Part I **Identification of the Living**

Chapter 1

Familiar face recognition

Vicki Bruce

1.1 Introduction

The task of the police officer investigating a crime, or of the forensic anthropologist working with human remains, is to establish an identity of a criminal or of a victim. This often involves working with images of faces – building a composite image from the memory of a witness to the crime, seeking CCTV images of the person or persons who might have committed the crime, or building a model of the face of an unknown person from their skull. Many of the chapters in this volume describe the processes involved in such reconstructions.

Once an image of a face is obtained, however, it needs to be identified. TV programmes such as Crimewatch in the UK (see www.bbc.co.uk/crimewatch) often display witness-generated facial composites or CCTV images of people associated with crimes which sometimes result in people volunteering new information; Richard Neave's constructed image of the unidentified victim 115 of the London Underground tube station fire at Kings Cross in 1987 (Chambers 2007) was displayed widely in newspapers and posters. CCTV images of the London 'nail bomber' in 1999 led to hundreds of responses, including one phone call from a man who named a work colleague, David Copeland, as resembling the images. This crucial lead quickly led to the arrest of David Copeland, unfortunately not soon enough to prevent the third devastating bomb attack on The Admiral Duncan, a London pub.

While the problems of building or using images of unfamiliar faces are the subject of many other chapters in this volume, here I focus on relevant issues about the identification of familiar faces, and also elaborate on the contrast between familiar and unfamiliar face recognition.

1.2 Familiar versus unfamiliar face recognition

When police investigate crimes the identity of the villain or villains is often unknown, and the police will question witnesses who may be asked in different ways to assist with establishing the identity of a person who was unfamiliar to them before the criminal incident. When a witness to a crime is asked to try to identify the criminal from photographs or a line-up, or to build an image of the face using a face composite system, this is typically difficult and error-prone. In contrast, most of us manage to recognise the faces of familiar people in our everyday lives or on the television or other media reasonably accurately, despite the enormous variation shown between different appearances of the same person. Error-prone unfamiliar face memory is transformed to a generally very reliable performance with known faces.

There is some neuropsychological evidence that somewhat different processes underlie the recognition of unfamiliar faces and the recognition of familiar ones. 'Prosopagnosic' people have difficulties recognising faces, sometimes as a result of brain injury. They may be unable to recognise even their closest relatives from their faces, and even their own face in the mirror may appear unfamiliar to them (see Harris and Aguirre, 2007, for a brief overview of this condition, and Duchaine and Nakayama, 2006 for one of many recent papers on congenital or developmental prosopagnosia). But some prosopagnosic patients can manage to match images of unfamiliar faces, albeit using laborious and

Craniofacial Identification, ed. C. Wilkinson and C. Rynn. Published by Cambridge University Press.

sometimes time-consuming processes. Young and his colleagues (1993) investigated residual deficits in face processing in a group of 34 brain-injured war veterans. Amongst this group they found one man who was impaired on familiar face recognition, but unimpaired at unfamiliar face matching, and another who was impaired on unfamiliar face matching but recognised familiar faces quite normally. This pattern of 'double dissociation' is consistent with the idea that different brain areas and/or processes are involved in the two tasks of recognising familiar faces and matching unfamiliar ones.

Further evidence for the dissociation between familiar and unfamiliar face processing is the observation that familiar face recognition seems to depend on rather different kinds of information from unfamiliar recognition. Lander and Butcher (Chapter 11) review work suggesting that representations of familiar faces capture something about their characteristic patterns of motion as well as static form. Even within static form, there appear to be different emphases in the visual representations of familiar and unfamiliar faces. Ellis *et al.* (1979) first demonstrated that unfamiliar face recognition was dominated by the external features of the face, including the hair. People find it very difficult to recognise once-viewed faces if their hair is concealed or changed. However, when faces are familiar their internal features are more important in their recognition. This dominance of the external over internal features for unfamiliar faces has been much-replicated (e.g. Young *et al.*, 1985*b*) and it is found in tasks that involve just matching unfamiliar faces without any memory load (Bruce *et al.*, 1999). O'Donnell and Bruce (2001), using newly familiarised faces, showed that it seemed to be the eyes in particular that became better represented in recently familiarised faces.

Why should representations of familiar faces become shifted towards internal and away from external features? For an unfamiliar face, at least in Western Europe where hair is generally visible and very variable in style and colour, the outer features of the face probably convey the most information that will be useful for matching and memory. As we see a face more frequently, however, we will attend more to its internal features, to see direction of gaze, expression and lip movements, which are all important for other social functions (see Bruce and Young, 1998 for an introductory overview). Moreover, hairstyle will vary from one occasion to another, and people may sometimes wear hats or scarves. Thus it is not surprising that the visual memory system begins to weight the internal features more strongly. Consistent with this, Megreya and Bindemann (2009) showed that when face recognition was tested in Egypt, adults, but not children, showed better memory for internal features for unfamiliar as well as familiar faces, probably because the commonplace wearing of head scarves means that attention is always oriented more towards internal than external features in that culture.

