This paper was accepted for publication in *Cognition* in November 2021.

Please cite as follows:

White, D., Wayne, T., & Varela, V. P. L. (2022). Partitioning natural face image variability emphasises within-identity over between-identity representation for understanding accurate recognition. *Cognition*, 219, 104966. <https://doi.org/10.1016/j.cognition.2021.104966>

Partitioning natural face image variability emphasises within-identity over between-identity representation for understanding accurate recognition

David White*, Tanya Wayne & Victor P. L. Varela
School of Psychology, UNSW Sydney

Word Count including references and figure captions: 9363

*Correspondence concerning this article should be addressed to David White, School of Psychology, UNSW Sydney, 2052, Australia. Email: david.white@unsw.edu.au

1 **ABSTRACT**

2
3 Accurately recognising faces enables social interactions. In recent years it has become clear
4 that people's accuracy differs markedly depending on viewer's familiarity with a face and
5 their individual skill, but the cognitive and neural bases of these accuracy differences are not
6 understood. We examined cognitive representations underlying these accuracy differences by
7 measuring similarity ratings to natural facial image variation. Natural variation was sampled
8 from uncontrolled images on the internet to reflect the appearance of faces as they are
9 encountered in daily life. Using image averaging, and inspired by the computation of
10 Analysis of Variance, we partitioned this variation into differences between faces (between-
11 identity variation) and differences between photos of the same face (within-identity
12 variation). This allowed us to compare modulation of these two sources of variation
13 attributable to: (i) a person's familiarity with a face and, (ii) their face recognition ability.
14 Contrary to prevailing accounts of human face recognition and perceptual learning, we found
15 that modulation of within-identity variation – rather than between-identity variation – was
16 associated with high accuracy. First, familiarity modulated similarity ratings to within-
17 identity variation more than to between-face variation. Second, viewers that are extremely
18 accurate in face recognition – 'super-recognisers' – differed from typical perceivers mostly in
19 their ratings of within-identity variation, compared to between-identity variation. In a final
20 computational analysis, we found evidence that transformations of between- and within-
21 identity variation make separable contributions to perceptual expertise in face recognition.
22 We conclude that inter- and intra-individual accuracy differences primarily arise from
23 differences in the representation of within-identity image variation.

24
25 (249 words)

26 **KEYWORDS**

27
28
29 Perception, Vision, Perceptual learning, Category learning, Face recognition, Individual
30 differences, Super-recognizers, Developmental Prosopagnosia.
31
32
33
34
35
36
37
38
39
40
41
42
43
44

1 INTRODUCTION

2
3 We must recognise objects in our visual world despite variations in their appearance, for
4 example due to lighting, hue and viewpoint. Variation in appearance is also informative,
5 signaling for example the time of day, ripeness of fruit or health of a loved one. Reading
6 visual information from people's faces – such as a person's identity, or their expression –
7 exemplifies this perceptual challenge. When recognizing faces, these two sources of
8 perceptual variation also give rise to an essential problem. First, variation in appearance
9 *between* faces is relatively constrained because faces are an especially homogeneous object
10 class, with every face sharing a common set of features and configuration. Second, variation
11 in appearance *within* a given face from one encounter to the next can be substantial, due to
12 changes caused by aging, expression, head angle, and so forth. As a result, images of the
13 same face can vary more than photos of different faces (e.g. Jenkins, White, Van Montfort, &
14 Burton, 2011; Cox, 2014).

15
16 To account for people's ability to solve this recognition problem, geometric models of
17 psychological similarity representation have been proposed. According to these models,
18 variation between faces is represented as deviations from the average appearance of faces that
19 a given viewer has experienced. This 'norm-based encoding' serves to tune perceptual
20 sensitivity to the properties – for example, the width of a person's nose – that distinguish
21 faces from one another in a given population of faces (see Rhodes & Leopold, 2011;
22 Valentine, 1991; c.f. Chang & Tsao, 2017). This provides an efficient coding of between-
23 identity variance, given the relative structural homogeneity of faces as an object class, and is
24 founded on a long tradition that characterises category and perceptual learning as the
25 discovery of differences between objects and object categories (e.g. Gibson & Gibson, 1955a,
26 1955b; see also Goldstone, Lippa & Shiffrin, 2001; c.f. Livingston, Andrews & Harnard,
27 1998; Collins & Behrmann, 2020).

28
29 But these models do not provide a full account of face representation for two main reasons.
30 First, geometric models of face representation – and object categorization more broadly –
31 only aim to capture representation of similarity across images of *different* objects (e.g.
32 Valentine, 1991; Valentine, Lewis, & Hills, 2016; Ashby & Perrin, 1988), and not perceptual
33 variation across different encounters with the *same* object (see Burton, Kramer, Ritchie, &
34 Jenkins, 2016; c.f. Tanaka, Giles, Kremen & Simon, 1998; Tanaka & Corneille, 2007). This
35 focus on differences between faces is reflected in the default experimental designs in this
36 field, which continue to be dominated by the use of highly standardized face images (see
37 Dawel et al. 2021). These default stimuli do not contain variability found in natural images
38 caused by lighting, expression, head pose and so on, causing face processing tasks performed
39 in the lab to be unrepresentative of the challenges of face processing tasks performed in
40 daily life (Burton. 2013). This methodological approach may also have important
41 downstream implications for theory development, leading to disproportionate focus on
42 cognitive processes that enable us to *tell faces apart* (i.e. between-identity variation), rather
43 than how we might *tell faces together* (i.e. within-identity variation, see Jenkins et al., 2011,
44 Burton, Jenkins, Hancock, & White, 2005).

The second reason that current models of face space are inadequate is that they do not distinguish familiar from unfamiliar faces, and by implication, do not offer an explanation of face learning. In simple perceptual matching tasks, familiar faces are identified much faster and more accurately than unfamiliar faces (e.g. Johnston & Edmonds, 2009; Jenkins et al., 2011; White, Burton, Jenkins, & Kemp, 2014), and increasing levels of familiarity with faces provide accumulating benefits for face identification (Clutterbuck & Johnston, 2005; see Johnston & Edmonds, 2009; Natu & O'Toole, 2011, Ramon & Gobinni, 2018, for reviews). Face space models do not account for these transformative effects of familiarity on performance, nor explain the processes by which incremental exposure to variation in the appearance of a face supports robust recognition (Jenkins et al., 2011; Balas & Pearson, 2017; Ritchie & Burton, 2017; Menon, Kemp, & White, 2018; Murphy, Ipser, Gaigg, & Cook, 2015; Kramer, Young, & Burton, 2018; Dunn, Kemp, & White, 2018; White et al., 2014; c.f. Chiachia, Falcao, Pinto, Rocha, & Cox, 2014; Blauch, Behrmann, & Plaut, 2020).

Recent computational studies have begun to address these shortfalls and are characterized by two main innovations. First, they use images containing natural image statistics that more closely model the perceptual exposure in daily life, thereby capturing the challenge of matching faces by identity outside of the lab (see Burton, 2013). Second, by using datasets that contain multiple images of the same identity, researchers have begun to model face learning as the accumulation of perceptual knowledge of how a given face varies from one image to the next (see Burton et al., 2016; Kramer et al., 2018; Blauch et al., 2020; c.f. Chiachia et al., 2014; Hill et al., 2019; Colón, Castillo, & O'Toole, 2020). This computational work shows that metric transformation of face image variation underlies familiarity-based improvements in supervised machine learning, and that these transformations are primarily found for representations of within-identity variation (e.g. Linear Discriminant Analysis: Kramer et al., 2018; DCNNs: Blauch et al., 2020, see section 3.6).

Whether these computational models capture human performance has not yet been tested. In the experiments reported here we directly compare representation of between- and within-identity variation in natural image statistics, to better understand the metric transformations of face image variation underlying high levels of face recognition accuracy. Interpreted within a signal detection framework (Hautus, Macmillan & Creelman, 20021; see also Ashby & Perrin, 1988), face recognition accuracy can be enhanced by two metric transformations to representational space, as illustrated in Figure 1.

