



Face perception: A brief journey through recent discoveries and current directions



Ipek Oruc^{a,b,*}, Benjamin Balas^c, Michael S. Landy^d

^a Ophthalmology and Visual Sciences, University of British Columbia, Canada

^b Neuroscience, University of British Columbia, Canada

^c Department of Psychology and Center for Visual and Cognitive Neuroscience, North Dakota State University, United States

^d Department of Psychology and Center for Neural Science, New York University, United States

ARTICLE INFO

Keywords:

Face perception
Face recognition
Other-race effect
Identification
Face diet

ABSTRACT

Faces are a rich source of information about the people around us. Identity, state of mind, emotions, intentions, age, gender, ethnic background, attractiveness and a host of other attributes about an individual can be gleaned from a face. When face perception fails, dramatic psycho-social consequences can follow at the individual level, as in the case of prosopagnosic parents who are unable to recognize their children at school pick-up. At the species level, social interaction patterns are shaped by human face perception abilities. The computational feat of recognizing faces and facial attributes, and the challenges overcome by the human brain to achieve this feat, have fascinated generations of vision researchers. In this paper, we present a brief overview of some of the milestones of discovery as well as outline a selected set of current directions and open questions on this topic.

1. Overview

The last few decades have seen the field of face perception firmly establish itself as a prolific area within vision research resulting in significant advances in our understanding at the conceptual, computational, neuropsychological, neuroscientific, developmental, and behavioral levels. These varied accomplishments across many distinct subdomains of face-recognition research are extensive and conveying the current state of knowledge and looking ahead to future challenges is difficult when there is this much material to draw upon. Our approach to this overview is thus to focus on a relatively small set of topics that we think are among those that point towards outstanding questions at the edge of what we know about how face recognition works. Following a broad overview on some of the major discoveries in the field, we elaborate on this set of selected topics: (1) face recognition in the wild, (2) the other-race effect and the impact of experience on face recognition, (3) critical features for various aspects of face recognition, and (4) social evaluation of face images. We provide a review of recent advances on these four topics followed by a list of suggested future directions for each. Our aim is to appeal to a seasoned as well as a novice audience. Thus, we provide a Glossary of important concepts in face recognition research at the very end. Terms that are defined in this glossary are printed in **bold** at the first appearance to alert the reader of an upcoming entry.

2. What we know

Conceptually, differences between various recognition tasks such as **detection**, **categorization**, **discrimination**, **individuation**, **memory** and **naming** have been better delineated, and the typical modes of **recognition** for various visual stimulus classes have been characterized (Palmeri & Gauthier, 2004; Tanaka & Gauthier, 1997). The computationally nontrivial challenge of balancing sensitivity (**telling faces apart**) and robustness (**telling faces together**) has been described (Andrews, Jenkins, Cursiter, & Burton, 2015; Jenkins, White, Van Montfort, & Burton, 2011). Utilization of multiple distinct cognitive strategies, such as holistic processing, and recognition by parts, has been explored (Maurer, Grand, & Mondloch, 2002; Rossion, 2008; Tanaka & Farah, 1993). Hemispheric specialization (predominantly right-sided), as well as a central role for localized intra-hemispheric structures (e.g., fusiform gyrus) for the processing of faces has been described based on neuropsychological impairments observed in acquired prosopagnosia and other acquired syndromes of high-level vision (Barton, 2008; Damasio, Tranel, & Damasio, 1990). These findings have resonated for the most part with neuroimaging work that mapped out face-responsive cortical regions (Grill-Spector & Weiner, 2014; Kanwisher, McDermott, & Chun, 1997), with specific modules as well as distributed pathways proposed for various perceptual tasks and

* Corresponding author at: University of British Columbia, Department of Ophthalmology and Visual Sciences, 818 West 10th Avenue, Vancouver, B.C. V5Z 1M9, Canada.

E-mail address: ipor@mail.ubc.ca (I. Oruc).

