

This is a repository copy of *Are we face experts?*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/125753/>

Version: Accepted Version

Article:

Young, Andrew William orcid.org/0000-0002-1202-6297 and Burton, Anthony Michael orcid.org/0000-0002-2035-2084 (2018) Are we face experts? Trends in Cognitive Sciences. pp. 100-110. ISSN 1364-6613

<https://doi.org/10.1016/j.tics.2017.11.007>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

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.

Authors' version, as accepted for publication in *Trends in Cognitive Sciences*. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article.

DATE OF ACCEPTANCE: 28.11.17.

Are we face experts?

Andrew W. Young and A. Mike Burton

Department of Psychology, University of York, YO10 5DD, UK

Keywords: Face recognition, face perception, face identity, expertise

Contact details: andy.young@york.ac.uk and mike.burton@york.ac.uk

Acknowledgements: We are grateful to our colleagues Tim Andrews and Rob Jenkins for helpful discussion of various issues. AMB's research received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n.323262.

Abstract

According to a widely used theoretical perspective, our everyday experiences lead us to become natural experts at perceiving and recognising human faces. However, there has been considerable debate about this view. We discuss criteria for expertise and show how the debate over face expertise has often missed key points concerning the role and nature of face familiarity. For identity recognition, most of us show only limited expertise with unfamiliar faces. Carefully evaluating the senses in which it is appropriate or inappropriate to say that we are face experts leads to the conclusion that we are, in effect, familiar face experts.

The appeal of the idea of face expertise

We spend much of our lives looking at faces. The idea that we gain naturally acquired expertise at perceiving and recognising the kind of faces we encounter therefore has a powerful intuitive appeal. In fact it has become one of the most widely used theoretical perspectives in studies of face recognition. According to this perspective, acquired expertise makes us very good at recognising faces. Expertise has also been claimed to account for **other-race effects** (see Glossary) in face recognition, for the development of **holistic processing**, for **inversion effects**, and for similarities and differences between the recognition of faces and other types of visual stimuli as diverse as cars, dogs, birds, butterflies, chess positions, and **greebles** [1-8].

Nonetheless, there has been a highly polarised debate about the issues surrounding this interpretation, with strong opinions expressed on both sides. Our aim here is to show how this debate has often missed key points concerning expertise, and especially the role of face familiarity. Although we are usually good at recognising familiar faces, most of us experience substantial problems with recognising the identities of **unfamiliar faces**. We will discuss different criteria for expertise and outline recent experimental and computational studies that help in understanding the ways in which it is appropriate or inappropriate to say that we are face experts. Understanding the underlying determinants of expertise and their implications leads to a different resolution of the issues that offers a novel way forward.

The facial expertise debate

The idea of naturally acquired expertise in recognising faces is grounded in the observation that faces form a class of visual stimuli with high inter-item similarity that require **within-category recognition** at a subordinate level; the superordinate category being 'faces', and the subordinate categories the individual faces themselves. This idea received support from studies of **acquired prosopagnosia**

which had shown that inability to recognise familiar faces was often associated with problems in recognising other highly similar visual stimuli [9]. Because faces will so often need to be individuated in everyday life, the claim is that we become expert at doing this and that this expertise explains such well-known properties of face recognition as inversion effects and holistic processing [1-2].

This approach has been extended by introducing the greebles class of artificial stimuli to study perceptual expertise created in the laboratory [3]. Learning to individuate upright greebles can lead to inversion effects and enhanced holistic processing for greebles, and these effects may involve changes in brain regions typically considered to show face-specific responses [3, 5]. On this basis it has been claimed that what appear to be face-specific brain regions are better considered as regions dedicated to within-category recognition that have become 'automatized by expertise' [4-5].

This interpretation has been disputed by others who argue there are grounds for postulating an evolved neural substrate to face recognition [10-12]. For example, contrary to the hypothesis that faces make use of a generic ability to distinguish items within any visually homogeneous category, some cases of acquired prosopagnosia do seem to involve a face-specific impairment [13-14]. The debate has been passionate because so much appears to be at stake. In essence, it seems to involve whether category selectivity (e.g. for faces) is a fundamental organising principle of visual cortex, or simply a by-product of acquired expertise. We will show that while this debate has undoubtedly sharpened our thinking and led to interesting findings, both conceptions miss key points.

What characterises acquired expertise?

Surprisingly, what it might actually mean to say that we are face experts is seldom discussed. Instead, the idea is usually directly linked to an assumed operational definition in which expertise is indexed through better recognition of upright than inverted stimuli or through enhanced holistic processing [1-3].

However, the idea of expertise clearly has a variety of different connotations. Most obviously, becoming expert requires extensive practice and perhaps some form of careful study, but also the consequence of this will involve reaching a high level of skilled performance. Criteria for expertise could therefore include a mixture of at least three things. First, expert performance should be based on substantial experience. Second, expert performance should lead to accurate responses. Third, as for other acquired skills, expert performance should involve a relatively high degree of automaticity [5, 15-17].

Of course, there are well-known examples beyond the domain of face perception where the idea of acquired perceptual expertise clearly does apply, both in real contexts such as radiology and for expertise acquired in the lab (greebles). Radiologists, for example, have learnt through years of practice and training to quickly and accurately detect abnormalities in scans [18]. The question at issue, though, is whether these forms of perceptual learning are actually comparable to face recognition?

