

In press in *Quarterly Journal of Experimental Psychology*

Accepted 6 Oct 2016

<http://www.tandfonline.com/toc/pqje20/current>

Evidence for view-invariant Face Recognition Units in unfamiliar face learning

David B. Eтчells, Joseph L. Brooks, and Robert A. Johnston

Centre for Cognitive Neuroscience and Cognitive Systems, School of Psychology, University
of Kent, UK

Corresponding author: David B. Eтчells, dbe4@kent.ac.uk

+44 (0)1227 827493

Acknowledgements: DBE was funded by a University of Kent 50th Anniversary PhD
studentship while conducting this work.

Abstract

Many models of face recognition incorporate the idea of a face recognition unit (FRU). This is an abstracted representation formed from each experience of a face. Longmore et al. (2008) devised a face learning experiment to investigate such a construct (i.e., view-invariance) but failed to find evidence of its existence. Three experiments developed Longmore et al.'s study further by using a different learning task, by employing more stimuli. One or two views of previously unfamiliar faces were shown to participants in a serial matching task (learning). Later, participants attempted to recognise both seen and novel views of the learned faces. Experiment one tested participants' recognition of a novel view, a day after learning. Experiment two was identical, but tested participants on the same day as learning. And experiment three repeated experiment one, but tested participants on a novel view that was outside the rotation of those views learned. Results revealed a significant advantage for recognising a novel view when two views had been learned, rather than a single learned view – for all experiments. The effect of view-invariance found when both views were learned is discussed.

Keywords: face recognition unit, face learning, face recognition, pictorial and structural encoding, view invariance.

Although our everyday experience of familiarity with faces of family members, close friends and colleagues is often taken for granted, the question of how this level of familiarity is achieved has been surprisingly difficult to investigate (Burton, 2013). It is clear that familiar and unfamiliar faces are processed in different ways. Changes in view, expression and context all impair unfamiliar but not familiar face recognition (e.g., Johnston & Edmonds, 2009). These differences tend to be explained in terms of the way that familiar and unfamiliar faces are represented in memory (e.g., Megroya & Burton, 2006). Although a face may have been seen before, it may still be unfamiliar and not as easily recognised in novel conditions because of its qualitatively different memory representation from that of a familiar face. Bruce and Young's (1986) influential face recognition model explicitly distinguished between qualitatively different codes that can be accessed when a face is seen: pictorial or structural. The difference between these can be employed to explain the differences between familiar and unfamiliar face recognition.

A 'pictorial' account distinguishes familiar and unfamiliar face representations primarily in terms of their frequency. Each episode or trace reflects only the stimulus properties of the experience and, in the stored representation, does not generalise beyond the properties of these experiences. Under this account, the primary difference between familiar and unfamiliar face representations is that familiar faces simply have far more traces stored - therefore there are no qualitative differences in the nature of familiar and unfamiliar representations. This greater frequency of traces for familiar faces increases the likelihood of a satisfactory level of similarity between a fresh encounter of that familiar face and a 'pictorial' trace already stored. Strictly speaking, with this sort of representation, familiar recognition should decline strongly as the difference between the novel view and previously seen views increases.

An alternative account proposes the use of structural representations. In this case, additional encounters with a face enhance its memory by updating an abstracted structural representation. This has been referred to as a Face Recognition Unit (FRU), and is at the core of several accounts of face processing (e.g., Bruce & Young, 1986; Burton, Bruce & Hancock, 1999; Burton, Bruce & Johnston, 1990). The creation of an FRU occurs through encounters with different views of a face and enables, “the perceiver to distil a powerful representation” (Burton, Jenkins, Hancock & White, 2005, p. 259). This structural representation emphasises important aspects of a face and de-emphasises non-diagnostic aspects, and extrapolates beyond just the information contained in the pictures seen to form a representation of the face that encompasses previously unseen views. This extrapolated representation, to the extent that it contains diagnostic information of the face, can aid the observer in recognising an identity from novel views. Tong and Nakayama (1999) found that, with many exposures, such representations can become so well learned that additional encounters confer no additional benefit on processing. However, in a review of the Bruce and Young (1986) model, Young and Bruce (2011) suggest that FRU representations may in fact change after being created, but that this process is slow compared to their initial period of establishment.

Although one might predict that encountering enormous variation across face views might make face recognition more difficult, Bruce (1994) has discussed how this variation can be an asset in forming a flexible representation that permits more effective processing of novel views. Encountering faces in, for example, different poses, lighting conditions, and across time, could increase the chances that a more robust FRU structural representation is formed, precisely because of the exposure to within person variability (e.g., Burton, 2013; Burton, Jenkins & Schweinberger, 2011).

The concept of an FRU structural representation has also found support from electrophysiological investigations. The ‘N250r’ event related potential (ERP) component ‘identity repetition effect’ occurs in response to repetition of a learned identity in two subsequent pictures. This occurs with pictures from the same view (i.e., a pictorial effect) as well as when a view change occurs. Importantly, the N250r appeared for view change trials only in the second half of the experiment as the faces were becoming familiar. This learned view invariance of the N250r is consistent with the formation of the FRU structural representation which arose from multiple view learning. Furthermore, this suggests that the FRU representation could be indexed electrophysiologically by the N250r ‘identity repetition effect’ (Zimmermann & Eimer, 2013).

Recently however, Longmore, Lui and Young (2008) found evidence inconsistent with the idea of an FRU structural representation being automatically generated through exposures to different views of the same face. In Experiment 3 of their study, participants learned to name a number of previously unfamiliar faces by encountering either front views only, profile views only, or both views. Learning continued until participants could reliably name the faces. Subsequent to (but on the same day as) this learning task the same participants completed an old/new recognition decision on these target faces intermixed with distractor identities. Each identity shown at test was a front view, profile view, or a novel three quarter view. The three quarter view images had not previously been seen by any participants although the target identities had been seen in one or both of the other views. If the learning phase training established a view invariant FRU structural representation from experience of multiple views, then participants should have been better at recognising the novel three quarter views of previously seen identities when they had learned with two views compared to when they had learned just one view. Experience with a single view would not be sufficient for creating a strongly view invariant FRU representation. Within-identity view

variation is critical for this. In contrast to this prediction of an FRU, recognition performance with three quarter views was no better when participants had learned both-views compared to when they had learned only one. Longmore et al. interpreted this as evidence for use of pictorial codes for face recognition of novel views in this task rather than view invariant structural representations.

The experiment conducted by Longmore and colleagues was elegant and their explanation well formulated. However, there are several additional factors that need to be explored before accepting their conclusion. First, their conclusion was based on a null effect. That is, the pattern of results predicted by a structural account of face representation failed to emerge. In addition, the nature of the learning paradigm implemented in their study only permitted the acquisition of very few faces, so the mean recognition scores that were compared (and failed to be distinguished statistically) were formed from a maximum of only four data points from each participant (Longmore et al., 2008, Exp. 3). Such a constraint is typical of face learning experiments where participants can only be expected to learn a limited number of items (e.g., Clutterbuck & Johnston, 2005). Additionally, their learning task involved a name to image association which may not reflect face familiarity learning as proposed by an FRU account (e.g., Bruce & Young, 1986; Burton, Bruce & Hancock, 1999; Burton, Bruce & Johnston, 1990). Indeed, it is possible that their name associative learning procedure may instead have promoted categorical perception processes that may harm unfamiliar face learning (e.g., Kikutani, Roberson & Hanley, 2008, 2010). Finally, Longmore et al. (2008) used images which contained extraneous features such as hair. Subsequent research using the same learning procedure, but with cropped extraneous features, did show evidence more consistent with FRU formation (Longmore, Lui and Young, 2015).

