

Canonical views of faces and the cerebral hemispheres

Bruno Laeng

University of Tromsø, Norway

Romke Rouw

Tilburg University, The Netherlands

Evidence is given for a special, canonical, status of one specific view in the identification of familiar faces. In the first experiment, subjects identified by name the fully frontal or profile poses of briefly familiarised individuals less efficiently than an intermediate pose. In addition, in a matching experiment using faces seen in different poses, it was found that one specific intermediate pose (corresponding to 22.5 degrees of angle from the full frontal view) was matched more efficiently in the right visual field (RVF) than in the left visual field (LVF). This finding supports the hypothesis of a superiority of the left hemisphere (LH) over the right hemisphere (RH) in processing a familiar face's canonical view. The other tested "non-canonical" views (i.e., full frontal, 45 degrees, and profile) of these same familiar faces were better matched in the LVF (i.e., the RH); especially at low levels of familiarity. We conclude that, for each familiar face, a viewer-centred representation of the canonical (22.5°) view is stored in the LH's memory system, whereas multiple views of familiar faces are stored in a memory system of the RH. With increasing levels of familiarity other views are increasingly more efficiently encoded by the LH, and in fact for facial self-recognition the full-front view is

Address correspondence to Bruno Laeng, PhD, Department of Psychology, University of Tromsø, Åsgårdveien 9, 9037 Tromsø, Norway. Email: bruno@psyk.uit.no

Bruno Laeng was supported by internal funds from the University of Tromsø (Norway) as well as a JSMF grant (No. 94-34) from the James McDonnell Foundation and The Pew Charitable Trusts while at the Department of Psychology, Harvard University. Romke Rouw, while at the Department of Psychology at Harvard University, was supported by the following Dutch agencies: "Het Bekker-laBastide Fonds", "de Vrijvrouw van Renswoude te Delft", "de Schuurman Schimmel van Outeren Stichting", "het Algemeen Studiefonds" and the "provileringconvenant" from the University of Amsterdam.

We thank Jackie Liederman, Michael Peters, and one anonymous reviewer for their comments on an earlier version. Thanks also go to Peak Chiu, Gunnar Jensen, Jan Are Johnsen, Ronny Mathisen, Ariane Park, and Marit Selnes (for collaborating in the preparation of the experimental material, for assisting in testing subjects, and discussions on the project during its several phases) as well as to many students and friends for lending their faces to the camera. We particularly thank Stephen Kosslyn for his reliable advice and substantial support to the initial part of the project, financially (grants NIA AG 12675 and NINDS 17778-09) as well as for allowing the use of his laboratory facilities.

superior to any of the other tested views. These findings taken together suggest that complementary lateralised memory subsystems in the two cerebral hemispheres store different sets, only partially overlapping, of view-centred face representations.

It would seem that faces, like other “objects”, are best seen and identified from a specific viewpoint. We suggest, in line with the proposals of some researchers (e.g. Troje & Bühlhoff, 1996) that there may be one specific view (or at least a narrow range of neighbouring perspectives) that is optimal for face perception. Several theorists have remarked that, whereas recognition can be largely viewpoint-invariant when the similarity among the objects in a class is low, strongly view-dependent recognition performance occurs when the similarity among the objects in a class is high (e.g., Edelman, 1995; Tarr & Bühlhoff, 1995). Because faces would seem to constitute one of the natural kinds with the greatest inter-object similarity, it would follow that face recognition should be strongly view-dependent. Considerations about what shape properties may be most relevant for the identification of a person from the face, plus evidence from previous studies, converge on the proposal that there is an optimal view of human faces (canonical view) that lies intermediate between the profile and full-front view. This view is conventionally or colloquially referred to as $\frac{3}{4}$ view. However, there is too much ambiguity about exactly what a $\frac{3}{4}$ perspective is (e.g., quarters between the two opposite profile views or between the front and the profile view?); consequently, we shall use the term only when discussing studies that have adopted such terminology. Indeed, some previous studies have clearly provided an operational definition of the intermediate (“ $\frac{3}{4}$ ” view) pose (e.g., Hill, Schyns, & Akamatsu, 1997; Troje & Bühlhoff, 1996). These recent quantitative operational definitions seem to reach a consensus on the perspective corresponding to the exact middle angular position between frontal and profile views; that is, the 45° perspective when we define head rotations as occurring around a vertical axis centred on the head with the full-face view as 0° . A clear advantage for skewed perspectives has been observed in previous visual matching studies (e.g. Troje & Bühlhoff, 1996, found optimal effects for views between 20° and 70° , when 0° is front and 90° is profile; O’Toole, Edelman, & Bühlhoff, 1998, found an advantage of the 45° over the 0° and 90° angles). With a different paradigm, in a study by Blanz, Tarr, and Bühlhoff (1999), subjects were asked to rotate digital images of objects and faces on the computer screen until they selected “the best possible impression of the object on the screen”; remarkably, the mean angular position of both objects and faces consisted in a non-frontal, intermediate, cluster of perspectives. According to Blanz and colleagues, the cluster of angles preferred for textured images of human heads was at about 30° of angle.

However, older studies with either recognition tasks (i.e. matching two face pictures) or identification tasks (i.e. naming a face picture), used the term $\frac{3}{4}$ view

but left it undefined (e.g. Baddeley & Woodhead, 1983; Bruce, 1982; Bruce, Valentine, & Baddeley, 1987; Krouse, 1981; Logie, Baddeley, & Woodhead, 1987). Thus, it would seem, the $\frac{3}{4}$ view can refer to a specified view or to any intermediate view. In this article, we will call all views between the full front and profile view "intermediate", and we will also attempt to gather evidence that one particular defined rotation (22.5°) is likely to approximate the "canonical view" status in the identification of human faces.

In addition, the present study intends to evaluate the effect of familiarity of an individual on the efficiency of different face views in establishing the individual's identity. Because multiple perspectives of highly familiar stimuli may have been overlearned, it is possible that whereas faces of individuals that are in the low range of familiarity (e.g., briefly familiarised faces, acquaintances) are best recognised when seen in a specific optimal view, highly familiar individuals (e.g., friends or your own face) may be recognised equally well from rather different perspectives. On the other hand, there is much evidence that even the recognition of highly familiar and common objects is view-dependent (e.g., Bartram, 1974; Jolicoeur, 1985; Lawson, Humphreys, & Watson, 1994; Newell & Findlay, 1997; Palmer, Rosch, & Chase, 1981). In this account, there would seem to be good theoretical grounds for expecting view-dependent performance in the recognition of highly familiar human faces.

Moreover, the present study will also investigate the role of each cerebral hemisphere in processing faces of varying familiarity when seen in different views. Few studies have specifically addressed the problem of orientation specificity in the perception of familiar faces (e.g., Bruce et al., 1987) and—to our knowledge—no study has investigated both of these factors in their relation to cerebral hemisphere differences. The standard view on the cerebral basis of face processing would seem to attribute such function exclusively to the right hemisphere (RH) and little or none to the left hemisphere (LH). This view has recently been reinforced by evidence for a specific face area in the fusiform gyrus of the RH (e.g., Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Gore, & Allison, 1997). However, several cognitive neuroscience studies on identifying familiar people from their faces suggest not only the involvement of both hemispheres (e.g., Damasio et al., 1996; Sergent, Ohta, & MacDonald, 1992; Sergent & Signoret, 1992) but even a special role of the LH (e.g., Kapur et al., 1995; Umiltà, Brizzolara, Tabossi, & Fairweather, 1978).

The theoretical basis for expecting a LH involvement in familiar faces identification may lie in the fact that the LH supports the semantic system and the lexicon (see Beeman et al., 1994; Caramazza & Shelton, 1998; Damasio et al., 1996; De Renzi, Scotti, & Spinnler, 1969; Gainotti, Silveri, Daniele, & Giustolisi, 1995; Gazzaniga, 1983; Warrington & Taylor, 1978). Biographical information about familiar people and their proper names may then strongly depend on LH structures (see Damasio et al., 1996). Therefore we would expect that this hemispheric specialisation for semantic/lexical information could

influence the functioning of the perceptual system in the same (LH) hemisphere (Van Kleeck & Kosslyn, 1991), including face-recognition processes. Specifically, stored semantic information on a specific person in the LH should facilitate or enhance the perceptual processing of the person's face within the hemisphere so that the identity of the person can be established more efficiently. That is, the matching of a familiar visual (face) pattern to semantic information would avoid the crossing of information via the corpus callosum and the consequent informational loss. This would seem a better design than one that exclusively localises visual processing areas in the hemisphere opposite to the one containing areas that support semantic knowledge.

In the study reported here, we shall first present evidence (Experiment 1) in favour of the optimal role of one intermediate view over two other representative views (the front and the profile) for matching a face to a person's name. In Experiment 2, we will (a) replicate the previous effect by using face stimuli of individuals who are actually familiar to the experimental subjects; (b) emphasise the importance of clear definitions indicating head angles, by comparing and contrasting two different intermediate views (22.5° view and 45° view); and (c) investigate the effects of view on recognition of the subject's own face. Finally, we shall present an account of hemispheric lateralisation for the identification of familiar persons that is somewhat different from the standard one.

