Fleeting images of shade

Identifying people caught on video

Vicki Bruce *gave the Presidents' Award lecture at the Society's Annual Conference in Brighton, March 1998.*

THE human face is an interesting object for psychologists to study. Interpersonal impression formation, effects of attractiveness or disfigurement, the perception and expression of emotions, neonatal imitation and neuropsychological impairments of social skills are amongst the many interesting and important topics where information from the face plays a vital role (see Bruce & Young, 1998 for an introduction). However during the past 25 years or so, the study of the identification of people from their faces has become an important topic for those interested in perception and cognition.

The problem of eyewitness testimony

My own interest in the problem of human memory for faces was stimulated in part by some rather paradoxical observations in the mid-1970s. On the one hand, experimental cognitive psychologists working within the laboratory were discovering the phenomenal capacity of human pictorial memory. If participants were briefly exposed to a series of pictures, each for just a few seconds, they were later able to discriminate studied from novel items highly accurately, typically scoring more than 90 per cent correct. When all the items in such experiments were pictures of human faces, recognition rates remained high, despite the similarity of the items one to another. Such observations led cognitive psychologists to be interested in why human face recognition was so good.

However, at about the same time, the Devlin Commission (Devlin, 1976) pointed to a number of serious and notorious errors of person identification made by eyewitnesses to crimes. One such case was that of Laszlo Virag, who was apprehended and convicted of an armed offence in Liverpool. This convic-

Figure 1. Laszlo Virag (left) and Georges Payen. Reproduced from Devlin (1976)

tion was supported by the evidence of several different eyewitnesses, one of whom proclaimed in court that Virag's face was 'imprinted' on his brain. Yet it turned out that Virag was not the 'Gunman of Liverpool'. Later, Georges Payen confessed to this crime. Payen and Virag bear a passing, but not striking, similarity to one another (see Figure 1). Cases such as these led Devlin's inquiry to recommend that eyewitness identification evidence was inherently unreliable, and that convictions should not generally be based upon this evidence alone. Such findings provoked a good deal of applied research into eyewitness testimony which was aimed at discovering means of overcoming the inherent fallibility of human person recognition.

Some aspects of this paradox are easily resolved once it was shown that memory for pictures of unfamiliar faces is good only when the picture is maintained between presentation and test. Changing the viewpoint, expression, lighting, hairstyle, paraphernalia such as glasses and facial hair, and contextual features such as clothing, all reduce recognition memory performance to a considerable extent (e.g. see Patterson & Baddeley, 1977; Bruce, 1982). The eyewitness to a crime may have to identify someone across variations of any or all these kinds — not surprising, then, that errors are made. There are further com-

Requests for reprints should be addressed to:
Professor Vicki Bruce
Department of Psychology
University of Stirling
Stirling FK9 4LA
Scotland
E-mail: vicki.bruce@stir.ac.uk

plications which may affect the accuracy of the witness. For example, people experience difficulties when identifying faces of people of another race, and line-up identification decisions may sometimes be biased, even unwittingly, by the line-up composition if the distractors in the line-up are not matched on relevant dimensions (e.g. if the witness remembers that the criminal was tall, or stocky, the suspect must not be the tallest or stockiest in the line), or by prior exposure to pictures of the suspect (a common feature in several cases reviewed by Devlin). Given all these factors, plus the task demands of the witness, it seems more surprising that such decisions can ever be accurate!

The problem of videowitness testimony

Given the fallibility of human memory for faces, there has understandably been considerable enthusiasm for the increasing use of videotape as a means of capturing identity. Now that closed-circuit television (CCTV) systems are common in banks, shops and city streets, there is increasingly the chance that an image of a criminal is captured on tape. This seems neatly to side-step human face memory. Once a suspect has been apprehended (using some combination of eyewitness and other forensic evidence), the person's identity can readily be confirmed by comparison with the security videotape. Or can it?

In 1988 police raided the home of Cherry Groce, searching for her son Michael who was suspected of a robbery at a building society. A police officer shot Mrs Groce in the raid and the incident attracted considerable publicity. Subsequently, the son was apprehended and prosecuted for the robbery, a prosecution based entirely on the evidence of a CCTV image which showed a young black man, alleged to be Groce. One witness for the prosecution claimed that he was able to prove that the identities matched by comparing the precise numbers of pixels (the very small elements that make up a picture) separating key features of the face. Now this was impressive technical stuff for a jury, until the defence called an expert in facial measurement as its own witness. Alf Linney, a medical physicist from University College London, pointed out that if any of his students had committed such elementary mistakes in face comparison and measurement (for example, not correcting for the viewpoints in two images to be compared or considering the resolution with which the face was depicted) they would certainly fail. Michael Groce was acquitted. However, this early and much-reported case made it clear that a CCTV image alone might prove rather little about the precise identity of a suspect.

