

# Predicting Human Performance for Face Recognition <sup>1</sup>

Alice J. O’Toole  
The University of Texas at Dallas  
Email: otoole@utdallas.edu

Fang Jiang  
The University of Texas at Dallas  
Email: fxj018100@utdallas.edu

Dana Roark  
The University of Texas at Dallas  
Email: danar@utdallas.edu

Hervé Abdi  
The University of Texas at Dallas  
Email: herve@utdallas.edu

## Abstract

*The ability of humans to recognize faces provides an implicit benchmark for gauging the performance of automatic face recognition algorithms. In this chapter we review the factors that affect human accuracy. These factors can be classified into facial structure constraints and viewing parameters. The former include factors such as face typicality, gender, and ethnicity. The latter include changes in illumination and viewpoint, as well as the perceptual complications introduced when we see faces and people in motion. The common thread of the chapter is that human experience and familiarity with faces can overcome many, if not all, of these challenges to face recognition. A goal of computational algorithms should be to emulate the ways in which humans acquire familiarity with faces. It may then be possible to apply these principles to the design of algorithms to meet the pressing challenges of face recognition in naturalistic viewing conditions.*

## 0.1 Introduction

Human face recognition abilities are impressive by comparison to the performance of many current automatic face recognition systems. Such is the belief of many psychologists and computer scientists. This belief, however, is supported more by anecdotal impression than by scientific evidence. In fact, there are relatively few systematic and direct comparisons between the performance of humans and automatic face recognition algorithms, (though see [84], for exceptions). Comparisons of these sorts can be carried out both at a quantitative and qualitative level. On the qualitative side, as the field of automatic face recognition develops, the assessment of working systems has expanded to consider, not only overall levels of performance, but performance under more restricted and targeted conditions. It may thus be important to know, for example, whether an algorithm performs accurately recognizing “male faces of a certain age”, “female faces of a certain ethnicity”, or faces that are “typical” in the context of the database to be searched. How secure is a security system that was developed to operate on Caucasian faces when it is put to the test in a context filled with diverse faces of many ethnicities?

By contrast to the study of automated face recognition, the study of human face recognition abilities has long focused on a range of qualitative factors that affect performance. The rationale for this endeavor has been to consider these factors as indices into the nature of human representations of faces. Do we encode a face in terms of its three-dimensional structure? Do we encode a face in terms of its similarity relationship to the other faces

---

<sup>1</sup>This work was supported by a contract from TSWG to A. O’Toole and H. Abdi.

we know? In fact, much is known about the qualitative factors that affect human accuracy recognizing faces. There is ample evidence to suggest that not all faces are equally recognizable. There is evidence also that the experiential history of individuals affects their ability to recognize faces.

Using the range of qualitative factors that are known to affect human performance, we can begin to tag individual faces with a probability of successful recognition. In addition to the properties that affect the recognizability of individual faces, much is known also about the viewing conditions that support good performance and about those that make humans prone to error. More recently, the study of human face recognition has begun to consider the factors that affect face and person recognition in more natural contexts. Accordingly, studies in recent years have considered the question of how humans recognize faces in motion. These studies open a window into understanding the mechanics of face recognition in the real world, where there is little or no control of illumination and where the relationship between the viewer and the person-to-be-recognized changes continuously. All of these factors impact the likelihood of successful human recognition in ways that are more or less knowable through the application of experimental studies of human memory.

In this chapter, we review the factors that affect human accuracy for face recognition. These factors can be classified into four categories: a.) face-based constraints, including the effects of typicality, gender, age, and ethnicity; b.) viewing constraints, including the effects of illumination, viewpoint, pose, and facial motion, and c.) experiential constraints, including how experience with individual faces and groups of faces affects recognition performance. Our purpose in each case is to define and discuss the nature of the predictors and the implications each has for understanding how humans represent and process faces. Our secondary goal is to begin to sketch out, for the field of automatic face recognition, a series of factors that might prove useful in deriving a more detailed and refined set of measures for the performance of face recognition algorithms. When all is said and done, any given face recognition algorithm must compete, not only against other automatic recognition algorithms, but against humans, who are currently performing the task in most applied situations.

## 0.2 Face-based Factors and the Face Space Model

A face is a complex three-dimensional object, with overlying surface pigmentation that specifies the reflectance properties of the skin, eyes, mouth and other features. To remember a face as an individual, and to distinguish it from other known and unknown faces, we must encode information that makes the face unique.

We begin by introducing the concept of a face space in psychological models of face processing [121]. The model has just a few components, but is ubiquitous in psychological and computational theorizing about face recognition. In a face space model of recognition, faces can be considered as points in a multidimensional space. The axes of the space represent the “features” with which a face is encoded. A face can be represented, therefore, by its coordinates on these axes, which specify the face’s value on each of the feature dimensions. At a psychological level, it is not necessary to specify the nature of features that form the axes of a space. It is generally enough to know that faces can be described using a set of features dimensions (e.g., width of face, distance between eyes, eye color, etc.) and that different faces vary in the combination of feature values they need to be re-created.

At the computational level, physical face spaces form the core of most automatic face recognition algorithms. The feature axes in computational models represent real physical features, often extracted statistically using principal or independent components analysis.

These kinds of analyses have been applied to face images (e.g., [115]), separated face shapes and image data (e.g., [34]), or three-dimensional data from laser scans (e.g., [12]). Face recognition then becomes a problem of projecting a “target” face into a space and comparing its location to other faces in the space. If the target face is sufficiently close to one of these other faces, then it is recognized. Otherwise, the face is rejected as unknown.

In the context of a face space model, it is relatively easy to see conceptually how the properties of individual faces and categories of faces may affect both human and machine recognition accuracy. We present the effects of each of these categories on human face recognition performance using the analogy of a face space as the supporting theoretical construct.

### 0.2.1 Typicality

A longstanding focus of research in human perception and memory centers on the importance of the “average” or “prototype” in guiding recognition and categorization of visual stimuli. The theory is that categories of objects, including faces, are organized around a prototype or average. The prototype is the best or most typical item in the category. Typical items, those similar to the prototype, are easier to categorize as exemplars of the category than are unusual items. For example, human subjects categorize sparrows as “birds” faster than they categorize penguins as “birds”, because sparrows are “typical” birds and penguins are “unusual” birds [101]. In general, the idea is that the closer an item is to the category prototype, the easier it is to “recognize” as an exemplar of the category.

The problem for faces, however, is to “recognize” each person as an *individual*. Thus, it is not sufficient to say that something is a face. Rather, we must determine: a.) if the face is known to us; and if so, b.) whose face it is. Here, the inverse relationship holds between typicality and recognizability. Typical faces, which are similar to the prototype, are more difficult to recognize. This occurs, presumably, because typical faces are plentiful and so there are more faces that could be falsely mistaken for any given typical face than for any given distinctive face [121].

Face typicality is one of the best known and most robust predictors of human face recognition performance. In a classic study, Light et al. (1979) [73] found that faces rated as “typical” were recognized less accurately than faces rated as “unusual”. In a follow-up experiment, Light et al. showed convincingly that the disadvantage found for typical faces was due to their high inter-item similarity. Faces rated by subjects as “typical” were found to be generally more similar to other faces, than faces rated by subjects as “unusual”. These findings have been replicated many times, with a variety of different measures of typicality (e.g., [122]).