One obvious potential parallel between face recognition and these examples of acquired expertise with stimuli other than faces lies in our first criterion for expertise; the role of experience. The importance of experience with faces is evident in many infant studies, including work on perceptual narrowing showing that infants become less adept at remembering pictures of monkey faces as they become more experienced with human faces [19]. Moreover, the extent of this perceptual narrowing can be reduced if non-human faces are more strongly present in the infant's environment, which again shows how the faces that are seen can shape perceptual organisation [20]. As adults we remain susceptible to effects of previous experience, as shown for example by other-race effects [21-23] whose origin can again be traced back to experience in early years [24-27]. Although it is debated whether social psychological factors may also contribute to other-race effects [21, 28], there is a clear contribution from perceptual factors [21, 29]. Likewise, the constraining effects of previous experience are revealed in the fact that it is very

difficult for adults to learn new faces from seldom-encountered formats, such as inversion or **contrast negation**[30].

Substantial, lifelong perceptual experience is therefore undoubtedly important to face recognition, and in this sense the first of our criteria for face expertise does seem applicable. However, other criteria for expertise, including our second criterion of accurate performance, are not always met. Rather surprisingly, most of us are less good at some types of face recognition than we may like to think. In laboratory recognition memory tasks that involve trying to learn new faces, for example, any change between the studied and test images of a face creates a notable decrement in performance [31-32]; we seem to learn more about the specific *photographs* than the *faces* that were studied. Importantly, this is not simply a memory problem; the same difficulties arise in purely perceptual matching tasks with unfamiliar faces, as shown in Figure 1 [33-35].

----- FIGURE 1 ABOUT HERE, PLEASE -----

In the face matching task shown in Figure 1 (a), average error rates are substantial (typically 30% with unlimited viewing time) even though the images are fairly standardised [33]; the upper and lower images in Figure 1 (a) have much the same pose and expression and they were all taken on the same day, but with a different camera.

When a wider range of image differences in pose, expression, lighting and the date when the photograph was taken is introduced, most people do even worse. Try for yourself the task shown in Figure 1 (b), which simply involves deciding how many different people are present in the array of 40 photos. When the faces used are unfamiliar to participants, they arrive at solutions involving between 3 and 16 different identities, with nine identities being the most common number [34]. In other words, participants typically think there are around nine different individuals in the set of 40 photos. However, there are actually only two. Most people tend to mistake differences between the images for differences in identity, leading them to

overestimate the number of faces in the display. If the faces are of another race, the level of performance drops still further [22-23].

This observation that most of us can be confused by the variability between images of the same unfamiliar face has been developed into the widely used Glasgow Face Matching Test, illustrated in Figure 1 (c). The test shows a remarkable range of abilities across different individuals in the population, ranging from near-chance to near-perfect performance [35]. Furthermore, performance is highly stable; some people are consistently better at matching unfamiliar faces than others.

In sum, whereas the idea of generic naturally acquired expertise for face identity implies that we should find little difficulty in seeing whether photographs show same or different individuals, studies using unfamiliar faces show that, for most of us, this is far from the case.

Image variability is central to understanding the phenomena shown in Figure 1. Photographs of faces differ in many ways that include pose, expression, lighting, camera and lens characteristics. This variability can result from within-person variability (differences between different images of the same face) or between-person variability (differences between images of different faces). As Figure 1 shows, within-person variability can be substantial and for unfamiliar faces it is easily confused with between-person variability.

Importantly, just as the photographs in Figure 1, real-life views of faces are also highly variable; this is true whether the faces are seen in person, in videos, or photographs. So potentially we should all have abundant experience of the consequences of image variability to draw upon. Despite this, our perception and recognition of unfamiliar face identity often does not meet standards of accuracy that could justify our second criterion for the 'expert' label, being instead prone to error in many circumstances. Even with lifelong experience, image variability can limit our ability to perceive identity in unfamiliar faces.

The same findings arise in highly realistic conditions. For eye witness testimony, although witness confidence may well be correlated with correct identification of the suspect across a sample of cases [36], it has long been known that individual witnesses can nonetheless be mistaken even when they are certain that they are correct [37]. Even when no event memory is required, such as when matching someone's face to a presented passport or an identity card, most people make many errors [38-40], as described in Box 1.

----- BOX 1 ABOUT HERE, PLEASE -----

As well as failing to account for problems most of us experience with unfamiliar face identity and the surprisingly high incidences of misrecognition in applied and forensic settings (Figure 1 and Box 1), there are other substantial limitations to the applicability of the natural expertise hypothesis to face recognition. It cannot readily explain the large and stable individual differences in unfamiliar face recognition ability in the population [35]. Moreover, recent studies have shown significant heritability and genetic specificity of this ability, demonstrating influences beyond the role of naturally acquired expertise [43-46].

In fact, attempts to train or enhance face expertise mostly end in failure [40, 47-48]. While there are substantial inter-individual differences in unfamiliar face recognition ability, these are little affected by experience. For example, despite their training and extensive practice, passport officers and others who use face recognition throughout their working lives show the same variability in performance as the rest of the population [40]. Only forms of training involving unusual and task-specific strategies [42] have shown any benefit (see Box 1).