EXPERIMENT 1

In this experiment, we attempt to gather support for the hypothesis that an intermediate view of the face has an advantage over the frontal and profile views. This investigation can bring us close to identifying a canonical, preferred, view of the human face. Palmer et al. (1981) showed different converging methods for individuating the canonical view of an object. However, according to Palmer and colleagues, the central criterion for individuating the canonical view is to find the perspective where the object is most identifiable. Therefore, we selected this method for our experiment. Subjects were asked to learn the names of a few individuals whose face were shown in black and white photographs, and subsequently they were asked to judge whether one of the names presented together with one of the pictures was a correct match or not. The faces could be seen in full frontal or profile view, or in an intermediate perspective. Our predictions are straightforward: Subjects should choose the correct response (either yes or no) more efficiently (i.e. faster) when perceiving the intermediate view than either of the other two perspectives.

As described earlier, older studies used the term $\frac{3}{4}$ view, but often no exact definition was provided. Therefore, failures to reveal an advantage of the intermediate view for familiar faces (or highly practised ones) in some previous studies (e.g. Bruce et al., 1987; Hill et al., 1997) may reflect the use of less than

optimal sets of perspectives. As Troje and Bühlhoff (1996) point out, if the angle of head rotation exceeds $30\text{--}40^\circ$ then an eye is obscured by the nose, and other features (e.g. the corner of the mouth) are visible only in one half of the face. Therefore, there are reasons to believe that the 45° perspective may obliterate information highly relevant to facial identity. In this study we used a specific intermediate view that is less skewed than the 45° angle. We defined this view as: *the perspective in which both eyes are fully visible but one of the ears is fully occluded from sight*. This view is approximately just as skewed from the full frontal pose as it is from the 45° pose; therefore, in terms of degree of angular rotation from the full frontal view, it is likely to correspond to a 22.5° perspective. This pose would seem to reveal clearly the internal portion of the face (including salient parts like the eyes, nose, and mouth) and at the same time the rotation of the head would seem to provide relevant surface depth information. Although in this case a range, or “cone”, of lines of perspective could satisfy the requirement, our definition nevertheless restricts the choice to a rather narrow set of angles of regard.

Finally, particular attention was given in this study to response time (RT) data. The majority of the studies mentioned earlier exclusively analysed error rates (one important exception is Bruce et al., 1987 study); it is not unlikely that indistinguishable accuracy levels between different poses are associated in actuality with rather different efficiency levels (i.e. speed of processing).

Methods

Subjects. A total of 12 female and 14 male subjects, all students and by self-report right-handed (mean age = 25, SD = 4.5) at the University of Tromsø (Norway) participated as volunteers in an experiment advertised as a study of face-recognition skills.

Stimuli and apparatus. The stimuli were grey-tone photographs of two young women and two men who volunteered to have their face photographed for a psychology experiment. All models were Caucasian (Figure 1 shows an example of one of the female faces and the three poses used in the experiment). Models were also asked to assume their most neutral expression during the photo-session. Three shots were taken for each model. First, the experimenter took a picture of the model's fully frontal pose, circa 1 m away from the model. Second, from this position, the model was asked to slowly rotate the head towards the right shoulder (keeping a constant chin elevation) and to stop when requested by the experimenter (that is, as soon the right ear disappeared from the experimenter's sight); this procedure yielded an intermediate view, which is reasonably close to a 22.5° view. Finally, the model was asked to continue rotating the head to the right and stop when reaching the position of the precise left-sided profile of the head. Photographs were taken indoors with a Nikon compact camera (zoom-lens set at



Figure 1. From top to bottom: Front, intermediate, and profile pose of one female model.

50 focus). In all pictures the illumination was artificial ceiling lighting with the addition of the camera's in-built frontal flash. Photos were printed on black and white paper and subsequently digitised using a Microtek Scanmaker 600ZS. Each image was edited with Adobe Photoshop software. In the editing, head size was normalised, the colour of the clothing was uniformed to a neutral black, any facial blemish was touched over, and the background was rendered in a neutral greyish-white. Opposite face views were obtained by digitally generating a mirror image. Each final image was 7.6×7.6 cm in size, with the face centred and occupying an area of approximately 5 cm in diameter.

A Macintosh Powerbook 1400cs/133 was used to present the stimuli and collect responses. The "B" key and the "N" key on the QWERTY keyboard were labelled with "Ja" and "Nei" respectively (i.e., "yes" and "no" in Norwegian). The names used were short proper names that are common both in Norway and English-speaking countries (namely, Anne, Tina, Frank, and Martin); these names appeared centred at the bottom of the screen in letter size 18 (Geneva). Other material used in the experiment included a print-out of each of the 12 pictures, in the same size as those seen on the computer's screen. The trials sequence was piloted by MacLab software (Costin, 1988).

Procedure. Subjects were randomly assigned to one of three experimental groups, balancing each experimental group for sex. In the "learning phase", subjects were shown pictures of each model with the model's name printed underneath. One group saw only one picture of each model taken from the front view, another group saw the intermediate views, and another only the profile views. Subjects were requested to learn the names of the models and were given unlimited time to inspect these pictures (no subject took more than three minutes). The picture-name matching task came directly after the learning phase. The subject began each trial by pressing the space bar on the computer's keyboard. This caused a written name to be presented for 1 second on the computer screen (centred and near the bottom). After a 50 ms delay (blank screen) a picture of one of the models, in one of three poses, appeared centred on the screen and remained visible until the subject made a key press. The subject's job was to judge (by pressing the appropriate key) as quickly and accurately as possible whether the name preceding the picture was correct or not. During the first half of the experiment a beep was heard each time the response was incorrect, thus providing subjects with feedback on their performance.

There was a total of 240 trials, and these were ordered so that all combinations of conditions would be presented according to a semi-randomised sequence; that is, they were ordered randomly except for the fact that more than three consecutive trials repeating the same condition (e.g. view, yes/no responses, model) were not allowed. All conditions were counterbalanced so that, for example, it was equally probable that the correct picture was presented after a name or that a model would appear in each pose for a certain amount of trials. A female name (Anne or Tina) was always followed by a picture of a female model; a male name (Martin or Frank) was always followed by the picture of a male. Each subject saw all views of all models, with an equal number of left-looking and right-looking faces.

Results

Means of RTs and percent errors were calculated for each pose condition (front, intermediate, profile). Data were pooled over sex of model (female, male) and type of match (correct versus incorrect), because preliminary analyses showed

no influence of these factors. Similarly, sex of subject (female, male) is also excluded from the analyses here because preliminary analyses showed no influence of this grouping factor.

Response times. Error trials were discarded. In addition, outliers among the RTs were trimmed prior to the analysis; outliers were defined as RTs greater than 3 SDs from a subject's mean RT in a particular cell (levels on the independent variables). RTs were analysed with a repeated-measures ANOVA, with Pose (frontal, intermediate, profile) as the within-subject factor. Learning condition (frontal pose, intermediate pose, profile pose) was not considered because a preliminary ANOVA showed no effects or interactions with this factor.

As expected, a significant effect of pose was found, $F(2, 44) = 23.6, p < .0001$ (see Figure 2). The response time for the intermediate pose (mean RT = 976 ms; SD = 174) was faster than the response time for the frontal pose (mean RT = 1002 ms; SD = 193) as confirmed by a t -test, $t(25) = 2.5, p < .02$. The response




Figure 2. Means and standard errors (bars) for the three different poses in Experiment 1.

time for the intermediate view was also faster than the profile view (mean RT = 1050 ms, SD = 187) as confirmed by a *t*-test, $t(25) = 6.8$, $p < .0001$. Responses to the profile pose were also reliably slower than those to the front pose, $t(25) = 4.3$, $p < .0001$.

Error rates. A separate repeated-measures ANOVA was performed with mean % error rates as the dependent variable and Pose (front, intermediate, profile) as the within-subject factor. Again, Learning condition (frontal pose, intermediate pose, profile pose) was not considered because a preliminary ANOVA showed no effects or interactions with this factor. There were no significant differences in error rate for different poses (front pose: mean % error = 6.0, SD = 4; intermediate pose: mean % error = 6.0, SD = 4; profile pose: mean % error = 6.7, SD = 2.4), $F(2, 44) = 2.0$. Clearly, there was no speed-accuracy trade-off for the pose condition.

Discussion

We hypothesised that an intermediate, close to 22.5° degree, view would be the optimal perspective for the recognition of faces when compared to other representative views such as the full frontal and profile. Studies with infants (Fagan, 1979) lend support to the idea that the intermediate view best provides information allowing perceptual distinctions between different individuals. Infants of 22 weeks can reliably discriminate two different persons' faces only when both faces are viewed in the intermediate pose, whereas full faces result in low and unreliable discriminations and profiles yield chance performance. Our findings clearly showed superiority in speed of processing of our putative canonical view. As suggested by Palmer et al. (1981), the canonical view could contain the maximum of information and/or salient attributes (see Blanz et al., 1999; Verfaillie & Boutsen, 1995). In turn, these properties would facilitate recognition performance. There are several possible reasons why a skewed perspective on the face would meet such requirements.