The relationship between typicality and recognizability accounts well for the psychological findings suggesting the enhanced recognizability of facial caricatures relative to veridical faces. In general, artists draw caricatures in a way that exaggerates or enhances facial features that are “unusual” for the person. As such, the lips of Mick Jagger become thicker than they are, and the eyes of Prince Charles are even closer together than they actually are. Despite the fact that caricatures are grotesque distortions of a face, they are often recognized more accurately and efficiently than actual images of the faces [9, 80, 102]. Computer generated caricatures likewise operate by comparing a face to the “average face”, and then by exaggerating facial dimensions that deviate from the average [18]. The enhanced recognizability of caricatures by comparison to veridical faces may be due to the fact that the exaggeration of unusual features in these faces makes the person less confusable with other faces, and somehow or other “more like themselves”.

It is worth noting that face typicality correlates well with other ratings of faces, including

perceived familiarity (i.e., “This person looks like someone I know”) [123] and surprisingly, facial attractiveness [68]. The odd, but replicable, finding that typical faces are attractive indicates that faces similar to the “prototype”, which lack distinctive features, are aesthetically preferred to more memorable, albeit less attractive, faces.

Metaphorically, the same face space analogy is applicable to understanding the effects of face typicality on recognition. For automatic face recognition systems, a number of findings about the relationship between typicality and recognizability may be useful. For example, we would expect that the recognizability of individual faces should be predicted by the density of faces in the neighboring “face space”. We might also expect that the face space should be most dense in the center near the average. The space should become progressively less dense as we move away from the average. If a computationally-based face space approximates the similarity space humans employ for face processing, we might expect that “typical” faces would be near the center of the space and that unusual or distinctive faces be far from the center. It follows, therefore, that computational models of face recognition will not perform equally well for all faces. These systems should, like humans, make more errors on typical faces than on unusual faces.

### 0.2.2 Sex

The assumptions we make about the density of face space being highest near the center of the space and progressively less dense as we move away from the average face are somewhat simplistic. A more likely version of human face space is one in which there are a number of local averages at the centers of clusters of faces that belong to different natural categories, e.g., sex, race, age. In fact, when we consider that faces can be grouped into natural categories based on their sex, race/ethnicity, and possibly age, there is little reason to expect that “typical faces” are close to the grand mean of the space. Rather, the average face at the center of a space that is based on a set of diverse faces is likely to be androgenous and unrecognizable as a member of a racial or ethnic group. The dynamics of face recognition in this space are rather different than those discussed previously, though we can expect that within any given face category, many of the same principles of typicality and confusability should apply. This brings us to the question of how natural categories of faces such as sex, race, and age, might affect human face recognition accuracy.

We begin with the sex of a face, which is arguably the largest source of variance in facial appearance. Indeed, faces are readily identifiable as male or female by both humans and machines (cf., [1, 27, 41, 48, 89]). Computationally-based contrasts between male and female faces are easily derived with analyses like PCA. Applied to the three-dimensional laser scan data obtained from a large number of faces [89], this contrast is illustrated in Figure 1. Here we see that the information specifying the sex of a face can be localized to a few eigenvectors. In Figure 1, we show the contrast in face shape that can be achieved by varying the weights of two individual eigenvectors (the first and sixth) in combination with the average head. The contrasts clearly change the gender appearance of the face in a perceptible way.

How does the sex of a face affect face recognition accuracy for humans? A number of studies, which posed this question in the sixties and seventies found effects of the sex of a face on recognition accuracy that were minimal and inconsistent (see review [105]). Shepherd (1981) concluded that the most consistent finding in these studies was a small recognition advantage for female faces. A second intriguing finding was an interaction between the sex of the subject and the sex of the face. This interaction generally indicated that women are particularly accurate at recognizing female faces. There has been little or no work since these early studies that alters their main conclusions.




Figure 1: Illustration of the computationally-derived information about the sex of a face from three dimensional head models from laser scans [89]. The first row of the figure shows the average face plus the first eigenvector (left) and minus the second eigenvector (right). The second row of the figure shows an analogous display for the sixth eigenvector. The coordinates of a face on both the first and sixth eigenvectors predict the sex of the face.

More recently the focus of the relationship between the sex of the face and recognition has been to consider whether the identity-specific and sexually dimorphic information in faces are independent. As noted previously, in order to recognize a face, we must encode the information that makes it unique or different from all other faces in the world. To categorize a face, for example by sex, we must encode information that the face share with an entire group of faces. To know that a face is male, we must find features common to male faces, but unlikely to be found in female faces.

According to Bruce and Young’s (1986) classic model of face recognition, the sex determining information about a face is accessed independently of the information about identity [25]. Computational assessments of the information necessary for sex classification of faces versus identification support the claim that there is reasonably good separability of these two kinds of information in faces [84].

A recent study of “speeded classification” of faces by sex and familiarity, however, is at odds with the conclusion of independent processing of these kinds of information in faces [44]. In this kind of study, subjects are asked to classify faces by sex (the target feature), for example, as quickly as possible. In some cases, other “distractor” features of the faces in the sequence change during the task (for example, facial expression or familiarity). If the change in the other irrelevant feature interferes with the processing of the target feature, one can conclude that the processing of the target and distractor features is not independent. In a speeded classification task of this sort, interference between familiarity and the sex of faces was found when subjects classified faces either by sex or by familiarity [44]. Specifically, irrelevant variations in familiarity slowed sex classifications and vice versa. They concluded that humans are unable to attend selectively to either identity or sex information in faces.

However, a more direct method for testing the co-dependence of sex and identity information in faces was used by Wild and colleagues [125]. They compared recognition and sex classification of children’s and adults’ faces. The rationale behind this method is based



Figure 2: Computationally-derived gender information in children’s faces. Left is the male prototype, made by morphing boys together (see text for details), and right is the female prototype.

on the assumption that although children’s faces have reliable information for determining their sex, this information is less salient and more difficult to access than the information that specifies the sex of adult faces. In this study, the authors began by trying to establish a more factual basis for the above assumption. Using simple morphing methods, they created a prototype “boy face” and a prototype “girl face” (see Figure 2). For example, the construction of a prototype boy face was done by morphing pairs of boy faces together, and then morphing pairs of the morphs together, etc., until the resulting morph stopped changing. The prototypes indicate that the faces of boys and girls should be physically distinguishable, but that the differences are indeed subtle and difficult to describe.[125].

Next, Wild and colleagues asked children and adults to classify the faces of children and adults by sex. After the sex classification task, the participants were asked to recognize the faces they had just classified. The results indicated that although children’s faces were more difficult to classify by sex, they were recognized as accurately as adult faces.

In short, one can conclude that the quality of categorical information in faces is not related to the quality of the identity information. This finding is consistent with principal components analysis models of face recognition and categorization, which localize categorical information about faces in eigenvectors with large eigenvalues and identity information in eigenvectors with smaller eigenvalues [84]. Combined, the computational and human results re-emphasize the importance of the local structure of face space in predicting recognizability.

For computational models, the best current data come from recent findings in the Face Recognition Vendor Test 2002 [95]. This test was conducted by DARPA on the performance of ten commercial, automated, face recognition systems. The results of this analysis indicate a quite consistent advantage for recognizing male faces over female faces. The results showed that the performance advantage for male faces declined with face age. Although highly consistent, it is difficult to pinpoint the source of this advantage to idiosyncratic properties of the set of faces or to the algorithms tested (for which implementation details are proprietary). More details about the algorithms and a verification of the finding with a different set of faces is needed to be certain of the advantage. The properties of both the set of faces and the types of algorithms employed may play a role in performance, but this may need to be evaluated on a case-by-case basis.

In summary, for humans there are only minimal differences in accuracy recognizing male

and female faces. For machines, there can be quite large differences, though at present it is unclear if these are differences that generalize across different sets of faces and algorithms. There is some evidence from both human and machine studies that the information specifying the sex of a face can be effectively isolated from the information useful for recognizing the face.