Nonetheless, despite these problems with perceiving individual face identity, there are things we can do relatively well with unfamiliar faces. Inferences concerning social characteristics such as gender, race and apparent age are reliable across different perceivers and often highly accurate [49-50]. Although also linked to individual differences in overall ability, interpretation of facial expressions elicits

substantial inter-observer agreement [51-53]. People even agree with each other when making inferences of dubious validity, such as deciding whether the face of someone they have never seen before looks trustworthy or intelligent [50, 54-56]. We need to understand how we can be good at perceiving some characteristics in unfamiliar faces when perception of unfamiliar face identity can be so problematic.

These social inferences often meet our second criterion of accurate (or at any rate, consensual) responses, suggesting that the idea of expertise for making them may have some merit, but with our third criterion of automaticity the evidence is mixed. The problems with perception of unfamiliar face identity already discussed make clear that this cannot be considered automatic. We are often thought to have become so expert at perceiving other social characteristics in faces that they will be seen automatically [57], and accurate or consensual judgements can be made even to relatively brief presentations that are in effect a single glance [58-59]. For example, first impressions of unfamiliar faces based on 100 ms presentation often do not differ from impressions based on unconstrained presentation times [58-59]. However, only a relatively limited form of automaticity is shown if more stringent criteria are applied [16, 60]. For example, although people are very good at judging the gender or race of a briefly presented unfamiliar face seen in isolation, their accuracy drops sharply when asked to match pairs of faces on these dimensions and they don't know in advance whether they will be asked to make the match on gender or on race before the trial begins [60]. This drop in performance due to uncertainty about what to look for would not be expected if both gender and race are simultaneously and effortlessly perceived. It shows a limit to automaticity in perceiving more than just the identity of unfamiliar faces.

In sum, on balance there are grounds for claiming some form of generic face expertise based on our first criterion of lifetime experience of faces, but with more demanding criteria these grounds are stronger for the aspects of unfamiliar face perception that do not involve recognition of identity.

The key role of face familiarity

Findings of the type summarised in Figure 1 and Box 1 all involve unfamiliar faces. For familiar faces, matters are completely different. Recognising familiar faces is of primary importance to humans, since it forms a key means of identifying a known individual and retrieving appropriate person-specific knowledge to support interaction based on past experience [49, 61-64]. Familiar face recognition is therefore embedded in a broader context of retrieval of semantic and emotional information that gives familiar faces rich layers of meaning [64-69].

These rich semantic and emotional properties of familiar faces may in themselves account for some of the differences between familiar and unfamiliar faces [70-71], but there is also a strong visual component. Although we occasionally make errors in recognising familiar faces in everyday life [72-73], recognition is usually fast and accurate across huge variations in image properties, even for degraded or briefly presented images [74-75]. Matching tasks with familiar faces seem to most participants almost trivially easy [34, 76] and familiar face recognition can meet criteria for automaticity [60, 77-78]. We seem almost not to notice the sheer scale of variability in our encounters with familiar faces, and it may be this that misleads us into expecting that we will also be good at recognising unfamiliar face identity [60]. Although there may well be individual differences in ability to recognise familiar faces, even those at the bottom of the range of ability in matching or learning unfamiliar faces (for example, in **congenital prosopagnosia**) may still be able to recognise a number of familiar faces [79-81]. Moreover, covert responses in acquired prosopagnosia show that some aspects of familiar face recognition can be remarkably resilient [65, 82-85].

Understanding how view-invariant recognition of familiar faces is achieved, and why recognition of unfamiliar faces across equivalent image changes is relatively poor, are therefore key theoretical tasks [61-62, 75]. Significant progress has been made in recent work.

----- BOX 2 ABOUT HERE, PLEASE -----

Box 2 summarises recent work that shows that familiar face-based expertise is necessarily idiosyncratic to each of the faces we know well [86]. By studying the variability across different images of the same face (for example, the statistical range of images of Angela Merkel across different poses, lighting and facial expressions) it has become clear that the ways in which one person's face varies are different from the ways in which someone else's face will vary. To recognise Angela Merkel from any image of her, then, our brains need to have learnt how to take into account this idiosyncratic, Merkel-specific variability. This means that it is an overstatement to claim people are experts at recognising all faces. Instead, it is more appropriate to say we are experts at recognising each of the faces we know; Angela Merkel, Tom Hanks, and so on. In fact our ability to learn the characteristics of a familiar face seems to surmount some of the limitations that are evident with unfamiliar faces, such as the other-race effect; there is no obvious decrement for recognition of other-race faces that have become familiar [87-88].

Interestingly, it turns out that being expert at recognising even just a handful of familiar faces across highly variable images is sufficient to create representations that can reliably distinguish social categories of gender and race. A standard computer engineering approach that can be used to train a computer model to identify highly varied images of different faces involves using PCA to capture the main dimensions of variability across the images and **LDA** to then cluster together images of the same face identity; in effect making these trained faces 'familiar' to the model [89]. A trained model shows properties analogous to human face recognition in that it performs well at recognising entirely new and highly variable images of familiar (trained) faces and performs less well at establishing whether different images of unfamiliar (i.e. untrained) faces are of the same person. Importantly, the space derived to cluster together disparate images of the faces of only a small number of people can also code sex and race very reliably [89]. Dimensions corresponding to these social categorisations simply emerge as a side-effect of solving the identification problem, despite the fact that neither sex nor race

are explicitly coded in the modeling process. Remarkably, these dimensions generalise to unfamiliar faces too – allowing very accurate categorization of large numbers of faces, never previously encountered.

