

MIT Press • Open Encyclopedia of Cognitive Science

Face Perception

A. Mike Burton¹

¹University of York, UK and Bond University, Australia

MIT Press

Published on: Jul 24, 2024

DOI: <https://doi.org/10.21428/e2759450.57de91a4>

License: [Creative Commons Attribution-NonCommercial 4.0 International License \(CC-BY-NC 4.0\)](#).

The human face sends many different signals. We can identify somebody, if we happen to know them, but we can also judge their emotional state from their expression, perhaps read their lips if they are speaking, or follow their gaze. We can make broad judgments about somebody's age, gender, and perhaps cultural background, all in signals from the face. The core question in face research is to understand how the perceptual system decodes all this information carried in the visual signal. Researchers have drawn from multiple techniques to understand how different parts of the perceptual system are organized. There is also a growing understanding of the neural structures involved in face recognition and the ways in which these are integrated into our more general knowledge about people. There remain some mysteries about how we recognize each other, but research in face recognition has generated some robust findings, now adopted into forensic and security settings.

History

A key research ambition in the study of face perception has been to understand how the many different sources of information in a face are processed. For example, when recognizing a smiling friend, how do we compute that we are seeing a friend and seeing a smile, and to what extent are the same processes involved when we see a smiling stranger? Does the analysis of facial information occur in a fixed sequence, or do some processes occur in parallel? Early functional accounts of face perception culminated in the [Bruce and Young \(1986\)](#) model, which set out to delineate the various cognitive processes involved in face perception and to place these in the context of word and object recognition—about which much more was known at the time. This same approach, differentiating perception of different signals, was later taken when uncovering the neural architecture supporting face perception, such as in the influential model of [Haxby et al. \(2000\)](#).

The degree to which we all experience similar perceptions of faces has been a longstanding research question in the field. For example, early research by [Ekman and Friesen \(1971\)](#) suggested that the generation of emotional expression, and consequently its perception, was to some extent shared across all human cultures, an idea that had earlier been expressed by Charles Darwin. Similarly, the idea that all human viewers might find the same faces attractive has been much studied—an idea that is also found in ancient civilizations, including early Egyptian and Greek cultures who sought formulaic descriptions of facial beauty. The extent to which these, and other, perceptions of faces are indeed universal, societal, or personal remains a topic of detailed study today.

Much of the early psychological research in face recognition focused on the problem of eyewitness memory. It has been known for many years that witnesses observing an event may subsequently make errors in good faith, either failing to recognize someone present at the event or misidentifying someone who was not present ([Wells & Olson, 2003](#)). From the 1960s onward, researchers have aimed to reduce these errors, for example, by making recommendations about lineup procedures or interview techniques.

While this research has been influential in jurisdictions around the world, the fallible nature of eyewitness testimony remains a significant problem.

At the same time that researchers began to study poor face recognition in the real world, an apparently contradictory body of evidence was accumulating in laboratory-based memory. A number of researchers reported extremely good memory for photographs; for example, people could identify several thousand photographs seen for the first time a few days earlier ([Standing et al., 1970](#)). This apparently excellent memory performance was especially marked for pictures of faces, which seemed much better than memory for images of other stimuli, such as houses or snowflakes ([Goldstein & Chance, 1971](#)). This apparent paradox, good memory for faces in the laboratory alongside frequent well-documented memory errors in the real world, attracted experimental psychologists to the field.

In modern times, face perception research has recruited findings from a wide variety of methods, including laboratory experiments, cognitive neuropsychology, developmental psychology, computational modeling, cognitive neuroscience, and, more recently, individual differences. As the sophistication of these techniques has developed, technical advances have been incorporated into experiment and theory. A coherent theoretical understanding of the processes underlying face perception has emerged slowly and remains incomplete. Progress has typically occurred through consensus based on converging evidence from numerous sources rather than discrete “breakthroughs,” though, of course, many theoretical debates are ongoing.

Core concepts

Familiarity and recognition

Most humans are extremely good at recognizing the people they know. With seemingly no conscious effort, we interact with our family, friends, and colleagues, rarely puzzling over their identity, and we can identify these people even in very poor images ([Burton et al., 1999](#)). However, this ability stands in stark contrast to our recognition of unfamiliar people. As noted, witnesses asked to recognize someone who was seen for the first time during some critical incident are often mistaken. In fact, while a great deal of research on eyewitness memory has focused on fallible memory, it turns out that people are poor at recognizing unfamiliar faces, even in situations that do not require memory.