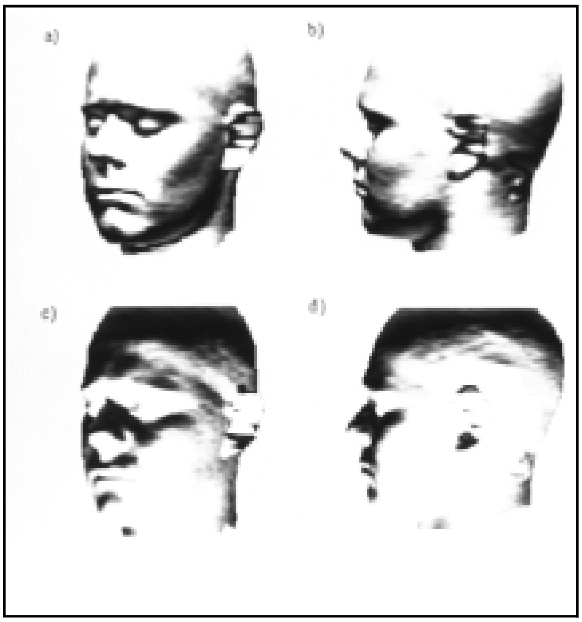

Figure 3. Examples of materials used by Hill and Bruce (1996). The same surface is shown from two different viewpoints (three-quarters and profile views) and in two different directions of lighting (from above and below the head). The participant's task when shown two such images was to decide if they showed the same or two different people.

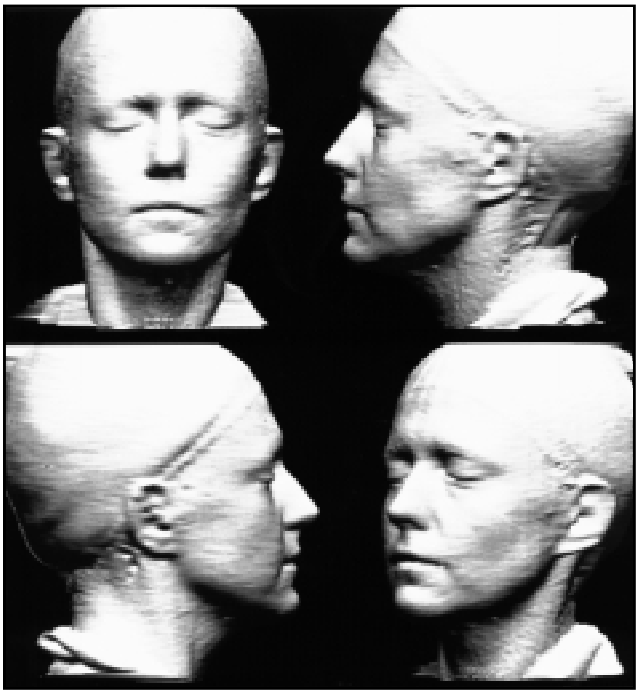

Figure 2. Surface images showing different views of a female head measured by laser at University College London. The database of 3D co-ordinates is used to build a 3D wire mesh computer model whose facets can then be displayed as a smooth surface by applying an artificial lighting model.

How well can people identify images of faces in such circumstances? There are a number of problems with typical CCTV footage which make the task inherently difficult in some circumstances — CCTV images are very variable in quality, and camera and lighting angles may conspire to produce no more than a poorly lit, messy image of the top or back of a person's head. Recent research findings suggest, however, that the process of matching identities across different images may be remarkably error-prone even when image quality is reasonably high.

Matching images of unfamiliar faces

One clue was provided by a research project by Harold Hill and myself (Hill & Bruce, 1996). We were interested in how well faces could be matched when they were lit from different directions. In our studies, participants were asked to decide whether two simultaneously presented images of faces showed the same or two different people. The two faces were either shown in the same or different viewpoints, and were lit from the same or different directions. To make the task impossible to do on the basis of superficial texture features such as hairstyle, the face images we used were pictures of three-dimensional surfaces of faces, obtained by measuring the shapes of faces with a laser (at Alf Linney's lab in London), and redisplaying the surfaces using computer-aided design software (see Figure 2, for an example of such a surface image derived from measurements made on my own face). Such heads can then be displayed with varying lighting (see Figure 3). In a series of experiments we showed that matching was more accurate when viewpoints matched, and better when lighting matched, than when either viewpoint or lighting was changed. A typical set of results is shown in Figure 4. An identical pattern of results was obtained when we