<https://doi.org/10.1016/j.visres.2019.06.005>

Received 11 April 2019; Received in revised form 9 June 2019; Accepted 11 June 2019

Available online 14 June 2019

0042-6989/ © 2019 Published by Elsevier Ltd.

processes typically associated with faces, such as detection (Kriegeskorte, Formisano, Sorger, & Goebel, 2007), individuation (Gauthier, Skudlarski, Gore, & Anderson, 2000; Natu et al., 2010), and computational strategies, such as holistic vs. part-based processing (Rossion et al., 2000) and the processing of changeable vs. invariant aspects of faces such as expression and gaze vs. identity (Haxby, Hoffman, & Gobbini, 2000).

Face recognition ability, considered by some to be a specific unitary dimension of human cognition (Wilmer et al., 2010), varies widely across the population. At the two extremes of this spectrum specialized populations have been described: super-recognizers, who show an extraordinary ability to remember faces (Noyes, Phillips, & O'Toole, 2017; Russell, Duchaine, & Nakayama, 2009; Tardif et al., 2019) and developmental prosopagnosics, who show a profound inability to recognize familiar faces (Duchaine & Nakayama, 2006b). Face recognition has come to be viewed as a type of visual expertise that requires a great deal of training and practice to develop. Indeed, face recognition skills continue to be honed with experience over much of childhood (Carey, Diamond, & Woods, 1980; de Heering, Rossion, & Maurer, 2012) reaching its peak in adulthood (Germine, Duchaine, & Nakayama, 2011). As with any expertise, face ability presumably emerges through the interaction of genetic predisposition and experience (Gauthier et al., 2014). Face diet, the collection of faces encountered by an individual on a daily basis, shapes the specific type of expertise one has with faces (Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003). Indeed, empirical studies of the face diet show that infants as well as adults have faces in view for a substantial proportion of their waking hours (Jayaraman, Fausey, & Smith, 2015; Oruc, Shafai, Murthy, Lages, & Ton, 2019; Sugden, Mohamed-Ali, & Moulson, 2014). In the other-race effect (ORE), a marked reduction is evident in the ability to discriminate and remember faces of ethnicities that are minimally present in the face diet compared to faces of commonly encountered ethnicities (Meissner & Brigham, 2001; Rhodes et al., 2009). In addition, lack of sufficient exposure to a rich diet of faces has also been found to curtail face abilities. For example, in sparsely populated regions, individuals who lack the opportunity to experience a rich variety of faces show lower face-identification ability (Balas & Saville, 2017). Low social motivation also limits encounters with faces (e.g., by reducing social interactions or face looking behaviors) and is associated with reduced face abilities in autism spectrum disorder (Oruc, Shafai, & Iarocci, 2018).

3. Some recent advances: A selection of active topics

3.1. Face recognition in the wild

One simple way to start thinking about face recognition is to address what seems like a straightforward question: How difficult is it? Though this question is simple to state, it turns out that relatively recent results have led to a renewed focus on understanding the conditions under which face recognition is easy and the circumstances that can make face recognition very difficult.

A great deal of face-recognition research has relied upon face stimuli that are highly controlled. That is, the face images we use for many of our experiments (including those designed by the authors!) have frequently been matched for mean luminance and contrast, or intensity histograms, had equal power spectra imposed, been aligned with one another based on eye position, and/or cropped with a uniform oval. On one hand, these manipulations provide some guarantees about what low-level visual processing can and cannot achieve when applied to these images, hopefully requiring higher-level processes to make the largest contribution to recognition. On the other hand, it isn't clear that the way we recognize such highly controlled images reflects the way we recognize faces in real visual environments. After all, the face images we see in natural settings have not been controlled at all, and thus incorporate all of the sources of variability mentioned above and then

some. Real-world faces can look different from moment to moment due to image variation (variability that results from differences in image-capture parameters like lighting or resolution) and appearance variation (variability that results from changes to the face itself, like changing facial hair or expression). In real settings, the image of an individual may look dramatically different across different exemplars due to both of these kinds of variability. What is face recognition like under those circumstances?

