



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/83702/>

Version: Submitted Version

Article:

White, David, Burton, A. Mike, Jenkins, Rob et al. (2014) Redesigning photo-id to improve unfamiliar face matching performance. *Journal of Experimental Psychology: Applied*. pp. 166-173.

<https://doi.org/10.1037/xap0000009>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Redesigning photo-ID to improve unfamiliar face matching performance

David White¹, A. Mike Burton², Rob Jenkins³ & Richard I. Kemp¹

¹ School of Psychology, The University of New South Wales, Australia

² School of Psychology, University of Aberdeen, UK

³ Department of Psychology, University of York, UK

ACKNOWLEDGEMENTS:

This research was supported by an ARC grant to Kemp (LP110100448), a bilaterally funded grant to Kemp (ARC: LX0083067), Burton and Jenkins (ESRC: RES-000-22-2519) and an ESRC Professorial Fellowship to Burton (ES/J022950/1). We thank Mark Howard for his help collecting data for Experiment 1, and also Graham Nisbett, Filippo Caranti and the Eva Renaldi for making the photographs in Figure 1 available for publication under Creative Commons licenses (CC BY 2.0).

(71 words)

CORRESPONDING AUTHOR:

Dr. David White

School of Psychology

The University of New South Wales

Kensington

Sydney

NSW 2052

david.white@unsw.edu.au

(+61) 2 9385 3254

RUNNING HEAD: Redesigning photo-ID

WORD COUNT (excluding abstract and figure captions): 5160 words

ABSTRACT

1
2
3
4
5 Viewers find it difficult to match photos of unfamiliar faces for identity. Despite this,
6
7 the use of photographic ID is widespread. In this study we ask whether it is possible
8
9 to improve face matching performance by replacing single photographs on ID
10
11 documents with multiple photos or an average image of the bearer. In three
12
13 experiments we compare photo-to-photo matching with photo-to-average matching
14
15 (where the average is formed from multiple photos of the same person) and photo-to-
16
17 array matching (where the array comprises separate photos of the same person). We
18
19 consistently find an accuracy advantage for average images and photo arrays over
20
21 single photos, and show that this improvement is driven by performance in match
22
23 trials. In the final experiment, we find a benefit of four-image arrays relative to
24
25 average images for unfamiliar faces, but not for familiar faces. We propose that
26
27 conventional photo-ID format can be improved upon, and discuss this finding in the
28
29 context of face recognition more generally.
30
31
32
33
34

35
36 (159 words)
37
38
39
40

KEYWORDS

41
42
43
44
45
46 Face Recognition; Unfamiliar Face Matching; Identity Verification; Facial Image
47
48 Comparison; Image Averaging.
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

INTRODUCTION

1
2
3
4 Photo ID documents are frequently used as proof of identity. Despite recent advances
5
6 in biometric technology and storage capacity of identity documents (e.g. passports),
7
8 facial appearance remains the most common means of checking identity at borders. In
9
10 addition, photo-ID is often required in everyday settings, for example when
11
12 purchasing age-restricted goods such as alcohol or tobacco. However, the widespread
13
14 use of photo-ID is at odds with psychological research, which consistently finds that
15
16 viewers perform poorly when matching unfamiliar people to their photos.
17
18
19
20
21
22

23
24 Estimates of human face matching performance vary depending on specifics of the
25
26 task. However even under optimal conditions people are surprisingly inaccurate at
27
28 identity verification from photographs. In an early study, Kemp, Towell and Pike
29
30 (1997) found that supermarket cashiers made over 30% errors when verifying the
31
32 identity of shoppers from Photo-ID cards, despite knowing they were taking part in a
33
34 trial. When an attempt was made to match foils to similar photos (same gender,
35
36 ethnicity, similar age and hairstyle), false acceptance rates rose to over 60%.
37
38
39
40

41 Laboratory-based matching studies have tended to use photo-to-photo matching tasks,
42
43 but also show high error rates (e.g. Bruce et al, 1999, 2001; Clutterbuck & Johnston,
44
45 2002; Megreya & Burton, 2006). However, when these studies are extended to
46
47 include live matches, equivalently poor performance is seen. For example, Megreya
48
49 & Burton (2008) reported an error rate of over 15% in a task requiring viewers to
50
51 match a person to a recent high quality photo, even though no time limit for decisions
52
53 was imposed.
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Importantly, face-matching performance is transformed by familiarity. Across a wide range of identification and matching tasks, viewers are consistently excellent at recognizing familiar faces, even under very poor viewing conditions (Burton et al, 1999; Jenkins, White, Van Monfort & Burton, 2011; Hole, George, Eaves & Rasek, 2002). Indeed, performance on matching tasks has been shown to be a good index of familiarity (Clutterbuck & Johnston, 2002, 2004, 2005). Furthermore, superficial image changes (e.g. viewing angle, expression) severely impair identification of unfamiliar faces, but do not impair identification of familiar faces (e.g., Bruce, Valentine & Baddeley, 1987; Hancock, Bruce & Burton, 2000; Hill & Bruce, 1996).