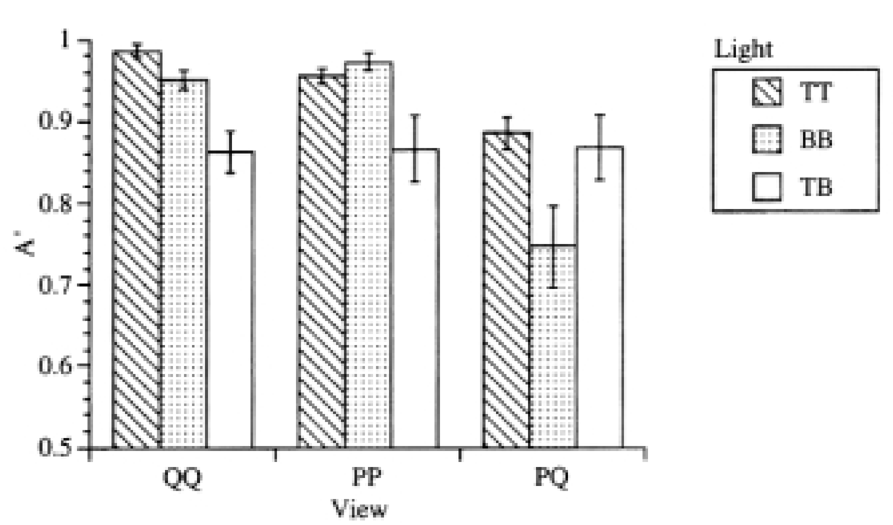

Figure 4. Results from Hill and Bruce (1996), showing reduction in A' (a measure of performance accuracy) when viewpoint changes, and when lighting changes. TT = two heads both lit from above (top); BB = both lit from below (bottom); TB = one top-lit, one bottom-lit. Heads were either shown in matched viewpoints (QQ= three-quarters; PP = profile) or with different viewpoints (PQ = one profile with one three-quarter view). Notice that PQ matches are more accurate when one or both heads are lit from above than when lit from below.

repeated the experiment using pictures of real faces with hairstyle concealed with a swimming cap, showing that the effects of lighting change did not simply arise because of the rather unnatural images used in most of our experiments. One additional finding illustrated in the data in Figure 4 was that people were much more accurate at matching different views of faces when these were lit from above rather than below. Control experiments showed that this was a genuine benefit of lighting direction, rather than a consequence of differential features which may be visible when light comes from different directions.

Richard Kemp and his colleagues at the University of Westminster (Kemp *et al.*, 1997) have used a more realistic everyday situation to explore how well people can verify the identities of people bearing photo identity cards (as in passports, or in photo-bearing credit cards which may be introduced to help prevent fraud). The participants in their study were supermarket cashiers, and they had either to verify or challenge the identity of a series of shoppers using mock-up photo-credit cards. Each card showed a small (2cm square) photograph of a person. In two conditions of the experiment, the identity of the photo matched that of the person carrying it: in one condition ('unchanged'), the photograph and the person carrying it matched on paraphernalia such as glasses, earrings, etc., but in the other some aspect of the paraphernalia had been changed ('changed' card condition). In two other conditions of the experiment the photograph on the card was of a person different from the bearer (a 'foil'), in one condition chosen to resemble the bearer (matched foil) and in the other condition simply matching in sex and racial group (unmatched foil). A summary of the accuracy observed in this experiment is shown in Table 1, which shows that these cashiers correctly detected the fraudulent identity on only a minority of the trials where the foils resembled the card-bearers, and even when the foils bore no particular resemblance to the bearer there were a large number of errors.

Unchanged	Changed	Matched foil	Unmatched foil
93.3	86.2	36.4	65.9

Table 1: Mean % of correct decisions by card type in Kemp* et al.*'s (1997) study.

My final example comes from preliminary data obtained in an ESRC-funded research project specifically targetted at the question of identifying people shown on videotape (Bruce *et al.*, in preparation). The project contains two distinct strands of research. In one, we are investigating the identification of people shown on relatively low quality CCTV security tape. In the other strand of research, we are exploring how well people can verify the identities of people shown on relatively high quality videotape. It is this latter part of the project which I will briefly describe here.

Our initial aim was to investigate how human matching of face images taken from reasonable quality video, against studio shots of the same individuals shown with foils, degraded as the viewpoint difference between the images to be compared was increased. This was so that we could compare the decline in human performance with the decline shown by typical computer face-matching algorithms which we have reason to believe share some properties with human face matching (e.g. see Hancock *et al.*, in press). Our expectation was that when people were asked to identify which of a set of full-face studio portraits matched the identity of a full-face 'target', with both the target and array shown in good quality pictures, and no time pressure exerted, that their performance would be highly accurate. Indeed, to ensure that performance in all conditions of our experiment was not at ceiling, we constructed each array of foils from faces which had been independently judged to bear some resemblance to the target allocated to the array.