One account may be that the portion of the face that is most diagnostic of personal identity is the internal region on the front (in particular, a small portion containing the eyes, nose, and mouth; Ellis, Shepherd, & Davies, 1979; Hosie, Ellis, & Haig, 1988; Young et al., 1985). It is possible that a perspective that is skewed from a fully frontal perspective could reveal more of the form attributes of internal parts (e.g. the length and protrusion of nose and lips, the curvature of cheek, forehead, and chin).

A second account of the informational relevance of the skewed view may have less to do with making visible parts or other local attributes of the face (e.g., pattern properties based on grey levels and edges), and more with providing better encoding of some global property of the face. One possibility is

that the intermediate, skewed view is optimal in conveying surface-based/depth properties of the whole facial structure. Indeed, several researchers have previously pointed out that a non-frontal view of the face is computationally more informative than a frontal view (e.g., Fawcett, Zisserman, & Brady, 1994; Vetter, Poggio, & Bülthoff, 1994). In particular, as Moses, Ullman, and Edelman (1996) remarked, a single non-frontal view of a bilaterally symmetrical object conveys much information about its 3D shape (which is limited in the case of its frontal view or for any single image of an asymmetrical object).

Another possibility is that an intermediate view provides the “generic” view of the face or, in other words, the view that provides the most stable information across minimal changes in orientation. As proposed by Koenderink and Van Doorn (1982), a view is stable if small transformations do not lead to significant qualitative changes (e.g., occluding or revealing significant features). Therefore, stable views can more easily generalise to novel views. In our domain, both the full frontal and profile view may be “accidental”, or “unstable”, compared to the intermediate view, because small changes in orientation can yield to loss of original information.

A final possibility is that there is a view that is the most frequently encountered perspective; e.g., in social interactions we may less commonly present or observe faces’ profiles. According to Troje and Kersten (1999), if a person is facing towards us, it is not unlikely that it would trigger attention towards that person’s face, thus decreasing exposure to the most skewed views (e.g., 45° and 90°). Moreover, the full head-on orientation (which does more or less coincide with the frontal plane of symmetry) is often avoided in social interactions (see Argyle, 1975), so the 22.5° intermediate view may indeed constitute the most frequently encountered perspective. Perrett, Oram, and Ashbridge (1998) have suggested that the time course of behavioural responses to specific views of the head and body may reflect the rate of accumulation of activity from neurons selective for particular viewing circumstances. Therefore, the speed of recognition of one view would depend on the extent that the object has been seen previously under the same particular viewing circumstances (that is, more cells will be tuned to the object’s configuration in the views most frequently experienced). All these accounts (i.e., informational relevance of local and/or global shape properties, informational stability, frequency of particular visual experiences) are not mutually exclusive and it is plausible that they may all contribute in yielding the same perspective, or a region of neighbouring views, as the facial canonical view.

For adults, who may be thought as “face experts” (Tanaka & Gauthier, 1997), sufficient information for successful recognition may be available from any particular view (if lighting is optimal and duration of inspection is not limited), as is demonstrated in several studies by the uniformly low error rate in face-matching tasks with different views of highly practised faces (e.g., Bülthoff, Edelman, & Tarr, 1995). Consistent with these findings, in our study,

the intermediate view advantage was also observed for RTs only, in the absence of differences among the (minimal) error rates of different poses.

Finally, there was no effect of the study view. Previous studies have shown a clear effect between the sample and match views, but the two views were presented in pairs sequentially (e.g., Bruce et al., 1987, Exp. 2; Troje & Bühlhoff, 1996). In this type of picture matching-to-sample task, we would expect strong effects and interaction between the two views. In our task, the study view was used to familiarise each subject with each face while learning names, and was used as a between-subjects factor in the experiment. Thus, the study views would seem unlikely to show effects over several subsequent picture–name matching trials.

EXPERIMENT 2

One issue that will be investigated in this experiment originates from the suspicion that much of the consensus on the 45° angle as the operational definition for a candidate optimal intermediate view (in previous studies often referred to as the $\frac{3}{4}$ view) may have been premature. First of all, we reasoned that the 45° perspective makes it more difficult to appreciate the overall outline/width of the face than a somewhat less intermediate view from the frontal plane of symmetry. Second, we considered that the 45° perspective could (depending on the observer's angle of elevation and distance as well as the head size of the model) present occlusions of relevant portions of the internal parts of the face. For example, one of the two eyes may not be fully visible, or from this pose the spatial relationship (distance) between the eyes and their surrounding regions may be less clear than from a view not as far rotated away from the frontal view. Other investigators, like Troje and Bühlhoff (1996), have pointed out similar weaknesses of the 45° perspective of the human face. Indeed, the method used in our previous experiment in selecting the intermediate view (i.e., ‘the view in which the ear just disappears from sight after a rotation of the head away from the full frontal position’) avoids these problems and yields a perspective that is less skewed than the 45° angle, approximating 22.5° .

The choice, in previous experiments, of the 45° angle as the optimal intermediate view carries the implication that both the 45° pose and the fully frontal pose (0° angle) would tend to be at an equal angular distance from what is instead proposed here to be the actual optimal perspective. Thus, we surmise, tests comparing recognition performance for the 45° and 0° views of familiar or learned faces may be prone to find little or no difference. Indeed, Bruce et al. (1987, Exp. 2) in a face-matching task failed to find a RT difference between 45° and 0° views of truly familiar faces (i.e., professors' faces as models and their students as subjects). In this same study (Exp. 1) when unfamiliar faces were instead used as stimuli, there appeared to be a RT advantage of the 45° over the 0° view. Somewhat consistent with the latter finding, in face-matching

studies of experimentally familiarised faces, Hill et al. (1997) found that the 45° view was matched better than the 0° view, whereas Troje and Bühlhoff (1996) showed that several different intermediate views, i.e., 22.5°, 45°, and 67.5° views, yielded a more accurate performance than the 0° (full-frontal) view. The latter study did not report RT data, therefore the question whether one of the intermediate views is more efficient than the other remains unresolved. Consequently, in Exp. 2 we attempt to compare directly performance for truly familiar and familiarised faces between two candidate intermediate poses, the 22.5° and 45° angles, as well as the 0° and 90°. Our prediction is that the 22.5° view should lead to better performance than the 45° view.

Another goal of this experiment regards the role of each cerebral hemisphere in face identification. The dominant or standard view concerning face recognition is that this is a specialisation of the right hemisphere (RH). Recent literature, especially on brain imaging of normal subjects, has proposed the existence of dedicated areas in the right temporal lobe for face-recognition tasks (e.g., Kanwisher et al., 1997; McCarthy et al., 1997). However, this evidence does not exclude the possibility that in some processing tasks, in particular the identification of familiar faces, visual areas of the left hemisphere (LH) are also involved. The goal of any system capable of identifying a face is that of establishing a one-to-one relationship between one visual structure and a unique and known person (i.e., provide the optimal entry-level for the face into a semantic knowledge database). Our proposal is that the knowledge database (semantic and verbal), containing both person-related facts and language information (e.g., proper names), is more dependent on processing in the LH than in the RH. Of particular relevance for this account is the evidence that the speech output control system is genetically constrained to develop in the LH (Corballis, 1991). Moreover, structures supporting associative memory also appear to be lateralised to the LH, as indicated by various evidence from lesions, split-brain patients, brain imaging and divided-fields studies (Beeman et al., 1994; Caramazza & Shelton, 1998; Damasio et al., 1996; De Renzi et al., 1969; Gainotti et al., 1995; Gazzaniga, 1983; Warrington & Taylor, 1978). In essence, the LH appears to possess a greater and finer representation, than the RH, of both the lexicon and amodal semantic (associative) memory. This asymmetry opens the possibility that the associative memory and/or lexical systems in the LH will exert better “feed-back training” (see Van Kleeck & Kosslyn, 1991) on the pattern recognition subsystems in the same (LH) hemisphere than the homologous pattern subsystem in the RH. The feed-back training from the associative memory and/or lexical systems would “tune” their input systems to provide them with the optimal input for entry-level access to their conceptual and lexical representations or, in other words, for face-identification purposes.

Support for a LH involvement in familiar face identification already exists. In divided-fields studies, a RVF/LH advantage in the recognition of famous faces has been repeatedly observed (e.g. Marzi & Berlucchi, 1977; Umiltà et al.,

1978). This finding has been typically explained away by noting that famous faces also tend to have famous names and, therefore, the RVF advantage may just reflect a LH advantage for the retrieval of such verbal labels. Moreover, several PET imaging studies indicate the involvement of LH areas in tasks that require explicitly the recognition of a familiar face and/or accessing information relative to the person. For example, Kapur et al. (1995) have shown that both recognising a particular face as repeatedly seen and judging whether a face belonged to a famous politician or not increased the cerebral blood flow in the left posterior hippocampus. Converging evidence has been collected with deep electrode recordings of the human hippocampus, which showed selective response in the left hippocampus to faces presented repeatedly (Heit, Smith, & Halgren, 1988). Also consistent with the role of both the LH and RH in the perception of familiar faces is the fact that bilateral activation of the temporal lobes has been observed in several brain-imaging studies that used familiar faces as stimuli (e.g., Damasio et al., 1996; Sergent et al., 1992; Sergent & Signoret, 1992) or required subjects to hold one or more faces in memory for long delays (Clark, Maisog, & Haxby, 1998; Haxby et al., 1995; Haxby et al., 1996; McIntosh et al., 1996).