The first transformation is expansion of between-identity variation (Figure 1B). By expanding the representational space containing different faces, individual faces are less likely to be confused with one another. This variance that has been strongly emphasised in existing geometric models of face recognition (e.g. Valentine, 1991; Valentine, Lewis & Hills, 2016). In addition, empirical studies that map metric changes in face stimuli to perceptual space have consistently focussed on changes in representation of between-face variation (Rhodes & Leopold, 2011; Faerber, Kaufmann, Leder, Martin, & Schweinberger,

2016; Chauhan, Kotlewska, Tang & Gobbini, 2020; McKone, Martini, & Nakayama, 2001). Recent studies have shown greater separation of familiar relative to unfamiliar faces in perceptual space (Faerber et al. 2016; Collins & Behrmann, 2020). However, stimuli in these studies only sample between-face variation and so it is not known how this transformation compares to other sources of variation.

Figure 1. What representational changes underlie expertise in face recognition?

(A) The core problem of face identity recognition is to cluster different images of the same identity in representational space, but this is challenging because variation in images of the same face can overlap with variation between different faces. In this example, the space devoted to Harrison Ford (orange shaded area), overlaps with a similar looking unfamiliar face (blue area), meaning that some encounters with the unfamiliar face would be mistaken for Harrison Ford. (B) One solution to this problem is to expand between-identity variation by increasing distance between individual faces in similarity space. (C) Another solution is to compress within-identity variation by reducing distances between images of the familiar face.

The second metric transformation that contributes to accurate face recognition is compression of within face variation (Figure 1C). Interest in within-identity variation is beginning to be acknowledged more widely, but it remains uncommon for stimuli in studies of face processing to include this type of variation (see Dawel et al. 2021). And while computational studies have begun to examine the source of metric transformations underpinning face recognition ability (e.g. Kramer et al. 2018), direct measurement of representational distance in artificial systems is more straightforward than in human participants. One reason for the lack of comparable studies on human participants therefore might be the methodological challenge associated with isolating these sources of variance.

To address this challenge, we introduce a new measure of geometric dispersal founded on natural image statistics, using face averages to represent centroids in physical similarity space. Face averages are routinely used to represent the centroid of face space (Bülthoff & Zhao, 2021), but their use as representations of person-specific subspaces has not been fully explored (see Burton et al. 2005, 2015). Here we use face averages for both of these purposes (see Figure 2). First, to measure perception of between-identity variation, we ask participants to rate the similarity of ‘gender averages’ – representing centroids of the overall facespace – to ‘identity averages’ that represent the average appearance of a single identity (Figure 2,

left). Second, to measure within-identity variation, we ask participants to rate the similarity of the identity averages – which represent centroids of identity-specific face *subspaces* – to exemplar images that were used to compute these averages (Figure 2, right).

This approach is analogous to the partitioning of between- and within-group variation in Analysis of Variance (ANOVA), enabling us to directly compare between- and within-identity perceptual variation for the first time. We compare the relative associations between representation of between- and within-group variation and high levels of accuracy in face recognition, in three ways. First, we test the relative extent to which differences in dispersal of these sources of variation can explain the transformative effect of face familiarity on recognition performance.

Second, we ask whether differences in representation of between- or within-identity variation can explain the large individual differences in people’s ability to identify faces. Over the past decade, research has shown that peoples’ ability to recognise faces varies substantially in the general population, with some people unable to recognise friends and family (Duchaine & Nakayama, 2005; Bowles et al., 2009) while others – super-recognisers – show extremely high accuracy in challenging face identity processing tasks (Russell, Duchaine, & Nakayama, 2009). Significant progress has been made in establishing the domain-specificity of these individual differences (e.g. Wilmer et al. 2012), their genetic basis (Wilmer et al., 2010; Shakeshaft & Plomin. 2015), and associations with other cognitive abilities (e.g. Jiahui, Yang, & Duchaine, 2018). However, mechanistic understanding of their cognitive and perceptual basis remains limited (see Wilmer, 2017 for a review). Third, we examine representations of between- and within-identity variation in a Deep Convolutional Neural Network optimised for face identity classification (VGG-Face, Parkhi, Vedaldi, & Zisserman 2015).

GENERAL METHODS

Key aspects of our experimental approach were similar across all four experiments and so these are presented together here. Methodological details that are specific to each experiment are presented in subsequent sections.

Stimuli

We conducted four experiments measuring perceived similarity between images of familiar and unfamiliar faces. Familiar faces were 20 US celebrities selected to be highly familiar to participants regardless of their country of residence, and unfamiliar faces were 20 Dutch celebrities selected to be unfamiliar outside of the Netherlands (see *Supplementary Materials* for list of celebrity names). All images used in the study were sampled from the internet to contain natural image variation. Images were downloaded by performing a Google Image search using the celebrity name as a search term and downloading the first 20 images that met inclusion criteria: (i) the face was approximately forward facing; (ii) the image was sufficient quality such that there were at least 60 pixels between the eyes; (iii) the face was

not occluded by another object. All other aspects of the images were free to vary including expression, lighting and pose.

Between-face variation

Within-face variation

Figure 2. Diagram showing example stimulus comparisons used to quantify dispersal of perceptual variation in psychological similarity space. Our approach was to partition natural image variation into between-face differences (left) and within face differences (right), following statistical convention in Analysis of Variance (ANOVA). At the level of between-identity variation (left) participants rated similarity of ‘identity average’ images, created from all images in the stimulus database of that individual face, to ‘gender average’ images created from all same-gender faces in the stimulus dataset. This constrained similarity judgments to the stable aspects of faces that are common to an identity, independent of effects of transient aspects of appearance caused by variations in lighting, expression etc. At the level of within-identity variation (right), participants rated similarity of single images of a given identity to the ‘identity average’ images created from all images of that identity. This constrained similarity judgments to changeable aspects of appearance that occur naturally between images of the same face.

Participant’s similarity ratings enabled us to quantify the dispersal of perceptual variation in psychological similarity space, at two different levels of representation. To quantify variation at the level of between-identity differences, we asked participants to rate the similarity of individual faces to ‘gender average’ images created from all of the same-gender faces in our stimulus dataset (representing the geometric centroid of the overall face space, see Figure 2 left). To quantify variation at the level of within-identity differences, we asked participants to rate the similarity of single images of an individual to an ‘identity average’ created from all of the images of that person’s face (representing the centroid of a person-specific face subspace, Figure 2 right).

Identity averages were created by manually aligning landmark anatomical features to a standard face template using image morphing software for each of the 20 images of that face. Previous research has shown that the appearance of identity averages converge on a stable appearance when 16 images or more contribute to an average (Burton et al. 2005) and so this number was deemed sufficient to derive a representative centroid. These resulting morphed images were then averaged to produce an average texture. These ‘shape-free’ average textures were then morphed back to the average shape coordinate template of that face to produce the averages (see Burton et al., 2005 for details of this procedure). Image exemplars

used in within-identity comparisons were morphed to an average shape template and back again, to ensure that artifacts of the morphing process did not affect our results. Gender averages used in between-identity comparisons were created by averaging all the same-gender identity averages using the method described above.

In a final processing step, all images were converted to greyscale and normalized by matching luminance histograms and Fourier amplitudes of individual images to the average of all experimental stimuli using a toolbox developed by Willenbockel et al (2010) to ensure as much uniformity in low-level image statistics as possible. All images were presented in greyscale and at a resolution of 190 x 285 pixels.

Design & Procedure

In all experiments, participants were asked to rate the similarity of images presented one after another. The task was self-paced, with participants viewing each stimulus for as long as they liked before viewing the next image. The second image in a pair was presented three seconds after the first image and participants were asked to ‘rate the similarity of this image to the previous image’ using a 9-point Likert scale (1 = ‘Very Dissimilar’, 9 = ‘Very Similar’).