Fortunately there has been a great deal of activity in this area for several years now, which has revealed some important aspects of how face recognition works “in the wild”. Much of this work has also used a simple testing paradigm based on sorting face images according to identity (Jenkins et al., 2011), which reveals many key aspects of how recognition of ambient images (Young & Burton, 2017) differs from that of highly controlled images. Note, however, that although we will discuss many studies based on this card-sorting paradigm here, it is the use of highly variable face images that we think is particularly important. While the testing paradigm has been the vehicle for much of this work, it's vital to remember that it's not card-sorting we're interested in, but rather how card-sorting may reveal properties of **ambient face recognition**.

The simplest observation about ambient face recognition is that it turns out to be quite hard in some circumstances. Presented with a set of cards depicting multiple images of an unknown number of unique and unfamiliar individuals, observers typically struggle to sort the images correctly. There are a number of ways to describe what goes wrong here, including simply counting the number of piles observers make (Jenkins et al., 2011), applying signal-detection descriptors to pairs of cards (Balas & Pearson, 2017), and using confusion matrices to quantify errors in sorting (Neil, Cappagli, Karaminis, Jenkins, & Pellicano, 2016; Yan, Andrews, Jenkins, & Young, 2016), but these approaches tend to agree on the basic phenomenology. Observers in these tasks tend to do the following things: (1) They often put images of the same person in different piles, and (2) They are less likely to put images of different people into the same pile. For example, a deck of cards depicting just two individuals may be sorted into a median value of 8 piles (Jenkins et al., 2011), with few “intrusions” – instances of a mostly homogeneous pile containing an image of a different person. This tendency to overestimate the number of unique identities in a deck of cards is further exacerbated when other-race faces are used (Laurence, Zhou, & Mondloch, 2016), reflecting further failures of face recognition when ambient image variability is incorporated into face stimuli. Importantly, when using signal detection measures to describe performance via the disposition of image pairs, it is evident that behavior in these tasks is characterized both by a strict criterion applied to accepting that pairs of images depict the same person, but also by low sensitivity to the difference between within-person differences and between-person differences (Balas & Pearson, 2017). To provide a short summary of face recognition abilities with such stimuli, it seems that ambient face recognition of unfamiliar faces can be characterized by fairly good abilities to *tell people apart*, but much poorer abilities to *tell people together* (Andrews et al., 2015). This pattern of results suggests that much of the difficulty in recognizing faces under natural viewing conditions is working out the many different ways that one person's face can look different from itself, both as a result of image variation and appearance variation.

Another important observation about performance with ambient images is that the recognition of familiar faces differs profoundly from the recognition of unfamiliar faces. By itself, this is not a new observation: There has been substantial evidence for a long time that unfamiliar-face recognition is poorer than familiar-face recognition in multiple contexts (Bruce et al., 1999; Bruce, Henderson, Newman, & Burton, 2001) and also that the manner in which familiar and unfamiliar faces are processed differs (Johnston & Edmonds, 2009). We suggest that what is important about the increased use of ambient images in recent face recognition research is that the use of naturalistic

