue. In our study, participants did not derive any extra movement advantage when characteristic movement patterns and structure-from-motion information were both present. This could be because the participants focussed their attention on characteristic movements at the expense of three-dimensional information, or it could be because the structure-from-motion information present in normal conversational head movements is insufficient to support matching. Alternatively, it is possible that it takes longer than 2 s to extract structure-from-motion cues: Pike et al. and Farivar et al. both used longer exposure durations than the current study (10 s and unlimited, respectively). While there is ample evidence to suggest that short clips (1-3 s) can carry sufficient characteristic movement information to give rise to a movement advantage (Lander et al., 1999; Lander & Bruce, 2000; Lander, Chuang, & Wickham, 2006; Rosenblum et al., 2002; Rosenblum et al., 2007), no research has investigated the duration of movement necessary to give rise to a structure-from-motion advantage.

The fact that people do not necessarily use structure-from-motion cues when other movement cues are present could also explain why previous studies of unfamiliar faces have found inconsistent results – many used static images at test, which does not allow participants to match or identify faces on the basis of characteristic movement patterns (e.g., Bruce et al., 1999; Bruce et al., 2001; Bruce & Valentine, 1988; Davis & Valentine, 2008).

Unlike the unfamiliar faces, famous faces showed a larger movement advantage for PLDs than for avatars. In general, this would indicate the use of structure-from-motion cues for familiar faces, but it is highly unlikely that the structure-from-motion explanation is applicable in this case.

If structure-from-motion improved matching performance for famous faces, we would expect famous PLDs to be matched better than avatars overall, which was not the case. Performance for famous moving PLDs and avatars was quite similar, and the main difference appeared in the S/S condition. Once again, this is not to say that participants were not using structure-from-motion cues to match the famous PLDs. However, from our results it is unclear whether they were relying primarily on structure-from-motion, characteristic movement patterns, or a combination of both. Future research may explore this question further, by finding stimuli that equate or minimize performance on static trials (e.g., shape-normalized PLDs).

In general, our findings suggest that the benefit of familiarity for moving faces is not as large or as clear-cut as has previously been assumed. Many studies using episodic memory and matching paradigms have found better performance for familiar than unfamiliar faces using static images (Johnston & Edmonds, 2009). Furthermore, several studies on movement-based face recognition have shown a movement advantage for familiar faces (e.g., Knight & Johnston, 1997). However, as with the various facial cues discussed above (structure and characteristic pose), our results suggest that the benefits of familiarity and of movement are not additive – rather, movement may help improve matching performance only up to a certain level. Since static processing of familiar faces is already very good – probably because people have access to more identifiable static information for familiar than unfamiliar faces (e.g., stored structural representations, characteristic poses or expressions) – familiar faces may only derive a small benefit from movement. On the other hand, static processing of unfamiliar faces is known to be less robust – for example, people are less sensitive to spatial changes in unfamiliar than familiar faces (Brooks & Kemp, 2007). As people are less proficient at using static cues in unfamiliar face matching, they may use movement as a compensatory cue to identity, and hence derive more benefit from movement than familiar faces. It would be interesting to examine whether a similar pattern of results is observed in other tasks – for example, a sorting task (Hill & Johnston, 2001) requires participants to compare multiple faces at once, which may make it harder for participants to rely on

characteristic head poses and could result in a movement advantage for both famous and unfamiliar faces.

Matching between formats vs. matching within formats.