Learning to distinguish the gender and race of any face can therefore be an incidental consequence of expertise in recognising a small number of familiar faces. It seems that the cues needed to distinguish gender or race covary with those needed to recognise familiar identities from the very different views we can encounter. This is interesting because learning to recognise a small number of familiar individuals closely approximates the task facing human infants in many societies and any evolved neural substrate for face recognition ability will have arisen within the context of relatively small social groups [90]. In contrast, a model trained only to classify gender offers no benefit to classifying identity [89]. What appear to be expert abilities in social categorisation may thus arise as an incidental consequence of expert recognition of personally familiar faces. This shows how what might seem like complex patterns of expertise can emerge from simple processes. Here we have a simple clustering model which leads to some (familiar) faces being well-recognised and other (unfamiliar) faces being poorly recognised while all can be categorised well on the untrained dimensions of sex and race.

Concluding remarks

Our message is that face expertise takes a different form than is usually thought. While there is undoubtedly a sense in which our huge experience of looking at the faces around us has created a kind of expertise, this does not take the form of the generic ability to recognise any face identity that has been so widely assumed. We do not of course seek to deny that experience with faces may have much to do with inversion effects, holistic processing and the like. Neither do we deny that there are aspects of unfamiliar face perception that are accurate, reliable or even close to automatic (see Outstanding Questions). However, recognition of unfamiliar face identity is not one of these. Although clearly shaped by experience in some ways, as shown for example by other-race effects, unfamiliar face recognition remains

generally vulnerable to the impact of the enormous variability in everyday images of faces.

To cope with the impact of this everyday image variability, our face recognition systems need to create representations that can deal with the different ways in which different known faces can vary. They are therefore driven to create an individually-tailored type of expertise for recognising each of the faces we know, and this explains why unfamiliar face identity can be problematic.

In effect, we are familiar face experts. This insight reframes the expertise debate in a way that has important implications for psychological theory and for a wide range of practical contexts such as passport control and witness testimony.

Glossary

Acquired prosopagnosia: An almost complete inability to recognise familiar faces following brain injury that cannot be explained by more general visual or intellectual difficulties. Usually, even the most familiar faces are not recognised. Studies of acquired prosopagnosia following brain injury have been very influential.

Congenital prosopagnosia: Problems in recognising faces that are present throughout life and are not linked to obvious neurological symptoms. There has been debate as to whether congenital prosopagnosia represents a distinct syndrome or the lower end of the continuum of ability in the population. Most studies define congenital prosopagnosia in terms of poor ability on tests of face matching or unfamiliar face learning. Familiar face recognition may also be weak, but in nearly all cases some familiar faces can still be recognised.

Contrast negation: Reverses the brightness values in an image, making light regions dark and dark regions light. This makes faces very hard to recognise.

Greebles: A class of artificial stimuli intended to create a category of novel exemplars that share a number of face-like properties [3].

Holistic processing: Refers to the idea that faces are processed as a whole, or Gestalt [2, 91].

Inversion effects: Comparisons of ability to recognise stimuli presented in the normally encountered upright orientation or presented upside-down. Face recognition usually shows a substantial inversion effect, with poorer performance for inverted faces.

LDA: Linear Discriminant Analysis. This statistical technique fits exemplars (here faces) to a space in which intra-class differences between images of the same face are minimised, while inter-class differences between images of different faces are maximised, so that images of the same person's face are clustered together. This is a technique which has been used in many computer models of face recognition.

PCA: Principal Components Analysis. A statistical technique that can reduce the dimensionality of a large data set to a relatively small number of principal components (PCs) that capture the main forms of variability in the data. It is often applied to face photographs.

Other-race effects: Refers to findings that most of us are better at learning and recognising unfamiliar faces that come from the own-race ethnic group with which we are most familiar. Similarly, we also show other-race effects for recognising facial expressions.

Unfamiliar faces: Faces of people we have not encountered before, or encountered very little. Understanding the perception and recognition of unfamiliar faces is critical to everyday concerns such as eyewitness testimony and passport control.

Within-category recognition: The idea that a type of visual stimulus has a sufficiently homogeneous overall appearance that we need to distinguish between items within the overall category to recognise individual exemplars. Faces, cars, and flowers are commonly used examples where within-category recognition is said to be needed.