### 0.2.3 Ethnicity

Another set of natural face categories can be found in the racial or ethnic differences between human faces. Large scale structural and reflectance differences exist between faces of different races. Psychologically, the categorical dimension of race differs from sex in that humans tend to have more experience with one race of faces, generally, their own race, than with other races. By contrast, most people have roughly equal amounts of contact and experience with males and females.

How does face race affect face recognition accuracy? Malpass (1969) was the first to demonstrate the “other-race effect” empirically [78]. He showed that human subjects recognize faces of their own race more accurately than faces of other races. This result has remained remarkably consistent over the years, with replications that have tested the effect with subjects and faces from a variety of races. In general, no effects of the race of the subject or the race of the face, *per se*, are found. The other-race effect consists solely of an interaction between the race of the face and the race of the subject.

The primary hypothesis for the cause of this effect is based on the differential experience individuals have with faces of their own race versus faces of other races. The so-called “contact” hypothesis predicts a relationship between the amount of experience we have with other-race faces and the size of the other-race effect. Several studies have tested this hypothesis, defining contact variously from simple questionnaires assessing previous exposure to members of other races [78] to the experience of living in an integrated neighborhood [40]. As noted previously [72], these studies have yielded inconsistent results, with some finding support for the contact hypothesis [30, 36, 33, 40, 104] and other studies failing to find support for this hypothesis [19, 70, 78, 83]. One reason for the lack of consistency among these studies might be linked to the diversity of the methods employed, and consequently, to the kinds of experience each may be measuring [42].

In reviewing evidence for the contact hypothesis many years ago, Shepherd [105] noted that among the few studies testing children and/or defining “contact” with other-race faces developmentally [36, 40], more consistent evidence for the contact hypothesis is found. These developmental studies examined children’s contact with other races in form of integrated versus segregated schools and/or neighborhoods [40, 36]. Shepherd suggests that this kind of early contact may be critical for developing an other-race effect for face recognition [105].

Complementing these studies, the developmental course of the other-race effect was tested further with Caucasian participants between the ages of 6 and 20 years old on a memory task for Caucasian and Asian faces [31]. In this study, the youngest participants, 6 year olds, recognized faces of both races equally well. By 10 years of age, however, there was a recognition accuracy advantage for Caucasian faces, which became successively larger for the older participants. Combined, these studies suggest the possibility that not all “contact” is equally effective in reducing/preventing an other-race effect. Contact early in life may be related to the magnitude of the other-race effect, whereas contact later on appears to be less consistently related to recognition skills for other-race faces.

The other-race effect for humans, therefore, represents an interaction between the properties of different categories of faces as well as the experiential history of the subject. In

particular, the age at which initial contact occurs may be particularly important in establishing the other-race effect.

At a glance, it is not obvious that this effect is relevant for face recognition algorithms. However, most recognition algorithms rely on computational and statistical descriptive tools, like PCA. These analyses are, by definition, sensitive to the statistical structure of the training set of faces. The basis vectors that define the face space are extracted from sets of faces that may vary in how representative they are of any particular population (e.g., the population for which one hopes to optimize a working automatic face recognition system). When face sets contain more faces of one race than of other races, it is reasonable to expect recognition accuracy differences for faces from the “majority” and “minority” race.

Furl et al. (2002) examined 13 face recognition algorithms for the presence of an other-race effect [42]. In all cases, the training set of faces was biased for the inclusion of Caucasian versus Asian faces. The algorithms fell into one of three categories in terms of their computational strategies and the extent to which a training set of faces could influence their performance. Two control algorithms based on image similarity comparisons (i.e., without low dimensional representations or complex distance measures), were not sensitive to the statistical structure of the training set. These algorithms showed no consistent difference for Asian versus Caucasian faces. Eight algorithms were based on PCA and varied in the distance measures used to assess matches. These algorithms performed best on the *minority* face race! In other words, they showed an “own-race effect”. The final three algorithms used a combination of PCA with a preselected biased training set and a second analysis that warped the space to optimize the separability of faces within the space (e.g., Fischer discriminant analysis, FDA). These algorithms performed most like humans, showing an other-race effect that favored recognition of the Caucasian faces.

Furl et al. (2002) speculated on the difference between the other-race advantage showed by the pure PCA models and the own-race advantage showed by the algorithms that mixed PCA with a space-warping algorithm like FDA. The authors suggest that the former are likely to perform better on the minority race faces due to the fact they these faces are “distinct” by comparison to the majority of faces in the face space. In other words, the minority race faces that were included in the PCA have representations in the sparsest part of the space and so are less likely to be confused with the minority distractor faces. By contrast, the combined models have the effect of warping the face space to provide the best and most distinctive representations for the majority faces. Thus, the learning of new minority race faces is limited to using a rather constrained and impoverished representation of these faces.

In summary, the interaction between experience and face race can affect performance for humans, and under more complicated metrics, for machines as well.

### 0.3 Viewing Constraints

In addition to the properties of faces themselves, nearly all components of the viewing parameters that have been tested with human subjects have been shown to affect face recognition accuracy. Most prominent are the effects of changes in viewpoint/pose and changes in the direction and intensity of the illuminant. These factors are also highly relevant for face recognition algorithms. At present, it would not be an exaggeration to say that viewing conditions constitute the single largest challenge to putting algorithms into the field. A common theme throughout this discussion is that as humans become familiar with faces, the apparent and strong effects of viewing parameters are lost or at least minimized. Common sense tells us that when we know someone well, it is possible to recognize them

under the poorest and least optimal of viewing conditions. We will return to the issue of familiarity after summarizing findings of viewing parameters.

### 0.3.1 Pose

To recognize a face from a novel view, we must be able to encode something unique about the face that distinguishes it from all other faces in the world *and* further must be able to access this unique information from the novel view. Studying the representations and processes that humans use to accomplish this task is difficult due to the complexity of the visual information observers experience in viewing faces from different viewpoints and due to the multitude of ways that such information can be encoded and represented. In general, humans can recognize faces from different views. However, they do this accurately only when the faces are well-known to them.

The nature of human representations of faces and objects is a long-standing issue that is still actively under debate. Two rather divergent theories of these representations have been put forth and tested. Structure-based theories suggest that the visual system constructs a three-dimensional representation of faces and objects from the two-dimensional retinal images on the left and right retinas [79]. Accordingly, objects can be recognized in a view-independent manner by their components (i.e., volumes or parts [11]).

The second type of theory is based on the direct analysis of images. These image-based theories assume a view-dependent representation of faces and objects. As such, experience with multiple views of an object or face is essential to be able to recognize the object or face from a novel viewpoint [97].

Both structural and image-based theories have been supported by psychophysical studies [110, 10]. It has been suggested, therefore, that the most viable model of object/face recognition should incorporate the most appealing aspects of both accounts [110].

To be able to recognize a face from a new view, we need to extract invariant information that can be used subsequently to identify the face. Because we know the structure of familiar faces, we can automatically encode them and therefore it is unlikely that we will find an effect of pose for recognition of familiar faces. Indeed, subjects do not benefit from seeing familiar faces in the same view at learning and at test [20, 117]. By contrast, subjects find it easier to recognize unfamiliar faces if the viewpoint of the face is the same at learning and at test. Thus, it seems that unfamiliar faces are sensitive to the change in viewpoint between the view of the face that was learned and the view that was tested.

Notwithstanding, it is also likely that some views of an unfamiliar face, those seen either at learning or at test, will make the extraction of the invariant information important for face recognition easier. Therefore, recognition of unfamiliar faces is likely to be influenced by the change in viewpoint between learning and test and by the information provided by the viewpoint of the face.