An overall examination of both experiments revealed that participants were much better at matching faces in Experiment 2 than Experiment 1. It is possible that presence of a non-degraded face image in Experiment 1 encouraged participants to pay more attention to form cues (e.g., individual features, texture, static shape) and structure, which would generally be the most informative cues to identity. The fact that the form cues were degraded, uninformative, or simply not present in the PLDs and avatars may have made it difficult for participants to complete the matching task. Of course, it also could be the case that the difference between the two experiments arose due to overarching problems generalizing from one type of image to another – a variation of the encoding specificity effect (Tulving & Thomson, 1973; Liu & Chaudhuri, 2000). Previous studies have found that static face learning is better when learning and test images have similar visual properties (e.g., overlapping spatial frequencies, Liu, Collin, Rainville, & Chaudhuri, 2000; both negated images, Liu & Chaudhuri, 1997), and our results suggest that the addition of movement information does not eliminate these effects. Given that participants are quite capable of matching identities from different moving sequences when the stimuli are in the same format (as in Experiment 2), future research on the movement advantage should use matching image formats, or take the encoding specificity principle into account when interpreting results.

Conclusions

Movement can play an important role in face processing, particularly face identification (Roark et al., 2003). Our findings suggest a general matching advantage for famous (compared to unfamiliar) faces, which arises primarily due to good performance with static images and avatars. However, we found no evidence that familiarity with a face leads to a larger movement advantage. Participants showed no movement advantage for either famous or unfamiliar faces when asked to generalize from a non-degraded image to a PLD or avatar, reinforcing previous findings that the

movement advantage is most prominent when face images are degraded (Knight & Johnston, 1997), and indicating that face recognition and matching may suffer when there is a mismatch of visual information at encoding and test (Liu & Chaudhuri, 2000).

In general, our results suggest that participants performed the task by using a variety of facial cues in a flexible, hierarchical manner. Participants appeared to rely primarily on static cues such as face shape and, when available, characteristic head pose and expression. When static cues were unavailable or unreliable, participants showed a movement advantage, driven by characteristic movement patterns (for unfamiliar faces, at least). Based on the current findings, we suggest that structure-from-motion information may only be used in face recognition or matching when there is no other movement information present (e.g., Farivar et al., 2009; Pike et al., 1997). Future studies may wish to investigate the effect of various factors – such as task, length of exposure to the face, and individual variation in face recognition ability – on the use of different facial cues for familiar and unfamiliar face recognition.

References

- Bassilli J N, 1979 "Emotion recognition: The role of facial movement and relative importance of upper and lower areas of the face" *Journal of Personality and Social Psychology* 37 2049-2058.
- Brooks K R, Kemp, R I, 2007 "Sensitivity to feature displacement in familiar and unfamiliar faces: Beyond the internal/external feature distinction" *Perception* 36 1646-1659.
- Bruce V, Henderson Z, Newman C, Burton M A, 2001 "Matching identities of familiar and unfamiliar faces caught on CCTV images" *Journal of Experimental Psychology: Applied* 7 207-218.
- Bruce V, Valentine, T, 1988 "When a nod's as good as a wink: The role of dynamic information in face recognition" *Practical Aspects of Memory: Current Research and Ideas* 1 169-174.
- Bruce V, Young A, 1986 "Understanding face recognition" *British Journal of Psychology* 77 305-327.
- 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, Jenkins R, Schweinberger S R, 2011 "Mental representations of familiar faces" *British Journal of Psychology* 102 93-958.
- Burton A M, Wilson S, Cowan M, Bruce V, 1999 "Face recognition in poor-quality video: Evidence from security surveillance" *Psychological Science* 10 243-248.
- Calvert G A, Campbell R, 2003 "Reading speech from still and moving faces: The neural substrates of visible speech" *Journal of Cognitive Neuroscience* 15 1 57-70.
- Christie F, Bruce V, 1998 "The role of dynamic information in the recognition of unfamiliar faces" *Memory and Cognition* 26 780-790.
- Davis J P, Valentine T, 2009 "CCTV on trial: Matching video images with the defendant in the dock" *Applied Cognitive Psychology* 23 482-505.

