Understanding professional expertise in unfamiliar face matching

David White, Alice Towler & Richard I. Kemp
School of Psychology
UNSW Sydney

*Forthcoming, to appear in:
M. Bindemann (Ed.), Forensic face matching: Research and practice. Oxford
University Press.*

Word Count (not including Abstract, References or Figure captions): 11,622

Author Note

Preparation of this chapter was supported by an Australian Research Council Linkage Project grants to White and Kemp (LP160101523), and an Australian Research Council Discovery Project grant to White (DP190100957). The authors thank Megan Papesh, Rebecca Heyer, Josh Davis and Fredrick Eckloff for providing additional information about the studies reviewed in this chapter, Gary Edmond for providing comments on an early version of the manuscript and Janice Yung for her support with finalising the manuscript.

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

Abstract (149 words)

Deciding whether or not two images are of the same unfamiliar face is an important task in many professions. These decisions are a critical part of modern identity verification processes with direct – and often profound – consequences for individual rights and the security of society. As a result, the public expect the people entrusted with these decisions to perform accurately. But do they? Here we review 29 published tests comparing face matching accuracy in professional and novice groups. Twelve of these tests show no significant differences between professional and novice groups, suggesting that merely performing the task in daily work is not sufficient to improve accuracy. However, specialist groups of Facial Examiners and Police Super-recognisers consistently outperform novices. Staff selection, mentorship, deliberate practice, motivation, feedback and training may all contribute to enhancing the performance of individuals working in these groups and future research is necessary to delineate their relative contributions.

Keywords

face perception; perceptual expertise; facial image comparison; super-recognizers; unfamiliar face matching; expertise; face recognition; learning; visual attention; visual processing.

EXPERTISE IN UNFAMILIAR FACE MATCHING

Deciding whether images of unfamiliar faces are of the same person or different people has been studied extensively in laboratory research since the early 1990s. Yet it is only very recently that researchers have begun to test the accuracy of professional experts on this task. The task is shown in Figure 1, and on first inspection it appears straightforward. We routinely identify faces in our everyday lives and so the task is superficially familiar, and the format of the task appears quite simple. It requires comparing images presented at the same time, and so does not rely on memory, giving the impression that this is an easy task.

The most important insight from laboratory studies of novice participants is that this initial impression is misguided. Unfamiliar face matching is substantially more difficult than it initially appears. When novice participants are given unlimited time to complete these tasks, error rates in pairwise face matching tasks range from around 20% with high quality standardised images (e.g. Bruce, Henderson, Newman, & Burton, 2001; Burton, White, & McNeill, 2010) to 30% or 40% in more challenging tests where images are captured in unconstrained environmental conditions (e.g. Davis & Valentine, 2009; Henderson, Bruce, & Burton, 2001; Phillips, Yates, et al., 2018). These high error rates are observed even when participants are matching images to people who are present in the room (Davis & Valentine, 2009), as is the case when checking Photo-ID documents (Kemp, Towell & Pike, 1997). Error rates are also higher when participants work under time pressure (Fysh & Bindemann, 2017), experience anxiety (Attwood, Penton-Voak, Burton & Munafò, 2013) or lack of sleep (Beattie, Walsh, McLaren, Biello, & White, 2016).

Figure 1: An example test item from the Glasgow Face Matching Test. This task requires participants to decide whether pairs of images show the same face or different faces. Despite the images being taken on the same day, with neutral expression, standardised pose and lighting conditions, novice participants make 20% errors on this task. The correct answer is that these images are of the same person.

Errors in face matching are not only of academic interest to cognitive psychologists. When these errors are made outside of the lab, they can have serious consequences. Face matching decisions are performed by many thousands of public and private sector staff across the world each day, such as passport officers, police officers, bank clerks, surveillance system operators, security personnel and retail workers. Identity verification decisions in these professional settings often require staff to compare two or more images of unfamiliar faces and decide if they show the same person or different people. For example when passport officers compare a traveller to their passport image, or when police officers compare CCTV evidence to mugshot images. As a result, the security of important identity management systems underpinning financial transactions, interactions with government, criminal investigations and legal processes all rely on the accuracy of unfamiliar face matching decisions.

Knowing whether the staff involved in these decisions are susceptible to the high levels of error observed in novices is therefore a critical question with immediate relevance to society. In professional settings, the task experimental psychologists refer to as ‘unfamiliar face matching’¹ is most commonly known as ‘facial image comparison’, or simply ‘facial comparison’. Despite the importance of this task and the substantial literature examining novice accuracy in laboratory studies, it is only within the past decade that researchers have begun to assess the accuracy of professional populations. The purpose of this chapter is to review this emerging literature.

PROFESSIONAL EXPERTISE IN UNFAMILIAR FACE MATCHING

Many people are required to be experts in unfamiliar face matching as a consequence of their position, professional training or experience. For instance, we require that border control officers can accurately decide if the person at a border crossing is the individual shown in their ID document. More specific requirements can also be imposed by legal requirements, for example in rules or guidelines that judges use to evaluate the admissibility and weight of expert evidence (see Edmond & Wortley, 2016). We have argued elsewhere that, on their own, a person’s position, qualifications, professional experience or training are not sufficient to qualify them as an expert in facial image comparison. Rather, a person’s expertise in this task should be established on the basis of proven and sustained levels of accuracy in face matching (e.g. Towler et al. 2018). This is consistent with the study of expertise more

¹ This term may be underspecified (see Stacchi, Huguenin-Elie, Caldara, & Ramon, 2019). Variants include tasks that require perceptually matching the identity of images of faces presented simultaneously, without the requirement that these images are committed to memory, and other ‘sequential’ versions that require holding an image in memory for a short duration. In this chapter we focus on a narrow definition of the task whereby images are presented simultaneously to the viewer, and so test perceptual discrimination ability with no requirement to store the face in memory (see Figure 1).

broadly, where a standard condition of being classified as an expert is “*consistent superior performance on a specified set of representative tasks for a domain*” (Ericsson & Lehmann, 1996).

The focus of this chapter is therefore on whether professional groups outperform novices on representative tasks in their domain. As researchers interested in expertise in unfamiliar face matching, we are in an enviable position compared to others working in the broader field of expertise (e.g. Anders, Charness & Feltovich, 2006), because it is relatively straightforward for us to define the *representative task*. Despite diversity in their job descriptions and roles, an important part of the professional duties performed by all professional staff reviewed in this chapter is captured by the task illustrated in Figure 1. The simplicity of this task, and its relative uniformity across a wide range of important applied contexts make it ideally suited to studying expert performance. In addition, the close correspondence between existing laboratory tests and tasks performed in the real world has enabled researchers working in this area to transition quickly from laboratory to field-based studies.

In addition to the clear practical importance, understanding professional expertise in unfamiliar face matching is also of theoretical interest, as it can potentially inform broader conceptual understanding of perceptual expertise. Notably, prominent theoretical accounts in the field of face perception have argued that people *in general* are ‘face experts’ (e.g. Carey, 1992; Diamond & Carey, 1986). According to this definition, perceptual expertise in face identification is acquired in the normal course of our everyday experience, as a consequence of the fundamental importance of accurate face perception in social interactions with other people. However, as Young and Burton (2018) have recently argued, processing the identity of unfamiliar faces is an important exception. For the average person, performing unfamiliar face matching tasks is error-prone and effortful so their performance does not satisfy recognized definitions of expertise, which emphasise accuracy and the presence of automatic cognitive processing (e.g. Ericsson & Lehmann, 1996; Shiffrin & Schneider, 1977). Instead, people in general are experts at recognising the *familiar* faces of people that they know — a task they perform routinely in daily life — but not *unfamiliar* face identification, a task that most people are rarely required to perform.

If the average person is not expert in unfamiliar face matching, and does not perform the task regularly, this raises the important question of whether expertise can be *acquired*. A large body of research shows that familiar face recognition and unfamiliar face matching rely on qualitatively different perceptual and cognitive processes (e.g. Bonner, Burton, & Bruce, 2003; Clutterbuck & Johnston, 2002; Megreya & Burton, 2006, 2007). Given these distinct processes, it is possible that performance is supported by qualitatively different learning trajectories. Face *recognition* is a task that everyone performs many times each day, and the learning curve that drives our expertise in this task is believed to asymptote around 30 years of age (Germine, Duchaine & Nakayama, 2011). However, it is not known whether a

separate trajectory of learning can lead to improved accuracy in unfamiliar face matching, in the special case where people are required to perform the task in their daily work.