The relative importance of the internal features of a familiar face, particularly the eyes, has implications for the effectiveness of different kinds of disguises. The true identities of Batman, and The Lone Ranger, were effectively concealed by their wearing of masks covering their eyes, in addition to hats that covered their hair. Sunglasses and hats can allow celebrities anonymity in a crowd.

Burton and his colleagues (2005) have suggested that representations of familiar faces are built up by simply averaging together individual instances of seen faces. The differential weightings of internal features during familiarisation, and for cultures where external features are not often viewed, will likely arise as a result of selective attention weighting areas of the face differently. An averaging mechanism with the addition of selective weighting through attentional mechanisms would allow the representation of a face to develop in a way which allows the face to be well-recognised, despite variations in expression and viewpoint which are extremely detrimental to recognition of unfamiliar faces (e.g. Bruce, 1982). Burton and his colleagues have shown that people are faster to recognise familiar faces when shown an average of 20 different photographic instances of them, than when shown individual instances alone. Moreover, in an interesting practical extension of this work, Jenkins and Burton (2008) showed that when an average of 20 instances of each of a set of famous face targets was used as a probe to match against a database of images of over 3000 celebrities, performance increased from 51% correct matching obtained from using individual image probes, to 100% correct when the averages were used as probes. From the evidence of this preliminary study Jenkins and Burton suggest that average images should be used on identity documents such as passports to improve both human and computer use of such images.

So, our representations of unfamiliar faces make it very difficult to generalise to novel views, expressions and contexts, while our representations of familiar faces transcend these limitations, perhaps because of consolidation of different instances to an average. We can recognise familiar faces well, even from low-quality images. Harmon (1973) showed good recognition of familiar faces in pixellated images where only very coarse-scale information about identity was preserved. And Burton and his colleagues (1999) showed that students could recognise with high accuracy very poor quality CCTV images of their lecturers *provided* that images of the face area were visible, even though these were very degraded. Indeed we are so good at recognising familiar people from moving degraded images that there must be real concern about the efficacy of attempts to conceal identities in TV clips showing witnesses or children. TV editors working on filmed footage are unfamiliar with the people depicted, and it is possible that the degree of blurring or pixellation applied, which seems satisfactory to them, will be insufficient to render these faces unrecognisable to familiar observers (Lander *et al*., 2001).

The representation of familiar faces emphasises internal features, particularly the eyes. Several recent studies have demonstrated that prosopagnosic people do not look at or use the eyes/upper face features when looking at faces and trying to recognise them (Caldara *et al*., 2005; Bukach *et al*., 2008). This is likely to be a symptom rather than a cause of their deficit, however, since prosopagnosic people are often completely normal at recognising facial expressions and perceiving eye gaze (e.g. Young *et al*., 1993; Duchaine *et al*., 2009), both of which also use information from the eye regions of the face.

In many other respects, though, the perception of familiar faces seems very similar to that of unfamiliar faces. Both familiar and unfamiliar faces appear to be processed 'configurally' (see Chapter 2) though more piecemeal processing may be beneficial in certain kinds of matching task (see later). The recognition of both familiar and unfamiliar faces is impaired by unusual transformations such as inversion and negation, or by changes in lighting direction which go beyond those usually experienced. Such extreme variations in orientation, contrast and luminance create changes well beyond the usual range of variation encountered for an individual face and which would therefore not be incorporated in the internal representation for a familiar one. Johnston and Edmunds (2009) provide a recent review of factors affecting unfamiliar and familiar face recognition.

1.3 Individual differences in face recognition

In recent years there has been renewed interest in the theoretical and practical implications of individual differences in face recognition. As with all other psychological abilities, tests of a reasonable sample of participants will reveal a range of performance scores. Recent interest has focused on understanding reasons for this variation (general ability? or something specific to face processing?). Woodhead and Baddeley (1981) noted a wide variation in abilities in face-recognition tests ranging from performance in discriminating familiar from novel faces at or near chance (d' = 0) to performance which was near perfect (d' = 6.8). Testing a group of relatively good and relatively poor recognisers again later, they found a significant difference in face recognition abilities remained, and the groups also differed slightly in their recognition memory ability for non-face pictures, but there was no difference in their verbal memory abilities.