- Farivar R, Blanke O, Chaudhuri A, 2009 “Dorsal-ventral integration in the recognition of motion-defined unfamiliar faces” *Journal of Neuroscience* 29 5336-5342 doi: 10.1523/JNEUROSCI.4978-08.2009.
- Haxby J V, Hoffman E A, Gobbini M I, 2000 “The distributed human neural system for face perception” *Trends in Cognitive Sciences* 4(6) 223–233.
- Henson R N A, Mouchlianitis E, Matthews, W J, Kouider, S, 2008 “Electrophysiological correlates of masked face priming” *NeuroImage* 40(2) 884-895 doi: 10.1016/j.neuroimage.2007.12.003.
- Hill H, Jinno Y, Johnston A, 2003 “Comparing solid-body with point-light animations” *Perception* 32 561-566.
- Hill H, Johnston A, 2001 “Categorizing sex and identity from the biological motion of faces” *Current Biology* 11 880-885.
- Jenkins R, Burton A M, 2011 “Stable face representations” *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 1671-1683.
- Johansson G, 1973 “Visual perception of biological motion and a model for its analysis” *Perception and Psychophysics* 14 201-211.
- Johnston R A, Edmonds A J, 2009 “Familiar and unfamiliar face recognition: A review” *Memory* 17 577-596.
- Knappmeyer B, Thornton I M, Bulthoff H H, 2001 “Facial motion can determine identity” *Journal of Vision* 3 337.
- Knight B, Johnston A, 1997 “The role of movement in face recognition” *Visual Cognition* 4 265-273.
- Lander K, Bruce V, 2000 “Recognizing famous faces: Exploring the benefits of facial motion” *Ecological Psychology* 12 259-272.
- Lander K, Bruce V, 2003 “The role of motion in learning new faces” *Visual Cognition* 10(8) 897-912.

- Lander K, Bruce V, 2004 “Repetition priming from moving faces” *Memory and Cognition* 32 640-647.
- Lander K, Bruce V, Hill H, 2001 “Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces” *Applied Cognitive Psychology* 15 101-116.
- Lander K, Christie F, Bruce V, 1999 “The role of movement in the recognition of famous faces” *Memory and Cognition* 27 974-985.
- Lander K, Chuang L, 2005 “Why are moving faces easier to recognize?” *Visual Cognition* 12 429-442.
- Lander K, Chuang L, Wickham L, 2006 “Recognizing face identity from natural and morphed smiles” *Quarterly Journal of Experimental Psychology* 59 801-808.
- Lander K, Davies R, 2007 “Exploring the role of characteristic motion when learning new faces” *The Quarterly Journal of Experimental Psychology* 60(4) 519-526.
- Liu C H, Collin C A, Rainville S J M, Chaudhuri A, 2000 “The effects of spatial frequency overlap on face recognition” *Journal of Experimental Psychology: Human Perception and Performance* 26 956-979.
- Liu C H, Chaudhuri A, 1997 “Face recognition with multi-tone and two-tone photographic negatives” *Perception* 26 1289 – 1296.
- Liu C H, Chaudhuri A, 2000 “Recognition of unfamiliar faces: Three kinds of effects” *Trends in Cognitive Sciences* 4 445–446.
- Macmillan N A, Creelman C D, 2005 *Detection Theory: A User's Guide* 2nd edition (Mahwah, N.J.: Lawrence Erlbaum Associates).
- Macmillan N A, Kaplan H L, 1985 “Detection theory analysis of group data: Estimating sensitivity from average hit and false-alarm rates” *Psychological Bulletin* 98 185-199.
- O’Toole A J, Roark D A, Abdi H, 2002 “Recognizing moving faces: A psychological and neural synthesis” *Trends in Cognitive Sciences* 6 261-266.