*Rotation effects.* In general, human observers' recognition performance is impaired when faces are rotated in depth between learning and test. When viewpoint change between learning and test is less than 30 degrees, performance declines in a roughly linear fashion as a function of the rotation offset. For changes of more than 30 degrees, the performance cost reaches a plateau. On the whole, a 45 degree rotation seems to impair performance significantly [5, 20, 57, 61, 76, 85]. Recent fMRI studies [4] seem to confirm this effect by showing that activity in the face fusiform area is sensitive to change in viewpoints. It is worth noting however that the fusiform responds to faces across viewpoints, as do other face sensitive brain regions like the superior temporal sulcus [3]. It has been possible as well to differentiate sub-areas of fusiform cortex and the lateral occipital complex that show varying degrees of invariance to changes in viewpoint and illumination versus changes in

the size and retinal position of objects and faces [50].

In human recognition studies, subjects could recognize unfamiliar faces as long as the change in viewpoint was smaller than 30 degrees [118, 119]. Interestingly, image-based simulations of face recognition such as PCA replicate these results. The strong correlation between simulations and human performance suggests that the deleterious effect of rotation on human performance is due to the decline of image similarity that occurs with rotation.

Contrary to the simulation results, however, human performance is still better than chance even after a 90 degree rotation. This could be due to the abstraction of a 3D model of the face (i.e., like one suggested by structural theories [79, 85] or by the abstraction of invariant features from the faces, e.g., moles, and other surface markings [120], or by both mechanisms [110]).

Even if the theoretical interpretation is still open to discussion, the empirical effect is clear. Unfamiliar faces are more difficult to recognize after a change of pose. Within limits, the larger the change the larger the decrement in performance. In addition, more distinctive faces are less sensitive to this effect; more typical faces are more sensitive to it [87].

*Recognition from the three-quarter pose.* In early work, several studies suggested that the three-quarter view of faces was “special” because it seemed to be recognized better than other viewpoints [5, 61, 20]. This finding was replicated several times [76, 21] and has been found also for young infants [39]. A neural network simulation of face recognition with an image-based code input to a radial basis function network spontaneously yields a 3/4 view advantage [118]. This suggests that the three-quarter view advantage can be attributed to the “more informative” nature of the view for recognition. In other words, there may be more unique information in face images taken from the three-quarter view than in face images taken from other views. This information-based explanation may also explain why the three-quarter view is favored for western portraits [6] as well as for cartoons [92].

Despite the fact that the three-quarter view advantage is replicable, the size of this effect is rather small and there are cases where the effect is not actually found [75]. In recent review of 14 studies testing for the three-quarter view advantage, it was reported that six of these failed to detect the effect. The authors concluded that the three-quarter advantage was not due the three-quarter view being more informative *per se* (i.e., being a “canonical view” or special view [91]). Rather they suggest that the three-quarter view is closer to most other views (i.e., profiles or full faces) and that it is therefore easier to transfer information to the three-quarter views (because the rotation angle to this view is smaller on average than to other views).

One counter-argument against this explanation is that the three-quarter view advantage has been found even when the view of the test stimulus matches the view that was learned [87]. Thus, learning and testing with three-quarter view faces produces better recognition performance than learning and testing with frontal or full-profile views.

In summary, there is good evidence that both the viewpoint, *per se*, and the change in viewpoint between learning and test impact human recognition performance.

### 0.3.2 Illumination

The effects of illumination variations on face and object recognition have been studied extensively and mirror the findings for viewpoint. The effects of illumination on the appearance of an object in an image can be every bit as dramatic as the effects of large changes in viewpoint. Illumination parameters can change the overall magnitude of light intensity reflected back from an object, as well as the pattern of shading and shadows visible in an image [110]. Both shading and shadows may provide cues about the three-dimensional

shape of a face. Indeed, varying the direction of illumination can result in larger image differences than varying the identity [2] or the viewpoint [110] of a face.

Although face and object recognition have been shown to be sensitive to illumination variations [15, 17, 55, 110, 114], considerable illumination invariance can be achieved under some novel illumination conditions [82, 16, 17]. In combination, these findings suggest that there are processes that attempt to discount the effects of illumination when recognizing faces, but that they are not perfect. Again, as for viewpoint, the effects of illumination seem to depend on the magnitude of change in illumination conditions between learning and test. The effects of illumination have been tested systematically only for unfamiliar faces. One might assume more robust processes for recognizing familiar faces over changes in illumination.

There is further evidence to suggest that humans prefer illumination that comes from above the face. At its extreme, this preference impacts the recognition of faces quite dramatically. For example, it is notoriously difficult to recognize faces in the photographic negative [43]. Photographic negatives of faces look like faces that are illuminated from below. It is difficult to recognize people, even those we know well, from photographic negatives. These data suggest that human recognition might be based on internal neural representations that are more like images than like structural descriptions of the face. Assuming the latter, recognition of photographic negatives should not be as difficult as it is.

The effects of illumination have been considered also at the level of neural encoding. These results are consistent with the findings from psychophysical studies. For example, although it is believed that illumination information can be extracted by lower visual areas, the results from a recent functional magnetic resonance imaging study suggest that the sensitivity to the direction of illumination is retained even in higher levels of the visual hierarchy [50]. In that study, the technique of functional magnetic resonance imaging adaptation was employed to test “invariance.” In general, the procedure works by assuming that the neural response adapts to repeated presentations of stimuli that are perceived identically. If an area of cortex is processing faces in a way that is invariant to illumination, then the neural response to a face varying in illumination, presented repeatedly, will continue to adapt. If the cortex response is sensitive to illumination, then adaptation will not occur with multiple presentations. Using this technique, Grill-Spector and colleagues found two subdivisions in the lateral occipital complex, one that showed recovery from adaptation under all transformations (size, position, illumination, and viewpoint) and another that showed recovery only for illumination and viewpoint [50]. This indicates the existence of a complex hierarchy of visual processes that ultimately contribute to human face recognition.

## **0.4 Moving Faces**

Recognition memory for moving faces is a new and growing area of research [88, 98, 128]. From an ecological perspective, the use of moving faces as stimuli in face recognition studies is useful for approximating the way in which people typically encounter faces in the real world – as dynamic objects. Moving faces provide the viewer with a wealth of social information [3] and, potentially, with a unique source of identity information.

### **0.4.1 Social Signals and Motion**

We begin with a brief taxonomy of facial movements. At the highest level of the taxonomy, facial movements are either rigid or non-rigid. Rigid motions of the head include nodding, shaking, tilting, and rotating about the vertical axis. It is worth noting that all of these movements change the view of the face available to a stationary observer. It is also important to note that each of these movements can convey a social signal or connotation. Nodding

the head can convey agreement, shaking back and forth can convey disagreement, and turns of the head, either toward or away from another person, can be used to initiate or break off a communication.

Non-rigid movements of the face occur during facial speech, facial expressions, and eye-gaze changes. Again, these movements produce highly variable images of the person that can distort many of the identifying “features” of faces, like the relative distances of the features. The difference between the image of a person smiling and one of the person with a surprised expression can be quite strong. The relative position of the eyebrows, for example, with respect to the mouth and the other features, changes radically between these images. Facial speech, expression, and eye gaze movements can also convey a social message. When a person speaks, they rarely do so with a static or neutral expression. Eye gaze changes can signal boredom or interest and facial expression can yield an nearly limitless amount of information about a person’s internal state.

Thus, we must bear in mind that facial movements may be challenging to the perceptual system. This is because they alter the nature of the “invariants” available for recognizing a face and because they must be monitored and interpreted constantly for the social information they convey.