stimuli tends to make these differences particularly stark, and the results observed across multiple recent studies raise important theoretical issues that remain to be explored fully. For example, when observers are presented with ambient images of people they know (e.g., celebrities) they make few sorting errors (Jenkins et al., 2011), if any. Face recognition mechanisms are thus not ineffective when confronted with high variability, but instead learn how to cope with that variability as faces become familiar. We emphasize that this includes both image variation and appearance variation: Familiar observers successfully cope with both kinds of variability. However, the lessons learned for familiar faces don't naturally generalize to other faces belonging to the same race/age, etc., which has important implications for models of face recognition. In particular, a strong hypothesis regarding the nature of face recognition based on these results is that face recognition may operate primarily in a person-specific manner. That is, rather than being based on general principles describing how all faces may vary as lighting or viewpoint changes, human face recognition may instead be based on representations that are learned for each familiar face and are thus not applicable to other individuals. Indeed, there is some evidence that variability across face images is idiosyncratic (Burton, Kramer, Ritchie, & Jenkins, 2016), meaning that the way a face looks different from itself both provides an independent cue for recognition and limits the generalizability of recognition strategies across individuals. Further exploration of these ideas is bound to be of considerable theoretical importance. The strong version of this hypothesis suggests that trying to model variability in general is likely to be of little use, but that learning models of variability per individual could dramatically improve the agreement between models and human performance. Of course, understanding how to model variability appropriately is a hard question (one of many!), but one that is likely to be of special importance as we work to understand, (1) How observers achieve robust recognition of familiar ambient face images, and (2) How unfamiliar faces become familiar, such that observers can cope with naturalistic variability.

3.2. The other-race effect and the impact of experience on face recognition

If we're thinking about face recognition in the context of ambient images in natural settings, we've seen that it turns out to be a difficult task. However, with experience, face recognition becomes substantially easier even when observers have to cope with a great deal of variability across different images of the same individual. Observers clearly *learn* to recognize faces, which leads to another question that is easy to state, but harder to answer: How do they do that?

One of the more robust behavioral effects in the face-recognition literature that is related to learning is the other-race effect (Malpass & Kravitz, 1969). Briefly, this refers to the relative impairment in recognizing and discriminating between faces belonging to racial categories that are underrepresented in an observer's experience when compared to performance for majority face categories. There are also some advantages that other-race faces enjoy, specifically with regard to categorization according to race (Sun, Zhang, Yang, Bentin, & Zhao, 2014) and visual search (Levin, 2000; Sun, Song, Bentin, Yang, & Zhao, 2013), but overall the phenomenology is that individuating other-race faces tends to be more difficult. This rich topic is one of the best examples of how varying experience with faces can lead to poorer recognition, and thus is critically important to consider if one is interested in how robust face-recognition abilities are established through maturation, through contact with a particular subset of faces in everyday interactions, or by virtue of social factors that influence the salience of various faces in the environment.

The other-race effect is known to depend critically on contact with faces belonging to different race categories. Infants develop an other-race effect following a developmental trajectory referred to as "perceptual narrowing" (Nelson, 2001), in which above-chance abilities to individuate own and other-race faces early in infancy eventually change such that other-race (or other-species) faces are no longer

distinguishable, but own-race faces are (Kelly et al., 2007; Pascalis, de Haan, & Nelson, 2002). Critically, the effects of experience in perceptual narrowing are generally negative, in that the end result is a loss of face recognition abilities for categories of faces not represented in the environment: Performance with frequently-seen faces does not improve, but performance with infrequently-seen faces deteriorates. That this trajectory is experience-dependent, rather than purely maturational, was demonstrated by Sugita (2008) who reared non-human primates in a face-free environment, followed by a period of time in which either own- or other-species faces were presented to the infant monkeys. Perceptual narrowing supporting either human or monkey face recognition followed as a function of what faces were made available in the environment. In human participants, the reversibility of the other-race effect as a function of the visual environment has also been reported in childhood (Bar-Haim, Ziv, Lamy, & Hodes, 2006; Sangrigoli, Pallier, Argenti, Ventureyra, & de Schonen, 2005), further demonstrating that it is neither the observer nor the stimulus alone that determines the other-race effect, but the observer's experience seeing the stimulus that matters. An important caveat, though, is that the timing of exposure also matters. Whether an observer has had early or late exposure to a category of faces during their lifetime determines their ability to acquire robust abilities with that face category later (Cassia, Kuefner, Picozzi, & Vescovo, 2009). Thus, while other-group face recognition effects can be reduced by experience (Cassia, Picozzi, Kuefner, & Casati, 2009), there is evidence that there are critical periods for exposure that support the potential for reactivation of dormant mechanisms for recognizing broader sets of faces.