Psychological theories of expert performance predict that deliberate and prolonged practice leads to improvements in performance (Ericsson, Krampe, & Tesch-Römer, 1993). In the remainder of this chapter, we evaluate support for this prediction in the task of unfamiliar face matching, by reviewing studies that have compared performance of professionals – engaged in the task as part of their employment – to novices.

First, we describe the types of professional groups engaged in unfamiliar face matching that have been tested. These practitioners come from diverse backgrounds, are exposed to different types of training and work experiences, and perform a range of different facial comparison tasks in a wide variety of environments (Heyer, MacLeod, Carter, Semmler, & Ma-Wyatt, 2011). Second, we present results from a recent meta-analysis of the studies reviewed in this chapter. This shows large variability in the standardized effect size comparing novice and expert performance across these studies, which appears to be attributable to the relative specialization of professional groups. Third, we examine the individual studies in some more detail to assess the factors that appear to promote superior accuracy in specialist groups, and the nature of the learning that might produce expert performance. Based on this detailed review, we propose that a combination of factors contribute to the higher accuracy of specialist groups.

We conclude that expertise in unfamiliar face matching can be acquired both *developmentally* and *professionally*. In our final section, we formalise these different types of expertise within a dual-route framework where two independent domains of expertise underpin an individual's accuracy in facial image comparison – one based on natural ability and rooted in the core face recognition system, and another that is acquired through professional experience and training in unfamiliar face matching, and is supported by more general perceptual and cognitive mechanisms.

SUBTYPES OF PROFESSIONAL GROUPS

There are three broad categorisations of professional roles that require people to perform unfamiliar face matching in their work. Following from recent tests involving diverse groups of face identification practitioners (Phillips, Yates, et al., 2018), we use the terms *Facial Reviewer*, *Facial Examiner* and *Police Super-recogniser* to capture the three types of professional experts that have been tested. We use these categories to organise the meta-analysis and literature review below, and provide further details of these roles in the following sections.

This classification is consistent with by the internationally recognised standards documents produced by the Facial Identification Scientific Working Group (FISWG, 2010). These are an authoritative reference for face identification practitioners worldwide. According to the FISWG classification there are two broad types of personnel that perform unfamiliar face matching in their daily work. *Facial Reviewers* process substantial volumes of face identification decisions and typically make these decisions quickly and alongside other administrative checks, such as checking that a traveller has a valid visa or has completed paperwork correctly. Examples include border control officers (e.g. Wirth & Carbon, 2017) and government staff who review candidate lists generated by face recognition software (e.g. Heyer, Semmler, & Hendrickson, 2018). These professionals typically receive only basic training, and do not receive mentorship on face matching. In contrast, *Facial Examiners* perform more rigorous comparison of facial images, may spend hours or days on each comparison, receive more extensive training and mentorship, and typically communicate their conclusions in formal reports. These individuals are often responsible for producing written reports that might be used in legal proceedings (e.g. Norell et al., 2015; Phillips, Yates, et al., 2018).

We also include a relatively new class of face identification professional, *Police Super-recognisers*, who have been selected on the basis of their high performance in face identification tasks to be part of specialist face identification teams to support criminal investigations (Davis, Lander, Evans, & Jansari, 2016; Robertson, Noyes, Dowsett, Jenkins, & Burton, 2016). There are relatively few studies on this type of professional, but there is a broader research literature examining non-professional populations that achieve high levels of accuracy on face identification tests (see Noyes, Phillips & O'Toole, 1997; Ramon, Bobak & White, 2019 for reviews). In general, natural variability in face identification ability in the general population is known to be relatively stable over repeated testing (Balsdon, Summersby, Kemp, & White, 2018; Wilmer et al., 2010; see Wilmer, 2017 for a review; c.f. Ramon et al., 2019). This finding has important implications for the selection and recruitment of individuals for professional roles. However, given the narrow focus of this chapter on performance of professional staff that perform unfamiliar face matching in daily work, we do not include studies of super-recognisers from the general population in our meta-analysis.

META-ANALYSIS OF COMPARISONS BETWEEN PROFESSIONALS AND NOVICES

We recently conducted a literature search to identify studies that compared professional experts and novice participants in unfamiliar face matching accuracy (White, 2020). Criteria for inclusion were that the professional groups performed face matching as part of their daily work, and at least one unfamiliar face-matching task was used to compare their accuracy to a comparison non-expert group. This meant excluding previous tests that either tested professional staff using memory-based face

identification tests (e.g. Burton, Wilson, Cowan, & Bruce, 1999; Davis et al., 2016; CFMT) or did not include a novice comparison group (e.g. Kemp et al., 1997; White, Kemp, Jenkins, Matheson, & Burton, 2014; photo-to-person test).

Although there was some variation in unfamiliar face matching tests used to compare novices to professionals across studies – for example in image quality or presentation format – they all shared the requirement that participants compared images of unfamiliar faces, presented simultaneously, and had to decide whether they showed the same person of different people. To ensure that the studies used tasks that were representative of real world tasks, we also removed tests from analysis that: (i) used test images that were digitally manipulated (e.g. Wirth & Carbon, 2017: ‘paraphernalia’, ‘distinctive feature’ and ‘hairstyle conditions’); (ii) involved other non-representative experimental manipulations (e.g. White, Phillips, Hahn, Hill, & O’Toole, 2015: Inverted and 2-second conditions); or (iii) showed clear evidence of ceiling effects (Towler et al., 2019: Experiment 2, high resolution-to-high resolution test).

We identified a total of 12 publications containing professional vs. novice comparisons. Many of these either contained multiple tests fitting our definition of unfamiliar face matching (e.g. White et al., 2014; Robertson et al., 2016) or multiple professional groups (e.g. White, Dunn, Schmid, & Kemp, 2015; Phillips, Yates, et al., 2018), giving a total of 29 comparisons that met our inclusion criteria. Over 1600 practitioners have been tested since 2013, a remarkable achievement given that our review did not uncover any published tests prior to this date. We find this especially encouraging because it suggests that professional organisations are beginning to recognise and address the problem identified by psychologists in the late 1990s – that novices perform badly at unfamiliar face identification tasks (Kemp et al., 1997; Bruce et al., 1999).

Effect sizes between professional and novice groups are shown in Figure 2, categorised by the three types of professional groups described above. Positive effect sizes signify superior performance of professionals compared to novices, with statistically significant comparisons indicated by confidence intervals that do not include zero. Visual inspection of this figure shows clear differences between the effects observed in the three professional groups. While facial examiners and super-recognisers consistently achieve higher accuracy than novices, facial reviewers generally have the same level of accuracy as novices (Mean *Cohen’s D* Facial reviewers = 0.34; Facial examiners = 1.43; Police super-recognisers = 1.66). In terms of aggregate accuracy differences, reviewers scored 1.5 percentage points higher than novice comparison groups, with examiners scoring 13 and police super-recognisers 14 percentage points higher.

EXPERTISE IN UNFAMILIAR FACE MATCHING

Figure 2. Effect sizes for 29 novice vs. professional comparisons in unfamiliar face matching tasks taken from 12 peer-reviewed scientific papers (data from White, 2020). A positive effect size indicates that the professional cohort outperformed novices, 95% confidence intervals that do not include zero indicate the comparison was statistically significant and the sample size is represented by the size of the marker. Details of the three types of expert and the tasks they perform are provided in the main text. For further details of individual tests please refer White (2020).

This meta-analysis indicates that the different practitioner groups have different levels of expertise. It is likely that the factors contributing to accuracy also vary across these professional groups. Police Super-recognisers are explicitly selected on the basis of their ability, whereas other groups are not. Differences between reviewer and examiner groups, on the other hand, may be due to differences in working practice in

these roles. While both of these groups perform face matching tasks regularly, there are important differences in the training and professional mentorship they receive, and the time available to undertake the respective tasks. In the following sections we examine the studies included in this meta-analysis in greater detail, to understand the conditions that lead to higher accuracy in some professional groups and not others. First, we review tests of facial reviewers, where evidence for expert performance is limited, before examining studies of more specialist groups – Facial Examiners and Police Super-recognisers – where evidence points to superior accuracy.