#### **0.4.2 Recognizing Moving Faces**

How do facial movements affect face recognition accuracy? To begin to answer this question, we must first ask, “In what ways might motion help face recognition?” and “In what ways might motion make face recognition more difficult?” For the former, there are two theories about how motion might improve recognition [88, 98]. The first theory posits that motion might provide facial identity signatures in the idiosyncratic patterns of dynamic information they undergo. If repeated regularly (e.g., as with characteristic gestures or facial expressions), these idiosyncratic patterns of movement might provide a reliable cue for the identity of the face. A second theory posits that motion could help face recognition by providing additional structure from motion information about the face. This can enhance the quality of the perceptual representation of a face [88, 98].

For the dynamic identity signature theory, the role of characteristic motions for learning to identify faces has been studied recently using animated synthetic three-dimensional head models [56, 59]. These studies have focused on the “learnability” of a dynamic signature when it is the most, or only, reliable cue to identity. For example, Hill and Johnston projected facial animations generated by human actors onto a computer-generated average head [56] (see Figure 3 for an example of their stimulus). Their participants learned to discriminate among four individuals based solely on the facial motion information. It is worth noting that non-rigid motion was less useful than rigid motion in this identity learning.

In another similar study, Knappmeyer and her colleagues trained participants to discriminate two synthetic faces that were animated with different characteristic facial motions [59]. When later viewing morphs between the two head models, the subjects’ identity judgments about the intermediate morphed heads were biased by the animated motion information they learned to associate with the faces originally. Both studies support the notion that inherently dynamic information about face movements can support recognition and can form a part of our representation of the identity of an individual.

More relevant for everyday face recognition, dynamic identity signatures are likely to be most helpful when faces are familiar. This is because it may take time, for any given individual, to learn the difference between characteristic movements and movements that are generated randomly. Indeed, the current literature suggests that face familiarity mediates the usefulness of facial motion as a recognition cue. The beneficial effects of facial motion




Figure 3: Illustration of stimulus creation from Hill and Johnson’s (2001) study. The motions of a human actor are projected onto the synthetic head. Subjects learn to identify the head by the motions.

are more robust and easier to demonstrate in recognition tasks with familiar/famous faces than with unfamiliar faces. For example, participants can recognize the faces of well-known politicians and celebrities more accurately from videotaped images than from static images. This finding is especially salient when the faces are presented in sub-optimal viewing formats (e.g., blurred, inverted, or pixilated displays; [60, 64, 66, 67]). Thus it seems that motion becomes more important as a cue to identity when the viewing conditions are sub-optimal. This probably occurs because the static features are less reliable cues to identity in these conditions, and so subjects are more likely to require additional information available in the identity signature to successfully recognize the person.

The second hypothesis about how motion might benefit face recognition posits that structure-from-motion processes can contribute to the quality of the face representation. This should apply most clearly in studying how motion affects our ability to learn new faces, (i.e. to create new face representations). In the case of newly-learned or unfamiliar faces, the data are not clear as to whether motion improves face recognition. Some studies report a motion benefit [65, 96, 98], whereas other studies find no benefit of motion [23, 24, 35, 54]. A closer inspection of these results suggest that the benefits of facial motion for unfamiliar face recognition tasks may be tied to the specific parameters of the learning and test conditions used across these different studies. Differences in the type of recognition tasks implemented and variations in the kinds of stimuli employed likely account in large part for the disparity in the results.

In a recent study, we attempted to control some of the confounds in previous experiments in order to assess the role of motion in face recognition [107]. In particular, we wanted to control for the extra views we see of a face when it is in motion. As noted, motion almost always provides subjects with extra views of a face in addition to whatever benefit may come from the motion, *per se*. In a single experiment, we compared recognition performance for subjects’ who learned faces in four conditions. One set of subjects learned each face from a

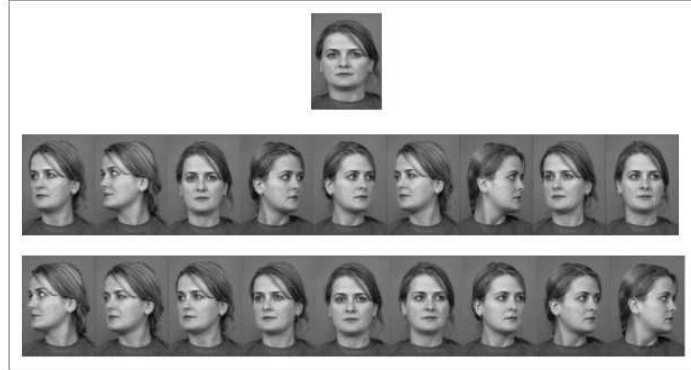


Figure 4: Stimuli for a study comparing static and dynamic learning conditions controlling for the extra views seen in the dynamic condition [107]. Subjects learned a single static image (row 1), or a random sequence of nine views (row 2), or an ordered sequence of nine views (row 3), or a video clip of the head rotating systematically through nine views.

static frontal image. A second set of subjects learned each face from nine images taken from different viewpoints and presented in an ordered sequence. A third set of subjects learned from the same nine images, but presented in random order to eliminate any implied motion signal. Finally, a fourth set of subjects learned faces from a video clip of the face rotating systematically through the nine viewpoints. This video looked like a person searching through a crowded room for someone. An example of the stimuli appears in Figure 4. Subjects were tested using the same kind of stimulus they learned.

The results of the experiment revealed absolutely no difference in recognition performance as a function of the learning condition (see Figure 5). This indicates that motion does not seem to benefit the learning of unfamiliar faces. To be sure that this effect was not based on the test conditions, which were matched between the learn conditions in the first experiment, we repeated the study. This time, however, we tested with a single frontal image in all conditions. Again, we found the same result. Motion provided no additional benefit for recognition.

## 0.5 Motion and Familiarity

The disparity of findings concerning the role of motion in face recognition hinges primarily around the familiarity of faces. Most studies have employed the faces of celebrities as “familiar” faces or the faces of people already well-known to the subjects (e.g., professors at their university). The major shortcoming of this method is that the researcher has no control over the learning conditions any given subject undergoes along the path to becoming familiar with a face. How many previous exposures have they had to the celebrity? Which views have they seen? Have the views occurred over years (e.g., Harrison Ford) or over months (e.g., some newly popular celebrity)?

One recent experiment employs a technique by which subjects gain familiarity with a




Figure 5: The results of a study comparing static and dynamic learning conditions, controlling for the extra views seen in the dynamic condition [107]. No significant differences were found among the learning conditions indicating that motion does not provide any benefit for face learning.

face under controlled conditions during the course of the experiment [99]. In that study, subjects learned people from a variable number of exposures of the same stimulus, either a close-up of the face (moving or static) or a distance view of the person walking. They were tested with the “other” stimulus. Thus people learned from facial close-ups were tested with distance videos and vice versa. The findings indicated, first, that pure repetition of the same stimulus, was remarkably effective in improving recognition over these substantial changes in viewing conditions.

Second, recognition was more accurate when participants viewed faces in motion *both* at learning and test than when participants viewed moving images only at learning or test. Again, this held even when the two stimuli were presented in different viewing formats (i.e., high-quality facial images versus whole-body surveillance-like videos). This finding, in light of other studies that have likewise reported inconsistent recognition benefits with moving faces (e.g., [24]), may suggest that unfamiliar face recognition tasks benefit from facial motion primarily in learning-test conditions that involve “motion to motion” transfers.

Tentative converging support for this motion match hypothesis comes from a recently advanced neural theory of face processing proposed by Haxby and colleagues [53]. In this theory, the variant and invariant aspects of faces are processed independently in the brain ([53]). In Haxby et al.’s model, the moving aspects of faces are processed in the superior temporal sulcus region (STS) of the dorsal visual stream, whereas the static aspects of faces (i.e., facial features) are processed in the fusiform face area of the ventral visual stream (FFA, [58]). The proposed separation of face processing between the two streams leaves open the possibility that recognition tasks that permit processing to remain in the same stream are more likely to support successful recognition performance. For now, however, it is clear that more research is needed to generate a coherent picture of how facial motion affects recognition.