In face matching tasks, viewers are asked to indicate whether faces presented simultaneously belong to the same person or not. These tasks might be presented as lineups (e.g., “Which of these people is the target face?”) or simple paired tasks (e.g., “Are these two photos the same person?”). These types of decisions are commonplace in everyday life, such as when proving one’s identity using a passport or other photo ID. Matching is also important in more specialist settings, such as when police officers attempt to identify someone captured on CCTV. However, despite the fact that face matching is commonplace, research consistently shows that people are generally poor at matching unfamiliar faces, making frequent errors, even with good quality photographs taken recently ([Bruce et al., 1999](#)). Perhaps

even more surprising is the fact that professional people who use face matching in daily life (e.g., passport officers, notaries, check-out staff) are typically no better at the task than anyone else ([White et al., 2014](#)).

The large discrepancy between our perceptual abilities with familiar and unfamiliar faces remains a critical aspect of face perception. One key observation is that we can recognize familiar faces over a wide range of photographs (e.g., changes in view, lighting, age, expression, and so forth), whereas representations of people we do not know are more tied to specific images. For example, in memory experiments, people are good at remembering faces if the same image is used at learning and test. But if a different photograph is used on the two occasions, familiar face recognition remains good while unfamiliar face recognition is much reduced ([Bruce, 1982](#)). This finding resolves the paradox raised in the History section: Why is laboratory-based face memory so good when memory in the field seems so poor? In short, *picture memory* is excellent for all faces, but *face memory*, generalizing across different pictures, is only good for the people we already know.

Modern studies of familiar face recognition incorporate generalizability by highlighting both between-person variability (people look different from each other) and within-person variability (the same person looks different on different occasions) ([Jenkins et al., 2011](#)). The problem of face recognition, then, incorporates both *telling people together* and *telling people apart*, processes which can be tracked as faces are learned and as children develop these skills ([Laurence & Mondloch, 2016](#)).

The inversion effect, holistic and configural processing

It is easy to demonstrate that recognizing upside-down faces is difficult—just present a friend with an inverted image of somebody famous. However, researchers studying the *face inversion effect* go beyond this demonstration to investigate the claim that faces are not only difficult to process upside down but *particularly* difficult—that is, perception of faces is affected more by inversion than are other visual stimuli. This is a longstanding idea ([Yin, 1969](#)) and has been a key component of many theoretical developments within the field, including the concepts of *configural* and *holistic processing*. There are many definitions of these terms, but the field has largely adopted those of [Maurer et al. \(2002\)](#): *Holistic processing* coheres facial features into a perceptual gestalt, *first order configural processing* involves detection of the basic layout of a face (eyes above nose above mouth), and *second order configural processing* is sensitive to the spatial layout of the facial features.

Holistic processing is often demonstrated using the *composite face effect* ([Young et al., 1987](#)). If we take the top half of one face and position it over the bottom half of another, then the two halves appear to fuse into an entirely novel identity—even if the two constituent faces are well-known. (Try it.) This seems to indicate that our perceptual system is somehow designed to perceive the identity of whole faces, as though some kind of template were triggered when a face shape appears. The illusion is quickly destroyed if the two half-faces are inverted or slightly misaligned.

Holistic and configural processing have often been contrasted with *featural processing*—a putative strategy in which individual features (eyes, noses, etc.) form the components of one's representation for a particular individual. In contrast, the *face superiority effect* ([Tanaka & Farah, 1993](#)) demonstrates the importance of a whole face. When asked to remember a particular nose, for example, viewers are more accurate when the nose appears within a whole face than when it appears in isolation or in a face in which the features have been jumbled up. The whole face context supports recognition, even though the participants are only required to remember one feature—good evidence for the role of a whole-face representation.

Configural processing is held to underlie the face inversion effect. When turned upside down, it becomes harder to discern configural changes (moving features slightly within a face) than featural changes (e.g., swapping one set of eyes for another) ([Le Grand et al., 2001](#)). This leads to the hypothesis that inversion disproportionately affects configural processing and, since face recognition is held to rely on configural processing, explains why faces are particularly hard to recognize upside down. This is a very popular theoretical position, but it remains somewhat unproven in the absence of a generally agreed measure of configural processing that can be taken independently of inversion. It also remains unclear whether configural processing underlies all face processing or only processing for certain tasks. For example, are configural processes recruited equally for familiar and unfamiliar faces? And what role do they play in judgments of expression or identity?