FACIAL REVIEWERS

A variety of facial reviewer cohorts have been compared to novices, including UK police officers (Towler et al. 2019), US bank tellers (Papesh, 2018), US notaries (Papesh, 2018), Australian passport issuance officers (White et al., 2014; White, Dunn, et al., 2015; Towler et al., 2019), Australian government staff (Calic, 2013; Heyer et al., 2018), German border control officers (Wirth & Carbon, 2017) and international cohorts of police and government staff (Phillips, Yates et al., 2018). All of these groups performed unfamiliar face matching in their daily work, and their decisions played a critical part of processes designed to protect against identity fraud and other criminal activity. Nevertheless, visual inspection of Figure 2 shows that the majority of these groups did not perform significantly better than novices, with roughly half showing numerically *lower* performance than people who do not perform the task regularly.

Across 18 tests, the aggregate percentage point advantage of reviewers over novice comparison groups was 1.6%. Just six of the 18 tests showed significant differences whereby reviewers outperformed novices. Effect sizes in tests of facial reviewers were generally small but show a large amount of variability across tests, ranging from 6% *less* accurate (*Cohen's D* = -0.61) to 17% more accurate (*Cohen's D* = 1.35). Because of this substantial variability, we first consider studies that did not find significant differences between novices and reviewers, before considering those reporting significant differences.

Non-significant differences between facial reviewers and novices

In total, 12 of 18 tests showed no difference between professional and novice groups. These null results include tests of police officers, bank tellers, notaries and passport officers – all staff that the public rely on to make accurate face matching decisions. Starting with studies showing weakest evidence of expert performance, police officers tested by Towler and colleagues (2019) had been working in roles that required them to make facial image comparison decisions for an average of one and a half-years, but only seventeen percent of these staff had received training. This study was designed as an evaluation of a professional training course, and so novice and professional groups performed tests before and after training to assess whether training improved

EXPERTISE IN UNFAMILIAR FACE MATCHING

accuracy (the test reported some limited evidence of improvement, see Towler et al. 2019 for details). For the purpose of our meta-analysis we analysed scores from the four unfamiliar face matching tests taken before training. Officers performed the GFMT and an additional three face matching tests, shown in Figure 3. Across all four of the tests, there were no significant differences between police officers and university students, and police were in fact *less* accurate than novice participants, as indicated by the negative effect size for these contrasts.

Figure 3. Police officers tested by Towler et al. (2019) completed four pairwise unfamiliar face matching tasks: the Good, Bad and Ugly test (A), Male Models test (B), Casework test (C); and the Glasgow Face Matching Test (see Figure 1).

Figure 4. (A) Photo-ID matching task used by Papesh (2018) to test bank tellers and notaries. (B) Photo-to-photo-ID matching test used by White et al (2014) to test passport issuance officers.

Bank tellers and notaries tested by Papesh (2018) – who also show negative effect sizes in comparison to novices – had many years experience verifying identity by checking photo-ID, to protect against identity fraud in financial and legal transactions. Despite the importance of these decisions, they made an average of 25% errors in the photo-to-photo-ID task matching shown in Figure 4 (A). Similarly, passport issuance

officers tested by White et al. (2014) and White, Dunn, et al. (2015) are responsible for ensuring that the image of the passport applicant matched their previous passport application, and had an average of eight and a half years experience in this role. When asked to decide whether pairs of images, like those shown in Figure 4 (B), are of the same person or different people, these staff made 20% errors on average – again, no better than university students.

In the three tests comparing the accuracy of passport officers to novices in White et al. (2014), one test did show significantly higher accuracy to novices (photo-to-photo-ID task, see Figure 4B). It is perhaps tempting to attribute this to the photo-ID checking task being more reflective of passport officer's daily work than the other tests, which used images that were not scanned from current photo-ID cards. However, photo-ID checking is not in fact the face-matching task performed most regularly by these passport issuance staff. Instead, the most common task they perform is shown in Figure 5, where they must decide if the person at the top appears anywhere in the gallery of faces below that have been returned by face recognition software. When White, Dunn, et al. (2015) tested their cohort on this precise task they found that passport issuance officers made errors in 50% of trials – the same as a comparison group of university students. This suggests that the superior performance on the photo-ID checking test was not simply due to the type of imagery being more familiar to the passport officers.

Figure 5. Example items from the face recognition software ‘candidate list’ task used by White et al. (2015) to test passport issuance officers. They found that facial reviewers who perform this task as part of their daily work made errors on 1 in every 2 decisions, and mistakenly selected the wrong face as matching the target 40% of the time. These types of false identifications could have serious consequences, especially in the context of law enforcement where this could result in wrongful conviction and imprisonment. The correct answer for A is top row, second from the left. The correct answer for B is target absent.

The task shown in Figure 5 is an increasingly common task in a variety of professional settings, where face recognition software is used to search for a face in a large database of facial images. The Australian Passport Office, for example, use face recognition technology to protect against identity fraud when issuing passports. When Australian citizens apply for a passport, their application image is used to search all other passport images. If a matching face is found under a different name, this may signal a fraudulent application. Critically, these types of ‘one-to-many’ database searches return a ‘candidate list’ of potential matching faces that a human practitioner must then review to decide if the target face is present (White, Dunn et al. 2015; White & Kemp, 2016). Face recognition systems used in criminal investigation and live surveillance produce similar arrays of images (see Noyes & Hill, this volume).

Heyer et al (2018) also tested a professional cohort using a candidate list task generated using face recognition software. While the task was substantially easier overall than in White et al (2015)², overall accuracy of staff from Australian Government and police did not exceed novices. However, this professional group was recruited from a range of Australian Government departments. These staff included passport issuance officers who compare passport images with the support of face recognition software (as in White et al. 2014; White, Dunn, et al., 2015), visa processing officers who compare applicants to passport images, and forensic artists who produce likenesses of suspects. As a result, there was substantial variability in the precise roles performed across these departments, and also the proportion of time staff devoted to unfamiliar face matching tasks in their work. When the researchers split this professional cohort by the average number of facial comparison decisions made in their daily work, they found that staff who made more than 50 decisions per day outperformed novices on the ‘candidate list’ task.

Heyer et al’s (2018) result may suggest that the number of decisions made by professionals each day is a key factor in the development of expertise. However, while this association appears to exist for the specific cohort they tested, it does not appear to be true in general. For example, German border control officers tested by Wirth and Carbon (2017) were professionals for whom “*person identification during immigration was one of the core tasks (if not the most frequent task) the participating*

² Heyer et al’s (2018) task was easier than White, Dunn, et al.’s (2015), because they constructed their task using images from a publicly available database of high-quality images (Phillips, Wechsler, Huang, & Rauss, 1998) where image capture conditions were under experimental control and images of the same person were taken on the same camera at roughly the same time. In contrast, White, Dunn et al. (2015) used real passport images where matching images were taken many years apart and in more diverse image capture conditions. Another important difference was that Heyer et al’s candidate lists were constructed using the top matches from a database of 1000 faces, but White et al. selected foil images from a dataset of over a million images. This meant that the non-matching ‘foil’ faces in White and colleagues’ study were likely to be much more similar in appearance to the target face than in Heyer et al.’s study.

police officers performed on their job”. Wirth and Carbon’s (2017) task was carefully constructed to be representative of the passport checking duties performed by their professional group, as shown in Figure 6. Despite making very frequent face matching decisions on this precise task in their daily work, their results show no overall difference between professionals and novices (*Mean accuracy difference* for unmodified images = 3.8%, *Cohen’s D* = 0.31, *CI* = [-0.04, 0.66]). Indeed, the relatively newer recruits performed significantly better than novices and staff that had been in the job for longer, suggesting that experience in the task does not necessarily lead to improved accuracy. This is consistent with both White et al (2014) and Papesh (2018) who also found that more professional experience did not predict higher accuracy in their cohorts.

Figure 6. Passport-to-photo matching task used by Wirth and Carbon (2017) to test German border security officers. The image on the left shows a participant performing the task, with examples of the images used in the test shown on the right.

Overall, the 12 tests showing non-significant differences between face matching accuracy of facial reviewers and novices appear to confirm that simply performing the task of unfamiliar face matching in daily work is not sufficient to produce expertise. Next, we consider the tests of professional groups that do show significant differences, to better understand which factors may give rise to expert performance.