## 0.6 Familiarity and Experience

We close this chapter by recalling our beginning claim that humans are the best available face recognition systems – superior to many, if not all, algorithms. It should be clear by now that this claim must be qualified in two ways. First, on the task of learning and remembering faces from a limited number of exposures, humans encounter the same kinds of difficulties that algorithms do, and likewise reflect these difficulties with impaired recognition performance. Second, although we generally assume that experience with faces we know well allows us to recognize them in sub-optimal viewing conditions (poor illumination, short glances, bad viewing angles), this problem has been little studied by psychologists. The reason for this is quite practical. In most cases, performance recognizing faces we know well is generally so good that few variables are likely to affect performance substantially.

Ultimately, we think that the issues surrounding how we *become* familiar with faces, and how the internal neural representation of faces changes with experience, are much neglected topics for psychologists studying human memory for faces. They are, however, the most valuable open questions for psychologists to study when they are interested in providing data for computational modelers of face recognition. Being able to recognize a friend, at a distance, in a dimly lit train station, from an odd angle, is currently a unique accomplishment of the human visual system. Certainly, algorithms can achieve this if enough is known about how we do it.

# Bibliography

- [1] H. Abdi, D. Valentine, B. Edelman, and A.J. O'Toole. More about the difference between men and women: evidence from linear neural networks and the principal-component approach. *Perception*, 24:539–562, 1995.
- [2] Y. Adini, Y. Moses and S. Ullman. *Face recognition: the problem of compensating for changes in illumination direction* (Report No. CS93-21). The Weizmann Institute of Science, 1995.
- [3] T. Allison, A. Puce and G. McCarthy. Social perception from visual cues : role of the STS region. *Trends in Cognitive Sciences*, 4:267–278, 2000.
- [4] T. J. Andrews, and M. P. Ewbank. Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *NeuroImage*, 1n press, 2005.
- [5] A. Baddeley, and M. Woodhead. Techniques for improving eyewitness identification skills. *Paper presented at the SSRC law and Psychology Conference*. Trinity College, Oxford, 1981.
- [6] A. Baddeley, M. Woodhead. Improving face recognition ability. In S. Llyood-Bostock, and B. Clifford (Eds.) *Evaluating witness evidence*. Chichester: Wiley, 1983.
- [7] J.C. Bartlett and J. Searcy. Inversion and configuration of faces. *Cognitive Psychology*, 25:281–316, 1993.
- [8] J.C. Bartlett and A. Fulton. Familiarity and recognition of faces in old age. *Memory and Cognition*, 19:229–238, 1991.
- [9] P.J. Benson and D.I. Perrett, 1991 Perception and recognition of photographic quality caricatures : implications for the recognition of natural images. *European Journal of Cognitive Psychology*, 3:105–135, 1993.
- [10] I., Biederman, and M. Bar, M. One-shift viewpoint invariance in matching novel objects. *Vision Research*, 39: 2885 – 2889, 1998
- [11] I. Biederman and P. Gerhardstein. Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance. *J. Exp. Psychol. Hum. Percept. Perform.*, 19:1162–1183, 1993.
- [12] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *SIG-GRAPH'99 Proceedings*, ACM: Computer Society Press, 187–194, 1999.
- [13] V. Blanz, A.J. O'Toole, T. Vetter, and H.A. Wild. On the other side of the mean: The perception of dissimilarity in human faces. *Perception*, 29: 885–891, 2000.

- [14] G.H. Bower and M.B. Karlin. Depth of processing pictures of faces and recognition memory. *Journal of Experimental Psychology*. 103:751-757, 1974.
- [15] W. L. Braje, D.J. Kersten, M.J. Tarr, and N.F. Troje. Illumination effects in face recognition. *Psychobiology*, 26:371–380, 1999.
- [16] W.L. Braje, G.E. Legge and D. Kersten. Invariant recognition of natural objects in the presences of shadows. *Perception*, 29:383–398, 2000.
- [17] W.L. Braje. Illumination encoding in face recognition: effect of position shift. *Journal of Vision*, 3:161–170, 2003.
- [18] S.E. Brennan. The caricature generator *Leonardo*, 18: 170–178, 1985.
- [19] J.C. Brigham and P. Barkowitz. Do "They all look alike?" The effects of race, sex, experience and attitudes on the ability to recognize faces. *Journal of Applied Social Psychology*, 8:306–318, 1978.
- [20] V. Bruce. Changing faces: visual and non-visual coding processes in face recognition. *British Journal of Psychology*, 73: 105–116, 1982.
- [21] V. Bruce, T. Valentine, A. Baddeley. The basis of the 3/4 view advantage in face recognition. *Applied Cognitive Psychology*, 1: 109–120, 1987
- [22] R. Brunelli and T. Poggio. Caricatural effects in automated face perception. *Biological Cybernetics*, 69:235–241, 1993.
- [23] V. Bruce, Z. Henderson, K. Greenwood, P.J.B. Hancock, A.M. Burton and P. Miller. Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, 5:339–360, 1999.
- [24] V. Bruce, Z. Henderson, C. Newman and A.M. Burton. Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied*, 7:207–218, 2001.
- [25] V. Bruce and A.W. Young. Understanding face recognition. *British Journal of Psychology*, 77(3):305–327, 1986.
- [26] D.M. Burt and D.I. Perrett. Perception of age in adult Caucasian male faces: Computer graphic manipulation of shape and colour information. In *Proceedings of the Royal Society London B*, 259:137–143, 1995.
- [27] A.M. Burton, V. Bruce, and N. Dench. What's the difference between men and women: Evidence from facial measurement. *Perception*, 22:(2)153–176, 1993.
- [28] A.M. Burton, V. Bruce, and P.J.B. Hancock. From pixels to people: a model of familiar face recognition. *Cognitive Science*, 23:1–31, 1999.
- [29] A.M. Burton, S. Wilson, M. Cowan, and V. Bruce. Face recognition in poor-quality video. *Psychological Science*, 10:243-248, 1999.
- [30] A.W. Carroo. Other-race face recognition: A comparison of Black American and African subjects. *Perceptual and Motor Skills*, 62:135–138, 1986.