Finally, neurological studies of brain-damaged patients (e.g., Benton & Van Allen, 1968; De Renzi & Spinnler, 1966; Hécaen & Angelergues, 1962; Milner, 1968; Warrington & James, 1967; Yin, 1970) and split-brain patients (e.g., Gazzaniga & Smylie, 1982; Levy, Trevarthen & Sperry, 1972) would seem to support the standard view that the RH is dominant for face perception. There are reported cases of prosopagnosia after single RH lesions, as revealed in brain imaging (e.g., De Renzi, 1986; De Renzi et al., 1994; Landis et al., 1986; Michel, Perenin, & Sieroff, 1986; Sergent & Villemure, 1989; Takahashi et al., 1995; Torii & Tamai, 1985; Whiteley & Warrington, 1977). However, these may be due either to (1) diaschisis effects over spared face-processing areas of the LH, or (2) individual differences in the degree of cerebral lateralisation (analogous to those observed for the cerebral organisation of speech in crossed-aphasia studies). Moreover, Tzavaras, Merieene, and Masare (1973) have reported one case of prosopagnosia following a *left* temporal lobectomy. Most interestingly, Hanley, Pearson, and Young (1990) reported a patient with a unilateral RH lesion who was unable to identify faces of celebrities who had become famous since her stroke, but was unimpaired in identifying faces of people who became famous before her stroke. If the LH does indeed participate in familiar face recognition, the most likely expectation would be that prosopagnosia would be observed in its most severe, frequent, full-blown expression after bilateral lesions; which is clearly the case (see Damasio, Damasio, & Van Hoesen, 1982; Hécaen & Albert, 1979; Meadows, 1974).

In the previous experiment faces were initially unknown to the experimental subjects and briefly familiarised. In this experiment, a small population of university students identified individuals who were truly familiar. Specifically,

in a name–face matching task, faces and names of other students in the same class or programme as the experimental subjects (that is, “moderately familiar” people condition) were presented, as well as those of students in the same class who were also personal friends or even partners in romance (that is, the “highly familiar” people condition). In addition, we also evaluated their recognition performance for faces that were only briefly familiarised before the test, as in the previous experiment (that is, the “low familiarity” condition). Half of the experimental subjects also viewed their own faces in the experiment. This additional manipulation should allow us to assess also whether the optimal view of the subject’s own face is affected by ecological constraints on visual experience (that is, visual experience of ourselves would seem to derive mainly from self-inspection in mirrors). This case presents an interesting and perhaps unique example of high familiarity (i.e., a life-time exposure and a wealth of autobiographical information) that is coupled with a strong constraint on dynamic perceptual experience. Intuitively, it would seem that the full frontal pose, not an intermediate pose, is the most frequently experienced in self-observations.

Thus, we may expect that whereas the 22.5° pose is canonical for “other” people’s faces, the front view may be canonical for recognising our “own” face. Perret et al. (1998) have explicitly proposed that the speed of recognition of a view of an object depends on the rate of accumulation of activity of neurons tuned to specific views and that, in turn, the amount of cells tuned to one view would depend on the statistical frequency of exposure to a particular perspective. Troje and Kersten (1999) have proposed that humans have a disproportionately high exposure to near-frontal views of one’s own face as seen in a mirror. In their study, they indeed found that frontal views of their subjects’ own faces were recognised better than their profiles. However, they did not compare self-recognition performance of frontal views with other views that, although possibly less frequently experienced, may be richer in shape-relevant information (e.g., the “canonical” 22.5° view).

Method

Participants. A total of 20 undergraduate students at the University of Tromsø (12 women and 8 men) volunteered for the experiment. All subjects except two were right-handed, according to self-report.

Stimuli and apparatus. All visual stimuli were grey-tone photographs of 18 Norwegian students (9 females) at the University of Tromsø. These models’ faces were photographed in four different poses: 0° (i.e. full-front), 22.5°, 45°, and 90° (i.e. profile), by asking each model to orient the head according to a template design, positioned on the floor, which indicated the positions corresponding to the four head rotations. Figure 3 illustrates these different poses.



Figure 3. Experiment 2: From top to bottom: 0°, 22.5°, 45°, and 90° views of one female model.

Only the left side of the face was photographed for each model and an image of the face oriented to the opposite side than the one photographed was again obtained by digitally generating the mirror image. Four of the female models were unknown to the experimental subjects as were four of the male models. Photographs were taken with a Kodak DC120 digital camera. Models wore no jewellery or eyeglasses and were photographed having a neutral facial expression. All stimuli were edited with Adobe Photoshop to reduce clothing to the same uniform black and were presented on the screen (set to 16 grey-tone) of a Macintosh Powerbook 1400cs/133 computer. Each picture was 15 cm ×

11.3 cm in size. The trials sequence was piloted by MacLab software (Costin, 1988). The verbal stimuli consisted of the true proper names of each model and they appeared close to the bottom edge of the screen in Geneva font, letter size 38.

Procedure. In the initial “familiarisation” phase, all subjects viewed paper printouts of all the models’ front view photographs (i.e. the 0° view). They were requested to memorise the name of each face (a task especially required for the eight novel, unfamiliar faces). No subject took more than 10 minutes in inspecting these photos. Afterwards, a Powerbook’s screen was positioned at a distance of 45 cm from a chin-rest, where the subject was requested to position the head so as to keep the viewing conditions constant. The subject started each trial by pressing the space bar on the keyboard. First, a name appeared centrally for 500 ms. Then the screen became blank for 50 ms, and next a fixation-cross (5 mm) appeared at the centre of the screen for an additional 600 ms. Subjects were instructed at the beginning of the experiment to gaze at the central cross whenever this appeared. Then a face picture was presented for 100 ms at either the left or right of fixation (the centre of the picture was at 9 cm from fixation). Subjects indicated whether the name and the following facial image correctly matched by pressing one of two keys labelled “Ja” or “Nei” (i.e., the “B” and “N” keys of a QWERTY keyboard). Subjects were asked to keep the index finger of each hand on one key and to respond as quickly and accurately as possible.

There were a total of 288 trials in the whole experiment. In half of the trials the name incorrectly matched the face and the foil name was always a randomly selected name among the persons within the same sex and familiarity models’ subset. All combinations of factors were randomly mixed. Each model’s view was presented four times, twice in each visual field. One of these lateralised views was left looking and the other right looking, so as to balance the advantage of presentations of salient facial features near the foveal position. However, one was paired to a true and one to a false label. In order to balance original/mirrored versions in the experiment with visual field and label (true/false), each original version shown in one visual field was paired for half of the models with the true label and for the other half of the models with the false label and vice versa for the mirrored versions. Moreover, this across-models balancing of original/mirror views and labels also occurred equally frequently for models of different sex and familiarity.

At the end of the divided-field study, subjects were requested to rate on a scale from 1 to 3 how familiar they would judge themselves to be with each model. Some guidelines were given in assigning each score. A score of 1 was to be assigned to an individual never seen before (i.e. the eight unfamiliar faces subset), a score of 2 to an individual who belonged to their student group but not in any especially close relationship to the subject, and a score of 3 was assigned

to individuals in the student group who were judged to be particularly close to the subject (i.e., doing some school project together, personal friends, partners in love, or, in some instances, the model was the subject). These scores were subsequently used to classify, subject-by-subject, the level of familiarity belonging to each individual trial as: low (i.e., faces familiarised immediately before the divided-fields task as well as during these experimental trials), moderate (i.e. fellow students in the class), and high (i.e. close friends, partners). However, for those who also had to match names to their own face, these trials were excluded from the main analysis and, as shown later, analysed separately (as the self-recognition task).

Results

Response times. Error trials were excluded from the analyses of the RTs and analysed separately (6.3% of all data). All responses longer than 2.5 SDs from each individual subject's mean RT were excluded (less than 1% of all RT data). In the analysis described here "yes" and "no" trials were collapsed together because a preliminary analysis showed no effects of this factor. A repeated measures ANOVA was performed on RTs to all the correct name-face matches, with Familiarity (low, moderate, high), View (0° , 22.5° , 45° , and 90°), and Visual Field (LVF, RVF) as the within-subject factors.

First of all, there was an effect of familiarity, $F(2, 38) = 84.5$, $p < .0001$. Responses were increasingly faster with increasing levels of familiarity (low: mean RT = 845 ms, SD = 199; moderate: mean RT = 660 ms, SD = 151; High: mean RT = 606 ms, SD = 141). As can be appreciated in Figure 5, with increasing levels of familiarity, the mean RTs did not only become faster, but the range of mean RTs for different poses became increasingly smaller.