Differences are observed in face-matching ability when there is no memory component at all. Megreya and Burton (2006) used a face-matching task introduced by Bruce *et al.* (1999), where participants must decide which, if any, of 10 faces in an array matches a good-quality target photograph. Performance on such tasks is surprisingly prone to error, but also quite variable (Megreya and Burton found mean performance of 82% correct and a standard deviation of 12%). An even simpler task, the Glasgow Face Matching Task (GFMT) requires participants to decide if two face images are the same or different people (see Figure 1.1). Using the short form of this test with 194 volunteers, Burton *et al.* (2010) report a mean of 81.3% and standard deviation of 9.7 on this task, which has been replicated in Newcastle.[1] In this sample (Burton *et al*., 2010), and in the samples tested in Newcastle about 10% of participants scored between 51% (around chance) and 70% correct, while the best 10% scored above 95% correct on the test.

Megreya and Burton (2006) examined what factors correlated with performance in the array-matching task. There were significant positive correlations between unfamiliar face matching and recognition memory for unfamiliar faces, and the matching task also correlated with a number of other tasks of visual memory and processing. The matching task did *not* correlate, however, with participants' abilities to

Figure 1.1 Are the two people shown the same or different? This can be a surprisingly difficult task. Figure courtesy of Mike Burton, Glasgow University, and also published in Burton *et al.* (2010).

recognise famous or recently familiarised faces – unless these were turned upside down! Indeed the best predictor of performance on matching unfamiliar faces when these were shown *upright*, was conducting the same or similar tasks with unfamiliar or familiar faces when these were *inverted*. Now we know that when faces are inverted, participants must rely more on analysing faces in a piecemeal way since the 'configural' processing that characterises upright face perception is impaired when faces are inverted (e.g. see Young *et al.*, 1987; Leder and Bruce, 2000). Megreya and Burton's finding suggests that matching unfamiliar faces well requires the ability to analyse local features and ignore the overall impression gained by more configural processing.

This suggestion that unfamiliar face matching requires a specific analytic processing style does not, however, account for the broader range of abilities to recognise faces from memory. Prosopagnosic people, who have difficulties recognising familiar faces in their everyday lives, seem to have specific deficits in processing 'configural', not local, face features. More recently, Russell *et al.* (2009) have studied volunteers who are particularly good at recognising faces: 'Super-recognisers'. These score extremely highly on the Cambridge Face Matching, as well as the Cambridge Face Memory tests (tests of unfamiliar face processing), but are not particularly skilled at tasks of inverted face recognition.

This suggests that skill with recognising familiar faces and remembering unfamiliar faces in some tasks relies on expertise in configural processing, but that other tasks involving faces, particularly matching unfamiliar faces, *also* require the ability to analyse local features well. Indeed these observations themselves reinforce earlier theoretical suggestions that there is a distinction between the processing of familiar and unfamiliar faces (Bruce and Young, 1986). If prosopagnosia is largely an impairment of the 'configural' processing of faces – on which most people are expert – then the observation that some prosopagnosics can perform well on unfamiliar tasks that might rely more on piecemeal processing is less surprising.

1.4 Stages in person recognition

When we meet someone we know well, and are expecting to see, then person recognition is usually effortless and complete. 'Hello Vicki' my father used to say, immediately, even though at the age of 90 he had very poor vision. But this facility in familiar person recognition belies a number of discrete stages in the process that can be revealed through systematic recording of everyday behaviour or more laboratory-based study.

Young *et al.* (1985*a*) collected a large number of instances of errors and difficulties in person recognition in a study that they imaginatively titled 'The faces that launched a thousand slips'. They asked 22 volunteers to note down any difficulties they experienced when recognising people for a period of 8 weeks, and they analysed the 922 recorded incidents from the 7 weeks following the initial familiarisation week. The first thing to note is that this equates to almost one incident per day, on average – showing that everyday person recognition is far from problem-free. Next was the grouping of incidents into distinct types. Most common (about a third of all the records) were when a person was misidentified – one person was confused with someone else, or a stranger was taken to be someone familiar, e.g. '. . . I saw a person with a dog and I thought it was a dog owner I sometimes see there. It was the wrong type of dog: I thought he must have got a new one!' Next most common – approximately 25% of all records – occurred where someone appeared to be familiar but the volunteer was unable to remember why they were familiar. Sometimes this situation was resolved by a protracted process, e.g. by the person in the bank who 'saw a person and I knew there was something familiar immediately. After a few seconds I realised she was from a shop on campus or a secretary of one of the departments. I eventually remembered by a process of elimination'. And the third most common type – 21% of records – was where it proved difficult or impossible to retrieve a full identification of an encountered person – often it was the name that proved elusive, e.g. someone recognised an actress from a poster and 'knew what films the actress was in and knew she does a lot of Beckett, but it was another minute before I could remember her name'.