- Pike G, Kemp R, Towell N, Phillips K, 1997 “Recognizing moving faces: The relative contribution of motion and perspective view information” *Visual Cognition* 4 409-437.
- Roark D A, Barrett S E, Spence M J, Abdi H, O’Toole A J, 2003 “Psychological and neural perspectives on the role of motion in face recognition” *Behavioral and Cognitive Neuroscience Reviews* 2 15-46.
- Roark D A, O’Toole A J, Abdi H, Barrett S E, 2006 “Learning the moves: The effect of familiarity and facial motion on person recognition across large changes in viewing format” *Perception* 35 761-773.
- Rosenblum L D, Niehus R P, Smith N M, 2007 “Look who’s talking: recognizing friends from visible articulation” *Perception* 36 157-159.
- Rosenblum L D, Yakel D A, Baseer N, Panchal A, Nordarse B C, Niehus R P, 2002 “Visual speech information for face recognition” *Perception and Psychophysics* 64 220-229.
- Saragih J, Lucey S, Cohn, J, 2011 “Deformable Model Fitting by Regularized Landmark Mean-Shifts” *International Journal of Computer Vision* 91 200-215.
- Saragih J, Lucey S, Cohn J, 2011 “Real-Time Avatar Animation from a Single Image” *Proceedings of the 9th IEEE International Conference on Automatic Face and Gesture Recognition*.
- Schiff W, Banka L, de Bordes-Galdi G, 1986 “Recognizing people seen in events via dynamic “mug shots”.” *American Journal of Psychology* 99 219-231.
- Shepherd J W, Ellis H D, Davies G M, 1982 *Identification evidence: A psychological evaluation* (Aberdeen, Scotland: University of Aberdeen Press).
- Thornton I M, Kourtzi Z, 2002 “A matching advantage for dynamic human faces” *Perception* 31 113-132.
- Tulving E, Thomson D M, 1973 “Encoding specificity and retrieval processes in episodic memory” *Psychological Review* 80 352-373.
- Vuilleumier P, Henson R N, Driver J, Dolan R J, 2002 “Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming” *Nature Neuroscience* 5 491–499.

Figure Captions

Figure 1: Positions of the points in the PLD stimuli. Participants were only shown the completed image on the right during the experiments.

Figure 2: Positions of the points tracked in the avatar stimuli. Participants were only shown the completed image on the right during the experiments.

Figure 3: Difference between d' for static-to-static condition and d' for other movement conditions in Experiment 1. Error bars represent ± 1 standard error.

Figure 4: Difference between d' for static-to-static condition and d' for other movement conditions in Experiment 2. Error bars represent ± 1 standard error.