The neural system supporting face perception

Advances in functional brain imaging have brought considerable progress in understanding the neural systems involved in face processing. Before imaging data became available, scientists had relied on information from neuropsychological patients whose deficits in face perception were associated with specific patterns of brain damage and on single-cell recordings made in monkeys. These sources had suggested that face perception involves the temporal cortex and is supported bilaterally in the brain but with a greater emphasis on processing in the right hemisphere.

Imaging data have revealed a *core network* ([Haxby et al., 2000](#)) of three brain areas that are particularly responsive to faces: the occipital face area (OFA), the fusiform face area (FFA), and the superior temporal sulcus (STS). The OFA is held to be involved in early visual processing of facial features and seems to act as an entry point to the distributed network for face processing. The FFA shows sensitivity to faces over many other classes of object (e.g., animals, houses, or jumbled faces) ([Kanwisher et al., 1997](#)). It is held to code those aspects of a face that are invariant, or at least change very slowly, such as a person's identity. The STS, on the other hand, is known to be involved in many social processes ([Pitcher & Ungerleider, 2021](#)) and appears to be closely associated with changeable aspects of the face, such as the perception of eye gaze, expression, and lip movements—all critical to our social interactions.

Of course, our interactions with others rely on many complex sources of information and not simply on visual analysis of faces (see [Theory of Mind](#)). For the people we know, there is a large store of individual

knowledge, and our relationship with the person also informs the way we behave—for example, with a family member versus a colleague. For people we do not know, we nevertheless make judgments about their attitude to us based in part on their appearance. This complex array of personal information is processed in the *extended face network*, which includes a more extensive range of brain areas, many of them not specialized for visual input. For example, it is now well-established that areas including the amygdala, anterior temporal cortex, and auditory cortex can all be recruited in our social interactions, and these are functionally connected to the core face areas FFA and STS in complex ways that are still being explored.

While neuroimaging techniques have enabled researchers to understand *where* in the brain faces are analyzed, they are poorer at revealing how this process unfolds over time. In contrast, the study of event-related potentials (ERPs), which have high temporal but low spatial resolution, has shown some responses that are sensitive to faces. The best-studied ERP component is the N170—a pattern of electrical negativity on the scalp that is observed roughly 170 ms after the presentation of a face, most clearly on electrodes mounted over posterior temporal sites ([Bentin et al., 1996](#)). The N170 is generally elicited by the presentation of all faces, whether they are familiar to the viewer or not. Later components, appearing around 250 ms and 400 ms following presentation, have been hypothesized to reflect visual and then conceptual knowledge about the person shown, respectively ([Wiese et al., 2023](#)).

Individual differences in face perception

There are widespread differences in people’s abilities in face perception. While there is some evidence for differing abilities in recognizing familiar people or identifying emotional expressions, the majority of work in individual differences focuses on unfamiliar face identity processing ([White & Burton, 2022](#)). This reflects the fact that there are some well-established tests of unfamiliar face recognition, such as the Cambridge Face Memory Test ([Duchaine & Nakayama, 2006](#)), that have well-documented psychometric properties and relatively high levels of test–retest reliability, making them good tools for work in individual differences.

There is considerable research focus on the extremes of face perception ability. Prosopagnosia refers to a deficit in which the patient cannot recognize the faces of others. In its acquired form, people may have suffered head injury or stroke. While it is often associated with other perceptual problems, some relatively “pure” cases exist, in which patients have no difficulty making other judgments about a face (i.e., they can point to features, estimate age, etc.) and can recognize familiar people from other cues (voices, clothes, etc.). Over recent years, there has been growing interest in *developmental prosopagnosia*, the condition in which people cannot recognize faces but have no known brain pathology ([Barton & Corrow, 2016](#)). This condition is sometimes known as *face blindness* and can result in social challenges associated with failing to recognize known people.

At the other end of the scale are people with extremely good face recognition abilities, labeled *super-recognizers* ([Russell et al., 2009](#)). These individuals consistently score much higher than the average

participant on tests of unfamiliar face identification (memory and matching) and have sometimes been recruited to specialist forensic and security organizations to support their operations. Nevertheless, this enhanced ability can come at some cost, and some super-recognizers report hiding their abilities in order not to stand out from the ordinary expectations of others.