Poor levels of performance lead one to ask how the problem might be addressed in practical settings. Perhaps it is possible to improve photo-ID by replacing the photograph with something that viewers find easier to match. One alternative might be to use video rather than photos on ID cards, as many cards now contain chips with sufficient storage for this. As it turns out, matching a person to a simultaneously presented high-quality video does not solve the problem (Davis & Valentine, 2009; Experiment 3). In Davis & Valentine's (2009) study, both hits and false alarms were unacceptably high - with error rates of 26% in match, and over 40% in mismatch trials - for video clips that were captured just one week earlier. Here we take a different approach, asking whether aspects of *familiar* face recognition, which is known to be highly accurate, can be built into the unfamiliar matching task.

Burton et al (2005) proposed a model of familiarity-based performance based on averaging together multiple images of the same face (Jenkins & Burton, 2011). According to this model, a stored representation is incrementally refined with each

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

encounter. The effect of adding more images to the average is thus to eliminate superficial differences, while preserving aspects of the images that are common across photos. By this process, the representation comes to emphasize unchanging features of the face that are diagnostic of the particular identity. An average image has been shown to be a useful representation for automatic computer-based face recognition systems, in the sense that matching new photos to an average gives much better performance than matching new photos to an existing photo (Jenkins & Burton, 2008a).

In this paper we ask whether a similar advantage for average images is seen in human performance. Using an unfamiliar face matching task, we tested whether viewers perform better when matching a photo to an average image than when matching two photos. For comparison, we also tested whether matching a photo to a photo array confers any advantage. Importantly, photo arrays preserve information about within-person variability in appearance (see Jenkins et al. 2011), whereas average images emphasize central tendency. Variance information could potentially boost performance by indicating the range of possible images that each face can project.

EXPERIMENT 1

In this experiment we test whether it is easier to match a face photo to another photo or to an average image. We compare performance for familiar and unfamiliar faces by testing participants in two locations (UK and Australia), and presenting images of national celebrities who are famous in only one of these locations. In this way, the

1 same stimuli can be used as both familiar and unfamiliar faces, eliminating any
2 potential confound between stimulus set and familiarity.
3
4
5

6 **Method**

7 *Participants*

8
9
10 Participants were 44 volunteers from The University of New South Wales, Australia
11
12 (28 females, mean age 19.5) and 44 volunteers from University of Glasgow, UK (27
13
14 female, mean age 23.7).
15
16
17

18 *Stimuli and Materials*

19
20
21 We constructed a stimulus set based on 40 UK national celebrities and 40 Australian
22
23 national celebrities. These celebrities were chosen to be known by participants in one
24
25 country, but not the other (for example, national TV presenters, sports personalities,
26
27 politicians). For each of the 80 celebrities, we collected 13 images using Google
28
29 Image search. The images thus sampled natural variability in facial, environmental,
30
31 and image-level parameters (Jenkins et al., 2011). We constrained image selection by
32
33 accepting only those that were of sufficient resolution (minimum 80 pixels between
34
35 the eyes), and where head-angle was no more than twenty degrees from full face. For
36
37 each celebrity, 12 photos were randomly selected to form the average image, and the
38
39 remaining photo was set aside for use as the target photograph in the matching task.
40
41 To construct the average image, we co-registered the twelve photos of each face by
42
43 aligning landmark anatomical features to a standard face template using in-house
44
45 image morphing software. This allowed us to calculate the average RGB values of
46
47 each pixel in a linear space. These ‘shape-free’ average textures were then morphed
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 back to the average shape of the twelve images to produce the final average (for
2 details see Burton et al., 2005).
3
4
5
6

7 For each celebrity, four stimulus-pairs were created: photo-photo and photo-average
8 pairs, in both match (same identity) and mismatch (different identity) combinations.
9