Importantly, the main effect of view was significant, $F(3, 57) = 4.6$, $p < .005$. As shown in Figure 4, the 22.5° view was faster than all the other views and these did not differ significantly from one another (as confirmed by t -tests). This effect confirms the overall superiority of an intermediate view over the front and profile view, as also seen in the previous experiment, but in addition it indicates the superiority of the 22.5° view over the 45° view. Moreover, these effects also replicate findings of other studies of no difference between the 45° and the 0° views when using familiar or highly practised faces as stimuli.

The factor of view also interacted with familiarity, $F(6, 114) = 3.4$, $p < .004$. As illustrated in Figure 5, whereas the efficiency of the 0° (frontal) view increased at each level of familiarity in relationship to the other views, the efficiency of the 90° (profile) view decreased (also relative to the other views) at each increasing step in familiarity. Specifically, the 0° view showed the lowest efficiency at low familiarity, being slower than the 22.5° , $t(19) = 12.2$, $p < .0001$, the 45° , $t(19) = 9.6$, $p < .0001$, and even the 90° view, $t(19) = 1.2$, $p < .0001$. However, with moderate familiarity its efficiency improved and it was not




Figure 4. Means and standard errors (bars) for the four different views in Experiment 2.

reliably different from the 45° and the 90°, although still inferior to the 22.5° view, $t(19) = 2.5$, $p < .01$. At high familiarity, matching of the 0° view further improved and it was now faster than both the 45°, $t(19) = 1.9$, $p < .07$, and the 90° view, $t(19) = 2.1$, $p < .01$, but no longer differentiable in speed from the 22.5°, $t(19) = 0.4$. The profile view (90°), at low familiarity was not significantly different from the 22.5°, $t(19) = 0.9$, and the 45°, $t(19) = 1.6$, and it was actually faster, $t(19) = 3.5$, $p < .0008$, than the 0°, frontal, view. Nevertheless, at the other, higher, levels of familiarity the 90° view was clearly the slowest in performance.

We now turn to the effects of visual field. The main ANOVA revealed no main effect of visual field and, indeed, the mean RTs for the LVF and RVF trials were practically identical (LVF: mean RT = 703 ms, SD = 194; RVF: mean RT = 704 ms, SD = 195). Visual field interacted with familiarity, $F(2, 38) = 3.6$, $p < .04$. This effect is due to the fact that whereas in the low and moderate familiarity condition there was a non-significant tendency for faster RTs in the LVF than in the RVF, in the high familiarity condition there was a significant




Figure 5. Experiment 2: Means and standard errors (bars) for four different views at each level of familiarity (low, moderate and high) and in each visual hemifield (LVF = left visual field; RVF = right visual field).

visual field advantage, $t(19) = 2.1$, $p < .05$, in the opposite direction (i.e., RVF faster than the LVF). Visual field also interacted with view, $F(3, 57) = 11.1$, $p < .0001$. As Figure 5 illustrates, there was a clear RVF/LH advantage for the 22.5° view only, $t(19) = 4.9$, $p < .0001$, whereas there was a LVF/RH advantage for the 45°, $t(19) = 2.1$, $p < .04$, and the 90°, $t(19) = 2.2$, $p < .04$, views and no visual field difference for the 0° view.

Importantly, the interaction of familiarity with view and visual field was also statistically significant; $F(9, 95) = 5.9$, $p < .0001$. Mean RTs and SDs for these conditions are shown in Table 1. The findings of this high-order interaction can




Figure 6. Experiment 2: Means and standard errors (bars) for four different views in self-recognition trials only.

be easily interpreted in the light of the previously exposed effects of visual field and familiarity that each had separately over view. In particular, it appears that the 0° view changed its role with increasing levels of familiarity from a low-efficiency view best matched in the LVF/RH to a high-efficiency view best matched in the RVF/LH. In contrast, the other views maintained a similar profile of efficiency and visual field effects (or lack of them) across familiarity levels.

Error rates. A repeated-measures ANOVA was performed after converting errors to percent error rates for each combination of the factors of Familiarity (low, moderate, high), View (0°, 22.5°, 45°, and 90°), and Visual Field (LVF, RVF). These same factors were used in the ANOVA as the within-subject factors. The analysis revealed only two significant effects. First there was a main effect of the familiarity factor, $F(3, 57) = 12.3, p < .0001$. As was to be expected, there were many more errors in the low familiarity condition (mean % error = 15.3, SD = 10) than in the other two (moderate: mean % error = 4.2, SD = 4; high: mean % error = 3.0, SD = 3). Second, the factor of familiarity also

TABLE 1

Experiment 2. Means and SEs of RTs for each familiarity level in each visual field and for each view

			00	22.5	45	90
Low	LVF	Means	835	904	808	790
		SEs	48	58	32	36
	RVF	Means	942	762	888	831
		SEs	54	31	51	25
Moderate	LVF	Means	676	639	648	652
		SEs	36	32	33	36
	RVF	Means	662	587	693	722
		SEs	31	26	35	37
High	LVF	Means	607	617	625	634
		SEs	37	33	37	32
	RVF	Means	571	541	615	637
		SEs	19	25	24	38

interacted with view, $F(6, 114) = 3.6$, $p < .05$. t -tests revealed that the 90° view had a larger error rate in the low familiarity condition than any of the other views, $2.3 < t(19) < 3.6$, whereas error rates did not differ among views at higher familiarity levels. This effect is interesting in the light of the RT data as it reveals the presence of a speed–accuracy trade-off. We speculate that when viewing profiles of low-familiarity faces, subjects in this experiment (unlike those in the previous experiment) may have chosen the strategy of matching names to faces on the basis of one or few salient and distinctive facial features that this view can offer (e.g. nose outline). This may yield a fast rate of response in correct trials but a rather low level of accuracy.

Analysis of self-recognition trials. For the 10 subjects who had also served as models in the task, a separate repeated-measures ANOVA was performed on the RTs to all (as there were no errors) of the self-recognition trials (i.e., trials in which the subject viewed her/his own face) with View (0°, 22.5°, 45°, and 90°) and Visual Field (LVF, RVF) as the within-subject factors. This analysis revealed only a significant effect of view, $F(3, 27) = 8.2$, $p < .0005$. As Figure 6 illustrates, RTs to the frontal 0° views were the fastest, and RTs were progressively slower with further rotations of the head from the optimal 0° of perspective. t -tests revealed that the 0° view was faster than the 22.5° view, $t(9) = 2.3$, $p < .03$, as well as the other views, the 22.5° was reliably faster than the 90°, $t(9) = 2.2$, $p < .04$, but not the 45° view and there was no difference between the 45° and 90° views.

Discussion

The experiment supported the hypothesis that there exists an optimal view for the recognition of familiar individuals. In a task of familiar face matching a 22.5° view, intermediate between 0° and 45°, appears to yield the fastest RTs. However, this conclusion is only generally true; in actuality, whether there is a single optimal view for the recognition of familiar individuals (and whether this corresponds to the 22.5° view) depends on the degree of familiarity between the viewer and the model. Specifically, when subjects viewed their own face, the 0° (frontal) view was in fact superior to the 22.5° view and, at high levels of familiarity of others' faces (i.e. faces of friends and partners, not just acquaintances or highly practised photos of unfamiliar faces), the 0° and 22.5° views could be recognised equally quickly.

Another novel conclusion from this experiment is that different intermediate views are not equally recognisable. In particular, we note that the two different intermediate views used here, i.e., the 22.5° and the 45° views, are to be considered qualitatively different perspectives of the human face, as the former is clearly superior overall for recognition. Interestingly, we also observed that at moderate levels of familiarity (i.e., faces of acquaintances) as well as high ones (i.e., faces of friends), there was clearly no difference in performance efficiency between the less efficient intermediate view (i.e. 45°) and the 0° frontal view. This phenomenon replicates the null findings of several previous studies when these two views were compared for familiar faces (e.g. Bruce et al., 1987, Exp. 2) or well learned photos of faces within the laboratory setting (Hill et al., 1997). Moreover, there was also a tendency for the 45° view to be better matched than the 0° view at the low familiarity level or, more precisely, when seeing faces that were entirely unfamiliar previous to the test (but practised to some extent through the task). Similarly, in Bruce et al.'s (1987; Exp. 1) study, a superiority of the 45° over the 0° view was found only when viewing unfamiliar faces and the same effect has been observed in other studies of matching views of unfamiliar faces (e.g., O'Toole et al., 1998).

Finding such clear view-dependent performance in facial recognition is remarkable considering that several studies have shown that difference in RTs between different object views are rather ephemeral and disappear with practice (Jolicoeur, 1985; Koriart & Norman, 1985; Tarr & Pinker, 1989). We only found some limited evidence that view-dependent performance disappears with practice (i.e., the indistinguishable speed between 0° and 22.5° views for highly familiar faces). Yet, also at very high level of familiarity, the two other tested views clearly yielded an inferior performance, and in self-recognition, which plausibly represents the highest level of familiarity tested in the study, there was a clear view-dependent difference in efficiency of identification.

A tentative explanation for these findings would be that there are physical constraints or preferences when viewing others' faces as well as our own face.