Significant differences between facial reviewers and novices

Across all 18 tests of facial reviewers, 6 showed significantly greater accuracy than novices. For 3 of these comparisons, the professional groups showed superior accuracy compared to novices in one test (Calic 2013, Experiment 2; Heyer et al., 2018, GFMT; White et al., 2014, photo-to-photoID) while showing equivalent accuracy in novices in other tests (Calic, 2013, GFMT; Heyer et al., 2018, FR Candidate List; White et al., 2014: GFMT, Photo-to-Photo). Moreover, where there were significant differences between groups in these studies, effect sizes were relatively small. Ericsson and Lehman’s (1996) definition of expertise, which we adopt here, states that a key criteria is ‘*consistent performance across a set of*

representative tasks', and so the inconsistent performance of these groups means that they do not meet this criteria.

The remaining three tests appear to show more compelling evidence of superior performance. Two of these tests were carried out on the same cohort of 204 passport issuance officers working for the Australian Passport Office (Towler et al. 2019, Experiment 2), who were an exhaustive sample of the passport issuance officers working for the Australian Passport Office. They outperformed students on both the Glasgow Face Matching Test (*Mean accuracy difference* = 9.3%, *Cohen's D* = 0.79) and a bespoke test that required comparing faces in low quality images to faces in high quality images (*Mean accuracy difference* = 4.2%, *Cohen's D* = 0.38). Interestingly, the passport officers tested by White et al (2014) were a subset of the larger sample tested by Towler et al (2019). Because the officers tested by White et al (2014) were recruited from a single regional office, they may not have been representative of the overall ability of officers in the organisation.

The final test showing a significant difference between facial reviewers and novices was conducted by Phillips, Yates, et al. (2018). Facial reviewers in this test – police and government employees primarily from the USA – were asked to complete a very challenging pairwise unfamiliar face matching task consisting of 20 image pairs (see Figure 6A). These reviewers were substantially more accurate than university students on this task (*Mean accuracy difference* = 17.5%, *Cohen's D* = 1.35), but facial reviewers had a window of up to three months to complete these face matching decisions, whereas the student comparison group completed them in a single experimental session. This may well have contributed to their higher accuracy, but the very large effect size – in comparison to the effects of providing novices with additional time in previous tests (e.g., White, Phillips et al., 2015; c.f., Özbek & Bindemann, 2011) – suggests that at least some of the effect can be attributed to the superior ability of this group. Nevertheless, given that details of their professional history, training and experience were unavailable – and that only a single test was used to verify their expertise – future research is necessary to determine the causes of their superior accuracy.

Overall, the evidence reviewed above provides very little support for *professionally acquired* expertise in facial reviewers. Importantly however, these professional groups did not receive any extensive training, feedback or mentorship in unfamiliar face matching tasks. All of these factors have been shown to play some role in the development of expertise in other fields (Ericsson et al., 1993), and so we next consider more specialist professional groups that do benefit from some of these learning opportunities.

Figure 7. Pairwise unfamiliar face matching tasks used to test facial examiners. (A) Example image pairs used to test examiners in White, Phillips et al. (2015), Towler et al (2017) and Phillip, Yates, et al. (2018). The left pair are of the same person and the right pair are of different people. (B) Image pair from the Person Identification Challenge Test used in White, Phillips et al. (2015). These images are of the same person. (C) An example of reference image (left) and three images of varying quality (right) used to create image pairs of varying quality in Norell et al. (2015).

FACIAL EXAMINERS

Facial Examiners are relatively homogenous group compared to Facial Reviewers. While they are employed by a variety of organisations, including government departments, police and private forensic service providers, they tend to perform a similar role in these professional settings. Typically, they are responsible for writing reports that provide a detailed and balanced assessment of the image comparison evidence. A key requirement of their role is to provide an estimate of the likelihood that two images are of the same individual or two different individuals using a ‘conclusion scale’, to provide justification for this conclusion and to potentially defend this in court (see Towler et al., 2018).

Facial Examiners typically receive facial comparison decisions that have been escalated to them via some initial identity screening process. This screening is

typically performed by the facial reviewers described above. For example, in a law enforcement context, the examiners might receive potential matches that have been uncovered via police investigation and perhaps by investigators using facial recognition software (see Klontz & Jain, 2013). At border control, examiners may be asked to compare an image of a traveller captured by an automated border gate to their passport image, when primary line officers are unable reach a decision.

All seven comparisons between Facial Examiners and novices show large effect sizes, with examiners outperforming novices. The aggregate percentage point advantage of examiners over novice comparison groups was 13%, with six of seven tests reporting *Cohen's D* greater than 1 and one of these was greater than 2 (Norrell et al. 2015, 2.55). They attained high levels of accuracy despite making decisions more quickly than they would normally be expected to in their daily work (e.g. 30 seconds in White, Phillips, et al. 2015; 18 seconds in White, Dunn, et al., 2015). Importantly, in all tests – except for one (Phillips, Yates, et al., 2018) – examiners also attained these high levels of accuracy in experimental conditions that did not permit use of digital tools or reference material that might be used in normal working practice. This suggests that the source of their expertise is in their perceptual and cognitive processing, rather than their ability to use specialist tools or decision-making support.

As further evidence of their expertise, when groups of Facial Examiners have been tested on multiple tests they have shown consistently superior performance across these tests. Facial examiners in the Australian Passport Office have shown superior performance in two separate studies conducted years apart (White, Phillips, et al., 2015; Towler, White & Kemp, 2017), and the international group of examiners tested by White, Phillips, et al (2015) showed superior accuracy across three representative tests performed in a single experimental session. Interestingly, the only test reporting an effect size less than 1 (White, Phillips, et al., 2015, Person Identification Test, *Mean accuracy difference* = 7.7%, *Cohen's D* = 0.79), included body and clothing information, and a subset of test items that were specifically selected because they did not include identifying information in the face (Figure 7B; see Rice, Phillips, Natu, An, & O'Toole, 2013 for details). Figure 7 shows examples of the different types of test items employed in the different tests of examiners.

The consistent superior accuracy of facial examiners compared to novices satisfies the criteria of expertise we adopted at the beginning of this chapter (Ericsson & Lehmann, 1996). There are many potential sources of learning above and beyond those available to facial reviewers, which may explain examiners' expertise. First, the training that facial examiners receive is much more detailed (see Towler, Kemp, & White, this volume; see also Towler et al, 2019). Second, most facial examiners work in small teams with many opportunities for mentorship and feedback on their performance. Both performing face matching tasks collaboratively with high performers (Dowsett & Burton, 2015) and receiving feedback on the accuracy of decisions (White et al., 2014; Papesh, Heisick & Warner, 2018; c.f., Alenezi &

Bindemann, 2013) have been shown to improve novice performance in laboratory studies. It is possible therefore that these factors enable effective learning in facial examiners.

However, it is important to clarify that the superior performance demonstrated by facial examiners is in fact *acquired in the job*, rather than arising in some other way. For example, it could be that increased motivation of professional groups contributed to their superior accuracy (Moore & Johnston, 2013; see Noyes et al., 2017). We see two main reasons to doubt this possibility. First, facial examiners have outperformed motivated groups of professional staff that lacked specialist experience in face identification (White, Phillips, et al., 2015; Phillips, Yates, et al. 2018). Second, it is presumably the case that all professional staff required to perform face identification tasks as part of their job description are motivated to perform well in these tests. However, as we have seen, this motivation does not translate into improved accuracy in facial reviewers. This is despite clear evidence of increased motivation relative to novice comparison groups. For example, in White et al (2014) passport issuance officers spent twice as long comparing images than the students, but did not achieve superior accuracy.

Another possible explanation for the superior accuracy of professional groups is that there may have been some element of selection for their role. It is now very well established that, as a consequence of the natural development of perceptual and cognitive systems, people vary in their ability to identify faces (see Wilmer, 2017 for a review). In the next section, we review studies of ‘Police super-recognisers’ – groups that have been selected on the basis of their performance in face identification tests to perform specialist roles. Facial examiners that have been tested were not explicitly selected in this way. However, there may be more subtle sources of selection bias, for example a tendency for people with higher levels of natural ability in face identification to gravitate towards these professional roles (e.g., Roy, 1951). This could either be self-selection of the participants themselves – for example, based on the intuition that they are ‘good with faces’ – or on the basis of their performance in the professional role, where they are believed to have made identification decisions correctly in the past.