References

1. Diamond, R. and Carey, S. (1986) Why faces are and are not special: an effect of expertise. *J Exp Psychol Gen*, 115, 107-117
2. Carey, S. (1992) Becoming a face expert. *Philos Trans R Soc Lond B Biol Sci*, 335, 95-103
3. Gauthier, I. and Tarr, M.J. (1997) Becoming a "greeble" expert: exploring mechanisms for face recognition. *Vision Res*, 37, 1673-1682
4. Gauthier, I. et al. (2000) Expertise for cars and birds recruits brain areas involved in face recognition. *Nat Neurosci*, 3, 191-197
5. Tarr, M.J. and Gauthier, I. (2000) FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nat Neurosci*, 3, 764-769
6. Tanaka, J.W. (2001) The entry point of face recognition: evidence for face expertise. *J Exp Psychol Gen*, 130, 534-543
7. Rhodes, G. et al. (2006) Expert face coding: configural and component coding of own-race and other-race faces. *Psychon Bull Rev*, 13, 499-505
8. McKone, E. et al. (2007) Can generic expertise explain special processing for faces? *Trends Cogn Sci*, 11, 8-15
9. Damasio, A.R. et al. (1982) Prosopagnosia: anatomic basis and behavioral mechanisms. *Neurol*, 32, 331-341
10. Kanwisher, N. (2000) Domain specificity in face perception. *Nat Neurosci*, 3, 759-763
11. Grill-Spector, K. et al. (2004) The fusiform face area subserves face perception, not generic within-category identification. *Nat Neurosci*, 7, 555-562
12. Kanwisher, N. and Yovel, G. (2006) The fusiform face area: a cortical region specialized for the perception of faces. *Philos Trans R Soc Lond B Biol Sci*, 361, 2109-2128
13. Busigny, T. et al. (2010). Acquired prosopagnosia as a face-specific disorder: ruling out the general visual similarity account. *Neuropsychologia*, 48, 2051-2067
14. Riddoch, M.J. et al. (2008) Are faces special? A case of pure prosopagnosia. *Cogn Neuropsychol*, 25, 3-26

15. Shiffrin, R.M. and Schneider, W. (1977) Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychol Rev*, 84, 127–189
16. Palermo, R. and Rhodes, G. (2007) Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia*, 45, 75-92
17. Moors, A. (2016) Automaticity: componential, causal, and mechanistic explanations. *Annu Rev Psychol*, 67, 263-287
18. Evans, K.K. et al. (2016) A half-second glimpse often lets radiologists identify breast cancer cases even when viewing the mammogram of the opposite breast. *Proc Natl Acad Sci U S A*, 113, 10292-10297
19. Pascalis, O. et al. (2002) Is face processing species-specific during the first year of life? *Science*, 296, 1321-1323
20. Pascalis, O. et al. (2005) Plasticity of face processing in infancy. *Proc Natl Acad Sci U S A*, 102, 5297-5300
21. Rossion, B. and Michel, C. (2011) An experience-based holistic account of the other-race effect. In *The Oxford Handbook of Face Perception* (A.J. Calder et al., eds.), pp. 215-243. Oxford: Oxford University Press.
22. Laurence, S. et al. (2016) The flip side of the other-race coin: they all look different to me. *Br J Psychol*, 107, 374-388
23. Yan, X. et al. (2016) Cross-cultural differences and similarities underlying other-race effects for facial identity and expression. *Q J Exp Psychol*, 69, 1247-1254
24. Tham, D.S.Y. et al. (2015) In infancy the timing of emergence of the other-race effect is dependent on face gender. *Infant Behav Dev*, 40, 131-138
25. Tham, D.S.Y. et al. (2017) The other-race effect in children from a multiracial population: a cross-cultural comparison. *J Exp Child Psychol*, 155, 128-137
26. Lee, K. et al. (2017) Face race processing and racial bias in early development: a perceptual-social linkage. *Curr Dir Psychol Sci*, 26, 256-262
27. Macchi Cassia, V. et al. (2014) How race and age experiences shape young children's face processing abilities. *J Exp Child Psychol*, 120, 87-101
28. Hugenberg, K. et al. (2011) Social categorization influences face perception and face memory. In *The Oxford Handbook of Face Perception* (A.J. Calder et al., eds.), pp. 245-261. Oxford: Oxford University Press.

29. Crookes, K. and Rhodes, G. (2017) Poor recognition of other-race faces cannot always be explained by a lack of effort. *Vis Cogn*, DOI:10.1080/13506285.2017.1311974
30. Kramer, R.S.S. et al. (2017) Natural variability is essential to learning new faces. *Vis Cogn*, DOI:10.1080/13506285.2016.1242522
31. Bruce, V. (1982) Changing faces: visual and non-visual coding processes in face recognition. *Br J Psychol*, 73, 105-116
32. Longmore, C.A. et al. (2008) Learning faces from photographs. *J Exp Psychol Hum Percept Perform*, 34, 77-100
33. Bruce, V. et al. (1999) Verification of face identities from images captured on video. *J Exp Psychol Appl*, 5, 339-360
34. Jenkins, R. et al. (2011) Variability in photos of the same face. *Cognition*, 121, 313-323
35. Burton, A.M. et al. (2010) The Glasgow Face Matching Test. *Behav Res Methods*, 42, 286-291
36. Wixted, J.T. et al. (2016) Estimating the reliability of eyewitness identifications from police lineups. *Proc Natl Acad Sci U S A*, 113, 304-309
37. Davies, G.M. (1996) Mistaken identification: where law meets psychology head on. *Howard J Crim Justice*, 35, 232-241
38. Megreya, A.M. and Burton, A.M. (2008) Matching faces to photographs: poor performance in eyewitness memory (without the memory). *J Exp Psychol Appl*, 14, 364-72
39. Kemp, R. et al. (1997) When seeing should not be believing: photographs, credit cards and fraud. *Appl Cogn Psychol*, 11, 211-222
40. White, D. et al. (2014) Passport officers' errors in face matching. *PLoS One*, 9(8), e103510
41. Robertson, D.J. et al. 2016 Face recognition by Metropolitan Police super-recognisers. *PLoS ONE*, 11(2), e0150036
42. White, D. et al. (2015) Perceptual expertise in forensic facial image comparison. *Proc R Soc*, B282, 20151292
43. Wilmer, J.B. et al. (2010) Human face recognition ability is specific and highly heritable. *Proc Natl Acad Sci U S A*, 107, 5238-5241