That is, if in principle we can experience other people's faces at every possible angle, there may exist a perspective that is simply most frequently seen. We can speculate that this most frequently encountered perspective corresponds to a range of angles in the neighbourhood of the 22.5° view. In contrast, when we visually experience our own faces, this must take place through the use of external tools like mirrors and, to some extent, static photographs and films. Mirrors clearly constrain the perspectives on our own face to the full front perspective and slight angular deviations from it. Thus, we speculate that the superiority of the 0° view would be caused by this narrow domain of facial self-inspection experience. On the basis of similar considerations, Troje and Kersten (1999) have conducted a face-naming study that showed a clear advantage of front view over profile views for facial self-recognition, whereas in this same experiment such front advantage was absent when viewing others' faces. Our results clearly confirm and complete their findings by showing that the front view is also superior to different intermediate views and for self-recognition only.

Accounts positing alignment processes of view-centred representations (e.g. Tarr & Pinker, 1989, 1990) predict that the addition of several stored specific views, incrementally to practice, would result in an increased ability to match novel views at equally high levels of efficiency. The persistence of optimal views, despite extended visual practice with multiple views, clearly indicates that greater familiarity with one or few particular views may not be the complete explanation. As Edelman and Bülthoff (1992) have demonstrated, canonical (most efficient) views can also exist for novel objects seen equally often from all the tested viewpoints. In other words, when practice on each perspective is carefully controlled there is still a view that can be matched optimally. Indeed, in Experiment 1, as well as for the low familiarity faces of Experiment 2, different poses of unfamiliar faces were presented for an equal number of trials and yet differences in efficiency between such poses were observed. This phenomenon confirms the original idea of Palmer et al. (1981) that one view may be more perceptually informative than others. Therefore we conclude that the 22.5° view of people's faces may contain more identity-relevant information (e.g., properties high in distinctiveness, therefore useful for subordinate, individual-level recognition) than other perspectives. Although we cannot exclude the possibility, entertained earlier, that the 22.5° view of other persons' faces may be the most frequently encountered perspective (or one in its neighbourhood), this explanation would seem destined to remain, at the current stage of knowledge, rather speculative.

Another important finding was the presence of hemispheric (visual field) differences for different poses of faces. This was modulated by the degree of familiarity with each specific face. First of all, at the lowest level of familiarity there was a clear RVF/LH advantage for the 22.5° view only, whereas there was a LVF/RH advantage for the other non-canonical views. One important

implication of these findings is that the verbal mediation hypothesis (i.e., LH advantage for the retrieval of verbal labels, like high-frequency proper names) can be ruled out as an explanation of the RVF/LH advantage in the recognition of familiar faces. Clearly, a verbal mediation hypothesis could predict the observed overall RVF/LH advantage but not view-dependent specific effects. Instead, the LH's advantage was view-dependent and, for low and moderate levels of familiarity, it was specific to the 22.5° pose. The view-specificity of the observed visual field effects also refutes an explanation of these results based on Goldberg and Costa's (1981) account that novel information is processed predominantly by the RH, while as information becomes more familiar, processing is dominated by the LH. An additional consideration to make is that previous contradictory findings (in the literature) indicating a LVF/RH superiority in matching famous faces (e.g. Young et al., 1984) could be reflecting a lack of control of pose of the faces. This is mostly a speculation, as in the majority of the studies the specific pose of the models was not reported explicitly. However, if not enough "optimal" canonical views were present in the stimuli used, and also depending on the level of familiarity (e.g., low) of the faces used as stimuli, we could actually expect the RH to become engaged more—or equally so—than the LH. Finally, with increasing levels of familiarity, the differences in efficiency between the poses changed. In particular the role of the 0° view changed from a low-efficiency, LVF/RH-advantaged view to a high-efficiency view and RVF/LH-advantaged. As discussed earlier, this changed role of the full-face view was clearest in self-recognition trials. These varying optimality effects of different views for both hemispheres may reflect the influence of "overlearning", practice, or frequency of encounters, with a particular facial pattern. One possibility is that with increasing familiarity the efficiency of response to non-canonical views also improves until an asymptotic level is reached, where further reductions on speed of identification are not possible and therefore differences in efficiency between the views tend to disappear. This effect could explain why, at high levels of familiarity with others' faces, the 0° view and the 22.5° view showed equal efficiency in the LH. However, despite the fact that the difference in efficiency between the canonical view and other non-canonical views, like the 45° and the 90° views, did increasingly reduce with increasing levels of familiarity, these more skewed views remained clearly below the optimal efficiency level. In conclusion, this experiment confirms the phenomenon of view-dependent performance in the recognition of highly familiar human faces.

GENERAL DISCUSSION

In Experiment 1 we provided some evidence in support of the idea that there is a view of the human face that is optimal for the identification of faces. In this experiment, subjects became familiarised with a small set of faces, through

repeated exposure to three different views, and learned the associated names. We found that the intermediate pose was better identified (in a face-name verification task) than the other two tested views (frontal and profile). The view tested in this experiment was operationalised and defined as a 22.5° angle view and in Experiment 2, we assessed directly the superiority of this view over others (i.e., 0° , 45° , and 90°) in another face-name verification task, where faces could also differ in familiarity (briefly familiarised, acquaintances, friends, and own face). We found evidence that when familiarity with a facial pattern increases, all views tested yielded increasingly efficient performance; although, as expected, the 22.5° was the view resulting in the most efficient performance. At high level of familiarity, when the faces viewed belonged to those of the observers' friends instead of simple acquaintances (or individuals familiarised by photograph), the 0° (full-front) view resulted in equally efficient performance to the 22.5° angle. When the observers identified their own faces, the best view was indeed this 0° (full-front) view. Therefore, we conclude that this view should be considered canonical for facial self-recognition. Finally, we found the 22.5° angle to be recognised best by the LH, when the faces were at least moderately familiar (e.g., acquaintances).

Following Palmer et al.'s (1981) definition, we conclude that this 22.5° view is, among the tested views, the best candidate for constituting the canonical view of the human face. Clearly, the findings of both Experiment 1 and 2 present an analogue, within the domain of face identification, to Palmer et al.'s (1981) finding that presentations of common objects and animals' $\frac{3}{4}$ (intermediate) views (e.g., a horse) yield a clear advantage in both RT and accuracy over other representative views of the same object. However, in the two experiments reported here, advantages or disadvantages in identifying faces from different views were revealed only in the speed and not in the accuracy measure of performance. We speculate that any view of a familiar face can present sufficient information to yield a successful match and that optimality of view can best be captured by efficiency measures such as the RT data.

Interestingly, the proposal that an intermediate view, close to or corresponding to the 22.5° , is canonical for the recognition of the human face finds clear correspondence in several cultural domains. It may be a familiar experience when looking at portraits in museums or books, that most of these depictions of real (and sometime fictional) faces take a preferred perspective. It would seem that, especially in the "realistic" portraiture of Western art, it is the so-called $\frac{3}{4}$ view that is the most frequently represented by painters and photographers (Brilliant, 1991; McManus & Humphrey, 1973). We believe that the portrait artists' preference for the intermediate pose is not coincidental but that it derives from its superiority in revealing the face's cues to its identity, therefore increasing the "likeness" of the portrait. Furthermore, the US Immigration & Naturalization Service stipulates clear specifications on what is an acceptable photograph for official identification documents (e.g., the "green card");

namely: ‘‘Photograph must show the subject in $\frac{3}{4}$ frontal portrait’’. Remarkably, in the sample picture provided by this US governmental office (see Brilliant, 1991, p. 42), the right cheek, both eyes, and the left ear are visible. Previous cognitive studies had already partially supported the idea that facial identity is better established after seeing an intermediate view of a face (e.g., Bruce et al., 1987; Hill & Bruce, 1996; Krouse, 1981; Logie et al., 1987). More recently, Troje and Bülthoff (1996) have reassessed the evidence.

Additionally, many classes of other real-world objects (e.g., either biological objects that possess a face, like horses or chickens, and artifacts that have a clear distinction between front and rear planes, like cars or clocks) show the same advantage for the ‘‘ $\frac{3}{4}$ ’’ perspective or a perspective intermediate of the front and side views (Palmer et al., 1981; Verfallie & Boutsen, 1995). In a study where subjects were asked to select the view that best represented an object, by manipulating computer-rendered 3-D images, Blanz et al. (1999) observed that their subjects chose intermediate views as the preferred views of many different classes of objects (natural kinds, artifacts) as well as of human faces (in this latter case, the mean preferred angle was approximately 30°). Finally, according to Nakayama, He, and Shimojo (1995), perceiving the surfaces of an object in intermediate ‘‘generic’’ views is always preferred by the visual system when perceiving the ‘‘accidental’’ view (of another object) would be also compatible with the input.

Since Palmer et al.’s (1981) study, several other studies have also investigated the issue of whether alternative measures of preference for views, such as inspection time, ‘‘goodness’’ ratings, and the estimated angle of a spontaneously generated visual image of an object, would also converge on the same skewed, non-frontal, perspectives. The results have so far been mixed. It would seem that goodness ratings do converge onto intermediate, stable perspectives as the preferred views, whereas mental images and inspection times may show great variability of responses among individuals and, if anything, a preference or longer inspection times for frontal views (see Blanz et al., 1999; Harries, Perrett, & Lavender, 1991; Niemann, Lappe, & Hoffman, 1996; Perrett, Harries, & Looker, 1992). However, if we consider that length of inspection may be directly proportional to the need and effort of perceptual processing, longer inspection times for certain accidental views would seem entirely consistent with longer RTs for the object’s recognition seen in these same views.