These three types – misidentification, familiar only, and difficulty in retrieving full details together account for almost 80% of the incidents. While not all the records involved face recognition (because there were some concerning, for example, failures to recognise voices), these three kinds of difficulty map well onto a broad 'flow model' of the stages involved in recognising someone from their face, shown in Figure 1.2.

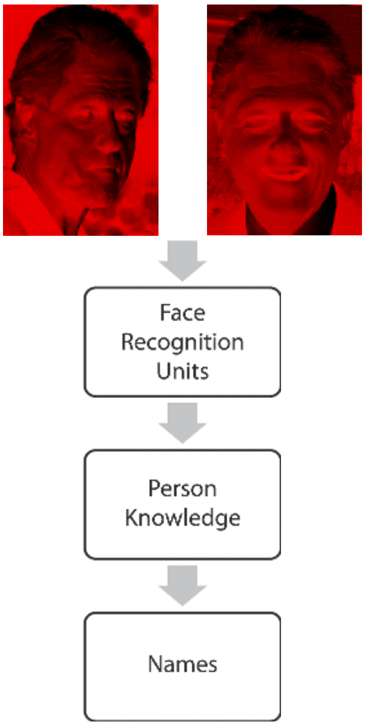

Figure 1.2 A simple three-stage model of person identification from the face.

This flow model describes three broad stages in person recognition. Step 1 is to match the incoming visual pattern against a stored visual representation of what that person looks like. If there is a match at this level then the pattern may be recognised as 'familiar'.[2] Step 2 involves retrieving information about why the person is familiar – where you know them from, what they do for a living and so forth. The final step, according to this simple 'stage' model, involves retrieving their name.

The flow model shows the steps occurring in this strict order. There is no route directly from the pattern-matching stage to the name, for example. There is a good deal of evidence that appears consistent with this proposal that name retrieval cannot occur before something else about the person is known. In the diary study (Young *et al.*, 1985*a*), there was not a single recorded error or difficulty in person recognition where a face was named but there was no knowledge of who they were. People never say things like 'I know that face – it's John Lennon – but I have no idea who John Lennon is or where I have seen this face before'. A more formal test of this strict sequence was conducted by Brennen *et al.* (1990) in an experiment which required that people answer a series of questions about celebrities from definitions – as they might if they were playing Trivial Pursuit, for example (e.g. What's the name of the person who played the nervous man in the shower scene in Hitchcock's *Psycho*?). For some questions, people knew the answers. Others, they knew they didn't know. But sometimes they felt they knew the answer but the name was 'on the tip of their tongue' and they just couldn't get to it. In this case, according to the sequence of steps, the participant is at Step 2, and stuck on the process of getting beyond that to the name. So, what happens if you see a picture of the person whose name you are trying to retrieve? Brennan *et al.* compared the correct names retrieved from tip-of-the-tongue states when a picture of the unnamed person was shown with when the original question was just repeated. Seeing the face gave no advantage at all. Indeed it appeared to annoy participants, who would say, through gritted teeth, such things as 'I know what he *looks* like, I just can't remember the name!' (Anthony Perkins, by the way.)

So it looks as though identifying a familiar person has (at least) three stages and problems can arise at any one of them. A person may be mistakenly not recognised at all, or they may be misidentified because the visual pattern gets matched to the wrong stored visual representation. (For example, I have recently discovered that I had a long-standing confusion in my own mind between two actresses – Zoe Wanamaker and Barbara Flynn – who I thought were both Zoe Wanamaker. Indeed it was only when my partner denied that we were watching Zoe Wanamaker in an old episode of the crime thriller *Cracker*, leading me to search google web images, that I was convinced that these really were two different women. So, in terms of our flow-chart, all the times I looked at Barbara Flynn and thought she was Zoe Wanamaker I was matching the pattern of her face against the wrong visual representation. Moreover, through doing this, I was creating an internal representation for these two women which was some kind of average of the two of them.) In a more important forensic context, there is plenty of scope for a police officer, or a member of the public viewing TV images, to mistakenly identify someone as a person they know, given a certain resemblance and, perhaps, some context (see later).