44. Zhu, Q. et al. (2010) Heritability of the specific cognitive ability of face perception. *Curr Biol*, 20, 137-142
45. Verhallen, R.J. et al. (2014) An online version of the Mooney Face Test: phenotypic and genetic associations. *Neuropsychologia*, 63, 19-25
46. Shakeshaft, N.G. and Plomin, R. (2015) Genetic specificity of face recognition. *Proc Natl Acad Sci U S A*, 112, 12887-12892
47. Towler, A. et al. (2014) Evaluating training methods for facial image comparison: the face shape strategy does not work. *Perception*, 43, 214-218
48. Towler, A. et al. (2017) Evaluating the feature comparison strategy for forensic face identification. *J Exp Psychol Appl*, 23, 47-58
49. Bruce, V. and Young, A. (2012) *Face Perception*. Hove, East Sussex: Psychology Press.
50. Todorov, A. (2017) *Face Value: the Irresistible Influence of First Impressions*. Princeton, NJ: Princeton University Press.
51. Ekman, P. (1972) Universals and cultural differences in facial expressions of emotion. In *Nebraska Symposium on Motivation*, 1971 (Cole, J., ed), pp. 207-283. Lincoln, Nebraska: University of Nebraska Press.
52. Biehl, M. et al. (1997) Matsumoto and Ekman's Japanese and Caucasian Facial Expressions of Emotion (JACFEE): reliability data and cross-national differences. *J Nonverbal Behav*, 21, 3-21
53. Yan, X et al. (2016) Cultural similarities and differences in perceiving and recognizing facial expressions of basic emotions. *J Exp Psychol Hum Percept Perform*, 42, 423-440
54. Santos, I.M. and Young, A.W. (2005) Exploring the perception of social characteristics in faces using the isolation effect. *Vis Cogn*, 12, 213-247
55. Oosterhof, N.N. and Todorov, A. (2008) The functional basis of face evaluation. *Proc Natl Acad Sci U S A*, 105, 11087-11092.
56. Sutherland, C.A.M. et al. (2013) Social inferences from faces: ambient images generate a three-dimensional model. *Cognition*, 127, 105-118
57. Martin, D. et al. (2015) The simultaneous extraction of multiple social categories from unfamiliar faces. *J Exp Soc Psychol* 60, 51-58

58. Willis, J. and Todorov, A. (2006) First impressions: Making up your mind after a 100-ms exposure to a face. *Psychol Sci*, 17, 592-598
59. South Palomares, J.K. and Young, A.W. (2017) Facial first impressions of partner preference traits: trustworthiness, status, and attractiveness. *Soc Psychol Personal Sci*, DOI:10.1177/1948550617732388.
60. Yan, X. et al. (2017) The automaticity of face perception is influenced by familiarity. *Atten Percep Psychophys*, 79, 2202-2211
61. Bruce, V. and Young, A. (1986) Understanding face recognition. *Br J Psychol*, 77, 305-327
62. Young, A.W. and Burton, A.M. (2017) Recognizing faces. *Curr Dir Psychol Sci* 26, 212-217
63. Yovel, G. and O'Toole, A.J. (2016) Recognizing people in motion. *Trends Cogn Sci*, 20, 383-395
64. Young, A.W. and Bruce, V. (2011) Understanding person perception. *Br J Psychol*, 102, 959-974
65. Bauer, R.M. (1984). Autonomic recognition of names and faces in prosopagnosia: a neuropsychological application of the Guilty Knowledge Test. *Neuropsychologia*, 22, 457-469
66. Ellis, H.D. et al. (1997) Reduced autonomic responses to faces in Capgras delusion. *Proc R Soc*, B264, 1085-1092
67. Gobbini, M.I. and Haxby, J.V. (2007) Neural systems for recognition of familiar faces. *Neuropsychologia*, 45, 32-41
68. Guntupalli, J.S. and Gobbini, M.I. (2017) Reading faces: from features to recognition. *Trends Cogn Sci*, <http://dx.doi.org/10.1016/j.tics.2017.09.007>
69. Natu, V. and O'Toole, A. J. (2011). The neural processing of familiar and unfamiliar faces: a review and synopsis. *Br J Psychol*, 102, 726-747.
70. Bonner et al. (2003) Meet the Simpsons: top-down effects in face learning. *Perception*, 32, 1159-1168
71. Schwartz, L. and Yovel, G. (2016). The roles of perceptual and conceptual information in face recognition. *J Exp Psychol Gen*, 145, 1493-1511
72. Young, A.W. et al. (1985) The faces that launched a thousand slips: everyday difficulties and errors in recognizing people. *Br J Psychol*, 76, 495-523