In Experiment 2, we found differences in the way each cerebral hemisphere can identify faces seen in different poses. However, these findings do not lend support to the standard view that the RH is dominant for all aspects of face processing. It would seem that, at least for face identification of familiar individuals, the LH is just as competent or even more efficient than the RH. This was shown in a divided-fields experiment, where at low and moderate levels of familiarity we found a superior performance in the RVF (i.e., LH) over the LVF (i.e., the RH) for the 22.5° angle of perspective. Also, at low levels of

familiarity, the RH appears superior for other non-canonical poses. For faces that are high in familiarity, the visual field differences specific to pose disappear and we observe an overall superiority of the LH.

These effects, although complex, are in fact consistent with several theoretical considerations we made at the outset. We assumed that highly familiar persons' faces are cues to a wealth of detailed person-related semantic information. Semantic knowledge is assumed to be generally more precise and rich in the LH. Thus, we would expect that the visual pattern subsystem in the LH would develop face representations of familiar individuals so as to directly access (i.e., by avoiding callosal transmission of information) semantic or biographical knowledge, and therefore establish their identity efficiently. In other words, the LH's semantic and verbal systems may exert "feed-back training" (Van Kleeck & Kosslyn, 1991) over the pattern recognition subsystem in this same hemisphere for the encoding of facial representations that can rapidly access facts and names. Other views that provide less information for the discrimination between identities (e.g. 45°, and 90°), would be identified initially by the LH at a slower rate than the canonical view. We suggest that the LH specifically profits from the canonical view, because it specialises in the rapid identification of known people. Instead, unlike the LH, the RH may specialise in multiple face-processing tasks besides person identification, which opens the possibility that its facial representations may be of a different "nature" from those of the LH. Consequently, the RH may not benefit from the same perspective in the same way. Indeed, we found that the LH could be more efficient than the RH in identifying people particularly from the intermediate 22.5° view. Although we no longer found a single canonical view at high levels of familiarity, there still remained view-dependent effects on identification (specifically, least efficient identification for the two most skewed views, i.e., 45° and 90°).

That object recognition is dependent on view has been often taken as evidence that the human brain's underlying perceptual representations are viewer-centred instead of object-centred (see Tarr & Pinker, 1989). As pointed out by Troje and Kersten (1999), the fact that subjects recognise the front view of their own face better than another view suggests that the visual system is unable, even for a highly familiar object such as your own face, to construct a viewpoint-independent representation. However, Troje and Kersten's study assessed self-recognition only for front views vs. profiles, therefore this finding may be confounded by the fact that the front view may simply be more informative than the profile view. Hence, an advantage of the front view over the profile may still be fully consistent with an object-centred representation of the human head. Nevertheless, our findings may help to dispel this doubt, as in this study the front view appeared superior in self-recognition not only to the profile (thus replicating Troje and Kersten's results) but also to the 22.5° view (which was best recognised when seeing familiar others' faces and theorised to

be the most informative view). Clearly, these differences in view-dependent performance between faces of self and others must reflect the disproportionately higher exposure of near-frontal views during self-observation than others-observation. Although, statistically, some views may be more frequent than others in social interactions, a restriction in range of visual experience does not apply for observing others' faces and, therefore, we would expect that the most informative view would emerge as the preferred view when we observe other persons' faces.

Manuscript received 11 April 2000

Revised manuscript received 10 August 2000

REFERENCES

- Argyle, M. (1975). *Bodily communication*. London: Methuen.
- Baddeley, A., & Woodhead, M. (1983). Improving face recognition ability. In S.M.A. Lloyd-Bostock & B.R. Clifford (Eds), *Evaluating witness evidence* (pp. 125–136). Chichester, UK: John Wiley & Sons.
- Bartram, D.J. (1974). The level of visual and semantic codes in object naming. *Cognitive Psychology*, 6, 325–356.
- Beeman, M., Friedman, R., Grafman, J., Perez, E., Diamond, S., & Lindsay, M.B. (1994). Summation priming and coarse semantic coding in the right hemisphere. *Journal of Cognitive Neuroscience*, 6, 26–45.
- Benton, A.L., & Van Allen, M.W. (1968). Impairment in facial recognition in patients with cerebral disease. *Cortex*, 4, 344–358.
- Blanz, V., Tarr, M.J., & Bülthoff, H.H. (1999). What object attributes determine canonical views? *Perception*, 28, 575–599.
- Brilliant, R. (1991). *Portraiture* (p. 192). Cambridge, MA: Harvard University Press.
- Bruce, V. (1982). Changing faces: Visual and non-visual coding processes in face recognition. *British Journal of Psychology*, 73, 105–116.
- Bruce, V., Valentine, T., & Baddeley, A. (1987). The basis of the $\frac{3}{4}$ view advantage in face recognition. *Applied Cognitive Psychology*, 1, 109–120.
- Bülthoff, H.H., Edelman, S.Y., & Tarr, M.J. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex*, 3, 247–260.
- Caramazza, A., & Shelton, J.R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience*, 10, 1–34.
- Clark, V.P., Maisog, J.M., & Haxby, J.V. (1998). fMRI study of face perception and memory using random stimulus sequences. *Journal of Neurophysiology*, 79, 3257–3265.
- Corballis, M.C. (1991). *The lopsided ape: Evolution of the generative mind*. New York: Oxford University Press.
- Costin, I.D. (1988). MacLab: A Macintosh system for psychology labs. *Behavior, Research Methods, and Computers*, 20, 197–200.
- Damasio, A.R., Damasio, H., & Van Hoesen, G.W. (1982). Prosopagnosia: Anatomic basis and behavioral mechanisms. *Neurology*, 32, 331–341.
- Damasio, H., Grabowski, T.J., Tranel, D., Hichwa, R.D., & Damasio, A.R. (1996). A neural basis for lexical retrieval. *Nature*, 380, 499–505.

- De Renzi, E. (1986). Current issues in prosopagnosia. In H.D. Ellis, M.A. Jeeves, K.F. Newcombe, & A. Young (Eds), *Aspects of face processing* (pp. 243–252). Dordrecht, Holland: Martinus Nijhoff.
- De Renzi, E., Perani, D., Carlesimo, G.A., Silveri, M.C., & Fazio, F. (1994). Prosopagnosia can be associated with damage confined to the right hemisphere—An MRI and PET study and a review of the literature. *Neuropsychologia*, 32, 893–902.
- De Renzi, E., Scotti, G., & Spinnler, H. (1969). Perceptual and associative disorders of visual recognition. *Neurology*, 19, 634–642.
- De Renzi, E., & Spinnler, H. (1966). Facial recognition in brain damaged patients. *Neurology*, 16, 145–152.
- Edelman, S. (1995). Class similarity and viewpoint invariance in the recognition of 3D objects. *Biological Cybernetics*, 72, 207–220.
- Edelman, S., & Bülthoff, H.B. (1992). Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, 32, 2385–2400.
- 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.
- Fagan, J. (1979). The origins of facial pattern recognition. In M. H. Bornstein & W. Keessen (Eds), *Psychological development from infancy: Image to intention*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Fawcett, R., Zisserman, A., & Brady, J. (1994). Extracting structure from an affine view of a 3D point set with one or two bilateral symmetries. *Image and Vision Computing*, 12, 615–622.
- Gainotti, G., Silveri, M.C., Daniele, A., & Giustolisi, L. (1995). Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory*, 3, 247–264.
- Gazzaniga, M.S. (1983). Right hemisphere language following brain bisection: A 20-year perspective. *American Psychologist*, May, 525–537.
- Gazzaniga, M.S., & Smylie, C.S. (1982). Facial recognition and brain asymmetries: Clues to underlying mechanisms. *Annals of Neurology*, 13, 536–540.
- Goldberg, E., & Costa, L.D. (1981). Hemisphere differences in the acquisition and use of descriptive systems. *Brain and Language*, 14, 144–173.
- Hanley, R., Pearson, N.A., & Young, A.W. (1990). Impaired memory for new visual forms. *Brain*, 113, 1131–1148.
- Harries, M.-H., Perrett, D.I., & Lavender, A. (1991). Preferential inspection of views of 3-D model heads. *Perception*, 20, 669–680.
- Haxby, J.V., Ungerleider, L.G., Horwitz, B., Maisog, J.M., Rapoport, S.I., & Grady, C.L. (1996). Face encoding and recognition in the human brain. *Proceedings of the National Academy of Science*, 93, 922–927.
- Haxby, J.V., Ungerleider, L.G., Horwitz, B., Rapoport, S.I., & Grady, C.L. (1995). Hemispheric differences in neural systems for face working memory: A PET-rCBF study. *Human Brain Mapping*, 3, 68–82.
- Hécaen, H., & Albert, M.L. (1979). *Human neuropsychology*. New York: John Wiley & Sons.
- Hécaen, H., & Angelergues, R. (1962). Agnosia for faces (prosopagnosia). *Archives of Neurology (Chicago)*, 7, 92–100.
- Heit, G., Smith, M.E., & Halgren, E. (1988). Neuronal encoding of individual words and faces by the human hippocampus and amygdala. *Nature*, 333, 773–775.
- Hill, H., & Bruce, V. (1996). Effects of lighting on the perception of facial structures. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 986–1004.
- Hill, H., Schyns, P.G., & Akamatsu, S. (1997). Information and viewpoint dependence in face recognition. *Cognition*, 62, 201–222.
- Hosie, J.A., Ellis, H.D., & Haig, N.D. (1988). The effect of feature displacement on the perception of well-known faces. *Perception*, 17, 461–474.