It is also possible that some period of memory consolidation (for a review of consolidation see Dudai, 2004) may be necessary for the FRU structural representation to reach a level of robustness which is able to withstand pictorial changes. The role of sleep in declarative memory formation is widely known (e.g., Ellenbogen, Payne & Stickgold, 2006; Tronson & Taylor, 2007). Research on novel word learning has suggested that consolidation does not necessarily occur immediately after learning, but instead may occur later only after certain conditions are met (e.g., sleep), despite no further encounters with the learned items (Dumay & Gaskell, 2007). In Longmore et al. (2008) the test phase occurred immediately following the learning phase, and it may be that pictorial codes were all that were available at such a short interval following learning.

The present experiments set out to test whether, under a different set of learning conditions, we could find evidence for development of an FRU during unfamiliar face learning. In Experiment 1 we used the same basic design as Experiment 3 in Longmore et al. (2008) which examined recognition hits at test as a function of learning view and test view. However, based on the considerations of previous work above, we established a new learning paradigm that we predicted would increase the chances of developing the FRU structural representation. Compared to Longmore et al. (2008), we included a larger number of 27 identities which led to more (nine) identities for each recognition accuracy estimate. We also cropped their external features (e.g., hair) because internal features are more diagnostic of familiar face processing (Clutterbuck & Johnston, 2002; 2004) and have been shown to aid view-invariant unfamiliar face learning (Longmore, Lui & Young, 2015).

Second, our learning procedure did not involve learning the names of the identities. Instead, we employed a one-back identity matching task in which participants saw a sequence of frontal and profile views of the 27 identities. Each identity was seen 42 times during the learning task and the sequence was structured such that, for all identities, participants

conducted a one-back matching judgement. Across trials this involved comparing each identity to itself as well as to other identities. As in Longmore et al. (2008), one third of learned identities were seen in the right profile view, another third in the full frontal view, and the final third in both views. The one-back matching task with cropped faces made direct comparison of internal face features task relevant. In particular, for identities seen in both views, the task required participants to directly compare features of faces from different views. We expected that this direct comparison between views should aid in development of the FRU structural representation and its associated view invariant recognition because it emphasized processing of within-identity feature variance.

At test, we presented the learned identities, intermixed with unseen distractor identities, in full frontal, profile, or the novel right three quarter view and assessed the percentage of hits (correct recognition) as a function of learned view. We addressed the issue of whether consolidation was necessary for FRU formation by conducting the test phase of the experiment either immediately after the learning phase (Experiment 2) or on the following day (Experiment 1).

To assess whether our new learning paradigm fostered creation of an FRU, the critical comparison in the experiments was how recognition performance on a novel three quarter view at test varied between identities which were learned using single views (frontal or profile) and those learned using both views (frontal and profile). We predicted that when an identity was learned from two different views of a face, recognition accuracy would be greater than when the identity was learned from a single view (Experiments 1 and 2). In other words, participants should recognise the novel view better than would be predicted if they just compared the novel view to the image-based pictorial traces in memory.

Instead, for identities learned from both views, a structural representation which represented diagnostic features of non-experienced views would have been created, and this representation would enhance recognition beyond that available through comparison to the pictorial traces alone (single view learning conditions). However, it is unclear whether the FRU advantage would occur only when the novel view fell within an internal rotation between the learned views (e.g., when the three quarter view falls between profile and frontal views), or whether all novel views benefit even if they fall outside of the range spanned by the learned views. In Experiment 3, we provide a rationale for this question and test how general the FRU structural representation is and provide one test of the range of novel views that it supports.

Experiment 1: Recognition accuracy after overnight consolidation

For this first experiment we wanted to test whether learning a single front view, single right profile, or both of these views, would result in better or worse recognition at test for the same view learned, the other view learned, or critically, a novel right three-quarter test view – after an overnight period of consolidation. Importantly, and relevant to experiment three, the novel test view for this experiment was an internal rotation, i.e., a right three quarter view which was a rotation between the two views seen at learning.

Method

Participants

Twenty-seven Caucasian undergraduates (22 females, 5 males) aged between 18 and 32 years (mean age, 19.7 years) participated in exchange for course credit. All participants had normal or corrected-to-normal visual acuity and no history of neurological illness (self-

report). All participants gave informed consent and the procedures were approved by the University of Kent, School of Psychology Ethics Committee.

Design

The experiment consisted of two phases: a learning phase which comprised a one-back face identity matching task with seven blocks of trials, and a recognition test phase (i.e., old or new identity judgement) that was administered the following day. In the learning phase, each of the 27 identities appeared in one of the three learned view conditions: front facing view only (FF), right-profile view only (RP), or both views (BV). To assign identities to viewing conditions, the 27 identities were randomly split into three groups (A, B, and C) of nine identities. For each participant, these three identity groups were then assigned to the learned view conditions according to a Latin square design (see Table 1). This ensured that, across participants, each identity occurred equally in each learned view and therefore identity could not account for any effects of the learning view factor. In the learning phase, the dependent variable was the percentage of correct one-back matches. This was measured separately for each of the seven blocks of trials. Thus, the overall design for the learning phase was a 3x7 within-subjects design with learned view and block number as factors.

For the recognition test phase, we measured the percentage of hit responses to target identities (i.e., saying 'yes' to faces previously encountered in the learning phase). We presented faces in one of three views: front facing (FF), a novel right three-quarter view (RTQ) or right profile (RP). Identities were assigned to test view conditions based on the counterbalancing laid out in Table 1. This ensured that each identity appeared equally often in each of the test view conditions and ensured that none of the effects of test view could be systematically accounted for by item-effects. Overall, the analysis of recognition test phase data comprised a 3x3 between-subjects design with learned view and tested view as factors. It

is worth pointing out that, despite us approaching this as a between-subjects design, each participant appeared in three cells of this 3x3 design (see Table 1 counterbalancing for detail). Thus, with only 27 participants we were able to achieve 9 participants per cell for the 3x3 design. The cells that each participant appeared in did not map on to a single factor. Therefore, we approached this as a between-subjects design in terms of the factors and analysis.

Materials and Apparatus

Images were presented on a 17-inch LCD monitor. Responses were made using a standard computer keyboard and the experiment was controlled with SuperLab 4 (Cedrus, Phoenix, Arizona, USA). All images were 15° (13.5 cm) vertically and ranged from 6.3° to 13.5° horizontally. The faces of 59 Caucasian men, taken from the Glasgow Unfamiliar Face Database (GUFDB: Burton, White & McNeill, 2010), were cropped to remove background detail and head hair, and all were free of non-face distinguishing features (e.g., tattoos, glasses and jewellery). The database contained two sets of greyscale photographs, representing the same identities taken with different cameras (camera sets 1 and 2) and from various viewpoints. For all identities, six types of image were prepared from each camera set (2 FF, 2 RP and 2 RTQ). The RTQ views were used only in the test phase. Five identities were used in the practice session and were not used again in the learning or test phases. Twenty-seven identities were randomly selected for use in the learning session and shown as images from camera set one. This set was the same for all participants. During the test phase these same identities were shown but with images from camera set two. The remaining 27 identities were not seen in the learning phase and were only encountered as distractors in the recognition test session.