In one of our experiments, participants were asked to decide which of an array of 10 clean-shaven, young male faces shown unsmiling in full-face, high quality photographs matched a 'target' face shown separately. A target was present on every trial and participants were aware of this. The target was a still image of one of these same 10 people, grabbed from a video taken of them in the same photographic session, so that hairstyle, visible beard stubble and so forth was identical. However, the video was taken under rather different lighting conditions to those used for the studio shots, and the very careful alignment of the head in the studio pictures was not achieved in the video. Each participant was asked to make this judgement for 40 separate target faces.

A typical array from this experiment is shown in Figure 5. (We compared the accuracy of matching when targets and arrays were monochrome images like those in Figure 5 with what happened when coloured images were used, but have so far found no clear benefits for colour.)

Three different matching conditions were compared. In the first, the target was grabbed from video in the best approximation to a full-face, neutral expression to match the array photographs as closely as possible (as in the example in Figure 5). In the second condition the target was grabbed with a full-face, smiling expression, and the third condition showed a neutral expression but the head was turned 30 degrees away from the camera. Forty participants were tested in each of these conditions. Performance was only 79 per cent correct in each of the two full-face (neutral expression and smiling expression) conditions and dropped to just 70 per cent correct when head angle differed by 30 degrees. Although performance is clearly much better than the chance rate of 10 per cent, nonetheless, given the absence of any memory component or any time pressure in the matching task, it reveals a remarkably high rate of errors. Moreover, this overall accuracy of 79 per cent correct conceals

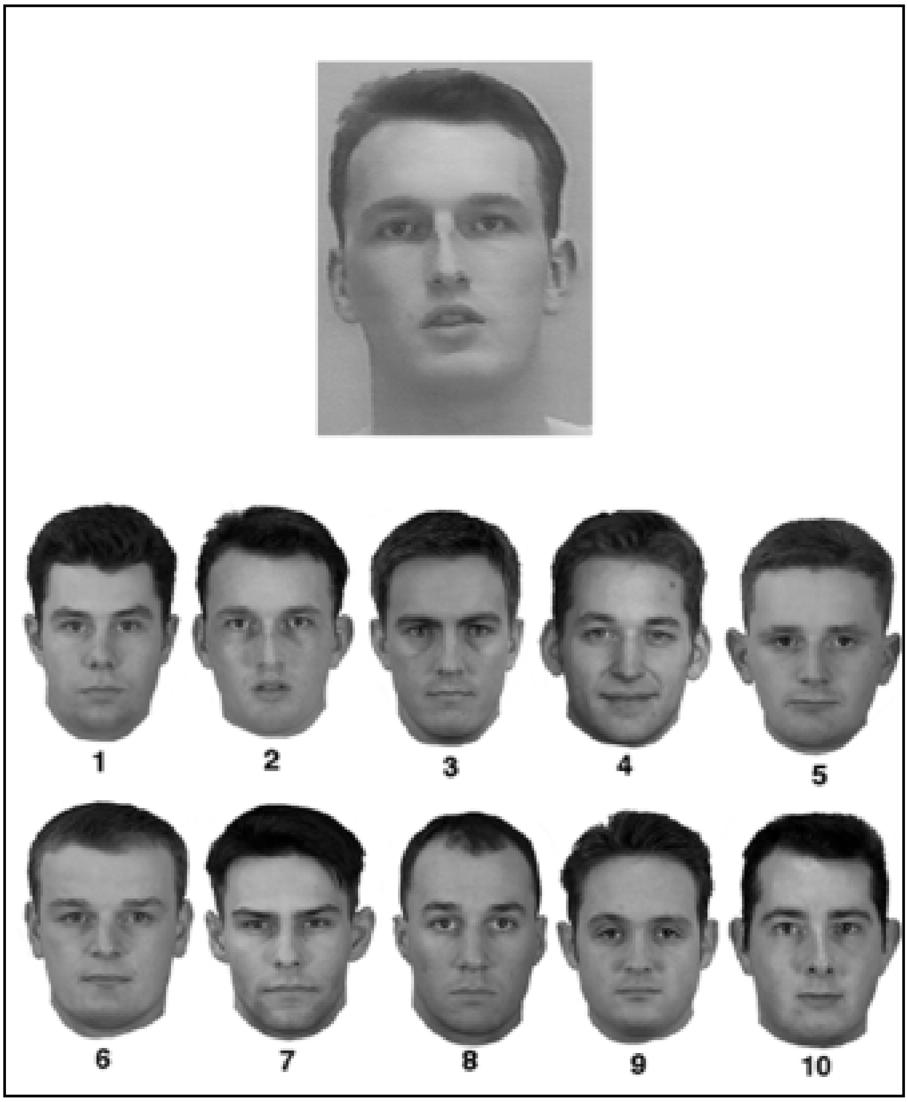

Figure 5. A typical array used in experiments by Bruce* et al. *(in prep). The task is to decide which of the faces in the bottom array is the same person whose image is shown (taken from video) at the top. Courtesy and Copyright PITO, the Home Office, reproduced with permission. (In this example, the correct answer is face number 2. This array is fairly typical of those used in the experiment, and, while many participants make the correct choice, we have also recorded choices of faces 1, 3, 4, 6, 7, 9 and 10 from this array)

quite a large variation between different faces, some of which were regularly and much more frequently confused with others in the arrays.