At the second stage, a person may seem familiar, correctly, but the details of why they are familiar may be elusive, or misremembered. One of the more intriguing things about the 'super-recogniser' study (Russell *et al.*, 2009) was their claimed ability not just to remember faces, despite a considerable lapse of time, but to know *why* they seemed familiar too. One of the super-recognisers remarked 'I've learned to stop surprising people with bizarre comments like, "Hey, weren't you at that so-and-so concert last fall . . . I recognise you"' and another said 'I do have to pretend that I don't remember, however, because it seems . . . that they mean more to me than they do when I recall that we saw each other once walking on campus four years ago in front of the quad!'. Remembering why someone is familiar may involve retrieving specific *episodic* information, such as this (I saw him then and there) or it may be more general '*semantic*' information – that's the person who reads the news on TV. The simple stage model outlined here does not adequately distinguish between these two different kinds of knowledge we may have of why a face is familiar to us. These differences are important in a forensic context. We might see a CCTV image of a person wanted in connection with a crime and 'know' that is our next-door neighbour, or the person we buy our papers from, or we may 'know' this is the same person we interviewed in connection with a similar incident last year. But when it comes to the provision of an alibi, for example, we may need to remember not just who a particular person is, when questioned, but that we definitely did see them on their regular train to work that morning at the same time they were supposed to be committing a crime somewhere else.

The most difficult thing about identifying familiar people is remembering their names, even when these are highly familiar to us. There's something rather strange about names, even when they are short, simple and concrete. McWeeney *et al.* (1987), for example, discovered that people found it harder to learn the same lexical items when presented as names ('Mr Baker') than when presented as occupations ('a baker'). The simple stage model places names as the final stage in a sequence of identification steps, and there is certainly plenty of evidence consistent with the 'last stage' account. First, as noted above, names are never retrieved without knowing something else about the person. Second, tasks which require names to be retrieved or verified are conducted more slowly than tasks that require other kinds of information to be used. For example, Burton, Jenkins and McNeill (2002; see Calderwood and Burton, 2006) repeatedly showed pictures of just four famous faces – two pop stars and two politicians, two called Peter and two called Paul. Even after extended practice participants were slower to say the name than the occupation of each face that they saw. A control study showed that other participants were no slower to read the names than the occupations out loud suggesting it is not because of difficulties articulating these labels. Johnston and Bruce (1990) asked participants to make speeded judgements about whether pairs of faces shared the same first name, or shared properties such as nationality or being dead or alive (e.g. if on a trial the participants saw pictures of John Lennon and James Dean they would respond 'no' if the question was whether they shared a name; and 'no' if they were asked if they matched nationality or not, but 'yes' to the question of matching on dead–alive). Again, despite the use of a set of just eight repeated, highly familiar faces, the decisions which involved names were made more slowly, even than decisions about whether the two depicted individuals were both dead, or both alive.

However, the slowness and error-prone nature of name retrieval need not mean that they are reached 'after' other person information, as in Figure 1.2. It may just be that names are generally available later than such information. Burton and Bruce (1992) demonstrated, using an 'interactive activation with competition' (IAC) model, that empirically observed effects could be modelled without the need to posit an extra stage for names. The IAC model of person recognition supposes that faces (or voices, or names) are recognised as familiar when sufficient activation is present within 'person identity nodes' (PINS) which pool activity from modality-specific pattern processors that respond separately to faces, voices and so forth. Personal information, including such things as occupations, nationalities and so forth, is retrieved via the PINS. Burton and Bruce demonstrated that if proper names were treated just the same way as other things known about people, and placed within the same pool of personal information, they naturally became activated more slowly than other kinds of information. This is because a name is a unique item of information about identity. We know only one Tony Blair and one Margaret Thatcher. Occupations, however, even rare ones like Prime Minister, are shared by several, sometimes many, people. Thus the PINS for Tony Blair and Margaret Thatcher (and Winston Churchill and Gordon Brown, etc.) will all have a link to the piece of information about 'Prime Minister'. These multiple links make activation in shared nodes rise more quickly than in unique nodes, simulating neatly the differential ease of determining occupations and other types of information compared with names.

While the IAC model incorporates activation between different pools there is inhibition between all units within a single pool (there is 'competition' as well as 'interactive activation'). In the case of names, though, the original Burton and Bruce (1992) proposal cannot be quite right. Bredart *et al.* (1995) demonstrated that there cannot be a *single* pool of undifferentiated information covering names, occupations, nationalities, favourite foods and so forth, because if this was the case then it would be harder to retrieve information about people we know a lot about compared with people we know rather little about, because additional information would inhibit the activation of any particular semantic unit, including that for the name. Bredart and his colleagues (1995) suggest instead that there are separate pools for distinct kinds of person identity – a pool for nationality, a pool for occupation, a pool for names – and inhibition occurs within each pool separately.

The suggestion that names are at the same stage as other personal information and retrieved more slowly due to their connection patterns implies that in certain contexts, if the connection to a name is strong enough, it might be retrieved more quickly than, say, occupation. Calderwood and Burton (2006) have demonstrated this with personally familiar people. In line with our intuitions, perhaps, it is indeed easier to recall your partner's name than his/her occupation.