Procedure

During both learning and test phases, participants were seated approximately 50cm from the screen and the face stimulus was presented at the centre of the screen against a white background. Before the learning phase commenced, participants completed a short practice session which had the same format as the learning phase (described below) but with only five identities (not seen in the rest of the experiment) and 19 trials. No feedback was given about accuracy. Upon successfully completing this, participants initiated the first experimental learning phase block with a button press. Participants were not explicitly informed that they would be tested on their memory for the faces they had been exposed to in the matching procedure.

Each of the seven blocks of the learning phase comprised 162 face stimuli. Each face appeared in the centre of the screen for 500ms and was followed by a blank screen for 500ms. This was then followed by a message (black text on a grey rectangle) asking participants whether the last identity they saw was the same as the one before (i.e., a one-back identity matching procedure), and to respond by means of a key-press: 'c' for yes and 'n' for no. Responses were only recorded once the message appeared (i.e., participants had to wait to make a response). No feedback on accuracy was provided.

As described in the design section above, of the 27 identities shown in the learning phase, nine identities were shown in frontal view, nine identities in right profile, and nine in both views, and the assignment of identities to view conditions was counterbalanced across participants (see Table 1). Within each block, each identity appeared six times. Both-view identities were presented as two triplets of the same identity in the different views (FF/RP/FF and RP/FF/RP). Single view identities (FF or RP) were presented as two pairs of trials with the same image (e.g., RP/RP or FF/FF) plus two additional single trials of that image

interspersed amongst the triplets and pairs to form a pseudo-random sequence of trials. The triplets and pairs structure ensured that the sequence would contain sufficient occurrences of match trials which would occur relatively infrequently if we had just randomly ordered the trials, i.e., one-back identity matches would have been less likely to occur if a randomised structure was imposed.

The trials were organized such that BV, FF and RP consecutive matches were alternately presented and separated by mismatches. This also ensured that each identity and each view type was seen equally often. The trial order was different between blocks for a given participant but the same across participants. However, for participants with the different assignment of identities to conditions (see identity counterbalancing in design section above) the exact identities for each trial would have been different but the pattern of responses identical across participants. Thus, again, particular assignment of identities to conditions is not confounded with manipulations of learned view and test view. Overall, each participant saw each identity a total of 42 times over the entire learning phase. Participants took breaks between blocks and proceeded when they were ready.

For each block of trials in the learning phase there were 36 match trials for the both views stimuli (i.e., two per triplet), and 18 match trials (i.e., one per pair) for each single view condition (FF & RP) - that is 36 matches in total across the two single view conditions. Thus, there were 72 match trials (36 BVL + 18FFL + 18 RPL) and 90 non-match trials. This sums to a total of 162 responses per block.

For the test phase, participants returned the following day (a strict 24-hour return was not required). The test phase consisted of 54 randomly presented face images. Twenty-seven were target identities (i.e., identities encountered on the previous day) and the other 27 were distractor identities that had not been encountered before. As mentioned in the design section

above, Table 1 clearly describes the assignment of identities to the 9 conditions resulting from the factorial crossing of learned view and test view factors and makes clear the counterbalancing of identities across these conditions and participants. Distractor identities were also split equally between FFT, RTQT and RPT test views, so that the test list had an equal number of each test view type for both targets and distractors and followed equivalent counterbalancing. With this design, each participants individual hit rate within a cell was calculated from their responses to nine target trials (correct rejections were calculated from nine distractor trials). Keep in mind that each participant contributed to only three cells in the 3x3 design (see Table 1). Thus, the set of 27 target identities presented to each participant, was split equally into those three cells for that participant (see Table 1). Then their hit rate for that cell was calculated from their responses to those nine trials. This design ensured that each target identity was only seen once by each participant during the recognition phase.

The 54 test phase images were presented in the centre of the screen at the same size as the learning images in a different random order for each participant. Images remained on the screen until the participant made a response via the keyboard to indicate whether the face matched an identity which they had seen in the learning phase ('y' for yes and 'n' for no). Participant response times were unlimited, and accuracy was emphasised over speed of response. A two second interval was provided between the participant's response and the next stimulus onset. Participants saw each identity only once and were not provided with any feedback. Upon completion the participant was thanked for their time and provided with a debriefing document.

Results

From the learning phase data, the percentage of correctly identified matches (hits) was analysed with a 3x7 repeated-measures ANOVA with learned view type and block as factors.

The Huynh-Feldt correction for departures from sphericity was used for the main effect of block in the learning analysis. We observed a main effect of view type, $F(2, 52) = 19.38$, $MSE = 185.73$, $p < .001$, $\eta_p^2 = 0.42$, (Observed power = 1), but there was no main effect of block, $F(3.21, 85.51) = 1.64$, $MSE = 380.78$, $p = .183$, $\eta_p^2 = 0.59$, (Observed power = .43). However, the view type x block interaction was significant, $F(12, 312) = 10.27$, $MSE = 49.46$, $p < .001$, $\eta_p^2 = .28$ (Observed power = 1).

We broke down the interaction by examining the simple main effect of learned view at each block. We found that the simple main effect of learned view was significant in blocks 1-5 (all p -values $< .002$; compare squares, triangles, and diamonds within each block, Figure 1). However, by blocks 6 and 7, performance became equivalent across the learned view conditions and there was no simple main effect of learned view in these blocks (block 6, $p = 0.53$; block 7, $p = 0.07$; Figure 1).

Having established that the learning phase produced an equivalent level of performance for each of the three view types by the end of the session, analysis of the test phase was carried out. A hit rate was calculated for each participant and condition by computing the percentage of targets which received a “yes” response within each condition. These values were processed with a 3x3 between-subject’s ANOVA, with learned view (both-views; Front Facing Learned view; Right Profile Learned view) and test view (Front Facing Tested view; Right Three-Quarter Tested view; or Right Profile Tested view) as factors. Analysis of recognition (hit) rates in the test phase revealed the main effect of learned view was significant, $F(2, 72) = 21.49$, $MSE = 448.89$, $p < .001$, $\eta_p^2 = 0.37$, (Observed power = 1), the main effect of test view was also found significant, $F(2, 72) = 3.68$, $MSE = 448.89$, $p = .030$, $\eta_p^2 = 0.09$, (Observed power = .65), and the critical interaction

between learned view and test view was found significant, $F(4, 72) = 14.58$, $MSE = 448.89$, $p < .001$, $\eta_p^2 = .44$] (Observed power = 1).

Further analysis of the significant interaction focused on the critical comparison to test for the FRU effect, that is, an advantage in recognition of the novel three quarter view for identities learned from both views over those learned from single views. To assess this, we conducted a one-way between-subject's ANOVA to test the effect of learned view for the right three quarter test view only (i.e., comparing the three data points in the central column of Figure 2). The result was significant, $F(2, 72) = 5.63$, $p = .005$, $\eta_p^2 = .13$, (Observed power = .84). Pairwise comparisons (adjusted for multiple comparisons) revealed that when both-views had been learned (diamond, centre column, Figure 2), performance was significantly greater on the three quarter view test than when only full frontal view ($p = .006$; square, centre column, Figure 2) or right profile view ($p = .004$; triangle, centre column, Figure 2) were learned. Moreover, there were no significant differences between FF and RP learned views when tested with the right three-quarter view ($p = .900$).