The other-race effect is just one example of a specific difference in face exposure manifesting as a difference in face-recognition ability. A broader question remains relatively open: How does experience with faces constrain face recognition ability in general? Experience does not have to be estimated solely in terms of in-group and out-group categories (which can be difficult to define clearly), but can be understood both more broadly as the full visual diet of faces in the environment, but also more specifically as the faces any one individual actually knows (Jenkins, Dowsett, & Burton, 2018). For example, observers from sparsely-populated communities tend to have poorer face-recognition abilities than observers from denser urban environments (Balas & Saville, 2015, 2017), even for own-race faces. This suggests that it is not just the category biases inherent to one's face experience that matter, but the raw number of faces available to be recognized that further constrains recognition abilities. This suggests that we need to do better at characterizing what individuals see in their environment beyond labeling faces according to race or age, perhaps by describing face homogeneity within an environment, estimating the size of the face network available to observers (Hill & Dunbar, 2003), and possibly by characterizing the depth of interaction between observers and other individuals in their environment.

3.3. Critical features for various aspects of face recognition

One aspect of face processing that may change as a function of experience, training, and/or development is the information used to recognize faces. Better face recognition may follow from observers establishing a more useful vocabulary for measuring facial appearance or more efficient and accurate application of mechanisms for measuring appearance. But what information is available to observers to recognize faces, and what measurements are the most useful for performing different face-recognition tasks? Faces convey a wide variety of attributes, such as identity, expression, gender, age, attractiveness, mental state, and personality traits (Todorov, Mandisodza, Goren, & Hall, 2005; Willis & Todorov, 2006). The perception of many of these attributes is influenced by the types of stimuli immediately preceding them. Termed adaptation aftereffects, these phenomena describe changes in the perception of a particular facial trait following exposure to another. For example, after viewing male faces, an androgynous face is often

perceived as more feminine than it would otherwise, whereas the same androgynous face tends to appear more masculine when seen after viewing female faces (Webster, Kaping, Mizokami, & Duhamel, 2004). While aftereffects in high-level perception, such as those of facial traits, necessarily inherit adaptation at lower levels of visual processing, face aftereffects are believed to be at least partly due to changes in the responses of high-level neural mechanisms that code for facial attributes, and thus have been used to study the properties of such mechanisms (Webster, 2015). Specifically, adaptation, as a methodology, has been used to examine how various aspects of faces have been coded in the human brain, resulting in evidence consistent with a norm-based representation for facial identity (Leopold, O'Toole, Vetter, & Blanz, 2001; Rhodes & Jeffery, 2006), and a multichannel model for the perception of eye gaze (Calder, Jenkins, Cassel, & Clifford, 2008). As a perceptual phenomenon, adaptation aftereffects have been observed for the perception of facial identity (Jiang, Blanz, & O'Toole, 2006; Leopold et al., 2001; Oruc & Barton, 2010b), expression, gender and ethnicity (Webster et al., 2004), age (O'Neil & Webster, 2011; Schweinberger et al., 2010), among others. It remains unclear whether the neural changes that give way to aftereffects serve a functional purpose (Clifford et al., 2007), but some evidence suggest that adaptation may sharpen tuning for various aspects of faces and improve discrimination ability for attributes such as identity and gender (Oruc & Barton, 2011; Yang, Shen, Chen, & Fang, 2011). In addition, it has been suggested that in the context of a norm-based model, where faces are coded in reference to deviations from a norm that lies at the origin of a multi-dimensional face space, adaptation may serve the functional purpose of dynamically recalibrating the norm based on prevailing face exposure (see, Webster & MacLeod, 2011, for a detailed review).