- [31] J.E. Chance, A.L. Turner and A.G. Goldstein. Development of differential recognition for own- and other-race faces. *Journal of Psychology*, 112:29–37, 1982.
- [32] Y. Cheng, A.J. O’Toole & H. Abdi. Sex classification of adults’ and children’s faces: Computational investigations of subcategorical feature encoding. *Cognitive Science*, 25:819–838., 2001.
- [33] P. Chiroro and T. Valentine. An investigation of the contact hypothesis of the own-race bias in face recognition. *Quarterly Journal of Experimental Psychology, A, Human Experimental Psychology*, 48A:879–894, 1995.
- [34] I. Craw and P. Cameron. Parameterizing images for recognition and reconstruction. In P. mowforth, editor *Proceedings of the British Machine Vision Conference* Springer, London, 1991.
- [35] F. Christie and V. Bruce. The role of dynamic information in the recognition of unfamiliar faces. *Memory & Cognition*, 26:780–790, 1998.
- [36] J.F. Cross, J. Cross and J. Daly. Sex, race, age, and beauty as factors in recognition of faces. *Perception & Psychophysics*, 10:393–396, 1971.
- [37] S. Edelman and H.H. Bühlhoff. Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, 32(12):2385–2400, 1992.
- [38] P.J. Ekman and W.V. Friesen. *The Facial Action Coding System: A Technique for the Measurement of Facial Movement*. Consulting Psychology Press, San Francisco, 1978.
- [39] J. Fagan. The origin of facial pattern recognition. In M. Bornstein and W. Keesen (Eds.) *Psychological development from infancy: Image to intention*. 1979
- [40] S. Feinman and D.R. Entwisle. Children’s ability to recognize other children’s faces. *Child Development*, 47(2):506–510, 1976.
- [41] M. Fleming and G.W. Cottrell G W. Categorization of faces using unsupervised feature extraction. *Proceedings of IJCNN-90* Vol 2, Ann Arbor, MI:IEEE Neural Networks Council, 65–70, 1990.
- [42] D.S. Furl, P. J. Phillips, and A.J. O’Toole. Face recognition algorithms as models of the other-race effect. *Cognitive Science*, 96:1–19, 2002.
- [43] R.E. Galper and J. Hochberg. Recognition memory for photographs of faces *American Journal of Psychology*, 84: 351–354, 1971.
- [44] T. Ganel and Y. Goshen-Gottstein. Perceptual integrality of sex and identity of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 28:854–867, 2002.
- [45] I. Gauthier, M.J. Tarr, A.W. Anderson, P. Skudlarski, and J.C. Gore. Activation of the middle fusiform face area increases with expertise recognizing novel objects. *Nature Neuroscience*, 2:568–573, 1999.
- [46] M.S. Gazzaniga(Ed.), 1995 *The Cognitive Neurosciences*., Cambridge, MIT.

- [47] G. Givens, J.R. Beveridge B.A. Draper, P. Grother and P/J. Phillips. How features of the human face affect recognition: A statistical comparison of three face recognition algorithms. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2004.
- [48] M. Gray, D.T. Lawrence, B.A. Golomb, T.J. Sejnowski. A perceptron reveals the face of sex. *Neural Computation*, 7:1160–1164, 1995.
- [49] D.M. Green and J.A. Swets. *Signal detection Theory and Psychophysics*. Wiley, New York, 1966
- [50] K. Grill-Spector, T. Kushnir, S. Edelman, G. Avidan, Y. Itzhak, and R. Malach. Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, 24: 187–203, 1999.
- [51] P.J.B. Hancock, V. Bruce, and A.M. Burton. Recognition of unfamiliar faces. *Trends in Cognitive Sciences*, 4:(9)263–266, 1991.
- [52] J.V. Haxby, M.I. Gobbini, M.L. Furey, A. Ishai, J.L. Shouten and J.L.Pietrini. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293:2425–2430, 2001.
- [53] J.V. Haxby, E.A. Hoffman, and M.I. Gobbini. The distributed human neural system for face perception. *Trends in Cognitive Sciences*, 20(6):223–233, 2000.
- [54] Z. Henderson, V. Bruce and A.M. Burton. Matching the faces of robbers captured on video. *Applied Cognitive Psychology*, 15:445–464, 2001.
- [55] H. Hill and V. Bruce. Effects of lighting on the perception of facial surface. *J. Exp. Psychol. : Hum. Percept. Perform.*, 4:(9)263–266, 1991.
- [56] H. Hill and A. Johnston. Categorizing sex and identity from the biological motion of faces. *Current Biology*, 11:880–885, 2001.
- [57] H. Hill, P. Schyns, and S. Akamatsu. Information and viewpoint dependence in face recognition. *Cognition*, 62: 201–202, 1997.
- [58] N. Kanwisher, J. McDermott, and M. Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.*, 17: 4302–4311, 1997.
- [59] B. Knappmeyer, I.M. Thornton and H.H. Bulthoff. The use of facial motion and facial form during the processing of identity. *Vision Research*, 43:1921–1936, 2003.
- [60] B. Knight and A. Johnston. The role of movement in face recognition. *Visual Cognition*, 4:265–273, 1997.
- [61] F. L. Krouse. Effects of pose, pose change, and delay on face recognition performance *Journal of Applied Psychology*, 66: 201–222, 1981.
- [62] P.K. Kuhl, K.A. Williams and F. Lacerdo. Linguistic experience alters phonetic perception in infants by 6 months of age. *Science*, 255:606–608, 1992.

- [63] J. Kurucz and J. Feldmar. Prosopo-affective agnosia as a symptom of cerebral organic brain disease, *Journal of the American Geriatrics Society*, 27:91-95, 1979.
- [64] K. Lander and V. Bruce. Recognizing famous faces: exploring the benefits of facial motion. *Ecological Psychology*, 12:259–272, 2000.
- [65] K. Lander and V. Bruce. The role of motion in learning new faces. *Visual Cognition*, 10:897–912, 2003.
- [66] K. Lander, V. Bruce and H. Hill. Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces. *Applied Cognitive Psychology*, 15:101–116, 2001.
- [67] K. Lander, K. Christie and V. Bruce. The role of movement in the recognition of famous faces. *Memory & Cognition*, 27:974–985, 1999.
- [68] J.H. Langlois, L.A. Roggman and L. Mussleman. What is average and what is not average about attractive faces? *Psychological Science*, 5:214–220, 1994.
- [69] A. Lanitis, C.J. Taylor, and T.F. Cootes. Automatic interpretation and coding of face imaging using flexible models. *IEEE Transactions Pat. Anal. Mach. Intell.*, 19: 743, 1997.
- [70] P.J. Lavarkas, J.R. Buri and M.S. Mayzner. A perspective on the recognition of other-race faces. *Perception & Psychophysics*, 20:475–481, 1976.
- [71] D. Leopold, A.J. O’Toole, T. Vetter, and V. Blanz. Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4:89-94, 2001.
- [72] D. Levin. Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, 129:559–574, 2000.
- [73] L. Light, F. Kayra-Stuart, and S. Hollander. Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory*, 5: 212–228, 1979.
- [74] D.S. Lindsay, P.C. Jack, and M.A. Christian. Other-race face perception. *Journal of Applied Psychology*, 76:587–589, 1991.
- [75] C. H. Liu, and A. Chaudhuri. Reassessing the 3/4 view effect in face recognition. *Cognition*, 83: 31–48, 2002.
- [76] R. Loggie, A. Baddeley, and M. Woodhead. Face recognition, pose and ecological validity. *Applied Cognitive Psychology*, 1:53–69, 1987.
- [77] N.K. Logothetis, J. Pauls, H.H. Bülthoff, and T. Poggio. Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5:552–563, 1991.
- [78] R.S. Malpass and J. Kravitz. Recognition for faces of own and other race faces. *Journal of Personality and Social Psychology*, 13:330–334, 1969.
- [79] D. Marr. *Vision*. Freeman, San Francisco, 1982.