Analysis of correct rejections at test (i.e., saying “no” to an identity which had not been seen at learning), indicated that when a distractor was a FFT test view, over 94% were correctly rejected; 88% for RTQT; and, 73% for RPT. A one-way ANOVA with test view as a factor showed that there was a significant effect of this factor, $F(2, 78) = 17.40$, $MSE = 229.32$, $p < .001$, $\eta_p^2 = 0.30$ (Observed power = 1). Pairwise analysis (adjusted for multiple comparisons) revealed that correct rejections of RPT views were significantly lower than FFT ($p < .001$) and RTQT ($p < .001$).

Discussion

The results of Experiment 1 show that participants were better at recognising a novel view of an identity that they had seen during the learning phase when they had learned

two views (frontal and profile), compared to when they had only learned one view of that face. This pattern of results is consistent with establishment of an FRU (abstracted structural representation) for identities which were learned with both views. An FRU integrates information from the two views and forms a structural representation. This representation includes information or predictions about face structure at intermediate, unseen views. This information can facilitate recognition of novel views. In contrast, a completely pictorial representation arising from single view learning alone does not provide much information about views other than those experienced. Thus, the FRU representation created by two views causes better novel recognition performance than the pictorial representation associated with single view learning.

We did find, unexpectedly, that the correct rejection rates at test differed significantly across the different test view types. In particular, correct rejection rates were lower for right profile test views than for the other views. This means that participants were more likely to say “yes” (i.e., they remembered seeing the identity in the learning phase; less likely to correctly say “no”) to distractors which were a right profile view than for the other views. This difference could reflect a response bias by test view type and could indicate that the hit rates in the right profile test conditions (right column, Figure 2) are inflated. Importantly though, this difference cannot be used to explain the FRU effect which is of primary interest here. The conditions associated with the FRU effect all have the same test view type (i.e., RTQ - centre column, Figure 2).

These results lead to a different conclusion than those of Longmore et al. (2008). They did not find support for an FRU structural representation involved in recognition of unfamiliar faces. Specifically, unlike our results, they did not observe the difference in novel view recognition between identities learned from both views and those learned from single view identities. This could be due to differences between our learning procedure and their

procedure. Our learning procedure was designed to maximise the learning of unfamiliar faces by having participants make identity matching decisions using a one-back matching procedure. This involved directly comparing pictures, in memory, of the same identity (as well as comparing to different identities). For the both views condition, this involved comparing two different views of the same face. This comparison process may have helped participants to focus on the critical differences and similarities between these views of each identity. This process could have fostered establishment of a structural representation of that identity which extrapolates information about unseen views. We also cropped our face stimuli, had no naming task, and conducted our test phase on the following day to allow a consolidation period. The results of Experiment 1 cannot determine which of these differences might explain the FRU effect that emerged in the test phase. Regardless of the exact process of formation of the FRU, it is clear that our procedure produced substantial learning of unfamiliar faces (mean accuracy at block seven of the learning phase for all view types was 82%) and a novel view FRU recognition advantage for identities learned from two views.

Experiment 2: Recognition accuracy without overnight consolidation

In Experiment 1, the learning and test phases were conducted on separate days with an overnight period between them. To test whether this delay and putative consolidation period is necessary for the transference effect (i.e., better performance with both views), Experiment 2 was conducted as a replication of Experiment 1 but with both phases on the same day, with only a short delay between them. If a period of overnight consolidation is required for the creation of an FRU, then we expect that we will not see the view-invariance effect that we observed in Experiment 1. Alternatively, if we do see similar results in Experiment 2 then this will suggest that an FRU can be set up immediately during learning and have an immediate impact on face familiarity.

Method

Participants

Twenty-seven Caucasian undergraduates (20 females, 7 males) aged between 17 and 23 years (mean age, 19.52 years) participated in exchange for course credit. This group was different from those in Experiment 1 but recruited from the same pool. All participants had normal or corrected-to-normal vision (self-report). The procedures were approved by the University of Kent's School of Psychology Ethics Committee.

Design, Materials and Apparatus

The design, materials, and apparatus were exactly the same as in Experiment 1.

Procedure

The procedure repeated that of Experiment 1, except that the learning and test phases were carried out on the same day. Participants completed the learning phase and were then provided with ten basic maths questions, which they were not required to complete fully, and were intended to act only as a filler task while the test phase of the experiment was set-up. This took on average, ten minutes. Participants then completed the test phase which was exactly the same as Experiment 1.

Results

As with Experiment 1, learning phase match accuracy was analysed with a 3x7 repeated-measures ANOVA. View type (BVL: both-views; Front Facing view: FFL or Right Profile view: RPL) and block (B1-B7) were independent variables. The Huynh-Feldt correction for departures from sphericity was used for the main effects of view type and block for the learning analysis.

The learning phase ANOVA revealed a main effect of view type, $F(1.20, 33.99) = 32.97$, $MSE = 428.11$, $p < .001$, $\eta_p^2 = 0.56$ (Observed power = 1), but not block, $F(3.59, 93.58) = 0.943$, $MSE = 266.71$, $p = .43$, $\eta_p^2 = 0.035$ (Observed power = .27). However, the view type x block interaction was significant, $F(12, 312) = 9.96$, $MSE = 60.80$, $p < .001$, $\eta_p^2 = .27$ (Observed power = 1). Again, as with Experiment 1, we broke down the interaction by examining the simple main effect of learned view at each block. We found that the simple main effect of learned view was significant for all blocks (all p-values $< .029$; compare squares, triangles, and diamonds within each block, Figure 3).

For the test phase, a percentage correct score was calculated for each participant based on the number hits (correctly saying “yes” to previously seen identity) achieved in the test phase (for means, see Figure 4 – formatted to match Longmore et al., 2008). These values were processed with a 3x3 between-subject’s ANOVA, with learned view (both-views; Front Facing Learned view; Right Profile Learned view) and test view (Front Facing Tested view; Right Three-Quarter Tested view; or Right Profile Tested view) as factors. Analysis of recognition (hit) rates in the test phase revealed the main effect of learned view was significant, $F(2, 72) = 27.20$, $MSE = 444.70$, $p < .001$, $\eta_p^2 = 0.43$ (Observed power = 1), the main effect of test view was also found significant, $F(2, 72) = 4.99$, $MSE = 444.70$, $p = .009$, $\eta_p^2 = 0.12$ (Observed power = .79), and the critical interaction between learned view and test view was found significant, $F(4, 72) = 32.65$, $MSE = 444.70$, $p < .001$, $\eta_p^2 = .64$ (Observed power = 1).

Further analysis of the significant interaction focused on the critical comparison to test for the FRU effect, that is, an advantage in recognition of the novel three quarter view for identities learned from both views over those learned from single views. To assess this, we conducted a one-way between-subject’s ANOVA to test the effect of learned view for the

right three quarter test view only (i.e., comparing the three data points in the central column of Figure 4). The result was significant, $F(2, 72) = 4.69$, $p = .012$, $\eta_p^2 = .11$ (Observed power = .77). Pairwise comparisons (adjusted for multiple comparisons) revealed that when both-views had been learned (diamond, centre column, Figure 4), performance was significantly greater on the three quarter view test than when only full frontal view ($p = .004$; square, centre column, Figure 4) or right profile view ($p = .038$; triangle, centre column, Figure 4) were learned. Moreover, there were no significant differences between FF and RP learned views when tested with the right three-quarter view ($p = .387$).