10 As our average face images are automatically cropped to remove extraneous
11 background, we cropped the comparison (i.e. non-target) image in the photo-photo
12 pairings in the same way (Jenkins & Burton, 2008b). All images were presented on a
13 computer monitor at a resolution of 200 by 300 pixels (see Figure 1 for example
14 stimuli).
15
16
17
18
19
20
21
22

23 ----- FIGURE 1 -----
24
25
26
27
28
29

30 *Design and Procedure*

31 Participants completed a 160-trial face-matching test (one match and one mismatch
32 trial per celebrity). Each trial comprised a (target) photo of a celebrity on the left side
33 of the screen and either a second (comparison) photo or an average image on the
34 right. Comparison photos were selected at random from the same set that had been
35 used to create the average images. For match trials, the target photo was of the same
36 celebrity, and for mismatch trials the target photo depicted a different unfamiliar face
37 that matched the same basic verbal description as the target face (e.g. young adult
38 male with dark hair). Participants indicated same identity or different identity
39 judgments via keypress. The task was self paced, and stimuli remained on screen until
40 a response was made. Trial order was randomised throughout.
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 The familiarity manipulation was then checked by showing participants the names of
2 all the celebrities that were presented in the experiment, and asking whether they were
3 familiar with each person's face. As expected, familiarity with home celebrities was
4 high (UK, $M = 34.4$ $SD = 5.5$; Australia, $M = 26.9$, $SD = 10.6$), and familiarity with
5 overseas celebrities was low (UK, $M = 2.0$ $SD = 2.8$; Australia, $M = 1.2$, $SD = 1.6$).
6 Home celebrities that turned out to be unfamiliar, and overseas celebrities that turned
7 out to be familiar, were excluded from analysis for each subject.
8
9
10
11
12
13
14
15
16
17
18

19 **Results**

20 For all experiments in this paper we present accuracy separately for match and
21 mismatch trials. Previous research has shown that face matching accuracy on match
22 trials is not predictive of accuracy on mismatch trials (e.g. Megreya & Burton, 2007).
23 For this reason, we chose not to rely on statistics that combine these measure of
24 performance. However, for the interested reader, we also provide analysis of non-
25 parametric Signal Detection Theory statistics in Supplementary Materials (A' and B'';
26 see Stanislaw & Todorov, 1999).
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41

42 ----- FIGURE 2 -----

43
44
45
46 Accuracy data for Experiment 1 is shown in Figure 2. For match trials, two-way
47 ANOVA with the within-subjects factors of Familiarity (Familiar, Unfamiliar) and
48 Image Type (Photo, Average) revealed significant main effects of both Familiarity, [F
49 (1,87) = 219, $p < 0.01$] and image type, [F (1,87) = 37.5, $p < 0.01$], as well as a
50 significant interaction between these factors [F (1,87) = 4.36, $p < 0.05$]. Simple main
51 effects showed an advantage for average images over photos in both the Unfamiliar
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

condition [$F(1,87) = 27.5, p < 0.01, \text{Cohen's } d = 1.124$], and the Familiar condition [$F(1,87) = 10.7, p < 0.01, d = 0.703$], with the interaction being driven by a larger effect for unfamiliar faces. For mismatch trials there was a significant effect of Familiarity, [$F(1,87) = 18.4, p < 0.01$], but no effect of Image Type and no interaction ($F_s < 1$).

Discussion

Consistent with all previous research (e.g. Clutterbuck & Johnston, 2005; Jenkins et al 2011), matching was more accurate for familiar than unfamiliar faces. More importantly for the current study, we found better face matching performance for average images than for single photos. Rather surprisingly, this was true for familiar faces as well as unfamiliar faces, despite high overall accuracy in the familiar condition. One possible interpretation of this finding is that familiarity was not asymptotic for these national (as opposed to global) celebrities.

Overall, the results demonstrate a performance boost for average images that may be of practical benefit. The averaging technique eliminates some of the transient characteristics of a photograph that profoundly affect appearance, but are irrelevant to identity (e.g. effects of lighting direction). Current technology would allow such digital images to be stored on photo-ID cards, potentially improving identification accuracy by human operators as well as for machines (Jenkins & Burton, 2008a). In the next experiment, we ask whether matching to a photo array might also yield performance benefits.