This experiment shows that significant numbers of errors are made in matching images of faces obtained from different media, even when viewing angle and expressions match, and when image quality is reasonably high. The similar, and slightly superior, performance shown in the full-face trials suggests that hairline is a significant feature used in matching.

Our investigations are still in progress, but the results to date have surprised me, even after 25 years' researching the topic of face matching and memory. I had always assumed that the difficulties that people had in matching two different views, expressions or lightings were because of major changes along these dimensions. We now have evidence that rather subtle pictorial differences are hard for human vision to deal with, a characteristic also of several of the more successful computer systems for face recognition (Hancock *et al.*, in press).

The courts understand that there are dangers posed by video evidence of identity, but not perhaps that these difficulties may extend to high quality rather than low quality video tape:

> 'some at least of the considerations underlying the safeguards built into the regulatory procedures laid down for identity parades and the showing of photographs come into play too with regard to the showing of video tapes. Particularly in cases where the quality of the video is poor or the opportunity it provides for recognition is limited … it would be desirable to regulate its showing so as to maximize the prospects of any recognition evidence being truly spontaneous … it is not for this court to suggest what the appropriate procedures should be; clearly regard will need to be had to the practicalities. We would, however, urge that some immediate thought be given to this matter at high level.' (R v. Caldwell & Dixon (1993))

The House of Lords Select Committee on Science and Technology (1998) recommended that the Turnbull Warning (on witness identification from line-ups and photographs) be adapted 'so that uncertainties inherent in images as evidence are made clear to the jury'; but its focus was on the authentification of the image as having been obtained on a particular date and time, and worries about effects of image processing and enhancement. On the basis of our own research, we feel that there must also be concerns that courts do not place too much weight on apparent visual resemblance (or lack of it) between a video image and a suspect even when image quality is high, and efforts should be increased to try to find automatic or semi-automatic face verification algorithms which can be used reliably to supplement human visual impressions of resemblance.

Effects of movement

One obvious way to enhance the identification of faces shown on video might be to exploit the fact that what is captured on a CCTV film often shows the person moving (this is not always the case, since some systems record sample stills only once or twice per second to economise on tape). Would people be better at verifying identities if shown the moving clips rather than stills? Theoretically, movement may help to establish a better 3D representation of the face, perhaps allowing better generalization to novel views, and in any case a moving clip contains many more 'instances' than a single static viewpoint. So far, however, there are mixed answers to the question of whether viewing faces in motion helps to match identities.

In an undergraduate project at Stirling this year, Karen Greenwood has compared the identification rates found in the target matching task described above when people are shown short clips from the original video, with those obtained when they study for the same total viewing time static snapshots selected from the same clip. The accuracy of matching obtained in these two conditions was identical. Similarly, in a recognition memory study, Fiona Christie and I found that people were no better at recognizing faces they had studied in animated sequences compared with study of the same image frames presented without animation (Christie & Bruce, in press). In contrast, however, Pike *et al.* (1997) report an interesting study in which there is a clear benefit to recognition memory for unfamiliar faces studied rotating in front of the camera compared with static comparison conditions in which different distinct viewpoints from the same video tape were shown statically.

One possible reason for these discrepancies may be differences in the lighting conditions used in the different studies, and the extent to which discrimination between different individual faces may rely more or less on 3D information. We know 3D surface images of faces alone are rather poor likenesses (did anyone spontaneously recognize the face shown in Figure 2 of this paper?), yet for certain types of discrimination (e.g. sex discrimination, see Bruce *et al.*, 1993, for example) 3D shape may be important. It may be that movement helps to derive or restore 3D shape information under certain lighting conditions, or that the 3D shape-from-motion that can be derived is more useful when matching across certain kinds of lighting changes than others. This will be an interesting topic for further exploration.

Whatever the theoretical resolution of these discrepant findings, clearly it is important to explore any possible advantages that there may be to verification by exploiting the movement that may be present in the original footage.