Univariate analysis of distractor identity responses showed that the percentage of correct rejections did not differ as a function of view, $F(2, 78) = 0.780$, $MSE = 52.759$, $p = .462$, $\eta_p^2 = 0.02$ (Observed power = .17) - FFT view, 96%; RTQT view, 93%; and, RPT view 95%. It was also noted that mean accuracy in experiment one when both views had been learned and tested on the novel RTQ view was 83% - in experiment two it was 66%. However, although this difference was notable, it was not significant, $F(1, 16) = 3.21$, $MSE = 418.423$, $p = .092$, $\eta_p^2 = .16$ (Observed power = .39).

Discussion

The results of Experiment 2 replicated the critical result of Experiment 1. That is, for identities tested in a novel right three-quarter view, participants were more likely to respond that the identity was familiar when it had been learned with both views compared to when it was learned with a single view alone. This provides further evidence that our learning procedure led to the development of an FRU structural representation. Furthermore, we can conclude that the FRU emerges immediately after the learning phase and does not require a period of overnight consolidation.

Experiment 3: Novel view is an external rotation

Experiments 1 and 2 provided evidence for creation of an FRU structural representation after learning from multiple views. The FRU structural representation presumably includes information about the expected structure at unseen views and this affords a benefit in recognising these novel views. In both experiments, the novel view that we tested was always a right three quarter view. This could be described as an internal (shortest distance) rotation between the two learning phase views (i.e., full frontal and right profile). It is clear from the results that the FRU supports recognition of at least one view along this internal rotation (around the head's vertical axis) between the learned views. We assume that other novel views along this internal rotation would show a similar benefit. This may have arisen because, in the two views learning condition, the learning phase task required participants to directly compare, in memory, the two views of each view identity. Speculatively, this could have been achieved by employing a mental rotation along the shortest path (rotating around the head's vertical axis) between the two views. If that process happened, it could have aided creation of an FRU representing information about these novel views along the internal rotation. Regardless of the exact mechanism of FRU creation from multiple views, our question in Experiment 3 is whether the benefit of the FRU representation is strictly limited to views along this internal rotation between learned views. Alternatively, the information in the FRU representation could also enhance recognition of rotations at a wider range of angles around the axis of rotation but outside of the set of views that fall between the learned views.

This question arises because, although recognition of non-face objects can also show a two view learning advantage (like our FRU effects), this occurs only for internal rotations between the two views. It does not generalise to external rotations (e.g., a left three quarter view would be an external rotation from frontal and right profile views) from the learned

views (Wong and Hayward, 2005; and see Hayward, 2003). These studies used non-symmetric Amoeboid and Geon stimuli, and their lack of symmetry may have played a critical role in the poorer recognition at external rotations from the learned view because distinguishing features of external rotation novel views may not have been visible from the learned views. In contrast, for symmetric objects such as faces (symmetric along the head's vertical axis when viewed from the front), distinguishing features along the internal rotation from frontal view to right profile view is likely to be highly similar (though not completely identical) to that along a symmetric external rotation between the frontal view and the left profile view. Furthermore, there is presumably a strong expectation of symmetry and regular structure of faces which could easily allow generalisation from internal rotations to other views which are expected to be near mirror symmetric based on facial structure expectations.

For example, assuming that faces are vertically symmetric when viewed from the front, the left three quarter view would be expected to have a large mirror symmetric overlap of information with a right three quarter view. Thus, for symmetric face stimuli, there is a strong case to predict that the FRU structural representation is not strictly limited to benefit only internal rotations but, at the very least, also external rotations which are mirror image views of the learned internal rotation. To be clear, the left three quarter view of a face is not strictly a mirror symmetric reflection of the right three quarter view because faces are not perfectly symmetrical. Nonetheless, given that there is substantial overlap in information between the two views, we would expect a flexible FRU representation to generalise its benefit to these very similar external rotations. It is important to note that it is not necessary for this to be the case if the FRU is strictly limited and does and is not mirror reflection invariant. Thus, we conducted Experiment 3 to test this aspect of the FRU representation.

To do this, in Experiment 3 we replaced the right three quarter novel view at test with a novel left three quarter view. This external rotation should have substantial, but not

complete, mirror symmetry with the internal rotation novel view used in Experiments 1 and 2. The database that we used for our face stimuli contained both left and right three quarter views, thus we were able to use true left three quarter views and did not simply mirror reflect the right three quarter view used in Experiments 1 and 2. Based on the expectation of mirror symmetry of faces, we expected the FRU representation effect that we observed previously to appear again in Experiment 3 but now for the left three quarter view when learned with both right profile and full frontal views.

Method

Participants

Twenty-seven Caucasian undergraduates (22 females, 5 males) aged between 18 and 24 years (mean age, 19.22 years) participated in exchange for course credit. This group was different from that in Experiments 1 and 2. All participants had normal or corrected-to-normal vision (self-report). The procedures were approved by the University of Kent's School of Psychology Ethics Committee.

Design, Materials and Apparatus

These were exactly the same as Experiments 1 and 2, with the exception that the critical test view was a left three quarter view (LTQT), selected from the same database as used in Experiments 1 and 2.

Procedure

The procedure was exactly the same as Experiment 1, meaning that the test phase occurred on the following day.

Results

Learning phase one-back matching accuracy was analysed with a 3x7 repeated-measures ANOVA with view type (BV: both-views; Front Facing view: FFL or Right Profile view: RPL) and block (B1-B7) as factors. The Huynh-Feldt correction for departures from sphericity was used for the main effects of view type, block and the interaction for the learning analysis.

We observed a main effect of view type, $F(1.39, 36.18) = 15.17$, $MSE = 205.36$, $p < .001$, $\eta_p^2 = 0.36$ (Observed power = .99), but not block, $F(4.66, 121.37) = 0.71$, $MSE = 63.88$, $p = .602$, $\eta_p^2 = 0.027$ (Observed power = .24). However, the view type x block interaction was significant, $F(9.61, 249.92) = 5.52$, $MSE = 47.73$, $p < .001$, $\eta_p^2 = .17$ (Observed power = 1). Again, as with Experiment 1 and 2, we broke down the interaction by examining the simple main effect of learned view at each block. We found that the simple main effect of learned view was significant for blocks 1 to 4 (all p-values $< .008$; compare squares, triangles, and diamonds within each block, Figure 5). However, by blocks 5 to 7 performance became equivalent across the learned view conditions and there was no simple main effect of learned view in these blocks (all p-values $> .125$; compare squares, triangles, and diamonds within each block, Figure 5).

For the test phase, a percentage correct score was calculated for each participant based on the number hits (correctly saying “yes” to previously seen identity) achieved in the test phase (for means, see Figure 6 – formatted to match Longmore et al., 2008). These values were processed with a 3x3 between-subject’s ANOVA, with learned view (both-views; Front Facing Learned view; Right Profile Learned view) and test view (Front Facing Tested view; Left Three-Quarter Tested view; or Right Profile Tested view) as factors. Analysis of recognition (hit) rates in the test phase revealed the main effect of learned view was

significant, $F(2, 72) = 32.39$, $MSE = 384.09$, $p < .001$, $\eta_p^2 = 0.47$ (Observed power = 1), but the main effect of test view was not found significant, $F(2, 72) = 1.44$, $MSE = 384.09$, $p = .243$, $\eta_p^2 = 0.039$ (Observed power = .30), however, the critical interaction between learned view and test view was found significant, $F(4, 72) = 20.66$, $MSE = 384.09$, $p < .001$, $\eta_p^2 = .53$ (Observed power = 1).