Participants compared pairs of images as illustrated in Figure 1. In experiments where participants made between-identity comparisons (Experiments 1, 3, & 4) they compared all identity averages to respective same-gender average. Where they made within-identity comparisons (Experiments 2, 3 & 4), they compared only one image exemplar-average pair per identity. This limited exposure to the average images, to ensure that viewing averages of familiar faces did not cause change in the existing face representations and so that participants did not become familiar with the unfamiliar faces¹. In Experiment 1, participants made between-identity comparisons only. In Experiment 2, participants made within-identity comparisons only. Experiments 3 and 4 combined the design of Experiments 1 and 2 into mixed factorial designs with Variation Type (Between-identity, Within-identity) as between-subjects factors and Familiarity (Unfamiliar, Familiar) as within-subjects factor.

After each experiment, participants were presented with names of each celebrity and were asked to confirm whether they were visually familiar with the persons’ appearance prior to

¹ Two ratings were collected for each image pair because, each time in the opposite sequence. This was because similarity ratings are asymmetric when comparing two stimuli that differ in prototypicality, with more atypical instances of an object being rated as more similar to typical instances than when the comparison is made in the opposite direction (e.g. Rosch, 1975; Op de Beeck, Wagemans, & Vogels, 2003). Indeed, we had originally designed Experiments 1 and 2 to test whether similarity asymmetries were also found in face image comparisons. We found consistent evidence for this predicted pattern of similarity asymmetries in face image comparisons for both between- and within-identity variation. While this was not of primary interest in the main paper, these results provide potentially important constraints for models of face recognition, so we have continued this investigation and will report results in a forthcoming manuscript.

1 taking part in the experiment. In addition, they completed a post-experiment questionnaire
2 which gauged their self-reported engagement with the task, and ‘catch trials’ showing pairs
3 of identical images were interspersed with experimental trials. We set *a priori* conditions of
4 inclusion that we expected would lead to subject attrition. Participants were excluded prior to
5 analysis if they: (i) did not complete the study, (ii) did not pass quality control criteria based
6 on a self-report questionnaire measuring level of engagement in the study and/or failed a
7 ‘catch trial’ test; (iii) reported recognising less than 50% of the US celebrities or reported
8 recognising more than 50% of the Dutch celebrities. For participants that were included in
9 our analysis, rates of familiarity were very high, ranging from 90 to 98% across studies.

11 **EXPERIMENT 1**

13 In Experiment 1 we asked whether familiar faces show more dispersal in psychological
14 similarity space compared to unfamiliar faces. Previous studies suggest that familiar faces
15 should show greater dispersal because famous celebrity faces are perceived as less typical
16 than unfamiliar faces (Faerber et al. 2016), and because faces become less similar to one
17 another after experimental familiarisation (Collins & Behrmann, 2020). We therefore
18 predicted that familiar faces would be rated as less similar to average faces representing the
19 center of facespace, indicative of greater dispersal in facespace.

21 **Method**

23 *Design and Procedure*

25 Participants each rated similarity of each of the 40 stimulus identities (20 familiar, 20
26 unfamiliar) to the average face (Figure 1A). For one rating, the norm was presented before
27 the prototype, and in the other trial the same stimuli were presented but in the opposite
28 sequence (see Footnote 1). The 80 trials were presented in a random order for each
29 participant, with familiar and unfamiliar faces presented in separate blocks.

31 *Participants*

33 Participants were all US citizens recruited via *Amazon Mechanical Turk*. After exclusions
34 there were 146 participants (79 female, $M_{\text{age}} = 36.6$ years, $SD_{\text{age}} = 11.7$ years).

36 **Results**

38 Results for Experiment 1 are shown in Figure 2 (top left panel). Contrary to our prediction,
39 the difference in similarity ratings between familiar and unfamiliar faces in between-identity
40 variation in Experiment 1 was non-significant, $t(290) = .994$, $p = .321$, 95% CI: $-.16, .49$,
41 *Cohen's d* = 0.12.

43 **EXPERIMENT 2**

1 In Experiment 2 we asked whether exemplar images of familiar faces are perceived as more
2 similar to identity-specific face averages – representing centroids of person-specific
3 subspaces – compared to unfamiliar faces.

4 5 **Method**

6 7 *Design and procedure*

8
9 The design of Experiment 2 required 40 stimulus comparisons to be made for each of the 40
10 stimulus identities, giving a total of 1600 comparisons. To avoid participants becoming
11 familiar with the experimental stimuli, these comparisons were split into various versions of
12 the task and distributed across a large number of participants. Each participant made 40
13 stimulus comparisons to either familiar US or unfamiliar Dutch celebrities.

14
15 Twenty versions of the test were created for the familiar stimuli, and 20 versions using the
16 unfamiliar stimuli, with participants allocated at random to one of the 40 versions. Image
17 pairs consisted of an image average of an individual celebrity, and a randomly selected
18 exemplar image that had contributed to that average. All other aspects of the methodology
19 were identical to the previous experiment.

20 21 *Participants*

22
23 All participants were US Citizens recruited via *Amazon Mechanical Turk*. After exclusions
24 there were 540 participants in the familiar group (273 female, $M_{\text{age}} = 34.0$ years, $SD_{\text{age}} = 11.5$
25 years) and 522 participants in the unfamiliar group (323 females, $M_{\text{age}} = 33.9$ years, $SD_{\text{age}} =$
26 12.0 years).

27 28 **Results**

29
30 Results for Experiment 2 are shown in Figure 2 (top right panel). Exemplar images of
31 familiar faces were rated as being more similar to their identity average than for unfamiliar
32 faces [$t(1060) = 8.729$, $p < .001$, 95% CI: .46, .73, *Cohen's d* = 0.541].
33

Figure 3. Familiarity modulates similarity of within- more than between-identity variation. In Experiments 1, 2 and 3 similarity did not consistently modulate dispersal of between-identity variation (left column). Individual images of faces were rated as more similar to the average appearance of that face when it was familiar compared to when it was unfamiliar (right column). Familiarity therefore appears to support recognition by compressing similarity space containing images of the same face, rather than by expansion of the broader similarity space containing different faces. For each plot, the circle markers are the mean similarity across participants and grey dots are average similarity rating given by individual participants. Dashed lines represent the median and interquartile range. Shaded regions are probability density functions.

EXPERIMENT 3

Results of Experiment 2 suggest that familiarity with a face causes images of the same face to appear more similar to the average appearance of that face, thereby reflecting a compression of image variance. This is in contrast to the results of Experiment 1, where familiarity did not appear to modulate the dispersal of individual faces relative to the average face. However, these experiments did not directly compare the modulating effects of familiarity on between- and within-face variation. In Experiment 3, we combined the design of Experiments 1 and 2 to examine the relative effects of familiarity on perception of between- and within-face variation.

Method

Design and procedure

Design of previous experiments were combined in a 2 by 2 mixed factorial design with between-subjects factor Variation Type (Between-identity, Within-identity) and within-subjects factor Familiarity (Unfamiliar, Familiar). Participants in this experiment were required to make two similarity judgments for each of 40 image pairs (20 familiar identities, 20 unfamiliar identities).

Details of the procedure were the same as previous experiments. As in Experiment 2, we again aimed to limit participants exposure to face prototypes and stimulus identities such that each participant rated a single image pair per face identity. This resulted in 20 versions of the test, so that across participants, ratings were collected for all average-exemplar pairings in the within-identity variation condition. Participants in the within-identity condition were randomly assigned to one of these 20 versions.

Participants

Participants were US citizens recruited from *Amazon Mechanical Turk*. The final sample after exclusions consisted of 527 participants who had been randomly assigned to either between-identity ($n = 271$; 135 female, $M_{\text{age}} = 37.41$, $SD_{\text{age}} = 11.01$) or within-identity rating conditions ($n = 255$; 120 female, $M_{\text{age}} = 36.37$, $SD_{\text{age}} = 10.49$).