While those lying at each end of the ability continuum are interesting, the full scale of abilities is utilized for research into fundamental aspects of face perception. For example, correlational studies can be used to explore the relationship between face recognition and other perceptual and cognitive abilities ([Yovel et al., 2014](#)). In fact, only relatively modest associations have been reported between face and general ability in object perception, whereas performance on different face perception tasks tends to correlate more strongly ([Verhallen et al., 2017](#)). Measures of individual differences can also be used to examine the genetic underpinnings of face recognition. Large-scale studies, including twin studies, conclude that face recognition is highly heritable ([Wilmer et al., 2010](#)). Furthermore, correlations have been observed between face recognition ability and aspects of an individual's neural anatomy ([Elbich & Scherf, 2017](#)), though this association remains controversial.

First impressions and social attribution

The idea that people's faces reveal aspects of their character is an ancient one that has little currency in modern science. However, contemporary research has been able to uncover robust associations between the physical structure of faces and our tendency to make *judgments* about them ([Todorov et al., 2015](#)). Researchers have sometimes asked participants to make free descriptions of faces or sometimes to rate them on multiple dimensions. In either case, their attributions can be summarized as lying along a small number of core dimensions that seem to capture social judgments. Ratings based on carefully controlled computer-generated faces—that is, without the visual “noise” introduced in natural photographs—suggest a two-factor model in which most of the variance in social attributions is described by *dominance* and *trustworthiness* ([Oosterhof & Todorov, 2008](#)). Later work, using naturally varying images gathered from the internet, has also reported a structure using these two factors but with a third important judgment dimension: *youthful attractiveness* ([Sutherland et al., 2013](#)).

Social judgments of faces are made very fast. For example, people apply attributes to photographs shown for 100 ms, or even less, and these tend to not change given longer exposures. This implies a quite fundamental, reflexive relationship between the structure of the face and the nature of the attribution—the display time is simply too short to allow for more complex decision-making. There is also a high level of agreement among participants on these attributions. The same faces tend to be viewed as highly trustworthy, for example, by most viewers. Once again, this suggests a strong association between physical face structure and social judgments.

One of the most puzzling aspects of social perception is that the judgments made to faces have poor validity. Despite the fact that most people will agree about whether or not a person seems agreeable (say), there is very little correspondence between the attribution and the person's actual agreeableness. This is

an important concern because there is clear evidence that people are willing to make serious real-world decisions on the basis of their first impressions. For example, judgments of competence, made solely on someone's face, have been shown to predict real election results ([Olivola & Todorov, 2010](#)), and judgments of facial trustworthiness predict real sentencing decisions ([Wilson & Rule, 2015](#)). It is therefore critical that we understand the nature of these biases in settings where they have the potential to influence important social outcomes.

Automatic face recognition systems

Alongside research focused on understanding human face perception, a large body of work based in engineering and computer sciences aims to recognize faces accurately, without necessarily trying to emulate the human system. Automatic face recognition systems are now commonplace at passport checks, secure access locations, and on personal technology such as phones and laptops.

The development of deep convolutional neural networks (DCNNs) has led to striking improvements in face recognition technology in recent years. These systems are now able to recognize faces over natural variations in viewing conditions (e.g., lighting, viewpoint, expression, etc.) with high levels of accuracy. They are based on multiple highly interconnected layers of simple processing units that are trained on very large numbers of faces—typically millions gathered over the internet. In simple face-matching tasks (“do these two images show the same person?”), these systems are now reliably more accurate than most human viewers ([O’Toole & Castillo, 2021](#)).

Automatic face recognition is normally used in settings in which a human viewer *unfamiliar* with the faces would traditionally make a decision (e.g., checking photo ID). While machines comfortably outperform most humans on such tasks, it remains unclear how they can perform against *familiar* viewers, who typically achieve very high accuracy rates when recognizing the people they know. Nonetheless, the high levels of accuracy in some face tasks have led some researchers to ask what is common between representations used in human cognition and in the DCNNs ([Nestor et al., 2020](#); [O’Toole & Castillo, 2021](#)), and this remains an active research question in cognitive science. In more practical terms, the technology for face recognition is now ubiquitous in many countries, being widely used in security and law enforcement. While accuracy is high, systems are not perfect, and serious concerns have been raised about the potential for biases, particularly racial biases, reflecting the nature of the training sets on which these systems are based ([Birhane, 2022](#)).