Although the issue about memory for names seems to be a rather theoretical one, there are implications of this for investigative work. It is important to note that remembering names, even of quite familiar people, is difficult and prone to error. While the police seek a name for a person seen at a crime, images from cameras or composites may trigger less specific cues to identity – 'the bloke I used to work with' – and these should be seen as important leads even when names are not known or remembered.

1.5 Contextual factors in person recognition

We are often surprised to meet people we know in unexpected places, and we may fail to recognise them completely under certain circumstances. Australian psychologist Don Thomson (1986), contrived a situation in which parents even failed to recognise their own daughter, by exploiting her unexpected presence in London when the parents were visiting there, and instructing her to stand near their hotel with an unfamiliar accomplice. The daughter was told to show not a single sign of recognition when her parents rushed up to greet her, and as a result the parents apologetically moved on. Although this seems unlikely, imagine you saw Paul McCartney in your local launderette – you would probably assume it wasn't him, but was just a striking resemblance. Don Thomson's friends clearly recognised their daughter, but then decided instead that she was someone who *looked like* (incredibly like) her. It is not quite clear, though, whether an inappropriate or unlikely context can make an otherwise familiar face seem completely unfamiliar and it would be difficult to contrive an experimental test of this.

Context can play some strange tricks, however. The same Don Thomson was once accused of a serious sexual assault and this seems to have been because he was appearing on the television while the unfortunate woman was being raped. His face therefore became associated with the incident. The current Labour MP for Neath and former government minister, Peter Hain, was accused of robbing Barclays Bank in 1975 – when he was a postgraduate student and leader of the Young Liberals and an anti-apartheid activist. He had been briefly in the vicinity of the bank at the time it was robbed, as he had been buying a new typewriter ribbon, and so it would have been possible for passers-by to note his presence out that day. There have even been suggestions that the robbery may have been deliberately staged using a look-alike of Peter Hain, whose political activities were not welcomed by the (then) apartheid regime in South Africa. Whatever the truth of such suggestions, newspapers published his picture with a headline about Hain's arrest *before* the main witness – the bank cashier – was shown a line-up in which he appeared, making it extremely likely that the witness could have found his face familiar, but not from the crime scene itself. Indeed even without the newspaper headline, someone who has some degree of celebrity from media coverage might be particularly susceptible to being mistakenly associated with a crime by virtue of their apparent familiarity.

The research by Young and his colleagues (1985*a*) on everyday errors and difficulties in person recognition, and other laboratory research by Hay, Young and Ellis (1991) and Hanley and associates (e.g. Hanley and Cowell, 1988) has demonstrated clearly that people quite often judge a face to be familiar without remembering why. Approximately 5–10% of attempts to identify a celebrity (Hanley and Cowell, 1988) led to judgements that the face was familiar, with no further information forthcoming. It is therefore very important that eyewitnesses asked to scrutinise faces in a line-up should not have been exposed to the person's face earlier, and that people who could be familiar to an eyewitness through general media exposure or from otherwise innocent contexts should not be placed in a line-up at all. In several cases reviewed by the Devlin committee (Devlin, 1976) witnesses were asked to observe a line-up having already been shown photographs of the suspect. For example, in the case of George Ince, who was tried in 1973 for murder following an armed robbery, the key witness who picked him out from an identity parade had already been shown photographs of him prior to the line-up.

Appropriate contextual information can also facilitate or 'prime' recognition of highly familiar faces in a way that speeds up familiarity decisions. Bruce and Valentine (1986) first demonstrated that faces from closely associated pairs – such as Stanley Laurel from 'Laurel and Hardy' were judged familiar more quickly if they immediately followed an image of their associated partner's face. This priming effect suggests that the representations of different people in memory are interconnected in some way. The IAC model introduced earlier provides one account of such priming effects.

Sinha and Poggio (1996) discussed a contextual priming effect possibly related to this. At the time it was published, Bill Clinton was President and Al Gore Vice-President of the USA. Viewers of the picture at that time quite readily recognised Al Gore standing beside Bill Clinton, even though the image shown actually has Bill Clinton's internal face features pasted into Al Gore's hairstyle. This demonstration depended upon the relative unfamiliarity of Al Gore at that time (hence the dominance of external face features given the appropriate context for his appearance). One way of explaining this would be to suggest that the recognition of Clinton's face produced activation of the recognition unit for Gore's which therefore required less evidence (and in this case, less accurate evidence) to become active too. Such demonstrations suggest that we might be convinced we have seen someone we expect to see at a particular time or place when actually we have only seen someone who resembles them.