Results and discussion

Results of Experiment 3 are shown in the bottom panel of Figure 3. Main effects of Variation Type and Familiarity were qualified by a significant interaction between these factors [$F(1, 525) = 268.155$, $p < .001$, $\eta p^2 = .338$]. Planned comparisons showed that familiarity affected both dispersal of between-identity variation and compression of within-face variation, however the effect was overwhelmingly carried by within-identity variation [$t(255) = 22.468$, $p < .001$, 95% CI: 1.21, 1.45, *Cohen's d* = 1.40], rather than between-variation [$t(270) = 4.398$, $p < .001$, 95% CI: .1, .25, *Cohen's d* = 0.26].

EXPERIMENT 4

In Experiments 1 to 3 we found that familiarity modulated similarity ratings of within-face variation substantially more than ratings of between-face variation. Changes in representation of within-identity variation therefore appear to underpin this major source of intra-individual variation in on recognition accuracy. Next, using the same approach as Experiments 1 to 3 we examined the extent to which differences in representation of between- and within-identity variation can explain *inter*-individual differences in face recognition accuracy. To do this, we compared similarity ratings made by two new participant groups. First, typical viewers – who were participants with average face recognition ability as measured by standardised pre-screening tests. Second, ‘super-recognisers’ – who had been selected as being amongst the very highest performing participants on pre-screening tests.

Method

Design and procedure

The design of Experiment 4 was the same as Experiment 3, but with an additional of factor participant group: 3 by 2 mixed factorial design with between-subjects factors Group (super-recognisers, typical viewers) Variation Type (between-identity, within-identity) and within-subjects factor Familiarity (unfamiliar, familiar).

Images used in between-identity comparisons were the same as in previous experiments. However, because the participant groups in Experiment 4 were necessarily smaller than previous experiments – and because we were interested in differences between participant groups – within-identity comparisons were made on a subset of the full stimulus set used in earlier experiments. We made 3 versions of the within-identity comparison experiment each with a different set of exemplar-prototype pairings. These pairs were selected to have roughly equivalent similarity based on data from the full set of 20 exemplar-prototype pairs per stimulus identity that had been used in previous experiments. Participants were randomly allocated to one of 3 versions of the within-identity comparison condition.

Participants

Super-recognisers were recruited from the UNSW Face Research Registry which contained 1157 participants that had performed three tests of face identity processing ability (UNSW Face Test: Dunn et al. 2020; Glasgow Face Matching Test: Burton et al. 2010; Cambridge Face Memory Test extended version: Russell et al. 2009). We invited 106 super-recognisers to participate that had exceeded average accuracy on each of these three tests by at least 1.7 standard deviations. This is a relatively stringent criteria relative to selection thresholds used in this field (see Ramon, 2021), and the requirement to show repeated high performance over subsequent tests limits any influence of chance and transient factors on performance. Summary statistics of super-recognisers screening tests performance relative to test norms are provided in Supplementary Materials (Table S2). From the 42 super-recognisers that completed the experiment, and after exclusions, 23 super-recognisers were assigned to between-identity comparisons (16 Female, $M_{\text{Age}} = 37.9$, $SD_{\text{Age}} = 7.92$) and 19 to within-identity comparisons (11 Female, $M_{\text{Age}} = 39.6$, $SD_{\text{Age}} = 9.4$).

Typical viewers in Experiment 4 were also recruited from a larger subset of the UNSW Face Research Registry that had completed the UNSW Face Test. Based on standard response rates, we emailed the experiment links to two separate random selections of 120 participants who had scored within one standard deviation of the mean on the UNSW Face Test. From the typical viewer respondents that completed the test, and after exclusions, 45 had completed between-face comparisons (29 Female, $M_{\text{Age}} = 47.0$, $SD_{\text{Age}} = 14.2$) and 31 within-face comparisons (10 Female, $M_{\text{Age}} = 39.8$, $SD_{\text{Age}} = 11.5$).

Results

Similarity ratings for super-recognisers and typical viewers are shown in Figure 4. We conducted a three-way mixed ANOVA with between-subjects factors participant Group (typical viewers, super-recognisers), Variation Type (between-identity, within-identity) and within-subjects factor Familiarity (familiar, unfamiliar). This analysis showed two significant two-way interactions. First, consistent with previous experiments, the effect of Familiarity was larger than for within-identity compared to between-identity variation [2-way interaction: $F(1, 114) = 198, p < .001, \eta p^2 = .635$; *Cohen's* $d_{(\text{between})} = 0.259$, *Cohen's* $d_{(\text{within})} = 2.042$]. Second, and more importantly for this study, the effect of Group was also markedly larger for within- than for between-identity variation [2-way interaction: $F(1, 114) = 52, p < .001, \eta p^2 = .314$; *Cohen's* $d_{(\text{between})} = 0.703$; *Cohen's* $d_{(\text{within})} = 1.467$]. This pattern of results shows that it is changes in representation of within-face variation, rather than between-face variation, that are primarily associated with individual differences in face recognition accuracy.

Figure 4. Face similarity ratings in super-recognisers and typical viewers. Between-identity similarity ratings for super-recognisers were slightly lower overall compared to typical viewers (left panel). As in previous experiments, modulation of similarity ratings by face familiarity was most pronounced for within-identity variation (right). Here, the difference between ratings by typical viewers and super-recognisers was also largest for within-identity ratings. Full analysis is in the main text.

Because the three-way interaction was significant [$F(1, 114) = 7.41, p = 0.007, \eta p^2 = .061$], and because visual inspection of Figure 4 shows a different pattern of data for between- and within-identity variation, we also analysed these data separately. For between-identity variation (Figure 4, left panel), analysis shows a small but significant main effect of Familiarity, $F(1, 66) = 5.50, p = .022, \eta p^2 = .077$, with lower ratings of similarity to familiar compared with unfamiliar faces ($M_{\text{Fam}} = 2.30, SD_{\text{Fam}} = 1.00$; $M_{\text{Unf}} = 2.46, SD_{\text{Unf}} = 1.02$,

Cohen's d = 0.259). The main effect of Group was significant, $F(1, 66) = 9.11, p = .004, \eta^2 = .121$, with super recognisers rating images as less similar than typical viewers ($M_{\text{SRs}} = 1.95, SD_{\text{SRs}} = 0.89; M_{\text{Control}} = 2.62, SD_{\text{Control}} = 0.92, \text{Cohen's } d = 0.703$). The interaction between Group and Familiarity was not significant for between-identity variation, $F(1, 66) = 1.02, p = .316, \eta^2 = .015$.

For within-identity variation (Figure 4, right panel), we found a main effect of Familiarity, $F(1, 48) = 237, p < .001, \eta^2 = .832$, and Group, $F(1, 48) = 49.353, p < .001, \eta^2 = .507$, with super recognisers rating images of the same face as more similar than typical viewers ($M_{\text{SRs}} = 7.7, SD_{\text{SRs}} = 1.2; M_{\text{Control}} = 5.87, SD_{\text{Control}} = 1.32$). The interaction between Group and Familiarity was also significant, $F(1, 48) = 6.592, p = .013, \eta^2 = .121$, driven by a larger effect of Familiarity for super-recognisers compared to typical viewers (super-recognisers: Mean difference = 1.895, $p < .001$, 95% CIs = 1.561, 2.228, *Cohen's d* = 3.05; typical viewers: Mean difference = 1.354, $p < .001$, 95% CIs = 1.093, 1.615, *Cohen's d* = 1.74). Thus, whereas familiarity-based modulation of between-identity similarity did not vary by group, familiarity-based modulation of within-identity variation was enhanced in super-recognisers.

COMPUTATIONAL STUDY

In a final study, we compared between- and within-identity variation using low-level image statistics (GIST descriptor: Oliva & Torralba, 2001) and in a Deep Convolutional Neural Network (DCNN) optimised for face identity classification (VGG-Face, Parkhi, Vedaldi, & Zisserman 2015).