Further analysis of the significant interaction focused on the critical comparison to test for the FRU effect, that is, an advantage in recognition of the novel left three quarter view for identities learned from both views over those learned from single views. To assess this, we conducted a one-way between-subject's ANOVA to test the effect of learned view for the left three quarter test view only (i.e., comparing the three data points in the central column of Figure 6). The result was significant, $F(2, 72) = 14.29$, $p < .001$, $\eta_p^2 = .28$ (Observed power = .99). Pairwise comparisons (adjusted for multiple comparisons) revealed that when both-views had been learned (diamond, centre column, Figure 6), performance was significantly greater on the left three quarter view test than when only full frontal view ($p = .006$; square, centre column, Figure 6) or right profile view ($p < .001$; triangle, centre column, Figure 6) were learned. In addition, and in contrast to the results from Experiments 1 and 2, right three quarter test performance was significantly higher for identities which were learned from full frontal views than those learned from right profile views ($p = .013$).

Univariate analysis of distractor identity responses showed that the percentage of correct rejections approached significance as a function of view, $F(2, 78) = 2.97$, $MSE = 139.05$, $p = .057$, $\eta_p^2 = 0.07$ (Observed power = .56), with pairwise (adjusted for multiple comparisons) analysis revealing that correct rejections of RPT views were significantly lower than FFT ($p = .017$) - FFT view, 95%; RTQT view, 91%; and, RPT view 87%.

Discussion

The results of Experiment 3 clearly demonstrate that the FRU effect that we observed in Experiments 1 and 2 for an internal rotation between the learned views also enhances recognition of at least one external rotation from the learned views. This demonstrates that the FRU advantage is not strictly limited to the exact range of rotations between the two learned views. Although we cannot determine the exact range of external rotations which are affected, one might assume that this is limited to those which could be considered to have mirror image symmetry with the internal rotations between the learned views. However, further work will be necessary to verify this and test the limits of the FRUs generalisation beyond the rotations used in this experiment. Our conclusion is that, in principle, the FRU does not strictly differentiate between internal and external rotations when affecting recognition.

General Discussion

We conducted three experiments to investigate whether we could find evidence congruent with development of Bruce and Young's (1986) notion of a face recognition unit (FRU) after substantial learning of unfamiliar faces using a one-back matching procedure. Across all three experiments, we found that participants had a recognition advantage for novel three quarter views of a face when they had learned them from two views rather than from just a single view in the preceding learning phase. This effect did not depend on a period of consolidation. It appeared immediately after the learning phase (Experiment 2) and persisted for at least one day after learning (Experiment 1). Furthermore, the representation could support recognition of novel views beyond those immediately between the learned views (Experiment 3), possibly due to the mirror reflection similarity between the other views and those learned. This suggests that with just 42 exposures to an unfamiliar identity from

only two different views, a robust and persistent view invariant representation can aid recognition immediately. This is likely very useful given the importance of faces in developing fruitful social relationships and avoiding negative interactions.

Our design largely followed that of Experiment 3 in Longmore, Lui and Young (2008) whose results did not show the FRU effect that we demonstrated here. They interpreted their results as supporting a largely pictorial account of learning unfamiliar faces even when learned from multiple views. Critically, we made some modifications to the learning procedure to further test this in our experiments. These included changing the task from a face naming procedure to a one-back identity matching task as well as cropping extraneous information such as hair from the face images. We also presented more identities which allowed us to have more trials at recognition, and hopefully more stable estimates for each hit rate in the test phase data.

Based on our data alone, we cannot conclusively determine which of these changes might account for the difference between our results and those of Longmore et al., (2008). However, a recent study by Longmore et al. (2015) did find data congruent with our conclusions about the existence of an FRU with a training paradigm similar to their 2008 paper, but by simply cropping their stimuli. This suggests that cropping may be a critical factor and that extraneous features such as hair may distract focus from critical diagnostic internal face features during learning. For example, it has been demonstrated that attention to the internal features are diagnostic of familiar face processing (Clutterbuck & Johnston, 2002; 2004). It could therefore be argued that by promoting attention to the internal features of unfamiliar faces in a learning paradigm will always promote a more direct route to an FRU representation being formed, rather than the lengthy real-world process of building a representation over time and after multiple varied episodes – where presumably implicit pairing down to the diagnostic features of a face takes place.

Furthermore, it is difficult to know how work on cropped faces generalises to face learning in naturalistic contexts. Natural external features such as hair are rarely cropped in naturalistic viewing except perhaps in limited contexts such as, for example, when wearing a hood to protect from rain or cold or religious or cultural attire such as the hijab which covers the hair and neck. Perhaps under most naturalistic circumstances which include extraneous features, development of the FRU representation takes longer than was measured in the 2008 Longmore et al. study. Alternatively, perhaps a different learning paradigm is required under these conditions. Our data, nor those of Longmore et al. (2015) can clarify this. Further work will be required to pin down the precise factors that constrain FRU formation under naturalistic conditions.

We have been able to demonstrate an alternative, purely visual, learning paradigm which generates clear evidence of FRU formation and that this does not depend critically on a procedure involving naming. Rather, FRU representations of unfamiliar faces can be learned without any reference to, or task requirement for, the name of the identity. Furthermore, we have demonstrated that this representation persists until at least the next day after learning. However, it is clear that in daily life, many of the faces that we learn well, and presumably would have FRUs for, are learned associatively with their names and much more biographical and contextual information. One of our motivations for using a purely visual paradigm was that associative name-face learning may encourage a representation focused on matching names to images and negatively impact structural abstraction process involved in forming an FRU. The results of Longmore et al. (2015) indicate that this is not the case because they showed evidence of FRU formation even when using a name-face learning procedure (but now with cropping of external features). Thus, our results can be taken to indicate that name-face learning is not absolutely necessary and that an FRU can be formed purely on the basis of learning based on visual matching. Although we do learn faces often in association with

names, it is not always the case as one often regularly encounters familiar strangers in their workplace, school, etc. without knowing their names.

It is worthwhile to point out that our conclusions may be limited to learning of static face stimuli presented as pictures on a screen. Use of such static pictures is common within the face processing literature. However, we acknowledge that different processes could be involved in learning dynamic faces which occur frequently in naturalistic social settings. A different set of structural codes, incorporating spatial as well as temporal information, may be relevant when viewing animated or moving faces. Nonetheless, in the modern world we often encounter static face pictures in virtual and online environments and our results are most certainly relevant to these settings.