Culminating in the extraction of such high-level features as identity, expression, gender, attractiveness and others from faces, the processing of face images, starts, like any other visual stimuli, with the front-end processing of low-level physical qualities of the face images, such as spatial frequency and orientation. Within the last few decades, we have come to learn some important properties of these processes in the context of face perception. We now know that despite the fact that face images are broadband and contain information over a wide range of spatial scales and orientations, observers utilize relatively narrow portions of the power spectrum available to them to recognize faces. For example, observers recognize faces mainly based on horizontal orientations (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Pachai, Sekuler, & Bennett, 2013) and specific scales (i.e., a narrow band of spatial frequencies) within the face image (e.g., Gold, Bennett, & Sekuler, 1999a; Nasanen, 1999). In addition, the range of spatial frequencies observers use depends on the task and viewing conditions (Bonnar, Gosselin, & Schyns, 2002; Oruc & Barton, 2010a; Schyns, Bonnar, & Gosselin, 2002; Shahangian & Oruc, 2014). One way that experience or training may affect face processing is that observers establish “information biases” that lead to a greater reliance on these critically important features relative to less useful visual information. It is an open problem to determine which specific measurements best support performance in any given task. Solving this problem will likely require complementary work along computational, psychophysical, and neural lines of inquiry.

3.4. Social evaluation of face images

Besides being able to recover variables like identity, gender, race, and age from face images, observers are also capable of estimating social properties from face images. These include a broad range of evaluations, including the estimation of trustworthiness (Rule, Krendl, Ivcevic, & Ambady, 2013), dominance (Bates, Re, & Perrett, 2015), aggression (Carre & McCormick, 2008), competence (Sussman, Petkova, & Todorov, 2013), extraversion (Borkenau, Brecke, Mottig, & Paelecke, 2009), and mental health (Fowler, Lilienfeld, & Patrick, 2009). In some circumstances, these judgments can be made given

extremely limited presentation time (~100 ms or so, Todorov, Pakrashi, & Oosterhof, 2009; Willis & Todorov, 2006) suggesting that observers can base complex social inferences on relatively impoverished input. However, it is important to note that these social judgments vary in terms of their validity: although observers may be consistent in their attributions of social properties to individuals based on facial appearance, those attributions do not necessarily reflect the real behavior or personalities of the individuals under consideration. Trustworthiness judgments, for example, are known to have poor validity (Rule et al., 2013). By comparison, judgments of aggression based on facial appearance are known to be diagnostic of real aggressive behaviors in some settings (Carre & McCormick, 2008). Other social attributions are frequently characterized as being above chance levels, but this often means that observers are only barely exceeding random guessing as a group. Regardless, the reliability of these judgments suggests that there are face-recognition mechanisms that use multiple aspects of facial appearance to infer complex social properties. How do these mechanisms work? What visual information is used to make such judgments? Furthermore, given the large space of candidate social evaluations that we *could* make from a given face, how should we characterize social face perception more generally?

The latter question has been the subject of a great deal of work over the past decade, largely based on applying various types of factor analysis to large datasets comprised of multiple social judgments made in response to many distinct faces. Specifically, rather than consider social face evaluation as a confederation of many different social inferences, these approaches attempt to model social evaluation in the context of a model in which each face occupies a position in a low-dimensional space defined by axes that capture most of the variance in social evaluations across images (Oosterhof & Todorov, 2008). Presently, the key conclusion of these studies is that social face evaluation is best characterized by a 2-dimensional space in which *valence* (or sometimes *trustworthiness*) and *dominance* are the primary axes of social evaluation. In some cases, a 3rd *attractiveness* dimension also captures a substantial amount of variance (Sutherland et al., 2013), but in all cases it seems that a wide range of social judgments can largely be accounted for by the coordinates assigned to faces along a small set of axes. This applies both to highly controlled face images as well as ambient face images that include substantial appearance variability (Sutherland et al., 2013), suggesting that this low-dimensional structure is highly robust to ecologically valid sources of image variability. The obvious benefit of this approach is that it transforms social face evaluation from a quest to understand how each of a large collection of highly specific social judgments are made from faces to a larger research enterprise in terms of understanding how a small set of social variables are estimated from face images.