EXPERIMENT 2

1
2
3
4
5 In this experiment we tested whether photo-ID might be improved if it contained
6
7 more than one photograph of the bearer. Identifying unfamiliar people from
8
9 photographs is difficult, because a person's appearance varies from one snapshot to
10
11 the next (Jenkins et al., 2011). We reasoned that incorporating such variation into a
12
13 photographic representation might make the task easier. Because photo arrays are
14
15 likely to require more elaborate processing than single images, we also included a
16
17 study duration manipulation, to test whether any benefit of photo arrays requires
18
19 extended study time.
20
21
22

Method*Participants*

23
24
25
26
27
28
29
30
31
32
33
34 Seventy-two undergraduates from University of New South Wales participated in the
35
36 study (36 female; mean age 19.7 years, SD = 2.8). None had participated in previous
37
38 experiments in our lab.
39
40

Stimuli

41
42
43
44
45
46 In this experiment we used photographs of 80 people who were unfamiliar to our
47
48 Australian participants (as verified at the end of the experiment). Thirty of these were
49
50 UK celebrities used in the previous experiment, and the remaining fifty were
51
52 consenting undergraduate psychology students who volunteered photos of themselves
53
54 from their Facebook accounts. From this set we selected six photos of each face at
55
56 random for use in the experiment. One of these was chosen at random to be the target
57
58
59
60
61
62
63
64
65

1 photo and the remaining five were used as array photos. For each identity we then
2 selected a similar looking person from existing databases to use as foils in mismatch
3 trials. All images were presented in full colour, cropped to a 2:3 aspect ratio and
4
5 scaled to 200 x 300 pixels.
6
7
8
9

10
11
12 ----- FIGURE 3 -----
13
14

15 16 17 *Design and Procedure*

18
19 Trial Type (match, mismatch) and Array Size (1, 2, 3, 4 photos) were manipulated
20 within-subjects, and Study Time (3 sec, 6 sec, 9 sec) was manipulated between-
21 subjects. Participants were allocated to one of the three Study Time groups at random.
22
23 Participants completed a 160 trial face-matching test (one match and one mismatch
24 trial per identity). Each trial consisted of a target image on the left side of the screen
25 and a photo array on the right. Array photos were selected at random from the five
26 available photos on a trial-by-trial basis, and were presented in a random order in a
27 predefined display configuration (see Figure 3 for an example display). For match
28 trials, the photo array was presented alongside the target photo. For mismatch trials,
29 the array was paired with the foil photograph.
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

46 On each trial, the participants' task was to decide whether the person on the left side
47 of the display was the same as the person on the right. We specifically instructed
48 participants that photos appearing on the right side (i.e. the array) would always show
49 the same person. As in Experiment 1, participants indicated same person or different
50 person decisions via keypress. The task was self-paced, and stimuli remained on
51 screen until response. After each decision, participants rated their confidence on a
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 scale from 1 to 100, so that we could relate objective performance to decisional
 2 confidence. Trials were presented in a random order. Counterbalancing was achieved
 3
 4 by rotating stimulus identities through Array Size conditions across participants, so
 5
 6 that each identity was presented in each condition an equal number of times¹.
 7
 8
 9

10
 11
 12 Participants responded by clicking on onscreen response buttons, and Study Time was
 13
 14 manipulated by delaying presentation of these buttons. Participants were instructed
 15
 16 that the delay should be used to study the faces, and were asked to respond quickly
 17
 18 and accurately once the response buttons appeared.
 19
 20
 21
 22

23 24 **Results**

25 26 27 28 29 ----- FIGURE 4 -----

30
 31
 32
 33 Accuracy data for Experiment 2 are shown in Figure 4. A three-way mixed ANOVA
 34
 35 with the within-subjects factors of Trial Type (match, mismatch) and Array Size (1, 2,
 36
 37 3, 4), and the between-subjects factor of Study Time (3 sec, 6 sec, 9 sec) revealed
 38
 39 significant main effects of Trial Type [$F(1, 69) = 14.4, p < 0.05$] and Array Size [$F(3, 69) = 10.6; p < 0.05$], but no main effect of Study Time [$F(2, 69) = 1.43, p < 0.05$].
 40
 41
 42 The three-way interaction between these factors was not significant ($F < 1$), and
 43
 44 neither were the two-way interactions between Study Time and Trial Type, and
 45
 46 between Study Time and Array Size ($F_s < 1$). Thus Study Time did not affect
 47
 48
 49
 50
 51
 52
 53
 54 performance in this situation.