In addition to the above, we also sought to explore whether there are strict limits on how far the FRU representation can generalise in terms of supporting recognition at novel views. From Experiments 1 and 2, it was clear that the FRU conferred recognition on a novel view which was an internal rotation along the shortest path between the frontal and right profile views seen during learning (e.g., right three quarter in this case). This suggests that other views along this path likely would benefit from a similar recognition advantage. However, it is not clear whether the FRU can support rotations beyond this internal rotation path. Previous work on non-symmetrical objects suggests that this is not the case (e.g., Wong and Hayward, 2005; and see Hayward, 2003). We reasoned that, given the symmetric and regular structure of faces, this would allow any FRU benefit found along one set of views to apply to a symmetric rotation of those views (see the full rationale in the introduction to Experiment 3) which we called ‘external rotations’. The results supported this and provided evidence that the FRU is not strictly limited to internal rotations between the learned views and, at the very least, can support recognition of external rotations as well. However, we

cannot be sure how far this generalisation occurs and further work will be required to determine this.

Although our primary focus was on the critical effects of learning view type for the novel three quarter test view (i.e., the FRU effect), it is worth pointing out that the pattern of performance for the rest of our conditions was very similar to that observed in Experiment 3 of the Longmore et al. (2008) paper. This can be observed in Experiments 1 and 2 (and to some extent, Experiment 3). For instance, for cases of learning from single views (dashed and dotted lines; Figures 2, 4, and 6) it is clear that recognition performance declined rapidly as a function of the viewing angle difference between learned and tested view. This suggests a lack of, or weaker, FRU formation (and structural encoding), and could be explained in terms of a pictorial effect. That is, performance depends primarily on whether the test image is visually/pictorially ‘more-similar’ to the learned view. As pictorial similarity between test and learning decreases, so does recognition accuracy at test. One may notice that performance on the views further from the test view are poor. In particular, recognition hits for full face views were only 5-25% (across the three experiments) for identities which were learned from a right profile view. Clearly recognition from single view learning is based on pictorial codes. The scale of the poor performance suggests that profile views, in particular, do not provide useful pictorial information for recognition of the full frontal view. This may be because one of the primary features of the profile is the silhouette outline of the face and this is not a feature clearly present in two dimensional frontal images of the face. Other work will need to examine the pictorial codes for different views and their relative contributions to recognition at other views after single view learning.

Robbins and McKone (2007) and others have suggested that face processing is ‘domain-specific’ while others suggest that face processes arise through accumulated ‘expertise’, sharing mechanisms with processing of non-face objects (e.g., Diamond & Carey,

1986). One open question is whether learning of a view invariant representation of faces, which we call the FRU here, is a domain-specific mechanism for faces. Given that faces have a reliable and expected structure, it is reasonable to speculate that a domain-specific learning mechanism could be efficient. One could even hypothesise an innate visual ability (e.g., reviewed in Nelson, 2001) or ‘face template’. Such a mechanism would be particularly useful if it incorporated weightings for structural codes which are particularly diagnostic of identity. Learning could then be focused on efficiently acquiring and storing this relevant information and avoiding noise introduced by non-diagnostic information. Assuming that information diagnostic of identity differs across different object classes, then it would be efficient to develop a domain-specific mechanism for learning. This is especially true for a functionally and evolutionarily relevant class of stimuli such as faces.

The issue of domain-specificity is a controversial one and not one which we can resolve with our data. Nonetheless, the results of Experiment 3 suggest that the FRU at least takes advantage of the mirror symmetry of faces and allows for invariance along this dimension. This shows some evidence that the FRU mechanism may hold expectations about face structure. However, it is unclear whether this may also apply to symmetric non-face objects as well. Future work will need to delineate a different line of experiments to address this contentious issue.

The concept of a Face Recognition Unit (FRU) proposed by Bruce and Young (1986), still provides conceptual utility in understanding how unfamiliar faces become familiar, and how they might be represented in memory. The current experiments sought to test the idea that FRUs could be identified, whether these representations required consolidation, whether they could persist a day after learning, and if the power of such a representation could generalise beyond the range of views strictly between those learned. On all counts, the results have been positively in favour of a flexible, immediate, and persistent FRU (Bruce &

Young, 1986). Future studies should not only use methods that obtain quantitative empirical evidence of the qualitative differences between unfamiliar and familiar representations; but perhaps more importantly, how accumulated learned visual information, expectation of a 'face template' or 'normative structure' (or similar conceptualisation), interact to produce the powerful effect of familiar face recognition.

Face learning research has been somewhat neglected in the past but is becoming more prominent, and while the current results indicate that the concept of an FRU still provides theoretical utility for the type of representation required for successful familiar face recognition, this will need to be extended and operationalised more fully by research methods that test its robustness. A purely 'pictorial' account of how unfamiliar faces become familiar also needs to be tested further, but it appears from the current results that this type of representation is qualitatively different from that of a FRU.

References

- Bruce, V. (1994). Stability from variation - the case of face recognition. *Quarterly Journal of Experimental Psychology Section A-Human Experimental Psychology*, 47(1), 5-28.
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77, 305-327.
- Burton, A. M. (2013). Why has research in face recognition progressed so slowly? the importance of variability. *The Quarterly Journal of Experimental Psychology*, , 1-19.
doi:10.1080/17470218.2013.800125
- Burton, A. M., Jenkins, R., & Schweinberger, S. R. (2011). Mental representations of familiar faces. *British Journal of Psychology*, 102, 943-958. doi:10.1111/j.2044-8295.2011.02039.x
- Burton, A., Bruce, V., & Hancock, P. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, 23(1), 1-31.
- Burton, A., Bruce, V., & Johnston, R. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology*, 81, 361-380.
- Burton, A., Jenkins, R., Hancock, P., & White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive Psychology*, 51(3), 256-284.
doi:10.1016/j.cogpsych.2005.06.003
- Burton, A. M., White, D., & McNeill, A. (2010). The glasgow face matching test. *Behavior Research Methods*, 42(1), 286-291. doi:10.3758/BRM.42.1.286

- Clutterbuck, R., & Johnston, R. A. (2005). Demonstrating how unfamiliar faces become familiar using a face matching task. *European Journal of Cognitive Psychology, 17*(1), 97-116. doi:10.1080/09541440340000439
- Clutterbuck, R., & Johnston, R. A. (2004). Matching as an index of face familiarity. *Visual Cognition, 11*(7), 857-869.
- Clutterbuck, R., & Johnston, R. (2002). Exploring levels of face familiarity by using an indirect face-matching measure. *Perception, 31*(8), 985-994. doi:10.1068/p3335
- Diamond, R., & Carey, S. (1986). Why faces are and are not special - an effect of expertise. *Journal of Experimental Psychology-General, 115*(2), 107-117. doi:10.1037/0096-3445.115.2.107
- Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram? *Annual Review of Psychology, 55*, 51-86. doi:10.1146/annurev.psych.55.090902.142050
- Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science, 18*(1), 35-39. doi:10.1111/j.1467-9280.2007.01845.x
- Ellenbogen, J. M., Payne, J. D., & Stickgold, R. (2006). The role of sleep in declarative memory consolidation: Passive, permissive, active or none? *Current Opinion in Neurobiology, 16*(6), 716-722. doi:10.1016/j.conb.2006.10.006
- Hayward, W. (2003). After the viewpoint debate: Where next in object recognition? *Trends in Cognitive Sciences, 7*(10), 425-427. doi:10.1016/j.tics.2003.08.004