Familiarity

It may be that CCTV images are of most help in the identification of criminals when the faces captured on tape are of someone who is known to the witness. One court case in which Alf Linney and Anne Coombes were acting as expert witnesses involved a policeman claiming that he could identify the person caught on the tape because he knew him well. Preliminary data of our own has certainly confirmed that people who are highly familiar are readily identified from CCTV images of a quality that makes it extremely difficult to identify unfamiliar faces, but we do not yet know whether expectations about identities might lead to bias in identification. For example, if you have reason to suspect that a particular local villain was involved in some crime or other, and a videotape shows someone who resembles him, is your familiarity and expectation more likely to lead to a correct discrimination (that looks like Bloggs, but is not) or to make a false identification? In our ongoing ESRC project we are conducting experiments investigating the effects of both familiarity and expectations on the discrimination of faces caught on low-grade CCTV, and examining specifically what may happen when some of the unfamiliar 'distractor' faces have been chosen to bear some resemblance to familiar ones.

In these studies we are examining also whether there is any benefit to identification of familiar faces from CCTV images when these are shown moving. Interestingly, there is some recent evidence that when it comes to the identification of famous faces shown in severely degraded images, movement certainly does enhance their identification. Knight and Johnston (1997) showed that movement helped the identification of famous faces shown in photographic negative, and Fiona Christie and Karen Lander at Stirling have replicated and extended this to show that movement enhances the identification of famous faces degraded through thresholding (transforming the image to black-on-white), pixellation (blocking) and blurring. Follow-up experiments by Karen Lander suggest that there are genuinely beneficial effects of dynamic information, rather than of the additional static information which is shown in the extra frames of a moving image (Lander *et al.*, submitted). At a theoretical level, these findings suggest that internal representations for person recognition