Method, Results & Discussion

Low-level image statistics

We first tested whether the patterns of similarity ratings by participants to familiar and unfamiliar face image pairs could be explained by the image statistics of our stimuli. Although we sampled and processed images of unfamiliar and familiar celebrities in the same way, it remained possible that inherent properties of the face image – rather than familiarity with the faces – influenced our results, because image-based face similarity can be closely related to face discrimination performance (Bülthoff & Zhao, 2021). To examine this, we first compared spatial frequency content of images using a standard approach developed in previous work (GIST descriptor: Oliva & Torralba, 2001). These statistics have been used to examine low-level similarity of faces in previous work and are correlated with neural response in the ventral stream (e.g. Rice, Watson, Hartley, & Andrews, 2014). Details of the computations involved are provided in Supplementary Material.

Similarity of low-level image statistics are shown in Figure 5A. We found no differences between familiar and unfamiliar faces for either between- or within-identity variation (between: $t(19) = 0.016, p = .987$, 95% CI: -.046, .045, *Cohen's d* = 0.004; within: $t(270) =$

4.398, $p < .001$, 95% CI: .1, .25, *Cohen's d* = 0.26) and so differences in these low-level image statistics do not explain our pattern of results. Notably, these data also show that low-level image similarity of between-identity variation was greater than within-identity variation – the opposite pattern to that observed for human similarity judgments. This is explained by the fact that averaging multiple images of the same face produces images that are unaffected by superficial variation in lighting, expression, pose etc (see Burton et al., 2005), yet humans perceive the images in the right panel of Figure 2 to look more similar than images in the left panel. This presumably reflects people's expertise in discriminating aspects of image variance that distinguish face identity.

Figure 5. Low level image statistics do not capture human similarity ratings. (A) Normalised cosine similarity of GIST descriptor vectors that represent global properties of image statistics. We found no differences between unfamiliar and familiar faces, and there was greater similarity in between-identity variation compared to within-identity variation. Together, these results show that low-level properties of images are not able to account for human ratings of similarity. (B) Normalised cosine similarity of feature vectors extracted from layers of a Deep Convolutional Neural Network (DCNN). The earliest layer of the DCNN represent the raw image data (Img), followed by convolutional layers (conv1, 2, 3, 4, 5) that process by pooling and convolving the pixel data in a similar way to early cortical processing in the brain. The final ‘fully connected’ layers (fc6, fc7) are trained via back propagation to extract patterns distributed across the local, image-bound representations in convolutional layers. Only fully connected layers are capable of producing accurate face recognition. As can be seen from this figure, changes to similarity representation in fully connected layers that enable accurate face recognition are driven by transformation of between- and not within-identity variation.

Deep Convolutional Neural Network

Low-level statistics show greater similarity in images containing only between-identity variation, compared to only within-identity variation. Therefore, in terms of raw statistics, the identity ‘signal’ – i.e. between-identity differences – is swamped by the ‘noise’ of within face differences. This raises the question of how perceptual expertise in recognising faces is achieved computationally: Is the ability to discriminate aspects of image variance that distinguish face identity achieved by (i) expanding between-face variance, or (ii) compressing within-face variance?

To address this question we ran another computational study that compared representation of between- and within-identity variation in a Deep Convolutional Neural Network (DCNN: VGG-Face, Parkhi, Vedaldi, & Zisserman 2015). Facial similarity scores produced by this

DCNN have been shown to discriminate between unfamiliar faces to a level that is comparable with human accuracy (Phillips et al. 2018). Therefore, examining representation of similarity at progressively higher layers of this network enabled us to test how representation of facial similarity changes in higher levels of visual processing that are capable of discriminating between unfamiliar faces. Because the network was not trained on the celebrities used in our experiments, and so was unfamiliar with all of these faces, we collapsed across familiar and unfamiliar faces for this analysis.

Average similarity of within- and between-identity variation in successive layers of the DCNN are shown in Figure 5B (see *Supplementary Materials* for further details of this analysis). Two aspects of these results are worthy of note. First, consistent with the GIST descriptor analysis – but inconsistent with human similarity judgments – we found that between-identity variation was less similar than within-identity variation in early network layers. Second, we observe a clear change in representational similarity of between-identity variation at higher ‘fully connected’ DCNN layers that are capable of discriminating between unfamiliar faces, but no accompanying change in representation of within-identity variation.

Importantly however, a previous study that compared representation of within-face variation for faces that were familiar and unfamiliar to this same DCNN did show compression of within-identity variation at fully connected network layers, specifically for faces that were familiar to the network (Blauch et al., 2020, section 3.6). This suggests that two separable sources of perceptual expertise with faces promote recognition accuracy by modulating distinct components of perceptual variation. On the one hand, our computational analysis shows that perceptual expertise in processing face identity leads to greater dispersal of unfamiliar face representations in face space. On the other, specific expertise with familiar faces causes compression of within-face variation for those faces, as shown in human experiments here and prior computational work (Blauch et al., 2020, section 3.6; see also Kramer et al. 2018). This may offer a useful resolution to recent debates about whether our perceptual expertise with faces is limited to the faces we know (see Young & Burton, 2018a; Rossion, 2018; Young & Burton, 2018b). Together, these results favour an account whereby people are experts in processing both unfamiliar and familiar face identity, but where these sources of expertise have distinct computational bases.

GENERAL DISCUSSION

In the experiments reported here we measured psychological similarity of two sources of natural image variation. First, between-face variation, defined as the structural differences between individual faces. Second, within-face variation, defined as the mutable differences that exist between different images of the same face. We made two comparisons that are known to account for substantial amounts of variability in human performance. First, between familiar and unfamiliar faces – the main cause of intra-individual differences in accuracy. Second, between typical viewers and super-recognisers – to examine representational differences underlying inter-individual differences in accuracy.

1 We were interested in the extent to which representations of similarity for between- and
2 within-identity variation were associated with these sources of human performance
3 differences and presented a novel method for measuring this directly.
4

5 Similar results were found for both comparisons. There were strong and systematic similarity
6 differences observed for within-identity variation, but weaker and less consistent differences
7 in between-identity variation. This places important theoretical constraints on understanding
8 of face recognition and visual expertise. A long history of theory in these fields has
9 characterised perceptual skill primarily as an ability to discriminate between individual
10 objects (Rhodes & Leopold, 2011; Valentine, 1991; Gibson & Gibson, 1955a, 1955b;
11 Goldstone, Lippa & Shiffrin, 2001; c.f. Livingston, Andrews & Harnard, 1998), or *telling*
12 *apart*. In contrast, results presented here suggest that the skill accounting for most variance in
13 human accuracy is in assimilating perceptual variance within a coherent identity
14 representation of familiar individuals – i.e., *telling together* (see Jenkins et al. 2011).
15

16 While modulation of within-identity variation was a product of both familiarity and
17 individual differences, these factors also interacted – with greater compression of within-
18 identity variation for familiar faces in super-recognisers. This suggests that people with high
19 levels of ability in face identification are especially adept at acquiring robust representations
20 of faces that they are familiar with. This result is surprising because super-recognisers in our
21 study were selected on the basis of test scores that were achieved on *unfamiliar* face
22 identification tasks, where participants receive very limited exposure to these faces. This
23 suggests that ability on unfamiliar face tasks predicts some of the differences in familiar face
24 identity processing (see also Russell et al. 2009; Devue, Wride & Grimshaw, 2019). But it
25 also implies that individual differences in face recognition cannot be fully understood without
26 understanding how high performing individuals acquire familiar face representations over the
27 course of more extensive exposure. Importantly, individual differences in familiar face
28 recognition ability have been studied only very rarely (but see Russell et al. 2009; Wilmer et
29 al. 2012; Bobak, et al, 2016; Devue, Wride & Grimshaw, 2019), and there has been no
30 systematic investigation of differences in familiar face representations across the ability
31 spectrum. Future work should therefore prioritise these knowledge gaps.
32

33 The differences we observed between similarity ratings to familiar and unfamiliar faces
34 provide the first direct support for predictions arising from recent computational models of
35 face representation. First, Kramer and colleagues (Kramer et al., 2018) modelled geometric
36 face space by deriving an unsupervised, bottom-up, representation of between- and within-
37 identity similarity using Principle Components Analysis (PCA). They then reshaped this
38 perceptual space using a supervised Linear Discriminant Analysis. Input to this model
39 included both bottom-up, PCA-derived representations of training images, and ‘top-down’
40 labels that grouped the individual training images by identity. The LDA reshapes the PCA
41 space such that differences between images of the same face are minimised, while differences
42 between faces are maximised. Consistent with other work, this top-down processing was
43 critical to enabling accurate face recognition (Mileva et al. 2020; Furl, Phillips & O’Toole,
44 2002). Consistent with our human data, this primarily resulted in images of the same face

1 being compressed towards the centre of their respective identity clusters, with more subtle
2 effects of reshaping on representation of similarity between different faces (Kramer et al.,
3 2018).