55
 56
 57 ¹ We did not counterbalance target image through array image positions, because it
 58
 59 was not clear how to achieve this for mismatch trials. However, we note that this
 60
 61 method of counterbalancing would provide a better model for the use of photo-ID in
 62
 63 real world situations, where the appearance of ‘targets’ would vary across encounters.
 64
 65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

There was a significant interaction between Trial Type and Array Size [$F(3,207) = 21.3, p < 0.05$]. Simple Main Effects revealed a significant effect of Array Size for match trials [$F(3,213) = 38.4, p < 0.05$], but not for mismatch trials, [$F(3,213) = 1.75, p > 0.05$]. The effect of Trial Type was non-significant for single-photo arrays ($F < 1, d = 0.063$), but significant for array of two photos [$F(1,71) = 12.6, p < 0.05, d = 0.840$], three photos [$F(1,33) = 29.3, p < 0.05, d = 1.284$], and four photos [$F(1,33) = 27.5, p < 0.05, d = 1.244$].

We also carried out planned comparison t-tests to break down the main effect of Array Size. Because there were no significant main effects or interactions involving study time, we collapsed across this factor before proceeding. Overall accuracy was 79.8% (SD = 8.1) for one-photo arrays, 83.0% (SD = 11.3) for two-photo arrays, 82.6% (SD = 13.2) for three-photo arrays, and 85.4% (SD = 9.2) for four-photo arrays. Planned comparisons revealed a significant difference between one-photo and two-photo arrays [$t(71) = 2.49, p < 0.05, \text{Cohen's } d = 0.325$], but no differences between two-photo and three-photo arrays ($t < 1, d = 0.033$) or between three-photo and four-photo arrays [$t(71) = 1.59, p > 0.05, d = 0.246$].²

Response times and confidence ratings were also collected in this experiment. These measures both corroborated the accuracy measure, showing that participants were more confident in their correct decisions when matching multiple-photo arrays, compared with single-photo arrays. As with accuracy data, this effect was found for match trials only, and saturated at array size two. Response time data confirmed that

² This pattern was also obtained in a separate experiment that excluded the Study Time factor (see Supplementary Materials, page 7).

1 the performance improvement was not due to a speed-accuracy tradeoff, as responses
2 in match trials were faster for multiple photo arrays than for single photo arrays (see
3
4
5 Supplementary Materials for full details of this analysis).
6
7

8 9 **Analysis of Similarity Ratings**

10 Our findings show that face matching performance can be improved by presenting
11
12 multiple comparison photos. We have previously argued that a single photographic
13
14 sample may not contain sufficient data for purposes of identification (Jenkins &
15
16
17 Burton, 2011). Evidently, additional samples go some way to solving that problem.
18
19
20
21
22

23
24 Although it is beyond the scope of this paper to specify the cognitive mechanism of
25
26 the observed performance enhancement, we note that there are at least two broad
27
28 processes that could account for a multiple-photo advantage. One possibility is that
29
30 the identity decision is dominated by the array photo that is most similar to the target
31
32 photo. Alternatively, viewing multiple images may lead the participant to construct a
33
34 more abstract representation of the face against which to match the target.
35
36
37
38
39
40

41 We attempted to distinguish between these accounts by collecting similarity ratings
42
43 for all target and array photos presented in Experiment 2. Our aim was to establish
44
45 whether trial performance was better predicted by the similarity between the target
46
47 photo and the best (most similar) array photo, or by the average similarity between the
48
49 target photo and the array photos. To this end, 28 participants (17 Female; Mean Age
50
51 = 19.3; SD = 2.1) each rated half of 800 comparisons. As it turned out, *best item*
52
53
54 *similarity* and *average similarity* were themselves very highly correlated (pooled
55
56
57
58
59
60
61
62
63
64
65

1 Spearman's $\rho = 0.91$), so it was not possible to distinguish between the two
2 accounts using this method (see Supplementary Materials for methods and analysis).
3

4 **Discussion**

5
6 Our results show that multiple-photo arrays can improve unfamiliar face matching
7 performance. As with Experiment 1, this advantage was observed for match but not
8 mismatch trials, so that the overall improvement was driven by increased accuracy in
9 detecting true matches. Requiring participants to spend more time on their decisions
10 did not improve performance. This might suggest that the critical information can be
11 extracted from multiple photographs rather quickly (i.e. within three seconds).
12
13
14
15
16
17
18
19
20
21
22
23

24 Alternatively, performance may be limited by the cognitive demands of processing
25 information from multiple face images, rather than by the information in the images.
26
27 Previous research has shown that in some circumstances, face identity processing can
28 be subject to strict capacity limits (see Bindemann et al. 2005, 2007). Thus the
29 information advantage of multiple photos may be partly offset by the increased
30 processing demands that they impose. If so, it is possible that a single average image
31 might be preferable to an array of separate photographs. We test this possibility in the
32 final experiment.
33
34
35
36
37
38
39
40
41
42
43
44
45