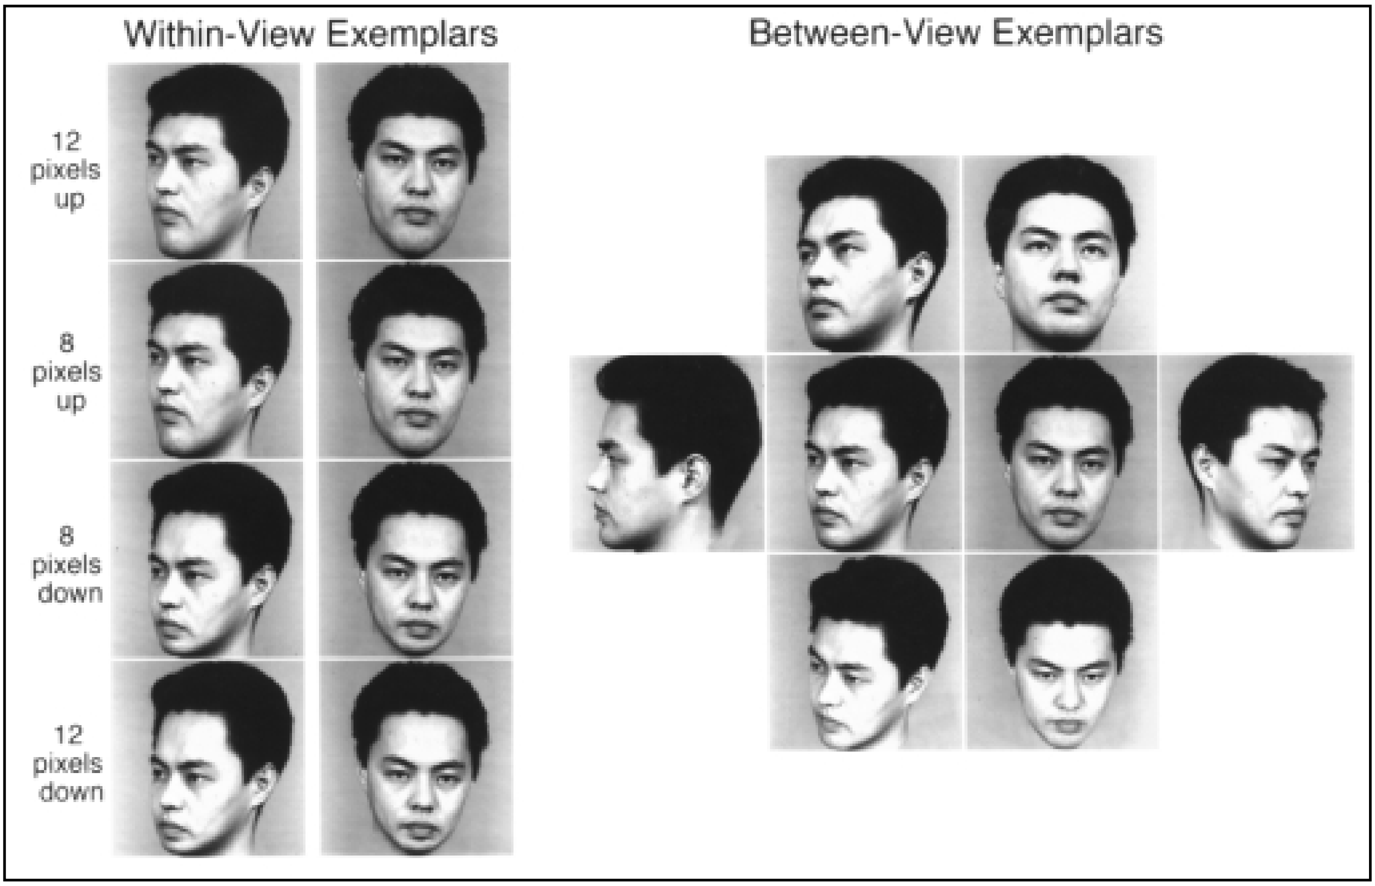

Figure 6. Variations in feature placement and head angle used by Cabeza* et al. *(in press). In our experiments, we find that participants exposed to variations in internal feature placement will tend to find previously unstudied central ('prototype') items as familiar at test as actually studied exemplars, but when participants are exposed to variations in head angle, they do not find intermediate (prototype) viewpoints familiar.

may include information about characteristic movements alongside the 'invariant' information that allows us to recognize a range of static views.

Theoretical implications

So far, I have reviewed evidence which points out how error-prone person identification is likely to be even when memory plays no role in the process at all. Different images of the same unfamiliar face can be extremely difficult to verify, and large numbers of errors are made when foils bear some resemblance to the targets. The precise details of specific images of faces dominate our impressions of likeness, and this can make it very difficult to match identity when, for example, the lighting differs. Indeed the importance of precise patterns of light and dark in face recognition was one reason behind my choice of title here.[1] However, recognition of familiar faces transcends pictorial effects, and recognition of videoclips of familiar persons can be highly accurate, even if the images are very poor, and perhaps especially if the clips are shown in motion.

What process allows the establishment of an internal representation for familiar faces that has such different properties from that of unfamiliar faces? In contrast to the image-specific representation for unfamiliar faces, familiar face recognition appears to be mediated by a representation which is abstracted across the numerous views, lightings and expressions in which the face has been encountered, and perhaps incorporates some information about characteristic motions as well.

How does a representation become established which shifts from a pictorial record to a flexible, structural one? We know that representations of familiar faces differ from those of unfamiliar ones by being more weighted towards the internal, expressive features of the face and away from the external features such as hairstyle (Ellis *et al.*, 1979). Such an internally weighted representation must be established over successive encounters as faces become familiar, though we do not know exactly how this may occur. In experiments of our own where participants study slightly different variants of a series of face identities, we have obtained evidence for the establishment of a 'prototype' representation of each distinct facial identity, abstracted in a way which seems to average over variations within the same viewpoint, but which may retain separate representations of different distinct viewpoints (Bruce *et al.*, 1991; Bruce, 1994; Cabeza *et al.*, in press, see Figure 6).

However, something must act to glue together viewpoint-specific representations of discrete faces. One possibility is that such viewpoints themselves may become bound together by the computational constraints present in image sequences, and there is some interesting recent computational work exploring such possibilities (e.g. Graham & Allinson, in press; Psarrou *et al.*, 1995).

Finally, a neglected contribution to the establishment of representations for face recognition may be the non-face context (time, place and so forth) which could also act to bind together the varying exemplars that characterize a face during familiarization. Of course there is considerable evidence for modularity of perceptual processes, with face-processing impairments of quite subtle types sometimes afflicting brain-injured people (Young *et al.*, 1993), so it is clear that face representations themselves are highly specific. Nonetheless, while a face moves and expresses, presenting its many views to the perceiver, the non-facial context conveniently signals that all these different images do belong to the same person. We have not yet considered the possible role that might be played in this way by non-facial information in 'supervising' or constraining the process of learning faces, but aim to explore this in future research.

Acknowledgements

The research on identification of faces shown on video is supported by a grant from the Economic and Social Research Council to Vicki Bruce and Mike Burton, and Zoe Henderson in Stirling has collected the data described here for this project. Harold Hill, Fiona Christie and Karen Lander have all been PhD students with me at the University of Stirling. ATR Laboratories, Japan, supported a collaboration with Stirling that allowed Harold Hill and Roberto

[1]Fleeting images of shade was a liberal adaptation of a line from P.B. Shelley's poem *The Witch of Atlas*:
For she was beautiful – her beauty made
The bright world dim, and everything beside
Seemed like the fleeting image of a shade.

Presidents' Award 1997

Professor Vicki Bruce received the 1997 Presidents' Award for her distinguished contribution to psychological knowledge. She was awarded the OBE in the Queen's Birthday Honours June 1997, for services to psychology, and has an international reputation for her work on face processing.

Her undergraduate and doctoral degrees were undertaken at the University of Cambridge where her PhD *Processing and Remembering Pictorial Information* was awarded in 1978. From there Professor Bruce moved, via Newcastle, to Nottingham University where she was awarded a Personal Chair in 1990. She became Professor of Psychology at Stirling in 1992 and is currently Deputy Principal there.