4
5 Blauch and colleagues (2020) show similar results to Kramer et al (2018) in a Deep
6 Convolutional Neural Network (DCNN). In the final layers of this network – which are
7 capable of accurate face recognition – they also found compression of within-identity
8 variation for familiar faces. Interestingly, we tested this same network in our study (VGG-16;
9 Parkhi et al., 2015) using faces that were not included in the training set used to develop this
10 algorithm. For these unfamiliar faces, we instead found that transformation of between-
11 identity variation – rather than within-identity variation – characterised differences between
12 earlier and later network layers, that are capable of accurate face recognition. We propose
13 that this modulation of between-identity variation reflects the basis of DCNN’s ability to
14 discriminate unfamiliar faces, whereas modulation of within- identity variation – reported by
15 Blauch and colleagues (2020) – explains the superior ability to discriminate the specific faces
16 that it has been trained on. Taken together with Blauch et al (2020), our computational
17 analysis therefore suggests that familiar face learning – which selectively modulates
18 representation of within-identity variation – is a computationally distinct source of perceptual
19 knowledge from ability to discriminate unfamiliar faces.

20
21 Differences between similarity ratings provided by super-recognisers and typical viewers
22 were mostly seen for within-identity comparisons and familiar faces. But super-recognisers’
23 ratings also reflected greater dispersion of between-identity variation and compression of
24 within-identity variation for unfamiliar faces. This might suggest that superior abilities of
25 super-recognisers are supported by both generic and person-specific components of expertise
26 outlined above. Naturally, deriving person-specific representation depends initially on the
27 successful discrimination of unfamiliar faces. According to the computational model by
28 Kramer et al (2018), initial success in deciding that an unfamiliar face is one that we have not
29 encountered before is necessary to facilitate clustering of the subsequent encounters with that
30 identity. So clear links exist between these sources of expertise exist in everyday face
31 learning.

32
33 Notwithstanding these links, our data show that differences in how individuals represent
34 within-identity variation appear to be the major determinant of face recognition ability.
35 Support for this is also seen in errors made by people with ‘face blindness’ and in young
36 children. Participants with Developmental Prosopagnosia – or ‘face blindness’ – appear to
37 have specific difficulty in matching ‘ambient’ images of the same unfamiliar face when they
38 contain naturalistic image variation (White, Palermo, Rivolta, Al-Janabi, & Burton, 2017;
39 Berger et al., 2021). Developmental research shows that school children are far more likely to
40 make errors of omission, by failing to recognise ambient images of their teachers, compared
41 to errors of inclusion where they misidentify their teacher as someone else (Laurence &
42 Mondloch, 2016). In addition, impairments in learning faces from variation have been
43 reported in Autistic Spectrum Disorder, which may be linked to a more general ability to
44 consolidate multiple exemplar images into integrated object representations (Ipser, Ring,

Murphy, Gaigg & Cook, 2016; see also Gastgeb, Wilkinson, Minshew, & Strauss, 2011). Together with our results, this evidence suggests that the process of assimilating variation in images of the same familiar face is critical to understanding individual differences in face recognition ability.

Our work is a first attempt at simultaneously studying the geometry of between-and within variation in psychological similarity space. Conceptually similar investigations have been carried out in computational systems (Kramer et al. 2018), but we know of no studies based on perceptual decision by humans. Because this is a first attempt, there are some limitations that should be addressed in future work. For example, our decision to compute variance relative to averages followed their extensive prior use as centroids in facespace (e.g. Leopold et al. 2001; Burton et al. 2005; Kramer et al. 2018), statistical convention and common practice in the analysis of variance². But some recent work suggests that statistical averages may not be regarded as the most typical appearance of faces (Ritchie, Kramer & Burton, 2018; see also Burton & Vokey, 1998), and so future work could examine whether statistical averages do indeed represent centroids of face spaces³. If not, there are various alternative ways to measure perceptual variance and partition sources of variance in natural image statistics. For one, variation could be measured as the sum of pairwise distances between individual face images, rather than distances from an average, although this would entail much larger volumes of human comparisons.

Moreover, an alternative to using photographs for within-identity comparisons, as we have here, is to model this image variance using freely available principal component analysis tools (e.g. InterFace; Burton et al. 2015). This would enable researchers to generate examples of individual faces that contain natural within-person variation without reliance on raw image data, enabling greater standardisation and experimental control over within-identity variation. Current popular generative frameworks for creating face stimuli model only between-face variation (e.g. Gerig et al. 2018), with generic modelling of variation in lighting, head angle and expression, rather than person-specific models captured from natural image statistics. Future frameworks that simultaneously represent variance at between-identity and within-identity levels can provide the basis for a more complete understanding of the representational geometry of facespace.

This type of approach might also enable researchers to address another important limitation. While our results clearly emphasise the relative importance of representing within-face variation in perceptual discrimination, they do not reveal much about the content of these

² A study published by Bülthoff and Zhao (2021), after we had completed our program of work, shows that averaging 40 or more same-gender faces produced stable ‘norm averages’. In light of this, future work should aim to use larger sets of identities to produce norm averages than we have used here.

³ We conducted a preliminary analysis of this question that is reported in Supplementary Materials. We asked 475 participants to rate how typical each individual image of familiar faces was of that person’s appearance and correlated the average typicality ratings for each image with the image-to-identity average similarity ratings collected in Experiments 2 & 3. This analysis shows a high correlation between similarity to averages and participants ratings of typicality, $r(400) = 0.685$, $p < 0.001$, suggesting that image averages do approximate centroids of person-specific subspaces. However, this was based on average ratings across participants and so does not capture idiosyncrasies in individual participant’s representations (see Ritchie, Kramer & Burton, 2018).

representations. Despite decades of research in face recognition, the perceptual content of familiar face representations remains poorly understood. And while there has been some recent progress – for example in understanding the extent to which representations contain abstractive (e.g. Burton et al. 2005, 2015) versus image specific information (e.g. White, Burton & Kemp, 2016; Armann, Jenkins & Burton, 2018; Dunn, Ritchie, Kemp & White, 2019; see also Hill et al. 2019; Colón et al., 2020; Ritchie, Kramer & Burton, 2017) – it is not clear how differences in the content of these representations is related to individual differences in accuracy. Neural and cognitive instantiation of the processes by which perceptual variation is assimilated into familiar face representations will be an important focus of future work in face recognition. This should also have implications for understanding object and category representation more broadly.