Indeed, the visual features supporting both valence and dominance estimates are understood fairly well, with some important caveats. Valence, for example, appears to depend heavily on the extent to which an individual's neutral facial expression tends to reflect positive or negative affect (Said, Sebe, & Todorov, 2009; Sutherland et al., 2015). Individuals who tend to look a bit happy are rated as more trustworthy or approachable, while individuals who tend to look a bit sad or angry are rated as less so. Dominance perception has similarly been characterized in terms of specific aspects of facial appearance: head tilt matters a good bit, for example, suggesting that dominance judgments are not invariant to variables like lighting or pose (Mignault & Chaudhuri, 2003). As for attractiveness (which, as we mentioned, is sometimes included as a 3rd dimension of social face space), it has been the subject of a great deal of work attempting to link attractiveness ratings of male and female faces to specific image properties like symmetry (Perrett et al., 1999) and contrast relationships within the face (Russell, 2003). There is thus good evidence that low-dimensional models have good explanatory power that accounts for a wide range of social evaluations based on facial appearance, and moreover that these social dimensions can be linked closely to specific visual features.

There is some trickiness here that remains to be understood. In particular, some recent results suggest that while low-dimensional models may be useful for characterizing social face evaluation in general, one may need different specific models to account for the way different populations of faces are evaluated, or the way different populations of observers make decisions about face stimuli. For example, recent results suggest that computer-generated (CG) faces (like those used in Oosterhof & Todorov, 2008) are not evaluated according to trustworthiness in the same way as real faces even when identity is matched across real and CG faces (Balas & Pacella, 2017). Further, it may matter a great deal who the observers are that make these judgments, as the specific dimensional structure that best accounts for social evaluations varies cross-culturally. Recent results suggest that Chinese observers, for example, appear to evaluate faces according to *valence* and *competence* rather than *valence* and *dominance* (Wang et al., 2019), which strongly suggests that the set of fundamental social judgments may vary with experience or culture. Further, Na, Kim, Oh, Choi, and O'Toole, (2015) demonstrated that while electability tends to be predicted well in Western populations using competence as a proxy, the same does not hold for Korean observers. The ongoing work of research groups participating in the Psychological Science Accelerator (Moshontz et al., 2018) aims to address this question by replicating the original study by Oosterhof and Todorov (2008) at a large scale at many sites around the world (Jones et al., 2019). Determining the extent to which there is variability in the dimensional structure that best accounts for variance across social judgments will be an important step towards understanding how social evaluations are learned and applied as a function of the face environment.

4. Future directions

There is a range of questions that follow naturally from recent progress in face-perception research, some of which have already been studied, and others that, to our knowledge, have not. While by no means an exhaustive list, we offer a few questions concerning how face images are recognized that we think will be important to address in the near future.

4.1. Face recognition in the wild

1. Are ambient face images processed holistically? There is a vast literature demonstrating in various ways that faces are processed holistically, rather than in a piecewise fashion. This includes many studies examining holistic face processing via the face-inversion effect (Yin, 1969), the composite-face effect (Young, Hellawell, & Hay, 1987), and the part-whole effect (Tanaka & Farah, 1993), the majority of which use highly controlled face images rather than ambient images. When ambient images are used, will the same results be obtained? If so, how is holistic processing applied to highly variable images? If not, what should we conclude about the relationship between holistic processing and face recognition in natural environments?
2. What visual features are used to recognize ambient images? In terms of specific image features that support face recognition, there are multiple studies suggesting that mid-range spatial frequencies (8–16 cycles per face or so, Ruiz-Soler & Beltran, 2006), horizontal orientation energy (Dakin & Watt, 2009), and descriptors of pigmentation (Russell, Sinha, Biederman, & Nederhouser, 2006) are particularly valuable for recognizing controlled face images. Again, how does this play out when we consider ambient face images instead? In general, how much do the lessons learned about the vocabulary of face recognition apply to faces appearing in natural environments?
3. How is ambient face recognition learned? Given that observers are capable of coping with high appearance variability when recognizing familiar faces, what are they learning that allows them to achieve this feat? This is an important question to consider both

developmentally (Matthews, Davis, & Mondloch, 2018) and in adult populations. To date, we know that high variability during training helps improve performance (Baker, Laurence, & Mondloch, 2017) during face learning, but more generally we would like to know more about how observers work out the way to “tell faces together” for new individuals.