- Jolicoeur, P. (1985). The time to name disoriented natural objects. *Memory and Cognition*, 13, 289–303.
- Kanwisher, N., McDermott, J., & Chun, M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for the perception of faces. *Journal of Neuroscience*, 17, 4302–4311.
- Kapur, N., Friston, K.J., Young, A., Frith, C.D., & Frackowiak, R.S.J. (1995). Activation of human hippocampal formation during memory for faces: A PET study. *Cortex*, 31, 99–108.
- Koenderink, J.J., & van Doorn, A.J. (1982). The shape of smooth objects and the way the contours end. *Perception*, 11, 129–137.
- Koriat, A., & Norman, J. (1985). Mental rotation and visual familiarity. *Perception and Psychophysics*, 37, 429–439.
- Krouse, F.L. (1981). Effects of pose, pose change, and delay on face recognition performance. *Journal of Applied Psychology*, 66, 651–654.
- Landis, T., Cummings, J.L., Christen, L., Bogen, J.E., & Imhof, H.-G. (1986). Are unilateral right posterior cerebral lesions sufficient to cause prosopagnosia? Clinical and radiological findings in six additional patients. *Cortex*, 22, 243–252.
- Lawson, R., Humphreys, G.W., & Watson, D.G. (1994). Object recognition under sequential viewing conditions: Evidence for viewpoint-specific recognition procedures. *Perception*, 23, 595–614.
- Levy, J., Trevarthen, C., & Sperry, R.W. (1972). Perception of bilateral chimeric figures following hemispheric deconnection. *Brain*, 95, 61–78.
- Logie, R.H., Baddeley, A.D., & Woodhead, M.M. (1987). Face recognition, pose and ecological validity. *Applied Cognitive Psychology*, 1, 53–69.
- Marzi, C.A., & Berlucchi, G. (1977). Right visual field superiority for accuracy of recognition of famous faces in normals. *Neuropsychologia*, 15, 751–756.
- McCarthy, G., Puce, A., Gore, J.C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 9, 605–610.
- McIntosh, A.R., Grady, C.L., Haxby, J.V., Ungerleider, L.G., & Horwitz, B. (1996). Changes in limbic and prefrontal functional interactions in a working memory task for faces. *Cerebral Cortex*, 6, 571–584.
- McManus, I.C., & Humphrey, N.K. (1973). Turning the left cheek. *Nature*, 243, 271–272.
- Meadows, J.C. (1974). The anatomical basis of prosopagnosia. *Journal of Neurology, Neurosurgery and Psychiatry*, 37, 489–501.
- Michel, F., Perenin, M.T., & Sieroff, E. (1986). Prosopagnosie sans hémianopsie après lésion unilatérale occipito-temporale droite. *Revue Neurologique*, 142, 545–549.
- Milner, B. (1968). Visual recognition and recall after right temporal lobe excision in man. *Neuropsychologia*, 6, 191–209.
- Moses, Y., Ullman, S., & Edelman, S. (1996). Generalization to novel images in upright and inverted faces. *Perception*, 25, 443–461.
- Nakayama, K., He, Z.J., & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S.M. Kosslyn & D.N. Osherson (Eds), *Visual cognition*, Vol. 2 (pp. 1–70). Cambridge, MA: MIT Press.
- Niemann, T., Lappe, M., & Hoffman, K.P. (1996). Visual inspection of three-dimensional objects by human observers. *Perception*, 25, 1027–1042.
- Newell, F.N., & Findlay, J.M. (1997). The effect of depth rotation on object identification. *Perception*, 26, 1231–1257.
- O'Toole, A.J., Edelman, S., & Bülthoff, H.H. (1998). Stimulus-specific effects in face recognition over changes in viewpoint. *Vision Research*, 38, 2351–2363.
- Palmer, S., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long & A. Baddeley (Eds), *Attention and Performance IX* (pp. 131–151). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Perrett, D.I., Harries, M.-H., & Looker, S. (1991). Use of preferential inspection to define the viewing sphere and characteristic views of an arbitrary machined tool part. *Perception*, 21, 497–515.

- Perrett, D.I., Oram, M.W., & Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: An account of generalisation of recognition without mental transformation. In M.J. Tarr & H.H. Bülthoff (Eds), *Object recognition in man, monkey and machine* (pp. 111–145). Cambridge, MA: MIT Press.
- Sergent, J., Ohta, S., & MacDonald, B. (1992). Functional neuroanatomy of face and object processing. *Brain*, 115, 15–36.
- Sergent, J., & Signoret, J.-L. (1992). Functional and anatomical decomposition of face processing: Evidence from prosopagnosia and PET study of normal subjects. *Philosophical Transactions of the Royal Society of London, B* 335, 55–62.
- Sergent, J., & Villemure, J.-G. (1989). Prosopagnosia in a right hemispherectomized patient. *Brain*, 112, 975–995.
- Takahashi, N., Kawamura, M., Hirayama, K., Shiota, J.-I., & Isono, O. (1995). Prosopagnosia: A clinical and anatomical study of four patients. *Cortex*, 31, 317–329.
- Tanaka, J.W., & Gauthier, I. (1997). Expertise in object and face recognition. In R. Goldstone, D. Medin & P. Schyns (Eds), *The psychology of learning and motivation, Vol. 36* (pp. 83–125). New York: Academic Press.
- Tarr, M.J., & Bülthoff, H.B. (1995). Is human object recognition better described by geon-structural-descriptions or by multiple-views? *Journal of Experimental Psychology: Human Perception and Performance*, 21(6), 1494–1505.
- Tarr, M.J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, 21, 233–282.
- Tarr, M.J., & Pinker, S. (1990). When does human object recognition use a viewer-centred reference frame? *Psychological Science*, 1, 253–256.
- Torii, H., & Tamai, A. (1985). The problem of prosopagnosia: Report of three cases with occlusion of the right posterior cerebral artery. *Journal of Neurology (Supplement)*, 232, 140.
- Troje, N.F., & Bülthoff, H.B. (1996). Face recognition under varying pose: The role of texture and shape. *Vision Research*, 36, 1761–1771.
- Troje, N.F., & Kersten, D. (1999). Viewpoint-dependent recognition of familiar faces. *Perception*, 28, 483–487.
- Tzavaras, A., Merieene, L., & Masare, M.C. (1973). Prosopagnosie, amnesie et troubles du langage par lesion temporale gauche chez un sujet gaucher. *Encéphale*, 62, 382–394.
- Umiltà, C., Brizzolaro, D., Tabossi, P., & Fairweather, H. (1978). Factors affecting face recognition in the cerebral hemispheres: Familiarity and naming. In J. Requin (Ed.), *Attention and performance VII* (pp. 363–374). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Van Kleeck, M.H., & Kosslyn, S.M. (1991). The use of computer models in the study of cerebral lateralization. In F.L. Kitterle (Ed.), *Cerebral lateralization: Theory and research. The Toledo Symposium*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Verfaillie, K., & Boutsen, L. (1995). A corpus of 714 full-color images of depth-rotated objects. *Perception & Psychophysics*, 57, 925–961.
- Vetter, T., Poggio, T., & Bülthoff, H. (1994). The importance of symmetry and virtual views in three-dimensional object recognition: Invariance to imaging transformations. *Current Biology*, 4, 18–23.
- Warrington, E.K., & James, M. (1967). An experimental study of facial recognition in patients with unilateral lesions. *Cortex*, 3, 317–326.
- Warrington, E.K., & Taylor, A.M. (1978). Two categorical stages of object recognition. *Perception*, 7, 695–705.
- Whiteley, A.M., & Warrington, E.K. (1977). Prosopagnosia: A clinical, psychological, and anatomical study of three patients. *Journal of Neurology, Neurosurgery and Psychiatry*, 40, 395–403.
- Yin, R.K. (1970). Face recognition by brain injured patients: A dissociable ability. *Neuropsychologia*, 8, 395–402.

- Young, A.W., Hay, D.C., McWeeny, K.H., Ellis, A.W., & Barry (1984). Familiarity decisions for faces presented to the left and right cerebral hemispheres. *Brain and Cognition*, 4, 439–450.
- Young, A.W., Hay, D.C., McWeeny, K.H., Flude, B.M., & Ellis, A.W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception*, 14, 737–746.