46 **EXPERIMENT 3**

47
48
49
50 In the final experiment, we directly compared unfamiliar face matching performance
51 for two different types of face representation - average images and photo arrays. In
52 the previous experiment we found that two-photo arrays improved performance over
53 single comparison photographs, and that increasing array size further yielded no
54
55
56
57
58
59
60
61
62
63
64
65

1 additional (statistically significant) benefit. However, since overall accuracy was
2 numerically highest for four-photo arrays (85.4%), we used four photos for the arrays
3 in this experiment.
4
5
6
7

8 9 **Method**

10 Participants were 28 volunteers from University of New South Wales, Australia (13
11 females; mean age = 20.7) and 28 volunteers from University of Glasgow, UK (17
12 females; mean age = 24.2).
13
14
15
16
17
18
19

20 The procedure was the same as Experiment 1, except that we replaced the photo
21 condition with a four-photo array condition. Participants completed a 160-trial face-
22 matching test (one match and one mismatch trial per celebrity). Each trial comprised a
23 photo of a celebrity on the left side of the display and either an average image or a
24 photo array on the right. As in Experiment 1, average images were constructed from
25 12 photos of the person. Four-photo arrays were generated on a trial-by-trial basis, by
26 selecting four of these twelve photographs at random.
27
28
29
30
31
32
33
34
35
36
37
38
39

40 For match trials, the target and comparison images showed the same person. For
41 mismatch trials, the comparison image was of a different unfamiliar face that matched
42 the same basic description as the target. Participants were asked to indicate whether
43 the face on the left (target) was the same as the face on the right (average image or
44 photo array). As in the previous experiment, it was made clear to participants that the
45 four photos in any array always showed the same person.
46
47
48
49
50
51
52
53
54
55
56
57

58 **Results**

59
60
61
62
63
64
65

----- FIGURE 5 -----

1
2
3
4
5 Accuracy data from Experiment 3 is shown in Figure 5. The overall pattern is similar
6
7 to that observed in Experiment 1, except that here we found an advantage for photo-
8
9 arrays over average images. For match trials, a two-way within subject ANOVA with
10
11 factors Familiarity (Familiar, Unfamiliar) and Image Type (Average Image, Photo
12
13 Array) revealed significant main effects of both Familiarity [$F(1,55) = 106, p < 0.01$]
14
15 and Image Type, [$F(1,55) = 6.97, p < 0.01$], as well as a significant interaction
16
17 between these two factors [$F(1,55) = 5.91, p < 0.05$]. Simple Main Effects confirmed
18
19 that the performance benefit for photo arrays was significant for unfamiliar faces [F
20
21 ($1,55) = 9.48, p < 0.01, \text{Cohen's } d = 0.415$] but not for familiar faces ($F < 1, d < 0.01$).
22
23
24
25
26
27
28

29 For mismatch trials there was a significant main effect of Familiarity only [$F(1,55) =$
30
31 $5.61, p < 0.05$], with no significant effect of Image Type ($F < 1$) and no interaction (F
32
33 < 1). Thus, matching performance using photo arrays exceeded performance using
34
35 average images, but this benefit was specific to same-person trials and unfamiliar
36
37 faces.
38
39
40
41
42
43

GENERAL DISCUSSION

44
45
46
47
48 In all three experiments, we found that alternatives to single-photograph
49
50 representations of faces can improve face matching performance. In Experiment 1,
51
52 matching a photograph to an average image was more accurate than matching two
53
54 photographs. In Experiment 2, matching a photograph to a multi-photo array was
55
56 better than matching two photographs. Finally, in Experiment 3, matching a
57
58
59
60
61
62
63
64
65

1 photograph to a multi-photo array was better than matching a photograph to an
2 average image. Our findings have important implications for face matching in
3 occupational settings. Foremost, they demonstrate that single-photo representations of
4 faces are suboptimal, and could be superseded by representations that incorporate
5 within-person variability. Either stabilizing the variability (by image averaging) or
6 increasing the number of samples (by presenting multiple photos) improves matters.
7
8
9