However, there is also reason to doubt this explanation. If selection alone were sufficient to explain the high performance of facial examiners, then facial examiners would be predicted to perform like high-performing novices. On the contrary, evidence suggests that examiners and novices perform the task in qualitatively different ways (White, Phillips et al. 2015; Towler et al., 2017). Two important qualitative differences were identified by White, Phillips, et al (2015). The first was that examiners only outperformed novices when given a long time to study the images (30 seconds), and not when images were presented for short durations (2 seconds). Novices were relatively unaffected by short study durations (see also Özbek & Bindemann, 2011), suggesting that the examiners extract the necessary information

for their decisions relatively slowly. The second difference was that examiners were relatively unimpaired by turning the images upside down. This ‘face inversion effect’ is thought to index the extent to which people engage in ‘holistic’ processing of faces (Maurer, LeGrand & Mondloch, 2002) – that is, is the tendency process identity information ‘at a glance’ via a single percept in which individual facial features are integrated. Importantly, the face inversion effect is larger in high performing novices – ‘super-recognisers’ (Bobak, Bennetts, Parris, Jansari, & Bate, 2016) – and parallel evidence from eye-tracking studies also point to enhanced holistic processing in this group (Bobak, Parris, Gregory, Bennetts, & Bate, 2017).

These qualitative differences suggest that facial examiners are not simply high performing novices, and instead that their expertise is acquired via professional experience and training. More specifically, their expertise appears to be related to a strategic shift in the way the task is performed, from a holistic analysis to a more piecemeal and analytic process of feature comparison. This is consistent with the training that these professional staff receive (see Towler et al., this volume; Towler et al., 2019; Heyer et al., 2011; Prince, 2012), which prescribes a careful feature-by-feature approach to comparison.

Other evidence of strategic differences in their performance of the task is provided by analysis of the way that examiners use the response scale relative to novices. Norell et al. (2015) tested the performance of 17 facial examiners from the European Network of Forensic Science Institutes, finding that the accuracy of the examiners was substantially better than novices (*Mean accuracy difference* = 12.9%, *Cohen’s D* = 2.55). Participants in this study made their responses on a 9-point conclusion scale that was chosen to reflect the response scales used by these practitioners in their daily work. Response options ranged from ‘*extremely strong support that the images are the same person*’ to ‘*extremely strong support that the images are of different people*’. Their analysis showed that, relative to novices, examiners were far more sensitive to the quality of images being compared when choosing the appropriate point on the response scale. When image quality was low, examiners tended to use midpoints in the scale reflecting low confidence in their decision. When image quality was high, examiners were more likely to use the end points of the scale. However, university students did not show this same systematic pattern suggesting that an essential part of the expertise of this group was in assessing the quality of the image evidence and incorporating this assessment into their response (see also Phillips, Yates, et al. 2018; Towler et al. 2018).

It could be argued that all the learning described above is non-perceptual. Differences in the use of the response scale are likely to reflect a greater understanding of how to use response scales effectively, and what the various response options mean within the context of their profession. So this may reflect a change in the way observations are converted to decisions, rather than reflecting a change in perceptual

representations underlying the decisions³. Similarly, strategic changes in the way that the visual information is parsed – reflected in a greater emphasis on features – alters the way that visual information is sampled. However, this change might be procedural and rule-based in origin, with instructions and training for example changing the way that examiners deploy their attention. Increasing focus on facial features may therefore be argued to reflect changes in cognitive control mechanisms, and not necessarily reflective of changes in perceptual representations (see Pylyshyn, 1999).

Notwithstanding the plausibility of this argument, there is preliminary evidence that these strategic changes are indeed accompanied by perceptual learning. Towler et al (2017) asked facial examiners to rate the similarity of certain facial features and examined the extent to which these feature ratings were diagnostic of whether the features were from the same face or different faces. The diagnosticity of examiners' feature ratings was far greater than novices' ratings for the same features. This appears to show greater perceptual sensitivity to diagnostic information contained at the level of individual features. In addition, examiners and novices differed with regard to which feature ratings were most diagnostic. For examiners, similarity ratings to ears and facial marks (scars and blemishes) were ranked as being amongst the most diagnostic, but novices' ratings to these features were not especially diagnostic. These observations raise the possibility that the strategic focus on facial features leads to the development of greater perceptual sensitivity to feature-level details that carry identity information. It also may lead to the discovery of features that novices do not regularly attend to in great detail when identifying faces in their daily lives – such as the ears or facial marks – but that are nevertheless very useful for identification when comparing images of unfamiliar faces.

POLICE SUPER-RECOGNISERS

In contrast to facial examiners, Police Super-recognisers are professionals who have been explicitly selected on the basis of their performance on face identity processing tasks. A substantial body of evidence now shows there is large variation in the general population between individuals' ability in face matching tasks, and that this variability is relatively stable over repeated testing (Wilmer et al., 2010; Baltho et al., 2018; see Wilmer, 2017 for a review; c.f. Ramon et al., 2019). Indeed, while we have focussed entirely on group differences reported in the studies described above, all of these studies also show large variability in the performance of individual staff, with some performing near perfectly and others as if they were randomly guessing (e.g. see White et al. 2014). Recruiting and selecting practitioners on the basis of their face matching ability is therefore a promising way to improve the accuracy of face

³ On the other hand, perceptual learning may support the assessment of image quality. Norell et al. (2015) finding that examiners are more sensitive to the quality of images may suggest they are more skilled at identifying image artefacts, such as lens flare, blurring or overexposure. Such perceptual expertise would support the ability to moderate the strength of the conclusions drawn from the image evidence.

identity decisions made in professional settings (White, Kemp, et al., 2014; Bobak, Dowsett, & Bate, 2016).

Both studies that have tested police super-recognisers recruited participants from the London Metropolitan Police Service (Robertson et al. 2016; Davis et al. 2016). In the study by Davis et al. (2016), 36 participants were tested on the Glasgow Face Matching Test and group performance was above novice accuracy (*Mean accuracy difference* = 5.3%, *Cohen's D* = 0.74). In the study by Robertson et al (2015), four participants were tested and the effect size for this group was substantially higher on the Glasgow Face Matching Test (*Mean accuracy difference* = 14.5%, *Cohen's D* = 1.98) and other matching tests (Male Models test: *Mean accuracy difference* = 16.7%, *Cohen's D* = 2.14; Celebrity lookalikes test: *Mean accuracy difference* = 14.5%, *Cohen's D* = 1.80).

Although these studies both test police super-recognisers from the same organisation, there are striking differences in the size of the professional groups in these studies and also the effect sizes that are reported. This is likely to be caused by the progressive specialisation of the super-recogniser unit in the London MET in the years since it was established. Robertson et al. (2016) tested a subset of participants from Davis et al's (2016) study, who that had been selected on the basis of their on-the-job performance and standardised tests of face identification ability (see Davis, 2019). This selection process is likely to explain the high levels of accuracy shown by the group tested by Robertson et al (2016), although the precise details of this process are not known.

Given uncertainty regarding how these police super-recognisers were selected, it is important to consider that other factors may have contributed to their higher accuracy in the task. On-the-job experience, mentorship, feedback, training and increased motivation may also contribute to the higher accuracy of these professionals, and because very few details relating to their professional history are available, it is not possible to rule these out. As we have discussed in previous sections, this is a general problem when aiming to establish the basis of expertise in professional groups. In many cases, there are a number of potential contributors to higher performance that are outside of the experimenters' control. To address this problem, it may be necessary to conduct longitudinal tests that track changes in an individual's performance over time, and to conduct studies that enable researchers to report details of training, experience and selection criteria used in organisations with more precision.

A final and important point relating to police super-recognisers is that the identity processing tasks these individuals perform are more diverse than those undertaken by the other professional groups in this chapter. While the primary task performed by other groups included in this review was unfamiliar face matching, super-recognisers are typically required to perform tasks that require memory, for example reviewing

CCTV recordings and verifying whether the people in the video are familiar to them or have appeared in other video sequences they have watched. A full discussion of the diverse set of skills and abilities that are potentially recruited in CCTV monitoring tasks is beyond the scope of the chapter. However, it has been argued elsewhere that this diversity poses substantial challenges when designing standardised methods for selecting and training these professionals (see Ramon et al., 2019 and commentaries).