REFERENCES

- Ashby, F. G., & Perrin, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review*, 95(1), 124.
- Armann, R. G., Jenkins, R., & Burton, A. M. (2016). A familiarity disadvantage for remembering specific images of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 42(4), 571.
- Balas, B., & Pearson, H. (2017). Intra-and extra-personal variability in person recognition. *Visual Cognition*, 25(4-6), 456-469.
- Blauch, N. M., Behrmann, M., & Plaut, D. C. (2020). Computational insights into human perceptual expertise for familiar and unfamiliar face recognition. *Cognition*, 104341.
- Berger, A. F., Fry, R., Bobak, A. K., Juliano, A., & DeGutis, J. (2021, November 9). Distinct abilities associated with matching same identity faces vs. discriminating different faces: Evidence from individual differences in prosopagnosics and controls. <https://doi.org/10.31234/osf.io/95bz6>
- Bobak, A. K., Bennetts, R. J., Parris, B. A., Jansari, A., & Bate, S. (2016). An in-depth cognitive examination of individuals with superior face recognition skills. *Cortex*, 82, 48-62.
- Bowles, D. C., McKone, E., Dawel, A., Duchaine, B., Palermo, R., Schmalzl, L., ... & Yovel, G. (2009). Diagnosing prosopagnosia: Effects of ageing, sex, and participant–stimulus ethnic match on the Cambridge Face Memory Test and Cambridge Face Perception Test. *Cognitive Neuropsychology*, 26(5), 423-455.
- Bülthoff, I., & Zhao, M. (2021). Average faces: How does the averaging process change faces physically and perceptually?. *Cognition*, 216, 104867.
- Burton, A. M. (2013). Why has research in face recognition progressed so slowly? The importance of variability. *The Quarterly Journal of Experimental Psychology*, 66(8), 1467-1485.
- Burton, A. M., Jenkins, R., Hancock, P. J., & White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive Psychology*, 51(3), 256-284.

- 1 Burton, A. M., Kramer, R. S., Ritchie, K. L., & Jenkins, R. (2016). Identity from variation:
2 Representations of faces derived from multiple instances. *Cognitive Science*, 40(1), 202-
3 223.
- 4 Burton, A. M., White, D., & McNeill, A. (2010). The Glasgow face matching test. *Behavior*
5 *Research Methods*, 42(1), 286-291.
- 6 Chang, L., & Tsao, D. Y. (2017). The code for facial identity in the primate
7 brain. *Cell*, 169(6), 1013-1028.
- 8 Chiachia, G., Falcao, A. X., Pinto, N., Rocha, A., & Cox, D. (2014). Learning person-specific
9 representations from faces in the wild. *IEEE Transactions on Information Forensics and*
10 *Security*, 9(12), 2089-2099.
- 11 Chauhan, V., Kotlewska, I., Tang, S., & Gobbini, M. I. (2020). How familiarity warps
12 representation in the face space. *Journal of Vision*, 20(7), 18-18.
- 13 Colón, Y., Castillo, C. D., & O'Toole, A. (2020). Facial expression is retained in deep
14 networks trained for face identification. <https://doi.org/10.31234/osf.io/dphsv>
- 15 Cox, D. D. (2014). Do we understand high-level vision? *Current Opinion in*
16 *Neurobiology*, 25, 187-193.
- 17 Clutterbuck, R., & Johnston, R. A. (2005). Demonstrating how unfamiliar faces become
18 familiar using a face matching task. *European Journal of Cognitive Psychology*, 17(1),
19 97-116.
- 20 Dawel, A., Miller, E. J., Horsburgh, A., & Ford, P. (2021). A systematic survey of face
21 stimuli used in psychological research 2000-2020. *Behavior Research Methods*.
22 <https://doi.org/10.3758/s13428-021-01705-3>
- 23 Devue, C., Wride, A., & Grimshaw, G. M. (2019). New insights on real-world human face
24 recognition. *Journal of Experimental Psychology: General*, 148(6), 994.
- 25 Duchaine, B., & Nakayama, K. (2005). Dissociations of face and object recognition in
26 developmental prosopagnosia. *Journal of Cognitive Neuroscience*, 17(2), 249-261.
- 27 Dunn, J. D., Kemp, R. I., & White, D. (2018). Search templates that incorporate within-face
28 variation improve visual search for faces. *Cognitive Research: Principles and*
29 *Implications*, 3(1), 37.
- 30 Dunn, J. D., Ritchie, K. L., Kemp, R. I., & White, D. (2019). Familiarity does not inhibit
31 image-specific encoding of faces. *Journal of Experimental Psychology: Human*
32 *Perception and Performance*, 45(7), 841.
- 33 Dunn J. D., Summersby S., Towler A., Davis J. P., White D. (2020). UNSW Face Test: A
34 screening tool for super-recognizers. *PLoS ONE* 15(11): e0241747.
35 <https://doi.org/10.1371/journal.pone.0241747>
- 36 Dunn, J. D., Kemp, R. I., & White, D. (2021). Top-down influences on working memory
37 representations of faces: Evidence from dual-target visual search. *Quarterly Journal of*
38 *Experimental Psychology*, 74 (8), 1368-1377.
- 39 Ester, E. F., Sprague, T. C., & Serences, J. T. (2020). Categorical biases in human
40 occipitoparietal cortex. *Journal of Neuroscience*, 40(4), 917-931.
- 41 Faerber, S. J., Kaufmann, J. M., Leder, H., Martin, E. M., & Schweinberger, S. R. (2016).
42 The role of familiarity for representations in norm-based face space. *PLOS One*, 11(5),
43 e0155380.

- 1 Furl, N., Phillips, P. J., & O'Toole, A. J. (2002). Face recognition algorithms and the other
2 race effect: computational mechanisms for a developmental contact
3 hypothesis. *Cognitive Science*, 26(6), 797-815.
- 4 Gastgeb, H. Z., Wilkinson, D. A., Minshew, N. J., & Strauss, M. S. (2011). Can individuals
5 with autism abstract prototypes of natural faces? *Journal of Autism and Developmental*
6 *Disorders*, 41(12), 1609-1618.
- 7 Gerig, T., Morel-Forster, A., Blumer, C., Egger, B., Luthi, M., Schönborn, S., & Vetter, T.
8 (2018). Morphable face models-an open framework. In *13th IEEE International*
9 *Conference on Automatic Face & Gesture Recognition (FG 2018)*.
- 10 Gibson, J. J., & Gibson, E. J. (1955a). Perceptual learning: differentiation or enrichment?
11 *Psychological Review*, 62, 32-41.
- 12 Gibson, J. J., & Gibson, E. J. (1955b). What is learned in perceptual learning? A reply to
13 Professor Postman. *Psychological Review*, 62, 447-450.
- 14 Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through
15 category learning. *Cognition*, 78(1), 27-43.
- 16 Hautus, M. J., Macmillan, N. A., & Creelman, C. D. (2021). *Detection Theory: A User's*
17 *Guide (3rd Ed.)*. Abington, UK: Routledge.
- 18 Hill, M. Q., Parde, C. J., Castillo, C. D., Colón, Y. I., Ranjan, R., Chen, J. C., Blanz, V., &
19 O'Toole, A. J. (2019). Deep convolutional neural networks in the face of
20 caricature. *Nature Machine Intelligence*, 1(11), 522-529.
- 21 Ipser, A., Ring, M., Murphy, J., Gaigg, S. B., & Cook, R. (2016). Similar exemplar pooling
22 processes underlie the learning of facial identity and handwriting style: Evidence from
23 typical observers and individuals with Autism. *Neuropsychologia*, 85, 169-176.
- 24 Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of
25 the same face. *Cognition*, 121(3), 313-323.
- 26 Jiahui, G., Yang, H., & Duchaine, B. (2018). Developmental prosopagnosics have
27 widespread selectivity reductions across category-selective visual cortex. *Proceedings of*
28 *the National Academy of Sciences*, 115(28), E6418-E6427.
- 29 Johnston, R. A., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A
30 review. *Memory*, 17(5), 577-596.
- 31 Kantner, J., & Tanaka, J. W. (2012). Experience produces the atypicality bias in object
32 perception. *Perception*, 41(5), 556-568.
- 33 Kramer, R. S., Young, A. W., & Burton, A. M. (2018). Understanding face
34 familiarity. *Cognition*, 172, 46-58.
- 35 Laurence, S., & Mondloch, C. J. (2016). That's my teacher! Children's ability to recognize
36 personally familiar and unfamiliar faces improves with age. *Journal of Experimental*
37 *Child Psychology*, 143, 123-138.
- 38 Livingston, K. R., Andrews, J. K., & Harnad, S. (1998). Categorical perception effects
39 induced by category learning. *Journal of Experimental Psychology: Learning, Memory,*
40 *and Cognition*, 24(3), 732.
- 41 McKone, E., Martini, P., & Nakayama, K. (2001). Categorical perception of face identity in
42 noise isolates configural processing. *Journal of Experimental Psychology: Human*
43 *Perception and Performance*, 27(3), 573.