73. Hay, D.C. et al. (1991) Routes through the face recognition system. *Q J Exp Psychol*, 43A, 761-791
74. Burton, A.M. et al. (2005) Robust representations for face recognition: the power of averages. *Cogn Psychol*, 51, 256-284
75. Burton, A.M. (2013) Why has research in face recognition progressed so slowly? The importance of variability. *Q J Exp Psychol*, 66, 1467-1485
76. Ritchie, K.L. et al. (2015) Viewers base estimates of face matching accuracy on their own familiarity: explaining the photo-ID paradox. *Cognition*, 141, 161-169
77. Young, A.W. et al. (1986). Face-name interference. *J Exp Psychol Hum Percept Perform*, 12, 466-475
78. Johnston. P. et al. (2016) Expectations about person identity modulate the face-sensitive N170. *Cortex*, 85, 54-64
79. Behrmann, M. and Avidan, G. (2005) Congenital prosopagnosia: face-blind from birth. *Trends Cogn Sci*, 9, 180-187
80. Barton, J.J.S. and Corrow, S.L. (2016) The problem of being bad at faces. *Neuropsychologia*, 89, 119-124
81. Bate, S. and Tree, J.J. (2017) The definition and diagnosis of developmental prosopagnosia. *Q J Exp Psychol*, 70, 193-200
82. De Haan, E.H.F. et al. (1987) Faces interfere with name classification in a prosopagnosic patient. *Cortex*, 23, 309-316
83. De Haan, E.H.F. et al. (1987) Face recognition without awareness. *Cogn Neuropsychol*, 4, 385-415
84. Young, A.W. et al. (1988) Cross-domain semantic priming in normal subjects and a prosopagnosic patient. *Q J Exp Psychol*, 40A, 561-580
85. Young, A.W. and Burton, A.M. (1999) Simulating face recognition: implications for modelling cognition. *Cogn Neuropsychol*, 16, 1-48
86. Burton, A.M. et al. (2016) Identity from variation: representations of faces derived from multiple instances. *Cogn Sci*, 40, 202-223
87. McKone, E. et al. (2007) Familiar other-race faces show normal holistic processing and are robust to perceptual stress. *Perception*, 36, 224-248

88. Zhou, X. and Mondloch, C.J. (2016) Recognizing "Bella Swan" and "Hermione Granger": no own-race advantage in recognizing photos of famous faces. *Perception*, 45, 1426-1429
89. Kramer, R.S.S. et al. (2017) Robust social categorization emerges from learning the identities of very few faces. *Psychol Rev*, 124, 115-129
90. Layton, R. et al. (2012) Antiquity and social functions of multilevel social organization among human hunter-gatherers. *Int J Primatol*, 33, 1215–1245
91. Maurer, D. et al. (2002) The many faces of configural processing. *Trends Cogn Sci*, 6, 255-260

TRENDS BOX**Trends**

- There is a wide range of ability to recognise the identities of unfamiliar faces in the population.
- This variability is influenced by genes but does not seem amenable to training. Even having a job that requires matching unfamiliar faces does not lead to much improvement.
- Images of faces seen in everyday life are highly variable, and much of the problem many of us experience with unfamiliar face identity comes from not being able to determine whether variability is image-related or identity-related.
- Although failures of familiar face recognition do sometimes occur, our ability to recognise familiar faces is mostly excellent, able to cope with degraded images and largely unaffected by image-related differences.
- Familiar face recognition meets broad criteria for expertise, but recognition of unfamiliar faces is expert only in the restricted sense that it is influenced by experience.

OUTSTANDING QUESTIONS BOX

Outstanding Questions

- Why are aspects of unfamiliar face perception other than identity recognition often accurate and/or reliable? Have we acquired specific expertise for these various tasks, or are some of these abilities incidental consequences of a system that is primarily oriented towards familiar face recognition?
- Although most of us are not very good at matching or recognising unfamiliar faces, a minority of individuals seem to be 'super-recognisers'. Are these super-recognisers simply the top end of the normal distribution of ability? What exactly is it that enables them to discount or disregard image differences that confuse most people? Does being a face super-recogniser entail any cost in some other type of perceptual task?
- If our recognition of unfamiliar faces is not, on average particularly good, why is it nonetheless the case that inversion effects and holistic processing are so pervasively found in studies of unfamiliar face recognition? What does experience with faces contribute to this?
- Is congenital prosopagnosia a qualitatively distinct type of problem, or just a way of labelling the bottom end of the normal distribution of ability? What implications does very poor ability to match and learn unfamiliar faces have for familiar face recognition? Most people described as suffering from congenital prosopagnosia offer anecdotal reports of everyday problems with familiar faces, yet they can usually still recognise a number of highly familiar faces. Is this number of recognisable familiar faces lower than might otherwise be expected, is the range of image variability that can be coped with reduced, or do people with congenital prosopagnosia simply need much more exposure before a face becomes familiar?
- To what extent do familiar and unfamiliar face perception rely on qualitatively different processes? Are unfamiliar faces just harder to recognise, or are they coded in a highly pictorial fashion, as some researchers have claimed?