Questions, controversies, and new developments

Are faces special?

One of the most controversial debates in the field of face perception has focused on the degree to which faces are “special” for perception. The key question is whether the human perceptual system has resources dedicated to face processing, including dedicated neural hardware, or whether the visual system is equipped with general mechanisms for making fine-scale discriminations between similar

objects that are recruited by face perception but are also available for other visual discriminations we might make. For example, if someone is an expert in judging dog shows or collecting stamps, that person makes fine-scale discriminations among similar items (e.g., dogs, stamps), which might make perceptual demands similar to face perception. Brain systems could have evolved for more general purposes, which are then used for face perception, but not exclusively.

The debate around this issue has elicited strong arguments on either side ([Gauthier et al., 2000](#); [Grill-Spector et al., 2004](#)). Evidence has been gathered from the study of visual experts in various specialties, neuroimaging, neuropsychological patients, and perception of artificial stimuli designed to capture some of the characteristics of faces. While the issue has not been conclusively resolved, the growing popularity of studies on individual differences suggests a rather isolated trait of face recognition, which does not correlate highly with other traits. However, since all such studies are inherently correlational, a final resolution seems unlikely. Many of the topics studied by researchers in face perception are orthogonal to the question of whether or not faces are special, and the field is now less focused on this question than in previous years.

The cognitive neuroscience of familiarity and learning

As discussed in the section [Familiarity and recognition](#), there are large differences between the processing of familiar and unfamiliar faces—such that viewers are typically good at recognizing multiple photographs of a familiar person but show more image-specific responses to unfamiliar faces. These effects can be observed over a range of behavioral measures, including memory for faces, face matching, and identity sorting (i.e., categorizing together multiple photographs of the same person). We might therefore expect equally large effects in neuroimaging, for example, showing differences in activation levels of key brain areas when viewing unfamiliar versus familiar faces. Interestingly, results of studies seeking simple location-based neural correlates of face familiarity have been equivocal.

One approach to this puzzle has been to propose that studies based on regions of interest in the brain are too simplistic and that instead, a whole brain approach is more likely to reveal key differences between processing familiar and unfamiliar faces ([Kovács, 2020](#)). Network analyses tend to reveal clearer effects of familiarity in the extended face processing network (see section [The neural system supporting face perception](#)) than the core visual areas. Of course, becoming familiar with new people involves much more than learning about their faces—we also come to learn information about them and perhaps develop an emotional response to them. One might therefore expect to see the involvement of extended network areas as we learn someone new, particularly in naturalistic settings rather than through tightly controlled laboratory-based exposure to repeated face images.

Studies of naturalistic learning are beginning to emerge, which track the perceptual changes that occur as a new person becomes familiar. For example, ERP evidence suggests that it can take many months to develop conceptual representations of new people but a shorter time to build robust visual representations of their faces ([Popova & Wiese, 2022](#)). Such studies are consistent with neuroimaging

research and reinforce the idea that there is no simple neural correlate of familiarity. Instead, what seems like a straightforward concept from behavioral evidence turns out to be more complex when studied physiologically.

Own- and other-group effects

It has been known for many years that memory for members of one's own group is better than memory for people outside this group. Most of the research in this field has concentrated on the *other-race effect*, in which members of different ethnic groups show a memory advantage for faces of their own ethnicity ([Meissner & Brigham, 2001](#)). Though generally demonstrated in memory experiments, this effect is also present for other tasks such as face matching or making perceptual decisions about composite faces. Explanations for the *other-race effect* have typically focused on perceptual experience; for example, the faces one sees while growing up will vary in particular ways, so viewers become sensitive to variation along those dimensions but not the dimensions along which faces from another ethnicity vary ([Valentine, 1991](#)). This is consistent with studies showing that children adopted between different cultural contexts show reversed *other-race effects*, taking on the same patterns of memory as the adoptive culture ([Sangrigoli et al., 2005](#)).

In fact, analogous face processing biases are also observed elsewhere. For example, both own-age and own-gender biases have been reported, each of which could plausibly arise from differential experience. However, a challenge to the perceptual account of these effects arises from studies that show similar biases across social groups. For example, when shown faces ostensibly from one's own or another university, students tend to remember the faces from their own institution, suggesting that in-group faces are encoded better than out-group faces, independently of any physical differences. Furthermore, the other-race effect can be eliminated by instructing viewers explicitly to individuate all faces when encountering them for the first time—suggesting that people may not dedicate as much attention to out-group faces, leading to subsequent poorer memory for them.