Because we have only focussed on unfamiliar face matching, the understanding of expertise proposed in this chapter is unlikely to generalise to the full range of face identification tasks that are performed in professional settings (e.g. face-in-crowd search, see Davis, Forrest, Tremblé, & Jansari, 2018). Indeed, as we outline below, our assumption is that expert performance on face *memory* tasks is unlikely to be acquired through professional experience, as appears to be the case for unfamiliar face matching. An implication of this is that different people are likely to be better suited to different expert face identification roles. For example, the slow, analytic approach of facial examiners (see White, Phillips et al., 2015) is not likely to be suited to live CCTV surveillance, where the requirements may be to quickly recognise faces of known suspects from poor quality imagery.

CAN EXPERTISE IN UNFAMILIAR FACE MATCHING BE ACQUIRED? A DUAL-ROUTE HYPOTHESIS

In reviewing studies of professional face matching performance, our aim was to provide some insight into the factors that might underlie superior performance in expert groups. Given that accuracy of facial reviewers does not differ reliably from novices, it appears that simply performing the task regularly in daily work is not sufficient to produce higher levels of performance. It is, however, very clear that two groups of professionals – Facial Examiners and Police Super-recognisers – meet our definition of expertise, by showing consistently higher accuracy than novices on tasks that are representative of their daily work. Moreover, effect sizes for comparisons between novices and these groups are large.

Importantly, these groups appear to achieve high levels of accuracy for different reasons. Police super-recognisers are selected on the basis of their accuracy on tests of face identification ability, and so it is perhaps unsurprising that they outperform novices. Facial examiners on the other hand do not appear to have been selected systematically. The critical distinction between these groups appears to be *how* expertise is acquired. In the case of super-recognisers, their superior ability appears to be predominantly acquired either congenitally or by the requirement to identify people during the course of their normal development – but in any case, prior to their deployment as face identification specialists. Conversely, facial examiners appear to acquire their expertise through their professional training and experience.

In our other chapter in this volume (Towler et al., this volume) we have proposed that these two routes to expertise map to distinct cognitive pathways. Expertise acquired through development exploits processes that have developed for the purpose of recognising faces – the ‘core’ face identity processing system described by Bruce and Young (1986). Professionally acquired expertise appears to exploit the alternative ‘directed visual processing’ route proposed in this model (see also Brunsdon, Coltheart, Nickels, & Joy, 2006). This is consistent with the view that recognition memory and unfamiliar face matching are reliant on different processes with the latter having a greater focus on piecemeal analysis of facial features (Megreya & Burton 2006, 2007). It is also consistent with the view that holistic and featural processing can offer parallel routes to face identification (Bartlett, Searcy, & Abdi, 2003; Brunsdon et al., 2006. Farah, 1991; c.f. Sergent 1984).

This account can explain why Facial Examiners do not meet some criteria that have been proposed as hallmarks of expert performance. Specifically, influential theories propose that with practice, cognitive processing becomes increasingly automatic and reliant on intuition (e.g. Chase & Simon, 1973; Kahneman & Klein, 2009; Kahneman, 2011). Yet facial examiners appear to be an exception to this rule. As White, Phillips, et al. (2015) show, their expertise is characterised by slow, analytic and controlled analysis of facial features. Super-recognisers on the other hand appear to rely on the fast, relatively intuitive mechanisms that enable them to recognise faces in daily life: they maintain high levels of accuracy when given relatively little time to study images (Russell, Duchaine, & Nakayama, 2009), and they appear to engage in more holistic processing than participants with normal face identification ability (Russell et al., 2009; Bobak, Bennetts, et al., 2016, Bobak, et al., 2017).

We assume that super-recognisers expertise is embodied in the core face recognition system and has been acquired through a combination of genetic (e.g. Wilmer et al. 2010) and environmental factors (Germine et al., 2011). Importantly, this expertise has been acquired for the purpose of recognition *memory* tasks that we typically perform in daily life when recognising family, friends and other acquaintances – and because it is essential to normal social functioning and survival, it has been subject to evolutionary pressure throughout human history. This necessity means that face recognition memory is overlearned, leading to asymptotic performance that is not amenable to further training, which is consistent with the finding that face memory ability in normal participants does not improve beyond around 30 years of age (Germine et al., 2011; Susilo, Germine, & Duchaine, 2013; see also Yovel et al., 2012).

On the other hand, facial examiners engage more general visual and top-down control mechanisms, that are not part of this core face recognition system, to bear on the task of unfamiliar face matching. This perspective has guided our research into the expertise that can be acquired via this route, and the mechanisms responsible for this learning (Towler et al., 2017; Towler et al., this volume). We reason that gains in

accuracy in unfamiliar face matching tasks are more likely to be found via this controlled processing route, through extensive practice in analytic, feature-by-feature comparison. This view is influenced by work showing that unfamiliar face matching recruits domain-general feature-based processing (Megreya & Burton, 2006), and that the development of visual expertise is associated with changes in frontal lobe activation, in areas associated with top-down attention control (see for example Martens, Bulthé, van Vliet, & de Beeck, 2018).

The greater reliance of facial examiners on facial features, and other strategic differences, suggests that such training plays some role in the development of their expertise. Paradoxically however, it does not appear that professional training courses have an immediate impact on accuracy, as tests conducted immediately before and after training sessions show no change in accuracy (Towler et al., 2019; Phillips, Heyer & Michalski, 2018). To account for this apparent discrepancy, we have argued elsewhere (Towler et al., 2019) that professionally acquired expertise is most likely to be developed through the subsequent application of the visual analysis protocols that are learned during training. By this view, the development of expertise occurs ‘on the job’ via *deliberate practice* of strategies prescribed in training courses.

The view that deliberate practice is essential to improvements in accuracy is consistent with the broader psychological literature on expertise. Deliberate practice is typically supervised by an instructor, and characterised by the delivery of immediate and informative feedback on the repeated performance of the same task (Ericsson et al., 1993). By this view, training does play an initial role in the development of expertise, by introducing practitioners to a new approach to the task that re-orientes their perceptual processing towards facial features, and discourages reliance on automatic and holistic processing. Subsequently, deliberate practice over an extended period of time enables them to discover which features are most useful for the purpose of identification, and to develop perceptual expertise in discriminating these features. It appears likely that effective mentorship and feedback are critical components of this learning process, given that laboratory studies have shown these to improve accuracy (Dowsett & Burton, 2015; White, Kemp, Jenkins, & Burton, 2014).

An important aspect of our hypothesis is that facial examiner’s learn to *override* the instinctive, automatic processing that enables us to recognise familiar faces in daily life. We propose that training in forensic facial comparison is akin to perceptual *unlearning*, as described in an influential study by Shiffrin and Schneider (1977). In their study, when participants were trained to find a set of letter targets amongst another set of letter distractors, their decisions became increasingly fast and automatic over training, reflective of perceptual learning of the target set. However, when experimenters then switched the target and distractor sets, their performance was impaired relative to initial performance at the start of training, an effect attributed to the fact that attention is captured by the original target set that are now being used as distractors. Importantly, in this study it took a very long period of learning to

overcome this impairment (2100 trials), a process that the authors refer to as *perceptual unlearning*. We propose that an analogous process occurs in facial examiners when they begin to perform unfamiliar face matching tasks: they must *unlearn* processes that support recognition memory, in order to discover features in the images that provide important cues to identity.

CONCLUSIONS

In this chapter we have reviewed the literature to determine whether professional groups outperform novices on unfamiliar face matching tasks. Over a third of comparisons between professional groups and novices show non-significant differences, confirming that mere practice is insufficient to develop expertise in unfamiliar face matching. Facial reviewers do not reliably outperform novices, but facial examiners and police super-recognisers do.

Greater opportunities for close mentorship and deliberate practice may partly explain why some groups show higher accuracy, and this appears associated with the relative specialisation of these groups. We have identified a number of potential sources of learning – strategic, decisional, and perceptual – that may drive improvements in accuracy in specialist groups. It is not possible based on the available evidence to make strong conclusions about the mechanisms that are involved in this learning. However, it appears likely that learning in each of these domains contributes to expert performance in facial examiners. More speculatively, this task-specific learning may rest on examiners ability to inhibit, or *unlearn*, more reflexive processes that support our common ability to recognise faces in daily life. Importantly, this is precisely the opposite trajectory than is typically emphasised to account for perceptual learning, where processing is argued to become more automatic, and less analytic, with practice (e.g. Kundel, Nodine, Conant & Weinstein, 2007; Thompson & Tangen, 2014).