10 The specific pattern of improvement for average images and photo arrays was also
11 consistent across experiments. In each experiment, improvement was observed only
12 in trials where target and comparison images showed the same person (match trials).
13
14 Apparently, providing more information about a person's appearance allows a viewer
15 more accurately to identify that person in a true match. Importantly, this benefit was
16 not accompanied by a general response bias to make 'same person' responses,
17 because accuracy on mismatch trials was always unaffected by extra visual
18 information. This is an important point, as it shows that the performance benefit
19 observed in match trials does not come at the cost of a performance decrement in
20 mismatch trials. Instead, we find a net gain in accuracy. In particular, it appears that
21 presenting multiple photographs of a face allows participants to be more
22 accommodating of within-person variance in appearance. Future research may
23 discover complementary methods for improving mismatch performance without
24 impairing match performance, as required for detection of identity fraud.
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52

53 It should be noted that this pattern of results is somewhat discrepant with our original
54 motivation. We sought to map aspects of familiar face processing onto unfamiliar face
55 processing to improve performance in the latter domain. Previous research
56
57
58
59
60
61
62
63
64
65

1 demonstrating familiarity-based improvement on matching tasks has found that
2 familiarity improves performance on both match *and* mismatch trials (e.g. Megreya &
3 Burton, 2006, 2007; Clutterbuck and Johnston, 2002, 2004, 2005). Although we also
4 found enhanced performance for familiar faces, the effect here was more pronounced
5 for match trials than for mismatch trials, perhaps due to the broader heterogeneity of
6 our stimulus images. Nevertheless, we observed the advantage for averages and
7 photo-arrays *only* in match trials, suggesting that these formats confer partial benefits
8 of familiarity. Previous studies typically report small effects of image-based
9 familiarization procedures on matching performance (e.g. Clutterbuck & Johnston,
10 2005; Osborne & Stevenage, 2008), or find that it does not improve accuracy at all
11 (Clutterbuck & Johnston, 2005). One direction for future research might be to develop
12 methods that accelerate the process of familiarization, and enhance the improvement
13 in matching performance seen here.
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30

31
32
33 Another fruitful direction for future studies would be to manipulate within-person
34 image homogeneity as a variable in its own right. Doing so should help to establish
35 whether it is the similarity of the closest matching photograph, or the similarity of the
36 entire array that drives improved performance in photo array conditions. In previous
37 work (Burton et al, 2005; Jenkins & Burton, 2011), we have proposed that familiar
38 face recognition is highly accurate *precisely because* it is based on representations
39 that summarise within-person variability in appearance. The resulting representations
40 are robust, in the sense that they can be matched to novel images of the same person,
41 provided that these vary in ways that are consistent with previous perceptual
42 experience. The idea behind the image formats tested here is to build variability into
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 the representation. If unfamiliar viewers are exposed to some variability in the target
2 person, they may benefit from some of the advantage of face familiarity.
3
4

5
6
7 Finally, we should note that all the experiments here use photo-to-image matching. In
8
9
10 real photo-ID settings, people are usually asked to make a match to a live person. In
11
12 fact, the relatively small literature comparing photo-to-image and live-to-live
13
14 matching has found surprisingly little difference in performance between the two
15
16 (Kemp et al, 1997; Davis & Valentine, 2009; Megreya & Burton, 2008). For this
17
18 reason, we expect that the performance benefits seen here would generalise to live
19
20 face-matching settings. Nevertheless, it is important to test this, and such experiments
21
22 will form the basis of future work. In that work it will be important to establish not
23
24 only whether the basic improvements in photo-ID format are observed, but also how
25
26 they might interact with characteristics of the observer (e.g. face recognition aptitude:
27
28 White, Kemp, Jenkins & Burton, 2013), and with realistic environmental factors such
29
30 as time constraints and cognitive load.
31
32
33
34
35
36
37

38
39 In summary, we have shown that traditional forms of photo-ID could be improved by
40
41 replacing individual photographs with representations derived from multiple photos of
42
43 the same face. Based on our current findings, we expect that this would have a
44
45 beneficial effect on identity verification procedures in occupational settings. Future
46
47 research should determine the optimal range of within-person variability, and how
48
49 best to summarize it. For now, it is clear that a single photograph is not the best way
50
51 to represent facial appearance.
52
53
54
55
56
57
58
59
60
61
62
63
64
65