It is clear that contextual factors can be very important in our recognition of familiar people. In a criminal context it may be extremely important that people are able to remember when or where they saw someone. A witness to a crime must know that the person in the line-up is familiar from the scene of the crime (and not from the newspapers, as in Peter Hain's case). And an alibi may be asked to testify that they saw a particular acquaintance in a place (not the crime scene) at a specified time, so they too must remember the context of a recent encounter with a known individual. While familiar people's identities are often bound up with places where they are seen ('are you on the telly?', or 'that's the man I see walking his dog') we know rather little about how well we are able to remember specific episodes involving particular people. I see the same elderly gentleman walking his greyhound most weekends, when my own walks are later than during the week. But was it last Saturday or Sunday that I saw him last? I couldn't say.

1.6 Conclusions

We are good at familiar face recognition. Members of the public, witnesses or police who claim to recognise a person in a CCTV image or other context as someone they know should be taken very seriously. This is why getting a good image of a face, via composite construction or CCTV, and broadcasting it widely, is so important to investigations. But we are not perfect. We may confuse two people, be unable to remember their names, or be misled by context. A witness must know that someone was familiar from the scene of the crime, rather than some other place. And the alibi must do more than recognise someone, they must be able to swear they were in a particular place at a particular time. This memory for time and place is a quite different ability, and one which we know less about, but may be absolutely critical in an investigation.

References

Bredart, S., Valentine, T., Calder, A. J. and Gassi, L. (1995). An interactive activation model of face naming. *Quarterly Journal of Experimental Psychology*, **48**, 466–486.

Brennen, T., Baguley, T., Bright, J. and Bruce, V. (1990). Resolving semantically induced tip-of-the-tongue states for proper nouns. *Memory and Cognition*, **18**, 339–347.

Bruce, V. (1982). Changing faces – visual and non-visual coding processes in face recognition. *British Journal of Psychology*, **73**, 105–116.

Bruce, V., Henderson, Z., Greenwood, K. *et al.* (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, **5**, 339–360.

Bruce, V. and Valentine, T (1986). Semantic priming of familiar faces. *Quarterly Journal of Experimental Psychology*, **38A**, 125–150.

Bruce, V. and Young, A. W. (1986). Understanding face recognition. *British Journal of Psychology*, **77**, 305–328.

Bruce, V. and Young, A. (1998). *In the Eye of the Beholder: The Science of Face Perception*. Oxford: Oxford University Press.

Bukach C. M., Le Grand, R., Kaiser, M. D., Bub, D. N. and Tanaka, J. W. (2008). Preservation of mouth region processing in two cases of prosopagnosia. *Journal of Neuropsychology*, **2**, 227–244.

Burton, A. M., and Bruce, V. (1992). I recognize your face but I can't remember your name – a simple explanation. *British Journal of Psychology*, **83**, 45–60.

Burton, A. M., Bruce, V. and Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology*, **81**, 361–380.

Burton A. M., Jenkins, R., Hancock, P. J. B. and White, D. (2005). Robust representations for face recognition: the power of averages. *Cognitive Psychology*, **51**, 256–284.

Burton A. M., White, D. and McNeill, A (2010). The Glasgow Face Matching Test. *Behaviour Research Methods*, **42**, 286–291.

Burton, A. M., Wilson, S., Cowan, M. and Bruce, V. (1999). Face recognition from poor quality video: evidence from security surveillance. *Psychological Science*, **10**, 243–248.

Caldara, R., Schyns, P., Mayer, E. *et al.* (2005). Does prosopagnosia take the eyes out of face representations? Evidence for a defect in representing diagnostic facial information following brain damage. *Journal of Cognitive Neuroscience*, **17 (10)**, 1652–1666.

Calderwood, L. and Burton, A. M. (2006). Children and adults recall the names of highly familiar faces faster than semantic information. *British Journal of Psychology*, **97**, 441–454.

Chambers, P. (2007). *Body 115: The Mystery of the Last Victim of the King's Cross Fire*. Chichester: Wiley.

Devlin, Lord (1976). *Report to the Secretary of State for the Home Department of the Departmental Committee on Evidence of Identification in Criminal Cases*. London: HMSO.

Duchaine, B., Jenkins, R., Germine, L. *et al.* (2009). Normal gaze discrimination and adaptation in seven prosopagnosics. *Neuropsychologia*, **47**, 2029–2036.

Duchaine, B. C. and Nakayama, K. (2006). Developmental prosopagnosia: a window to content-specific face processing. *Current Opinion in Neurobiology*, **16**, 166–173.