- Johnston, R. A., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A review. *Memory*, *17*(5), 577-596. doi:10.1080/09658210902976969
- Kikutani, M., Roberson, D., & Hanley, J. R. (2010). Categorical perception for unfamiliar faces: The effect of covert and overt face learning. *Psychological Science*, *21*(6), 865-872. doi:10.1177/0956797610371964
- Kikutani, M., Roberson, D., & Hanley, J. R. (2008). What's in the name? categorical perception for unfamiliar faces can occur through labeling. *Psychonomic Bulletin & Review*, *15*(4), 787-794. doi:10.3758/PBR.15.4.787
- Longmore, C. A., Liu, C. H., & Young, A. W. (2015). The importance of internal facial features in learning new faces. *Quarterly Journal of Experimental Psychology*, *68*(2), 249-260. doi:10.1080/17470218.2014.939666
- Longmore, C. A., Liu, C. H., & Young, A. W. (2008). Learning faces from photographs. *Journal of Experimental Psychology-Human Perception and Performance*, *34*(1), 77-100. doi:10.1037/0096-1523.34.1.77
- Megreya, A. M., & Burton, A. M. (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition*, *34*(4), 865-876. doi:10.3758/BF03193433
- Nelson, C. A. (2001). The development and neural bases of face recognition. *Infant and Child Development*, *10*(1-2), 3-18. doi:10.1002/icd.239
- Robbins, R., & McKone, E. (2007). No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition*, *103*(1), 34-79. doi:10.1016/j.cognition.2006.02.008

Tong, F., & Nakayama, K. (1999). Robust representations for faces: Evidence from visual search. *Journal of Experimental Psychology-Human Perception and Performance*, 25(4), 1016-1035. doi:10.1037//0096-1523.25.4.1016

Tronson, N. C., & Taylor, J. R. (2007). Molecular mechanisms of memory reconsolidation. *Nature Reviews Neuroscience*, 8(4), 262-275. doi:10.1038/nrn2090

Wong, A., & Hayward, W. (2005). Constraints on view combination: Effects of self-occlusion and differences among familiar and novel views. *Journal of Experimental Psychology-Human Perception and Performance*, 31(1), 110-121. doi:10.1037/0096-1523.31.1.110

Young, A. W., & Bruce, V. (2011). Understanding person perception. *British Journal of Psychology*, 102, 959-974. doi:10.1111/j.2044-8295.2011.02045.x

Zimmermann, F. G. S., & Eimer, M. (2013). Face learning and the emergence of view-independent face recognition: An event-related brain potential study. *Neuropsychologia*, 51(7), 1320-1329. doi:10.1016/j.neuropsychologia.2013.03.028




Figure 1. Experiment 1 Learning Phase Results. Mean percent correct one-back matching responses are plotted as a function of learned view (Both-views, Front Facing view & Right-Profile view) at each block of learning. Error bars represent standard error of the mean.




Figure 2. Experiment 1 Test Phase Results. Mean percent recognition hits of previously seen faces as a function of learned view and test view. The results indicate the overall effects of learning both-views or one view on recognition accuracy for three test views. Error bars represent standard error of the mean.




Figure 3. Experiment 2 Learning Phase Results. Mean percent correct one-back matching responses are plotted as a function of learned view (Both-views, Front Facing view & Right-Profile view) for each block of learning. Error bars represent standard error of the mean.




Figure 4. Experiment 2 Test Phase Results. Mean percent recognition hits of previously seen faces as a function of learned view x test view. The results indicate the overall effects of learning both-views or one view on recognition accuracy for three test views. Error bars represent standard error of the mean.




Figure 5. Experiment 3 Learning Phase Results. Mean percent correct one-back matching responses are plotted as a function of learned view (Both-views, Front Facing view & Right-Profile view) at each block of learning. Error bars represent standard error of the mean.

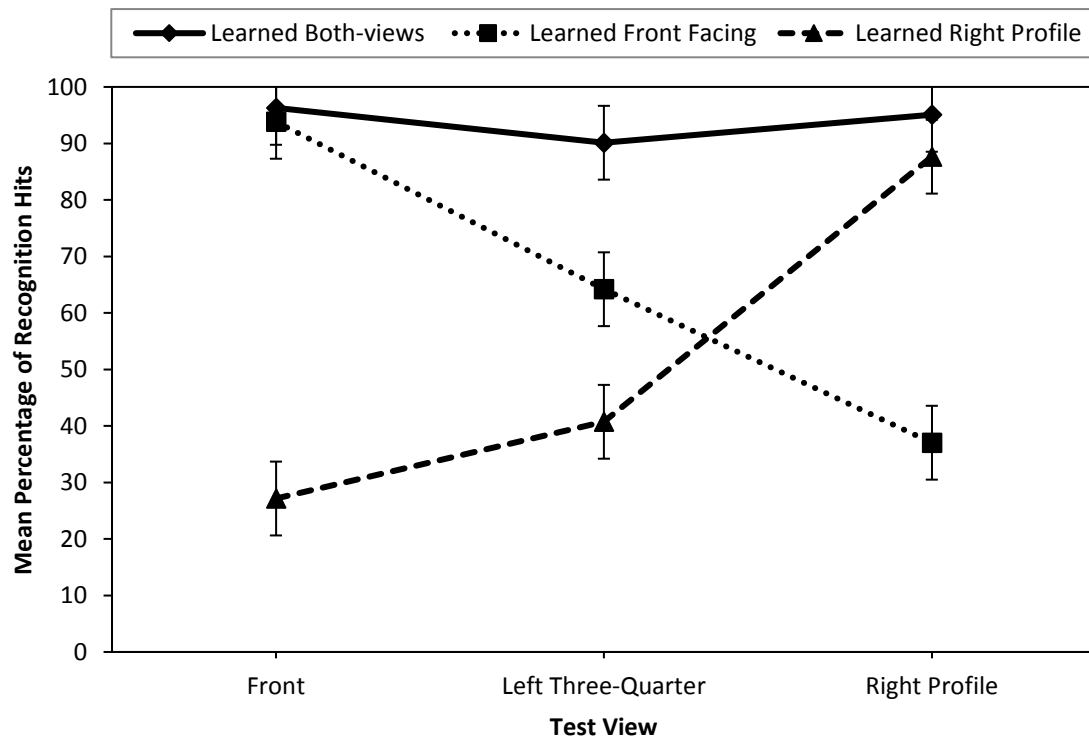


Figure 6. Experiment 3 Test Phase Results. Mean percent recognition hits of previously seen faces as a function of learned view and test view. The results indicate the overall effects of learning both-views or one view on recognition accuracy for three test views. Error bars represent standard error of the mean.

Table 1

Generic Learning Phase and Test Phase matrix - indicating learned view type and test view type, as well as identities used in each phase – applicable to all experiments.

Participant	Learning Group	Learn Both Identity	Learn Front Identity	Learn Profile Identity	Test Three-quarter Identity	Test Front Identity	Test Profile Identity
1-3	A1	1-9	10-18	19-27	1-9	10-18	19-27
4-6	A2	1-9	10-18	19-27	10-18	19-27	1-9
7-9	A3	1-9	10-18	19-27	19-27	1-9	10-18
10-12	B1	19-27	1-9	10-18	1-9	10-18	19-27
13-15	B2	19-27	1-9	10-18	10-18	19-27	1-9
16-18	B3	19-27	1-9	10-18	19-27	1-9	10-18
19-21	C1	10-18	19-27	1-9	1-9	10-18	19-27
22-24	C2	10-18	19-27	1-9	10-18	19-27	1-9
25-27	C3	10-18	19-27	1-9	19-27	1-9	10-18