The explanation for these effects remains contentious, and there have been some attempts to integrate theories ([Sporer, 2001](#)), though *other-race effects* have been easier to replicate than some *other-group effects*, suggesting they may be more robust. Furthermore, an analogous problem in automatic face recognition has recently arisen (see section [Automatic face recognition systems](#)). The accuracy of algorithms in everyday use is affected by the training set faces used to build face recognition systems, and this can lead to biased outcomes—an issue of pressing concern to current developers of these systems ([Cavazos et al., 2021](#)).

Broader connections

While the study of face perception has its beginnings in witness reliability, there remain strong ties between those attempting to understand how human viewers process faces and those working in judicial and security settings. Despite advances in other biometrics, including fingerprint, iris, and voice

recognition, faces are still an important cue used to identify people in many settings. The psychological study of individual differences has supported the development of recruitment tools for people working in these fields, and the use of combined human/computer decision-making is under active investigation.

A common criticism of laboratory-based research in face recognition is that the tasks used in experiments typically do not reflect recognition in the real world adequately. When we recognize someone in the workplace, for example, we might simultaneously have access to information about their face, voice, gait, and clothing and contextual information about where we expect to meet particular people. There is a substantial body of research in voice recognition that shares many of the same problems with face recognition. For example, to recognize speakers, we must be able to tell people apart (to the extent that different people sound different) and tell them together (to the extent that the same person can sound different on different occasions). How information is combined across audio and visual channels, and the potential use of real-life redundancy across identity cues, is an active topic for current research ([Lavan et al., 2022](#); [Young et al., 2020](#)).

More broadly, the study of face recognition is beginning to move out of the laboratory and into natural settings. As with other areas of cognition, there is a growing acknowledgment that the advantages of laboratory experiment (e.g., the ability to control stimuli precisely) can sometimes obscure key aspects of the problem under study. In the engineering domain, face recognition “in the wild” ([Huang et al., 2007](#)) has now become the standard problem to be solved. This is not yet the case for studies of human cognition, where demands of experimental control, and particularly replicability, make real-world studies difficult to carry out. Nevertheless, there is a growing acknowledgment that to understand how we recognize people, it will be necessary to combine our knowledge from multiple sources, including detailed study of real-life recognition.

Further reading

- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77, 305–327. <https://doi.org/10.1111/j.2044-8295.1986.tb02199.x>
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6), 223–233. [https://doi.org/10.1016/S1364-6613\(00\)01482-0](https://doi.org/10.1016/S1364-6613(00)01482-0)
- Maurer, D., Grand, R. L., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255–260. [https://doi.org/10.1016/S1364-6613\(02\)01903-4](https://doi.org/10.1016/S1364-6613(02)01903-4)
- Todorov, A., Olivola, C. Y., Dotsch, R., & Mende-Siedlecki, P. (2015). Social attributions from faces: Determinants, consequences, accuracy, and functional significance. *Annual Review of Psychology*, 66, 519–545. <https://doi.org/10.1146/annurev-psych-113011-143831>

References

- Barton, J. J. S., & Corrow, S. L. (2016). The problem of being bad at faces. *Neuropsychologia*, 89, 119–124. <https://doi.org/10.1016/j.neuropsychologia.2016.06.008>

↩

- Bentin, S., Allison, T., Puce, A., Perez, E., & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. *Journal of Cognitive Neuroscience*, 8(6), 551–565. <https://doi.org/10.1162/jocn.1996.8.6.551>

↩

- Birhane, A. (2022). The unseen Black faces of AI algorithms. *Nature*, 610(7932), 451–452. <https://doi.org/10.1038/d41586-022-03050-7>

↩

- Bruce, V. (1982). Changing faces: Visual and non-visual coding processes in face recognition. *British Journal of Psychology*, 73(Pt 1), 105–116. <https://doi.org/10.1111/j.2044-8295.1982.tb01795.x>

↩

- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77, 305–327. <https://doi.org/10.1111/j.2044-8295.1986.tb02199.x>

↩

- Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology. Applied*, 5(4), 339–360. <https://doi.org/10.1037/1076-898X.5.4.339>