Ellis, H. D., Shepherd, J. W. and Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features – some implications for theories of face recognition. *Perception*, **8**, 431–439.

Hanley, J. R. and Cowell, E. S. (1988) The effects of different types of retrieval cues on the recall of names of famous faces. *Memory and Cognition*, **16**, 545–555.

Hanley, J. R. and Damjanovic, L. (2009). It is more difficult to retrieve a familiar person's name and occupation from their voice than from their blurred face. *Memory*, **17**, 830–839.

Harmon, L. D. (1973). The recognition of faces. *Scientific American*, **229** (5), 70–82.

Harris, A. M. and Aguirre, G. K. (2007). Prosopagnosia. *Current Biology*, **17**, R7–R8.

Haslam, C., Kay, J., Hanley, J. R. and Lyons, F. (2004). Biographical knowledge: modality-specific or modality-neutral. *Cortex*, **40**, 451–466.

Hay, D. C., Young, A. W. and Ellis, A. W. (1991). Routes through the face recognition system. *Quarterly Journal of Experimental Psychology*, **43A**, 761–791.

Jenkins, R. and Burton, A. M. (2008). 100% accuracy in automatic face recognition. *Science*, **319**, 435.

Johnston, R. A. and Edmunds, A. J. (2009). Familiar and unfamiliar face recognition: A review. *Memory*, **17**, 577–596.

Johnston, R. A. and Bruce, V. (1990). Lost properties? Retrieval differences between name codes and semantic codes for famous people. *Psychological Research*, **52**, 62–67.

Lander, K., Bruce, V. and Hill, H. (2001). Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces. *Applied Cognitive Psychology*, **15**, 101–116.

Leder, H. and Bruce, V. (2000). When inverted faces are recognised: the role of configural information in face recognition. *Quarterly Journal of Experimental Psychology*, **53**, 513–536.

McWeeney, K. H., Young, A. W., Hay, D. C. and Ellis, A. W. (1987). Putting names to faces. *British Journal of Psychology*, **78**, 143–149.

Megreya, A. M. and Bindemann, M. (2009). Revisiting the processing of internal and external features of unfamiliar faces: the headscarf effect. *Perception*, **38**, 1831–1848.

Megreya, A. M. and Burton, A. M. (2006). Unfamiliar faces are not faces: evidence from a matching task. *Memory and Cognition*, **34**, 865–876.

O'Donnell, C. and Bruce, V. (2001). Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception*, **30**, 755–764.

Russell, R., Duchaine, B. and Nakayama, K. (2009). Super-recognisers: people with extraordinary face recognition ability. *Psychonomic Bulletin and Review*, **16**, 252–257.

Sinha, P. and Poggio, T. (1996). I think I know that face.... *Nature*, **384**, 404.

Thomson, D. M. (1986). Face recognition: More than a feeling of familiarity? In H. D. Ellis, M. A. Jeeves, F. Newcombe and A. Young (Eds.), *Aspects of Face Processing*, pp. 118–122. Dordrecht: Martinus Nijhoff.

Woodhead, M. M. and Baddeley, A. D. (1981). Individual differences and memory for faces, pictures and words. *Memory and Cognition*, **9**, 368–370.

Young, A. W., Hay, D. C. and Ellis, A. W. (1985*a*). The faces that launched a thousand slips: Everyday difficulties and errors in recognising people. *British Journal of Psychology*, **76**, 495–523.

Young, A. W., Hay, D. C., McWeeny, K. H *et al.* (1985*b*). Matching familiar and unfamiliar faces on internal and external features. *Perception*, **14**, 737–747.

Young, A. W., Hellawell, D. J. and Hay, D. C. (1987). Configural information in face perception. *Perception*, **16**, 747–759.

Young, A. W., Newcombe, F., De Haan, E. H. F., Small, M. and Hay, D. C. (1993). Face perception after brain injury: selective impairments affecting identity and expression. *Brain*, **116**, 941–959.

Notes

1. The individual differences in face matching appear to be highly robust. In unpublished undergraduate projects supervised by me in Newcastle, Lizzie Smith and Charles Okell in 2009 tested 100 participants on the short form of the GFMT and found a mean of 82.1 and standard deviation of 10.0. Claire Chandler and Joanne Sweeney tested a new sample of 77 participants in 2010 and found a mean of 81.3 and a standard deviation of 8.3.
2. Burton *et al.* (1990), elaborated the original model schematised here, and suggested familiarity results from a stage of pooling information across different input routes such as face and voice. However, some observations of differential access to semantics from faces, voices or names otherwise equated for familiarity are challenging for this model (see Haslam *et al.*, 2004; Hanley and Damjanovic, 2009).