- [80] R. Mauro and M. Kubovy. Caricature and face recognition. *Memory & Cognition*, 20:433–440, 1992.
- [81] W.H. Merigan. P and M pathway specialization in the macaque. In A. Valberg and B.B. Lee, editors. *From Pigments to Perception*. Plenum, New York. pages 117–125, 1991.
- [82] Y. Moses, S. Edelman, and S. Ullman. Generalization to novel images in upright and inverted faces. *Perception*, 25:(4)443–461, 1996.
- [83] W. Ng and R.C.L. Lindsay. Cross-race facial recognition: Failure of the contact hypothesis. *Journal of Cross-Cultural Psychology*, 25:217–232, 1994.
- [84] A.J. O’Toole, H., Abdi, K.A. Deffenbacher, and D. Valentin. Low dimensional representation of faces in high dimensions of the space. *Journal of the Optical Society of America A*, 10, 405–410, 1993.
- [85] A. J. O’Toole, H. Bülthoff, N. Troje, and T. Vetter. Face recognition across large viewpoint changes. In *Proceedings of the International Workshop on Automatic Face and Gesture Recognition*. Zurich, 1995.
- [86] A.J. O’Toole, K.A. Deffenbacher, and D. Valentine. Structural aspects of face recognition and the other-race. *Memory & Cognition*, 22: 208–224, 1994.
- [87] A.J. O’Toole, S.E. Edelman, and H.H. Bülthoff. Stimulus-specific effects in face recognition over changes in viewpoint. *Vision Research*, 38: 2351–2363, 1998.
- [88] A.J. O’Toole, D. Roark, and H. Abdi. Recognition of moving faces: a psychological and neural perspective. *Trends in Cognitive Sciences*, 6:261–266, 2002.
- [89] A.J. O’Toole, T. Vetter, N.F. Troje, and H.H. Buelthoff. Sex classification is better with three-dimensional head structure than with image intensity information. *Perception*, 26:75–84, 1997.
- [90] A.J. O’Toole, T. Vetter, H. Volz, and E.M. Salter. Three-dimensional caricatures of human heads: distinctiveness and the perception of facial age. *Perception*, 26:719–732, 1997.
- [91] S. Palmer, E. Rosch, and P. Chase. Canonical perspective and the perception of objects. In J. Long, and A. D. Baddeley, (Eds.) *Attention and performance IX*. Hillsdale: Erlbaum, 1981.
- [92] D. N. Perkins. A definition of caricature and recognition. *Studies in the Anthropology of Visual Communication*, 2: 1–24. 1975.
- [93] D. Perrett, J. Hietanen, M. Oram, and P. Benson. Organization and function of cells responsive to faces in temporal cortex. *Phil. Trans. Roy. Soc. Lond. B Biol. Sci*, 335: 23–30, 1992.
- [94] P.J. Phillips, H.J. Moon, S.A. Rivzi and P.J. Rauss. The FERET Evaluation methodology for face-recognition algorithms. *Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1090–1104, 2003.

- [95] P.J. Phillips, P. Grother, R.J. Michaels, D.M. Blackburn, E. Tabassi and M. Bone. Face recognition vendor test 2002. *NISTIR*, 2003.
- [96] G.E. Pike, R.I. Kemp, N.A. Towell and K.C. Phillips. Recognizing moving faces: The relative contribution of motion and perspective view information. *Visual Cognition*, 4:409–437, 1997.
- [97] T. Poggio and S. Edelman. A network that learns to recognize 3D objects. *Nature*, 343:263–266, 1991.
- [98] D.A. Roark, S.E. Barrett, M.A. Spence, H. Abdi and A.J. O’Toole. Psychological and neural perspectives on the role of motion in face recognition. *Behavioral and Cognitive Neuroscience Reviews*, 2:15–46, 2003.
- [99] D.A. Roark, S.E. Barrett, H. Abdi and A.J. O’Toole. Learning the moves: the effect of facial motion and familiarity on recognition across large changes in viewing format. (under review)
- [100] D.A. Roark, H. Abdi and A.J. O’Toole. Human recognition of familiar and unfamiliar people in naturalistic videos. *Proceedings of the IEEE International Workshop on Analysis and Modeling of Faces and Gestures*, 36–43, 2003.
- [101] E. Rosch and C. Mervis. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7:573–605, 1975.
- [102] G. Rhodes. *Superportraits: Caricatures and Recognition*. Psychology Press, Hove, UK, 1997.
- [103] M. Riesenhuber and T. Poggio. Models of object recognition. *Nature Neuroscience Supplement*, 3:–1204, 2000.
- [104] J.W. Shepherd, J.B. Derogowski and H.D. Ellis. A cross-cultural study of recognition memory for faces. *International Journal of Psychology*, 9:205–212, 1974.
- [105] J. Shepherd. Social factors in face recognition. In (Eds. G. Davies, H. Ellis, & J. Shepherd) *Perceiving and remembering faces*, Academic Press: London, 55–78, 1981.
- [106] M. Spiridon and N. Kanwisher. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*, 35:1157, 1991.
- [107] S. Snow, G. Lannen, A.J. O’Toole and H. Abdi. Memory for moving faces: Effects of rigid and non-rigid motion. *Journal of Vision*, 2:7, Abstract 600, 2002.
- [108] J. W. Tanaka and M.J. Farah. Parts and wholes in face recognition. *Quarterly Journal of Psychology*, 46A:(2)225–245, 1993.
- [109] K. Tanaka. Neuronal mechanisms of object recognition. *Science*, 262:685–688, 1991.
- [110] M. Tarr, and H. Bülthoff. Image-based object recognition in man, monkey and machine. *Cognition*, 67: 1–20, 1998.
- [111] M.J. Tarr, D. Kersten and H.H. Bulthoff. Why the visual recognition system might encode the effects of illumination. *Vision Research*, 38:2259–2275, 1998.

- [112] P. Thompson. Margaret Thatcher: A new illusion. *Perception*, 9:483–484, 1980.
- [113] N.F. Troje and H.H. Bühlhoff. Face recognition under varying pose: The role of texture and shape. *Vision Research*, 36:1761–1771, 1996.
- [114] N.F. Troje and H.H. Bulthoff. How is bilateral symmetry of human faces used for recognition of novel views? *Vision Research*, 38:79–89, 1998.
- [115] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3:71–86, 1991.
- [116] S. Ullman. *High-Level Vision*, MIT Press, Cambridge, MA, 1996.
- [117] D. Valentin. *How come when you turn your head I still know who you are: evidence from computational simulations and human behavior*. Ph.D. Dissertation, The University of Texas at Dallas. 1996.
- [118] D. Valentin, H. Abdi, and B. Edelman. What represents a face: A computational approach for the integration of physiological and psychological data. *Perception*, 26: 1271–1288, 1997.
- [119] D. Valentin, H. Abdi, and B. Edelman. From rotation to disfiguration: Testing a dual-strategy model for recognition of faces across view angles *Perception*, 28: 817–824, 1999.
- [120] D. Valentin, H. Abdi, B. Edelman, and M. Posamentier. 2D or not 2D? that is the question: What can we learn from computational models operating on 2D representations of faces? in M. Wenger, J. Townsend (Eds.), *Computational, geometric, and process perspectives on facial cognition*. Mahwah (NJ): Erlbaum. 2001.
- [121] T. Valentine. A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology*, 43A:161–204, 1991.
- [122] T. Valentine and V. Bruce. The effects of distinctiveness in recognising and classifying faces. *Perception*, 15:525–536, 1986.
- [123] J. Vokey and D. Read. Familiarity, memorability, and the effect of typicality on the recognition of faces. *Memory & Cognition*, 22:208–224, 1992.
- [124] J.F. Werker, J.H. Gilbert, K. Humphrey and R.C. Tees. Developmental aspects of cross-language speech perception. *Child Development*, 52:349–355, 1981.
- [125] H. A. Wild, S.E. Barrett, M.J Spence, A.J. O’Toole, Y. Cheng, and J. Brooke. Recognition and categorization of adults’ and children’s faces: examining performance in the absence of sex-stereotyped cues. *Journal of Experimental Child Psychology*, 77:269–291, 2000.
- [126] R.K. Yin. Looking at upside-down faces. *Journal of Experimental Psychology*, 81: 141–145, 1969.
- [127] A.W. Young, D. Hellawell, and D.C. Hay. Configurational information in face perception. *Perception*, 16:747–759, 1987.
- [128] W. Zhao, R. Chellappa, P.J. Phillips and A. Rosefeld. Face Recognition: A Literature Survey. *ACM Computing Surveys*, 35:399–459, 2003.