Figure 1

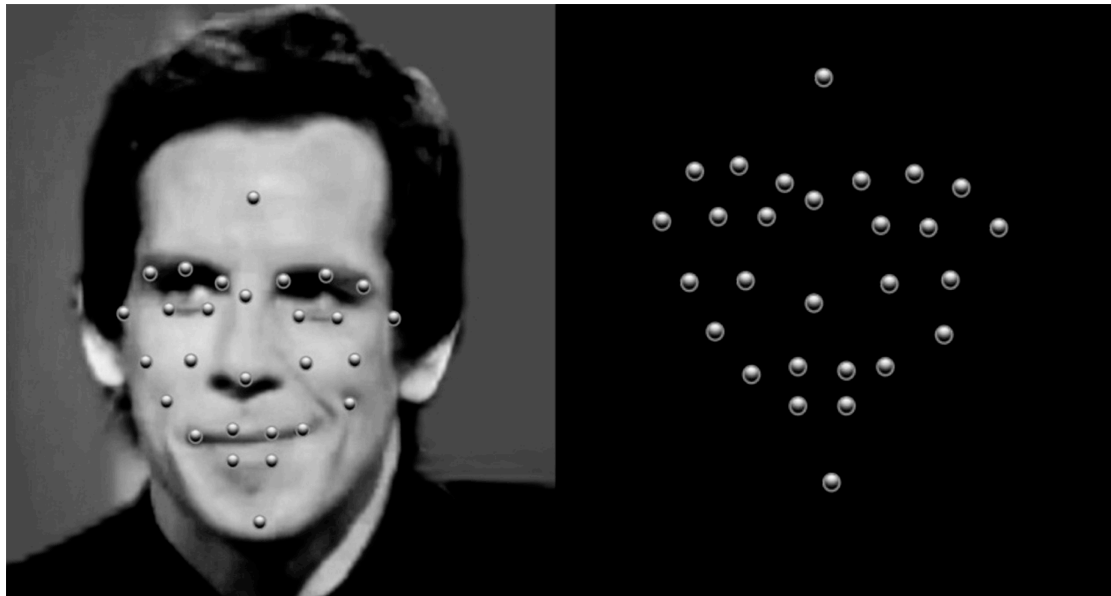


Figure 2

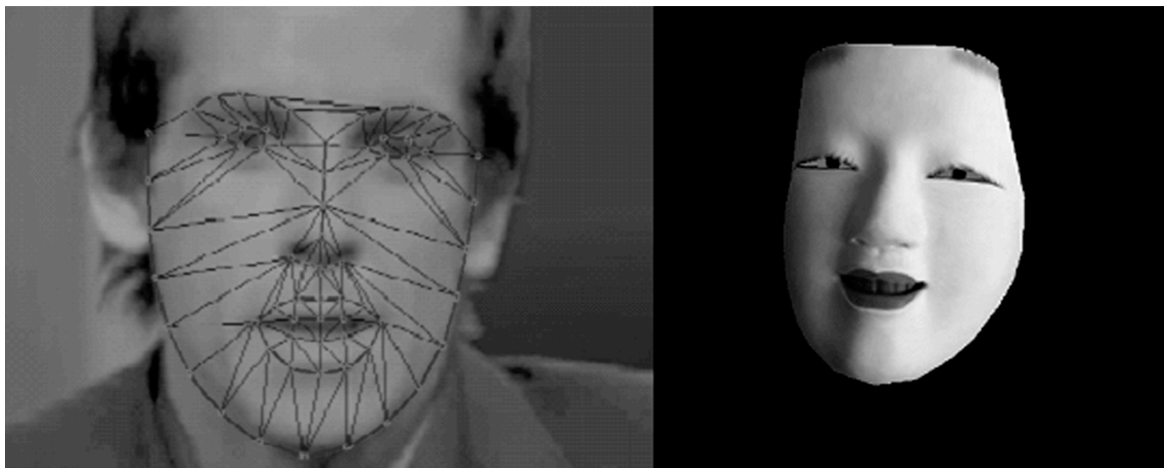


Figure 3

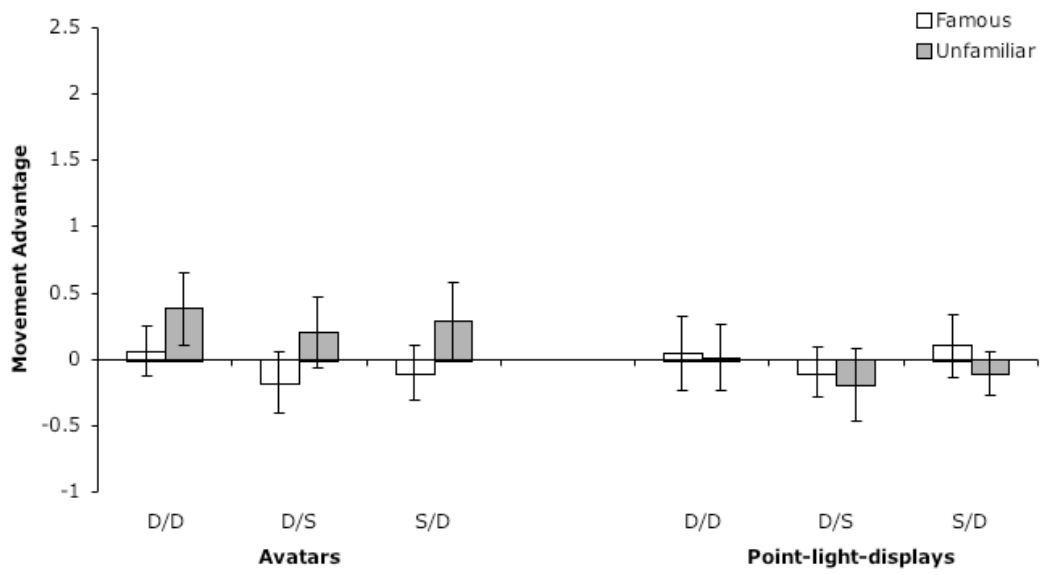


Figure 4

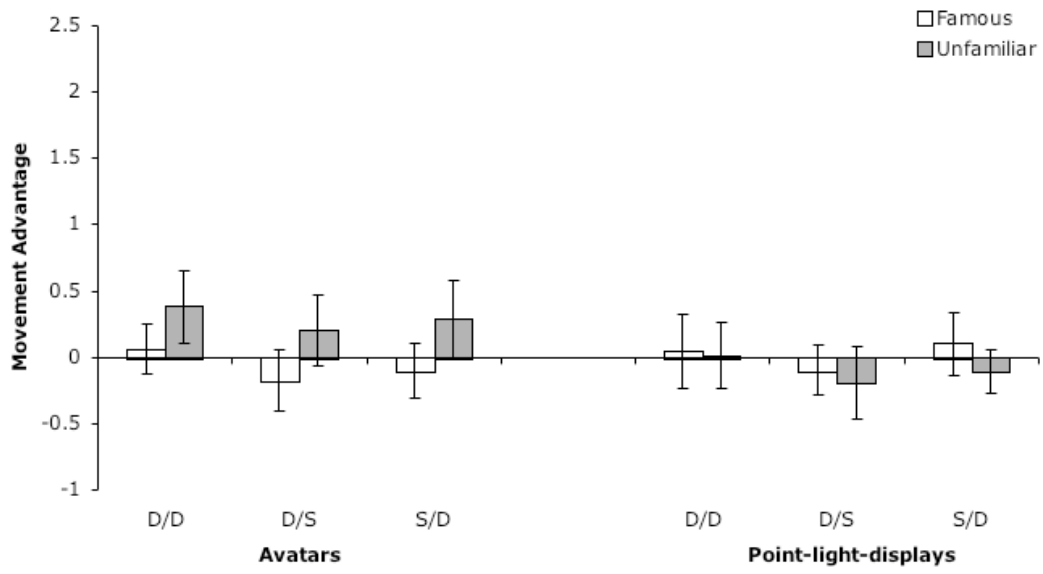


Table 1

d' for Experiment 1

	Avatars				Point-light-displays			
	D/D	D/S	S/D	S/S	D/D	D/S	S/D	S/S
Famous Mean	0.85**	0.62*	0.69*	0.79**	0.80**	0.66*	0.86**	0.75*
SD	1.13	0.92	1.08	1.13	1.13	1.03	0.94	1.23
Unfamil Mean	0.47*	0.29	0.38	0.09	0.80**	0.59*	0.68**	0.78**
SD	1.16	1.23	1.11	1.08	1.10	1.10	0.93	0.94

Note: All *d'* values were compared to chance. * $p < .05$ ** $p < .0005$

Table 2

d' for Experiment 2

	Avatars				Point-light-displays			
	D/D	D/S	S/D	S/S	D/D	D/S	S/D	S/S
Famous Mean	1.24**	1.08**	0.86*	1.16**	2.01**	1.68**	1.67**	0.57
SD	0.72	0.75	1.16	0.92	0.65	0.90	1.08	1.44
Unfamil Mean	1.01**	0.78*	0.48	-0.15	1.72**	1.53**	1.27**	0.11
SD	0.90	0.98	1.15	1.30	0.74	0.57	0.78	1.34

Note: All *d'* values were compared to chance. * $p < .05$ ** $p < .0005$