↩

- Burton, A. M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science*, 10(3), 243–248. <https://doi.org/10.1111/1467-9280.00144>

↩

- Cavazos, J. G., Phillips, P. J., Castillo, C. D., & O'Toole, A. J. (2021). Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(1), 101–111. <https://doi.org/10.1109/TBIOM.2020.3027269>

↩

- Duchaine, B., & Nakayama, K. (2006). The Cambridge Face Memory Test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants. *Neuropsychologia*, 44(4), 576–585. <https://doi.org/10.1016/j.neuropsychologia.2005.07.001>

↩

- Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology*, 17(2), 124–129. <https://doi.org/10.1037/h0030377>

↩

- Elbich, D. B., & Scherf, S. (2017). Beyond the FFA: Brain-behavior correspondences in face recognition abilities. *NeuroImage*, 147, 409–422. <https://doi.org/10.1016/j.neuroimage.2016.12.042>

↩

- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3(2), 191–197. <https://doi.org/10.1038/72140>

↩

- Goldstein, A. G., & Chance, J. E. (1971). Visual recognition memory for complex configurations. *Perception & Psychophysics*, 9(2), 237–241. <https://doi.org/10.3758/BF03212641>

↩

- Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, 7(5), 555–562. <https://doi.org/10.1038/nn1224>

↩

- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 4(6), 223–233. [https://doi.org/10.1016/S1364-6613\(00\)01482-0](https://doi.org/10.1016/S1364-6613(00)01482-0)

↩

- Huang, G., Ramesh, M., Berg, T., & Learned-Miller, E. (2007). Labeled faces in the wild: A database for studying face recognition in unconstrained environments.

↩

- Jenkins, R., White, D., Van Montfort, X., & Mike Burton, A. (2011). Variability in photos of the same face. *Cognition*, 121(3), 313–323. <https://doi.org/10.1016/j.cognition.2011.08.001>

↩

- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 17(11), 4302–4311. <https://doi.org/10.1523/JNEUROSCI.17-11-04302.1997>

↩

- Kovács, G. (2020). Getting to know someone: Familiarity, person recognition, and identification in the human brain. *Journal of Cognitive Neuroscience*, 32(12), 2205–2225. https://doi.org/10.1162/jocn_a_01627

↩

- Laurence, S., & Mondloch, C. J. (2016). That's my teacher! Children's ability to recognize personally familiar and unfamiliar faces improves with age. *Journal of Experimental Child Psychology*, 143, 123–138. <https://doi.org/10.1016/j.jecp.2015.09.030>

↩

- Lavan, N., Smith, H. M. J., & McGettigan, C. (2022). Unimodal and cross-modal identity judgements using an audio-visual sorting task: Evidence for independent processing of faces and voices. *Memory & Cognition*, 50(1), 216–231. <https://doi.org/10.3758/s13421-021-01198-7>

↩

- Le Grand, R., Mondloch, C. J., Maurer, D., Brent, H. P., Le Grand, R., Mondloch, C. J., Maurer, D., & Brent, H. P. (2001). Neuroperception. Early visual experience and face processing. *Nature*, 410, 890–890. <https://doi.org/10.1038/35073749>

↩

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

↩

- Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law. An Official Law Review of the University of Arizona College of Law and the University of Miami School of Law*, 7(1), 3–35.

↩

- Nestor, A., Lee, A. C. H., Plaut, D. C., & Behrmann, M. (2020). The face of image reconstruction: Progress, pitfalls, prospects. *Trends in Cognitive Sciences*, 24, 747–759. <https://doi.org/10.1016/j.tics.2020.06.006>

↩

- O'Toole, A. J., & Castillo, C. D. (2021). Face recognition by humans and machines: Three fundamental advances from deep learning. *Annual Review of Vision Science*, 7, 543–570. <https://doi.org/10.1146/annurev-vision-093019-111701>

↩

- Olivola, C. Y., & Todorov, A. (2010). Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of Nonverbal Behavior*, 34(2), 83–110. <https://doi.org/10.1007/s10919-009-0082-1>

↩

- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences of the United States of America*, 105(32), 11087–11092. <https://doi.org/10.1073/pnas.0805664105>

↩

- Pitcher, D., & Ungerleider, L. G. (2021). Evidence for a third visual pathway specialized for social perception. *Trends in Cognitive Sciences*, 25(2), 100–110. <https://doi.org/10.1016/j.tics.2020.11.006>