In future, longitudinal testing of specialist professional groups may be necessary to confirm that deliberate practice is effective in improving performance, and to provide greater insight into this learning trajectory. Combined with carefully designed experimental work, this approach can help delineate the learning mechanisms and knowledge outcomes that underpin professionally acquired expertise in this important task.

REFERENCES

- Alenezi, H. M., & Bindemann, M. (2013). The effect of feedback on face-matching accuracy. *Applied Cognitive Psychology*, 27(6), 735-753. doi:10.1002/acp.2968
- Anders, K. E., Charness, N., & Feltovich, P. (2006). *Cambridge handbook of expertise and expert performance*. Cambridge University Press.

- Attwood, A. S., Penton-Voak, I. S., Burton, A. M., & Munafò, M. R. (2013). Acute anxiety impairs accuracy in identifying photographed faces. *Psychological Science*, 24(8), 1591-1594. doi:10.1177/0956797612474021
- Balsdon, T., Summersby, S., Kemp, R. I., & White, D. (2018). Improving face identification with specialist teams. *Cognitive Research: Principles and Implications*, 3(1), 25. doi:10.1186/s41235-018-0114-7
- Bartlett, J. C., Searcy, J. H., & Abdi, H. (2003). What are the routes to face recognition. *Perception of faces, objects, and scenes: Analytic and holistic processes*, 21-52.
- Beattie, L., Walsh, D., McLaren, J., Biello, S. M., & White, D. (2016). Perceptual impairment in face identification with poor sleep. *Royal Society open science*, 3(10), 160321. doi:10.1098/rsos.160321
- 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. doi:10.1016/j.cortex.2016.05.003
- Bobak, A. K., Dowsett, A. J., & Bate, S. (2016). Solving the border control problem: Evidence of enhanced face matching in individuals with extraordinary face recognition skills. *PloS one*, 11(2), e0148148. doi:10.1371/journal.pone.0148148
- Bobak, A. K., Parris, B. A., Gregory, N. J., Bennetts, R. J., & Bate, S. (2017). Eye-movement strategies in developmental prosopagnosia and “super” face recognition. *The Quarterly Journal of Experimental Psychology*, 70(2), 201-217. doi:10.1080/17470218.2016.1161059
- Bonner, L., Burton, A. M., & Bruce, V. (2003). Getting to know you: How we learn new faces. *Visual Cognition*, 10(5), 527-536. doi:10.1080/13506280244000168
- Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, 5(4), 339-360. doi:10.1037/1076-898X.5.4.339
- Bruce, V., Henderson, Z., Newman, C., & Burton, A. M. (2001). Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied*, 7(3), 207-218. doi:10.1037/1076-898X.7.3.207
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British journal of psychology*, 77(3), 305-327. doi:10.1111/j.2044-8295.1986.tb02199.x

EXPERTISE IN UNFAMILIAR FACE MATCHING

- Brunsdon, R., Coltheart, M., Nickels, L., & Joy, P. (2006). Developmental prosopagnosia: A case analysis and treatment study. *Cognitive Neuropsychology*, 23(6), 822-840. doi:10.1080/02643290500441841
- Burton, A. M., White, D., & McNeill, A. (2010). The Glasgow face matching test. *Behavior Research Methods*, 42(1), 286-291. doi:10.3758/BRM.42.1.286
- Burton, A. M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science*, 10(3), 243-248. doi:10.1111/1467-9280.00144
- Calic, D. (2013). *From the laboratory to the real world: Evaluating the impact of impostors, expertise and individual differences on human face matching performance* (Doctoral dissertation). Retrieved from <https://hekyll.services.adelaide.edu.au/dspace/bitstream/2440/91444/3/02whole.pdf>
- Carey, S. (1992). Becoming a face expert. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 335(1273), 95-103.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55-81. doi:10.1016/0010-0285(73)90004-2
- Clutterbuck, R., & Johnston, R. A. (2002). Exploring levels of face familiarity by using an indirect face-matching measure. *Perception*, 31(8), 985-994. doi:10.1068/p3335
- Davis, J. P. (2019). The worldwide public impact of identifying super-recognisers for police and business. *The Cognitive Psychology Bulletin*, 4, 17-22.
- Davis, J. P., Forrest, C., Treml, F., & Jansari, A. (2018). Identification from CCTV: Assessing police super-recogniser ability to spot faces in a crowd and susceptibility to change blindness. *Applied cognitive psychology*, 32(3), 337-353. doi: 10.1002/acp.3405
- Davis, J. P., Lander, K., Evans, R., & Jansari, A. (2016). Investigating predictors of superior face recognition ability in police super-recognisers. *Applied Cognitive Psychology*, 30(6), 827-840. doi: 10.1002/acp.3260
- Davis, J. P., & Valentine, T. (2009). CCTV on trial: Matching video images with the defendant in the dock. *Applied Cognitive Psychology*, 23(4), 482-505. doi: 10.1002/acp.1490

EXPERTISE IN UNFAMILIAR FACE MATCHING

- Diamond, R., & Carey, S. (1986). Why faces are and are not special: an effect of expertise. *Journal of Experimental Psychology: General*, 115(2), 107-117. doi:10.1037/0096-3445.115.2.107
- Dowsett, A. J., & Burton, A. M. (2015). Unfamiliar face matching: Pairs out-perform individuals and provide a route to training. *British Journal of Psychology*, 106(3), 433-445. doi:10.1111/bjop.12103
- Edmond, G., & Wortley, N. (2016). Interpreting Image Evidence: Facial Mapping, Police Familiars and Super-Recognisers in England and Australia. *Journal of International and Comparative Law*, 3(2), 1-50.
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological review*, 100(3), 363-406. doi:10.1037/0033-295X.100.3.363
- Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, 47(1), 273-305. doi:10.1146/annurev.psych.47.1.273
- Facial Identification Scientific Working Group. (2010). *Guidelines and recommendations for facial comparison training to competency*. Retrieved from https://fiswg.org/FISWG_Training_Guidelines_Recommendations_v1.1_2010_11_18.pdf
- Farah, M. J. (1991). Cognitive neuropsychology: Patterns of co-occurrence among the associative agnosias: Implications for visual object representation. *Cognitive Neuropsychology*, 8(1), 1-19. doi:10.1080/02643299108253364
- Fysh, M. C., & Bindemann, M. (2017). Effects of time pressure and time passage on face-matching accuracy. *Royal Society open science*, 4(6), 170249. doi:10.1098/rsos.170249
- Germine, L. T., Duchaine, B., & Nakayama, K. (2011). Where cognitive development and aging meet: Face learning ability peaks after age 30. *Cognition*, 118(2), 201-210. doi:10.1016/j.cognition.2010.11.002
- Henderson, Z., Bruce, V., & Burton, A. M. (2001). Matching the faces of robbers captured on video. *Applied Cognitive Psychology*, 15, 445-464. doi:10.1002/acp.718
- Heyer, R., MacLeod, V., Carter, L., Semmler, C., & Ma-Wyatt, A. (2011). *Profiling the facial identification practitioner in Australia: Report on the human operator capability project survey. Report for the Department of Prime Minister and Cabinet: Grant PR09-0078*. Retrieved from

https://www.researchgate.net/publication/318206064_Profiling_the_Facial_Comparis_on_Practitioner_in_Australia

Heyer, R., Semmler, C., & Hendrickson, A. T. (2018). Humans and Algorithms for Facial Recognition: The Effects of Candidate List Length and Experience on Performance. *Journal of Applied Research in Memory and Cognition*, 7(4), 597-609. doi:10.1016/j.jarmac.2018.06.002

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: a failure to disagree. *American Psychologist*, 64(6), 515-526. doi:10.1037/a0016755

Kemp, R., Towell, N., & Pike, G. (1997). When seeing should not be believing: Photographs, credit cards and fraud. *Applied Cognitive Psychology*, 11(3), 211-222. doi:10.1002/(SICI)1099-0720(199706)11:3<211::AID-ACP430>3.0.CO;2-O

Klontz, J. C., & Jain, A. K. (2013). A case study on unconstrained facial recognition using the boston marathon bombings suspects. *Michigan State University, Tech. Rep*, 119(120), 1.