Box 1. Face matching in everyday life.

The tasks shown in Figure 1 might appear to be laboratory bound, and so the question arises whether unfamiliar face matching is a problem in real life. For example, checking someone's photo-ID involves interacting with a live, moving, three-dimensional individual, not just comparing two images. Despite the widespread requirement to use photos to prove our identity (in passports, drivers' licences and other documents), it turns out that the same problems arise for photo identification in everyday life.

In a study which copied the general design of the face matching task shown in Figure 1 (a), the target was sometimes a real person, to be matched against photos. This did not make the problem easier – viewers made just as many errors matching to a real person as matching to a photo [38].

In fact, this problem has been shown to persist even in highly natural settings. In a study conducted in a supermarket, using specially created identity cards that showed either a valid photograph (i.e., an actual photo of the person) or an invalid photograph (a photo of another person), the supermarket's own cashiers were told to look out for the fake identity cards and told that a bonus payment would be made to the cashier with the best performance [39]. Yet the cashiers only challenged about 10% of people presenting a valid card and accepted many of the invalid cards if there was some similarity in appearance. Thus, even comparing a photo of a face to the person standing in front of you is not as easy as we usually take it to be.

It turns out that unfamiliar face matching can also be a problem for professional groups that rely on this ability. For example, working passport officers make errors when asked to verify passport photos of volunteers during an identity check [40]. Just like members of other populations, the passport officers display large individual differences, with some performing extremely well (almost perfectly) but others much more poorly. What is interesting is that these differences seem unrelated to experience; officers with as much as 20 years experience can still perform poorly

[40]. Despite the fact that experienced officers have lengthy training histories, there are just as many high and low scores from people with all levels of experience.

The difficulty of unfamiliar face matching in operational settings is beginning to be addressed in some organisations. One approach is used by the London Metropolitan Police, who select a small number of officers with high natural ability for a 'super-recogniser' unit, specializing in identification. This group reliably outperforms other officers and the general population on unfamiliar face identification tasks [41].

An alternative to recruiting individuals with high aptitude is to provide highly specific training. Professional forensic face identification examiners, who are trained in how to make detailed comparisons across photographs, have also been studied [42]. They can outperform untrained participants and computer algorithms, but make use of careful feature comparison strategies that require substantial time and are little affected by inversion. In other words, forensic examiners have been trained to use methods quite unlike those used in normal face recognition. While they have been shown to do well on unfamiliar face identification tasks, they remain far from perfect and their performance doesn't reach levels which are commonplace in familiar face recognition.

Box 2. Face variability is idiosyncratic.

Different images of the same face can vary due to viewpoint, expression, lighting and other properties. Given this substantial variability, how can we characterise the invariant properties of each face that make recognition possible? One approach is to analyse how images of the same face vary by applying PCA to everyday photographs of the same person [86]; PCA is a statistical technique that can reduce the dimensionality of photographs to a relatively small number of principal components (PCs). This is a widely used technique in the computer science literature, but it is mostly applied to images of different faces. That is, researchers usually run PCA across images of many different faces to establish dimensions on which faces differ from one another. In order to avoid coding superficial differences, the images used in PCA are usually photographed under standard conditions, to eliminate changes in pose, expression, lighting and so on.

However, instead of asking what image properties can distinguish between all faces, as in standard applications of PCA, we can also ask what properties characterise the highly variable images that correspond to a specific face identity? By analyzing the statistics of naturally varying images, sampled from internet search and deliberately not controlled for low level characteristics, it is possible to describe the way that images of one person vary, and to compare this to the pattern for a different person [86].

It turns out that different faces have different dimensions of variation. Put simply, the ways in which one person's face varies from image to image is different from the ways in which someone else's face will vary. The characteristics that vary or remain relatively consistent across images are idiosyncratic; they differ between one person and another.

The implication is that we have to learn separately the relatively variant and invariant characteristics of each of the faces we know, because the variability across images is to some extent identity-specific [86]. We need to learn which

characteristics of a particular face are relatively consistent and which are variable. This explains why unfamiliar face matching is so hard. Confronted with two images of an unfamiliar person, we often do not know whether or not the images fall within the range of variation for that person, making it hard to judge whether they are the same individual. For familiar people, we have learned how their faces can change, making a matching decision much easier.

FIGURE

Figure 1: Examples of how tricky it can be to recognise the identities of unfamiliar faces, showing that this is in large part a perceptual problem.

Figure 1 (a): Is the top person present in the lower array, or are they missing?

Average error rates are substantial, yet there is no event-memory component to the task [33].

Figure 1 (b): Can you sort the 40 images into the different face identities? If the faces are unfamiliar, most people overestimate the number of identities present [34].

Figure 1 (c). Two example pairs of images from the Glasgow Face Matching Test [35].

This test simply asks people to decide whether or not two simultaneously presented images show the same person. The top row shows two images of different identities, while the bottom row illustrates a 'same identity' image pair. Note that all faces in the GFMT are unfamiliar and that all test items involve pairs of photographs with substantial superficial differences.

Figure 1 (a)




Figure 1 (b)



Figure 1 (c)