REFERENCES

- 1
2 Bindemann, M., Burton, A. M., Hooge, I. T. C., Jenkins, R. & De Haan, E. H.
3
4 (2005). Faces retain attention. *Psychonomic Bulletin & Review*, *12*, 1048-1053.
5
6
7 Bindemann, M., Jenkins, R., & Burton, A. M. (2007). A Bottleneck in Face Identification.
8
9 *Experimental Psychology*, *54*(3), 192–201.
10
11
12 Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M., & Miller, P.
13
14 (1999). Verification of face identities from images captured on video. *Journal of*
15
16 *Experimental Psychology: Applied*, *7*, 207-218.
17
18
19 Bruce, V., Henderson, Z., Newman, C. & Burton, A.M. (2001). Matching identities of
20
21 familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental*
22
23 *Psychology: Applied*, *7*, 207-218.
24
25
26 Bruce, V., Valentine, T., & Baddeley, A. (1987). The Basis of the 3/4 View Advantage in
27
28 Face Recognition. *Applied Cognitive Psychology*, *1*, 109–120.
29
30
31 Burton, A.M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-
32
33 quality video: Evidence from security surveillance. *Psychological Science*, *10*(3),
34
35 243–248.
36
37
38 Burton, A. M., Jenkins, R., Hancock, P. J. & White, D. (2005). The power of averages:
39
40 Robust Representations for Face Recognition. *Cognitive Psychology*, *51*, 256- 284.
41
42
43 Clutterbuck, R., & Johnston, R. A. (2002). Exploring levels of face familiarity using
44
45 an indirect face-matching measure. *Perception*, *31*(8), 985-994.
46
47
48 Clutterbuck, R., & Johnston, R. A. (2004). Matching as an index of face familiarity.
49
50 *Visual Cognition*, *11*, 857-869.
51
52
53 Clutterbuck, R., & Johnston, R. A. (2005). Demonstrating how unfamiliar faces
54
55 become familiar using a face matching task. *European Journal of Cognitive*
56
57 *Psychology*, *17*, 97-116.
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
- Davis, J. P., & Valentine, T. (2009). CCTV on trial: Matching video images with the defendant in the dock. *Applied Cognitive Psychology*, 23(4), 482–505.
- Hancock, P., & Bruce, V. & Burton, A. M. (2000). Recognition of unfamiliar faces. *Trends In Cognitive Sciences*, 4, 330-337.
- Hill, H. & Bruce, V. (1996). Effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 986–1004.
- Hole, G., George, P., Eaves, K. & Rasek, A. (2002). Effects of geometric distortions on face-recognition performance. *Perception*, 31, 1221-1240.
- Kemp, R., Towell, N. & Pike, G. (1997). When seeing should not be believing: Photographs, credit cards and fraud. *Applied Cognitive Psychology*, 11, 211-222.
- Jenkins, R., & Burton, A. M. (2008a). 100% Accuracy in Automatic Face Recognition. *Science*, 319, 435–435.
- Jenkins, R., & Burton, A.M. (2008b). Response to comment on “100% accuracy in automatic face recognition”. *Science*, 321, 912d.
- Jenkins, R., & Burton, A. M. (2011). Stable face representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1571), 1671–1683.
- Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of the same face. *Cognition*, 121, 313-323.
- Megreya, A, M. & Burton, A. M. (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition*, 34, 865-876.
- Megreya, A. M., & Burton, A. M. (2008). Matching faces to photographs: Poor performance in eyewitness memory (without the memory). *Journal of Experimental Psychology: Applied*, 14, 364–372.
- Osborne, C. D., & Stevenage, S. V. (2008). Internal feature saliency as a marker of

familiarity and configural processing. *Visual Cognition*, 16(1), 23-43.

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory

measures. *Behavior research methods, instruments, & computers*, 31(1), 137-149.


White, D., Kemp, R. I., Jenkins, R. & Burton, A. M. (2013). Feedback training for

facial image comparison. *Psychonomic Bulletin & Review*. Advance online

publication. [doi: 10.3758/s13423-013-0475-3](https://doi.org/10.3758/s13423-013-0475-3)


1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 1. Example image pairs used in Experiment 1. In each trial a comparison image (left) was paired with either an average image (top row) or a single photograph (bottom row). Image pairs in the left column show the same person (match). Image pairs in the right column show different people (mismatch).




1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 2. Mean accuracy (percent correct) for the face matching task in Experiment 1
 (\pm standard error).



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 3. Example stimulus displays from Experiment 2, showing each of four Array Size conditions. Displays in the left column (one-photo and three-photo arrays) show match trials, and arrays in the right column (two-photo and four-photo arrays) show mismatch trials.



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 4. Mean accuracy (percent correct) for the face matching task in Experiment 2
 (\pm standard error).





Figure 5. Mean accuracy (percent correct) for the face matching task in Experiment 3
 (± standard error).



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Supplemental Materials

[Click here to download Supplemental Materials: XAP_white_1032_supp.pdf](#)