↩

- Popova, T., & Wiese, H. (2022). The time it takes to truly know someone: Neurophysiological correlates of face and identity learning during the first two years. *Biological Psychology*, 170, 108312. <https://doi.org/10.1016/j.biopsycho.2022.108312>

↩

- Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16(2), 252–257. <https://doi.org/10.3758/PBR.16.2.252>

↩

- Sangrigoli, S., Pallier, C., Argenti, A.-M., Ventureyra, V. A. G., & de Schonen, S. (2005). Reversibility of the other-race effect in face recognition during childhood. *Psychological Science*, 16(6), 440–444. <https://doi.org/10.1111/j.0956-7976.2005.01554.x>

↩

- Sporer, S. L. (2001). Recognizing faces of other ethnic groups: An integration of theories. *Psychology, Public Policy, and Law. An Official Law Review of the University of Arizona College of Law and the University of Miami School of Law*, 7(1), 36–97.

↩

- Standing, L., Conezio, J., & Haber, R. N. (1970). Perception and memory for pictures: Single-trial learning of 2500 visual stimuli. *Psychonomic Science*, 19(2), 73–74. <https://doi.org/10.3758/BF03337426>

↩

- Sutherland, C. A. M., Oldmeadow, J. A., Santos, I. M., Towler, J., Michael Burt, D., & Young, A. W. (2013). Social inferences from faces: Ambient images generate a three-dimensional model. *Cognition*, 127(1), 105–118. <https://doi.org/10.1016/j.cognition.2012.12.001>

↩

- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, 46(2), 225–245. <https://doi.org/10.1080/14640749308401045>

↩

- Todorov, A., Olivola, C. Y., Dotsch, R., & Mende-Siedlecki, P. (2015). Social attributions from faces: Determinants, consequences, accuracy, and functional significance. *Annual Review of Psychology*, 66, 519–545. <https://doi.org/10.1146/annurev-psych-113011-143831>

↩

- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, 43(2), 161–204. <https://doi.org/10.1080/14640749108400966>

↩

- Verhallen, R. J., Bosten, J. M., Goodbourn, P. T., Lawrance-Owen, A. J., Bargary, G., & Mollon, J. D. (2017). General and specific factors in the processing of faces. *Vision Research*, 141, 217–227. <https://doi.org/10.1016/j.visres.2016.12.014>

↩

- Wells, G. L., & Olson, E. A. (2003). Eyewitness testimony. *Annual Review of Psychology*, 54(1), 277–295. <https://doi.org/10.1146/annurev.psych.54.101601.145028>

↩

- White, D., & Burton, A. M. (2022). Individual differences and the multidimensional nature of face perception. *Nature Reviews Psychology*, 1(5), 287–300. <https://doi.org/10.1038/s44159-022-00041-3>

↩

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

↩

- Wiese, H., Schipper, M., Popova, T., Burton, A. M., & Young, A. W. (2023). Personal familiarity of faces, animals, objects, and scenes: Distinct perceptual and overlapping conceptual representations. *Cognition*, 241, 105625. <https://doi.org/10.1016/j.cognition.2023.105625>

↩

- Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., Nakayama, K., & Duchaine, B. (2010). Human face recognition ability is specific and highly heritable. *Proceedings of the National Academy of Sciences of the United States of America*, 107(11), 5238–5241. <https://doi.org/10.1073/pnas.0913053107>

↩

- Wilson, J. P., & Rule, N. O. (2015). Facial trustworthiness predicts extreme criminal-sentencing outcomes. *Psychological Science*, 26(8), 1325–1331. <https://doi.org/10.1177/0956797615590992>

↩

- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81(1), 141–145. <https://doi.org/10.1037/h0027474>

↩

- Young, A. W., Fröhholz, S., & Schweinberger, S. R. (2020). Face and voice perception: Understanding commonalities and differences. *Trends in Cognitive Sciences*, 24(5), 398–410. <https://doi.org/10.1016/j.tics.2020.02.001>

[↵](#)

- Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, 16(6), 747–759. <https://doi.org/10.1068/p160747>

[↵](#)

- Yovel, G., Wilmer, J. B., & Duchaine, B. (2014). What can individual differences reveal about face processing? *Frontiers in Human Neuroscience*, 8(August), 562. <https://doi.org/10.3389/fnhum.2014.00562>

[↵](#)