- 1 Menon, N., Kemp, R. I., & White, D. (2018). More than a sum of parts: Robust face
2 recognition by integrating variation. *Royal Society Open Science*, 5(5), 172381.
- 3 Mileva, M., Young, A. W., Jenkins, R., & Burton, A. M. (2020). Facial identity across the
4 lifespan. *Cognitive Psychology*, 116, 101260.
- 5 Murphy, J., Ipser, A., Gaigg, S. B., & Cook, R. (2015). Exemplar variance supports robust
6 learning of facial identity. *Journal of Experimental Psychology: Human Perception and*
7 *Performance*, 41(3), 577.
- 8 Natu, V., & O'Toole, A. J. (2011). The neural processing of familiar and unfamiliar faces: A
9 review and synopsis. *British Journal of Psychology*, 102(4), 726-747.
- 10 Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation
11 of the spatial envelope. *International Journal of Computer Vision*, 42(3), 145-175.
- 12 Op de Beeck, H., Wagemans, J., & Vogels, R. (2003). Asymmetries in stimulus comparisons
13 by monkey and man. *Current Biology*, 13(20), 1803-1808.
- 14 Parkhi O. M., Vedaldi A., & Zisserman, A. (2015) Deep face recognition. In X. Xie, M.
15 Jones, & G. Tam. *Proceedings of the British Machine Vision Conference*, , 1-12.
16 Available at www.bmva.org/bmvc/2015/index.html.
- 17 Phillips, P. J., Yates, A. N., Hu, Y., Hahn, C. A., Noyes, E., Jackson, K., ... & O'Toole, A. J.
18 (2018). Face recognition accuracy of forensic examiners, superrecognizers, and face
19 recognition algorithms. *Proceedings of the National Academy of Sciences*, 115(24),
20 6171-6176.
- 21 Ramon, M. (2021). Super-recognizers—A novel diagnostic framework, 70 cases, and
22 guidelines for future work. *Neuropsychologia*, 158, 107809.
- 23 Ramon, M., & Gobbini, M. I. (2018). Familiarity matters: A review on prioritized processing
24 of personally familiar faces. *Visual Cognition*, 26(3), 179-195.
- 25 Ramon, M., Dricot, L., & Rossion, B. (2010). Personally familiar faces are perceived
26 categorically in face-selective regions other than the fusiform face area. *European*
27 *Journal of Neuroscience*, 32(9), 1587-1598.
- 28 Rhodes, G., & Leopold, D. A. (2011). Adaptive norm-based coding of face identity. In A.J.
29 Calder, G. Rhodes, M.H. Johnson, & J.V. Haxby (Eds.), *The Oxford Handbook of Face*
30 *Perception* (pp. 263-286). Oxford: Oxford University Press.
- 31 Rice, G. E., Watson, D. M., Hartley, T., & Andrews, T. J. (2014). Low-level image properties
32 of visual objects predict patterns of neural response across category-selective regions of
33 the ventral visual pathway. *Journal of Neuroscience*, 34(26), 8837-8844.
- 34 Ritchie, K. L., Kramer, R. S., & Burton, A. M. (2018). What makes a face photo a 'good
35 likeness'?. *Cognition*, 170, 1-8.
- 36 Ritchie, K. L., & Burton, A. M. (2017). Learning faces from variability. *Quarterly Journal of*
37 *Experimental Psychology*, 70(5), 897-905.
- 38 Rossion, B. (2018). Humans are visual experts at unfamiliar face recognition. *Trends in*
39 *Cognitive Sciences*, 22(6), 471-472.
- 40 Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with
41 extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16(2), 252-257.
- 42 Shakeshaft, N. G., & Plomin, R. (2015). Genetic specificity of face recognition. *Proceedings*
43 *of the National Academy of Sciences*, 112(41), 12887-12892.

- 1 Tanaka, J., Giles, M., Kremen, S., & Simon, V. (1998). Mapping attractor fields in face
2 space: the atypicality bias in face recognition. *Cognition*, 68(3), 199-220.
- 3 Tanaka, J. W., & Corneille, O. (2007). Typicality effects in face and object perception:
4 Further evidence for the attractor field model. *Perception & Psychophysics*, 69(4), 619-
5 627.
- 6 Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race
7 in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 43(2),
8 161-204.
- 9 Valentine, T., Lewis, M. B., & Hills, P. J. (2016). Face-space: A unifying concept in face
10 recognition research. *The Quarterly Journal of Experimental Psychology*, 69(10), 1996-
11 2019.
- 12 Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural
13 facial categories. *Nature*, 428(6982), 557-561.
- 14 White, D., Burton, A. L., & Kemp, R. I. (2016). Not looking yourself: The cost of self
15 selecting photographs for face identification. *British Journal of Psychology*, 107, 359-
16 373.
- 17 White, D., Burton, A. M., Jenkins, R., & Kemp, R. I. (2014). Redesigning photo-ID to
18 improve unfamiliar face matching. *Journal of Experimental Psychology: Applied*, 20(2),
19 166-173.
- 20 White, D., Palermo, R., Rivolta, D., Al-Janabi, S. & Burton, A. M. (2017). Face matching
21 impairment in developmental prosopagnosia. *Quarterly Journal of Experimental*
22 *Psychology*, 70, 287-297.
- 23 [Dataset] White, D., Wayne, T. & Varela, V. P. L. (2021). Supporting image and rating data
24 for “Partitioning natural face image variability reveals precedence of within- over
25 between-face variation in explaining high accuracy in face recognition”. Available via:
26 https://github.com/UNSWfacelab/WhiteWayneVarela_PrototypeSimilarityData
- 27 Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010).
28 Controlling low-level image properties: The SHINE toolbox. *Behavior Research*
29 *Methods*, 42(3), 671-684.
- 30 Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., ... &
31 Duchaine, B. (2010). Human face recognition ability is specific and highly
32 heritable. *Proceedings of the National Academy of sciences*, 107(11), 5238-5241.
- 33 Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Gerbasi, M., & Nakayama, K.
34 (2012). Capturing specific abilities as a window into human individuality: The example
35 of face recognition. *Cognitive Neuropsychology*, 29(5-6), 360-392.
- 36 Wilmer, J. B. (2017). Individual differences in face recognition: A decade of
37 discovery. *Current Directions in Psychological Science*, 26(3), 225-230.
- 38 Young, A. W., & Burton, A. M. (2018a). Are we face experts? *Trends in Cognitive*
39 *Sciences*, 22(2), 100-110.
- 40 Young, A. W., & Burton, A. M. (2018b). What we see in unfamiliar faces: a response to
41 rosson. *Trends in cognitive sciences*, 472-473.

SUPPLEMENTARY MATERIALS

The supplementary materials document referred to in the main text contains: (i) names of the celebrities used as stimuli; (ii) supporting details of methodology used to select super-recognisers in Experiment 4; (iii) details of the computational analysis; (iv) details of a supplementary analysis that tested whether participants viewed identity averages as centroids of person-specific space. In addition, images and associated similarity rating data from all studies presented in this paper are accessible online via:

https://github.com/UNSWfacelab/WhiteWayneVarela_PrototypeSimilarityData
[dataset] (White, Wayne & Varela, 2021).

AUTHOR CONTRIBUTIONS

DW developed the study concept and design. DW, TW and VV performed experimentation, data collection and analysis. DW wrote the manuscript, TW and VV provided critical revisions.

ETHICS STATEMENT

This research was approved by the Deputy Vice-Chancellor (Research) of the University of New South Wales on the advice of the Human Research Ethics Advisory Panel.

ACKNOWLEDGEMENTS

This research was supported by an Australian Research Council Discovery Project grant (DP190100957), an Australian Research Council Future Fellowship (FT200100353), a UNSW Faculty of Science Research grant and a UNSW Scientia Fellowship to DW. The authors thank Nadia Menon for assistance in Experiment 1 and 2 data collection.