Kundel, H. L., Nodine, C. F., Conant, E. F., & Weinstein, S. P. (2007). Holistic component of image perception in mammogram interpretation: gaze-tracking study. *Radiology*, 242(2), 396-402. doi:10.1148/radiol.2422051997

Martens, F., Bulthé, J., van Vliet, C., & de Beeck, H. O. (2018). Domain-general and domain-specific neural changes underlying visual expertise. *NeuroImage*, 169, 80-93. doi:10.1016/j.neuroimage.2017.12.013

Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255-260. doi:10.1016/S1364-6613(02)01903-4

Megreya, A. M., & Burton, A. M. (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition*, 34(4), 865-876. doi:10.3758/BF03193433

Megreya, A. M., & Burton, A. M. (2007). Hits and false positives in face matching: A familiarity-based dissociation. *Perception & psychophysics*, 69(7), 1175-1184. doi:10.3758/BF03193954

Moore, R. M., & Johnston, R. A. (2013). Motivational incentives improve unfamiliar face matching accuracy. *Applied Cognitive Psychology*, 27(6), 754-760. doi:10.1002/acp.2964

Norell, K., Låthén, K. B., Bergström, P., Rice, A., Natu, V., & O'Toole, A. (2015). The effect of image quality and forensic expertise in facial image comparisons. *Journal of Forensic Sciences*, 60(2), 331-340. doi:10.1111/1556-4029.12660

Noyes, E., Phillips, P. J., & O'Toole, A. J. (2017). What is a super-recogniser. *Face processing: Systems, disorders and cultural differences*, 173-201.

Özbek, M., & Bindemann, M. (2011). Exploring the time course of face matching: Temporal constraints impair unfamiliar face identification under temporally unconstrained viewing. *Vision research*, 51(19), 2145-2155. doi:10.1016/j.visres.2011.08.009

Papesh, M. H. (2018). Photo ID verification remains challenging despite years of practice. *Cognitive research: principles and implications*, 3(1), 19. doi:10.1186/s41235-018-0110-y

Papesh, M. H., Heisick, L. L., & Warner, K. A. (2018). The persistent low-prevalence effect in unfamiliar face-matching: The roles of feedback and criterion shifting. *Journal of Experimental Psychology: Applied*, 24(3), 416-430. doi:10.1037/xap0000156

Phillips, P. J., Heyer, R., & Michalski, D. (2018). Effectiveness of a Facial Forensic Training Course. *Journal of Vision*, 18(10), 560-560. Poster presented at the Vision Sciences Society Eighteenth Annual Meeting.

Phillips, P. J., Wechsler, H., Huang, J., & Rauss, P. J. (1998). The FERET database and evaluation procedure for face-recognition algorithms. *Image Vision Computing*, 16(5), 295–306. doi:10.1016/S0262-8856(97)00070-X

Phillips, P. J., Yates, A. N., Hu, Y., Hahn, C. A., Noyes, E., Jackson, K., ... & Chen, J. C. (2018). Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms. *Proceedings of the National Academy of Sciences*, 115(24), 6171-6176. doi:10.1073/pnas.1721355115

Prince, J. P. (2012). *Report on emerging use of facial recognition systems and facial image comparison procedures*. Retrieved from https://www.churchilltrust.com.au/media/fellows/2012_Prince_Jason.pdf

Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, 22(3), 341-365. doi:10.1017/S0140525X99002022

Ramon, M., Bobak, A. K., & White, D. (2019). Super-recognizers: From the lab to the world and back again. *British Journal of Psychology*. doi:10.1111/bjop.12368

Rice, A., Phillips, P. J., Natu, V., An, X., & O'Toole, A. J. (2013). Unaware person recognition from the body when face identification fails. *Psychological Science*, 24(11), 2235-2243. doi:10.1177/0956797613492986

Robertson, D. J., Noyes, E., Dowsett, A. J., Jenkins, R., & Burton, A. M. (2016). Face recognition by metropolitan police super-recognisers. *PloS one*, 11(2), e0150036. doi:10.1371/journal.pone.0150036

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2), 135-146.

Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic bulletin & review*, 16(2), 252-257. doi:10.3758/PBR.16.2.252

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology*, 75(2), 221-242. doi:10.1111/j.2044-8295.1984.tb01895.x

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2), 127-190. doi:10.1037/0033-295X.84.2.127

Stacchi, L., Huguenin-Elie, E., Caldara, R., & Ramon, M. (2019). Normative data for two tests of face matching under ecological conditions. doi:10.31234/osf.io/swjrq

Susilo, T., Germine, L., & Duchaine, B. (2013). Face recognition ability matures late: Evidence from individual differences in young adults. *Journal of experimental psychology: human perception and performance*, 39(5), 1212-1217. doi:10.1037/a0033469

Thompson, M. B., & Tangen, J. M. (2014). The nature of expertise in fingerprint matching: experts can do a lot with a little. *PloS one*, 9(12), e114759. doi:10.1371/journal.pone.0114759

Towler, A., Kemp, R. I., Burton, A. M., Dunn, J. D., Wayne, T., Moreton, R., & White, D. (2019). Do professional facial image comparison training courses work?. *PloS one*, 14(2), e0211037. doi:10.1371/journal.pone.0211037

Towler, A., White, D., Ballantyne, K., Searston, R. A., Martire, K. A., & Kemp, R. I. (2018). Are forensic scientists experts?. *Journal of Applied Research in Memory and Cognition*, 7(2), 199-208. doi:10.1016/j.jarmac.2018.03.010

Towler, A., White, D., & Kemp, R. I. (2017). Evaluating the feature comparison strategy for forensic face identification. *Journal of Experimental Psychology: Applied*, 23(1), 47. doi:10.1037/xap0000108

White, D., Dunn, J. D., Schmid, A. C., & Kemp, R. I. (2015). Error rates in users of automatic face recognition software. *PLoS One*, 10(10), e0139827. doi:10.1371/journal.pone.0139827

White, D., Kemp, R. (2016). Studying human performance to improve accuracy of biometric facial recognition. *Keesing Journal of Documents & Identity, Annual Report Identity Management 2016-2017*. Retrived from <https://www.linkedin.com/pulse/studying-human-performance-improve-accuracy-biometric-david-white/>

White, D., Kemp, R. I., Jenkins, R., & Burton, A. M. (2014). Feedback training for facial image comparison. *Psychonomic bulletin & review*, 21(1), 100-106. doi:10.3758/s13423-013-0475-3

White, D., Kemp, R. I., Jenkins, R., Matheson, M., & Burton, A. M. (2014). Passport officers' errors in face matching. *PLoS One*, 9(8), e103510. doi:10.1371/journal.pone.0103510

White, D., Phillips, P. J., Hahn, C. A., Hill, M., & O'Toole, A. J. (2015). Perceptual expertise in forensic facial image comparison. *Proceedings of the Royal Society B: Biological Sciences*, 282(1814), 20151292. doi:10.1098/rspb.2015.1292

White, D. (2020). Do professional groups outperform novices in unfamiliar face matching tasks? A meta-analysis. *PsyArXiv*.

Wilmer, J. B. (2017). Individual differences in face recognition: A decade of discovery. *Current Directions in Psychological Science*, 26(3), 225-230. doi: 10.1177/0963721417710693

Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., ... & Duchaine, B. (2010). Human face recognition ability is specific and highly heritable. *Proceedings of the National Academy of sciences*, 107(11), 5238-5241. doi:10.1073/pnas.0913053107

EXPERTISE IN UNFAMILIAR FACE MATCHING

Wirth, B. E., & Carbon, C. C. (2017). An easy game for frauds? Effects of professional experience and time pressure on passport-matching performance. *Journal of Experimental Psychology: Applied*, 23(2), 138-157. doi:10.1037/xap0000114

Young, A. W., & Burton, A. M. (2018). Are we face experts? *Trends in Cognitive Sciences*, 22(2), 100-110. doi:10.1016/j.tics.2017.11.007

Yovel, G., Halsband, K., Pelleg, M., Farkash, N., Gal, B., & Goshen-Gottstein, Y. (2012). Can massive but passive exposure to faces contribute to face recognition abilities? *Journal of Experimental Psychology: Human Perception and Performance*, 38(2), 285–289. doi:10.1037/a0027077