Professor Bruce has published articles in a large number of prestigious journals including: *Journal of Experimental Psychology*; *Proceedings of the Royal Society of London*; and *The Quarterly Journal of Experimental Psychology*.

Professor Bruce has an ability to present complex material in an accessible form for her many students. Her books include the highly successful textbook *Visual Perception: Physiology, Psychology and Ecology*, co-authored with Patrick Green in 1985, now in its third edition (with Patrick Green and Mark Georgeson). Her most recent book, *In the Eye of the Beholder: The Science of Face Perception,* written with Andy Young and published by Oxford University Press (1998), was produced to coincide with the exhibition on The Science of the Face at the Scottish National Portrait Gallery in 1998.

Professor Bruce has been awarded research grants and contracts, totalling over £1million, most of them by the Research Councils. She has acted as a consultant to the Royal Mint for projects involving the design of new coinage for the UK.

In 1996 Professor Bruce became a Fellow of the Royal Society of Edinburgh, and in 1997 was elected to an Honorary Fellowship of The British Psychological Society.

Within The British Psychological Society, Professor Bruce has served as Chair of the Scientific Affairs Board (1989–92) and as Editor of the *British Journal of Psychology* since 1995.

Cabeza to conduct postdoctoral work which is also mentioned in this article.

References

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

Bruce, V. (1994). Stability from variation: The case of face recognition. The M.D. Vernon memorial lecture. *Quarterly Journal of Experimental Psychology*, **47A**, 5-28.

Bruce, V., Burton, A.M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R. & Linney, A. (1993). Sex discrimination: How do we tell the difference between male and female faces? *Perception*, **22**, 131-152.

Bruce, V., Doyle, T., Dench, N. & Burton, A.M. (1991). Remembering facial configurations. *Cognition*, **38**, 109-144.

Bruce, V., Henderson, Z., Greenwood, K., Hancock, P.J.B., Burton, A.M. & Miller, P. (in preparation). Matching identities of faces caught on video.

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

Cabeza, R., Bruce, V., Kato, T. & Oda, M. (in press). The prototype effect in face recognition: Extension and limits. *Memory & Cognition*.

Christie, F. & Bruce, V. (in press). The role of movement in the recognition of unfamiliar faces. *Memory & Cognition*.

Devlin, Lord Patrick (1976). Report to the Secretary of State for the Home Department of the Departmental Committee on Evidence of identification in criminal cases. London: HMSO.

Ellis, H.D., Shepherd, J.W. & 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.

Graham, D.B. & Allinson, N.M. (in press). Characterising virtual eigensignatures for general purpose face recognition. To appear in H. Wechsler, P.J. Phillips, V. Bruce, F.F. Soulie & T.S. Huang (Eds), *Face Recognition: From Theory to Applications*. Springer-Verlag.

Hancock, P.J.B., Bruce, V. & Burton, A.M. (in press). A comparison of two computer-based face identification systems with human perceptions of faces. *Vision Research*.

Hill, H. & Bruce, V. (1996). Effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, **22**, 986-1004.

House of Lords Select Committee on Science and Technology. (1998). Digital Images as Evidence, Session 1997-8, 5th report. London: HMSO.

Kemp, R., Towell, N. & Pike, G. (1997). When seeing should not be believing: Photographs, credit cards and fraud. *Applied Cognitive Psychology*, **11**, 211-222.

Knight, B. & Johnston, A. (1997). The role of movement in face recognition. *Visual Cognition*, **4**, 265-273.

Lander, K., Christie, F. & Bruce, V. (submitted). The role of movement in the recognition of familiar faces.

Patterson, K.E. & Baddeley, A.D. (1977). When face recognition fails. *Journal of Experimental Psychology: Human Learning and Memory*, **3**, 406-417.

Pike, G.E., Kemp, R.I., Towell, N.A. & Phillips, K.C. (1997). Recognizing moving faces: The relative contribution of motion and perspective view information. *Visual Cognition*, **4**, 409-438.

Psarrou, A., Gong, S. & Buxton, H. (1995). Modelling spatio-temporal trajectories and face signatures on partially recurrent networks. *Proceedings of the IEEE International Conference on Neural Networks (ICNN'95)*, **1-6**, ch. 620, pp.2226-2231.

R. v. Caldwell & Dixon (1993). 99 Crown Appeal Reports, 73, 78 per Brown Wilkinson L.J., as cited by the House of Lords Select Committee on Science and technology. *Digital Images as Evidence*, Session 1997-8, 5th report. London: HMSO.

Young, A.W., Newcombe, F., De Haan, E.H.F., Small, M. & Hay, D.C. (1993). Face perception after brain injury. *Brain*, **116**, 941-959.