Of course there are many more things to think about in this domain, but we hope we have conveyed the richness of this topic and suggested some exciting directions for future work.

4.2. The other-race effect and the impact of experience on face recognition

Several factors have been shown to contribute to the diminished recognition ability for other-race faces as seen in the ORE, including degree of contact, implicating lack of experience with other race faces, and socio-cognitive biases implicating attitudes towards members of “out-groups” (e.g., Rossion & Michel, 2011). Previous studies have examined what changes these factors engender in the neural processes and mechanisms of recognition with evidence variably pointing to qualitative (Michel, Caldara, & Rossion, 2006; Michel, Rossion, Han, Chung, & Caldara, 2006; Tanaka, Kiefer, & Bukach, 2004) and quantitative differences (Harrison, Gauthier, Hayward, & Richler, 2014; Mondloch et al., 2010; Shafai & Oruc, 2018). How does experience mold recognition processes for faces in general, as well as in the context of the ORE? Below, we list several open questions on this topic.

1. What are the limiting factors for achieving expert native performance for faces? Infants' face exposure is predominantly to a remarkably homogenous set of faces (e.g., own-race, female, few identities) (Jayaraman et al., 2015; Sugden et al., 2014) compared to the richer face diet of adults (Oruc et al., 2019). It has been suggested that perceptual narrowing and specialization to own-race faces may be facilitated by the homogeneity and consistency in the infant face diet (Sugden et al., 2014). By extension of this logic, exposure to heterogeneous input may be viewed as detrimental to the development of expertise and specialization to faces. One prediction that follows is that sustained and substantial exposure to faces of more than one race during early development would lead to a generalist face-recognition system – a ‘*jack of all trades, master of none*’ system, that does not achieve expertise. Alternatively, exposure to a rich variety of faces may spur richer representations and better recognition abilities. Finally, both factors may play a role in the context of a temporal modulation of exposure statistics, such as with gradually enriched exposure.
2. What are the necessary and sufficient conditions of face experience that allow a typical individual to achieve their full potential of face ability? Are there any sensitive periods for the acquisition of those experiences?
3. Can adult face-recognition ability be improved? There is significant variation in face ability across the population with normal vision – some of us are good at it, others struggle. Quality of life can be negatively impacted even by mild face-recognition impairment. In addition, certain occupations implicitly or explicitly require competence in face recognition (as in the case of, e.g., teachers, health-care professionals, law-enforcement and border-services officers). Assuming some of the variation stems from experience-related factors, can we improve face recognition via targeted training?
4. What can we learn from super-recognizers (Noyes et al., 2017; Russell et al., 2009)? Regardless of genetic predisposition and experience factors, are there high-level strategies that can be purposefully adopted at a cognitive level (looking at a certain location in the face, preferring a certain viewpoint) to improve face recognition? Can we devise a kind of “face coaching” by training individuals to mimic the strategies used by people who are good at recognizing faces?

4.3. Critical features for various aspects of face recognition

Visual processes of low- and mid-level vision take advantage of environmental statistics to facilitate perception. For example, in the dome-crater illusion, a 2D shading pattern that is equally consistent with a convex dome lit from below or a concave crater lit from above is perceived as the latter. The human brain weighs in on the ambiguity by inserting the knowledge that scenes are typically lit from above (Adams, Graf, & Ernst, 2004). Visual scene statistics bear regularities in many other attributes relevant to low- and mid-level vision such as orientation (Girshick, Landy, & Simoncelli, 2011) and motion (Weiss, Simoncelli, & Adelson, 2002). The visual system uses these to resolve perceptual ambiguities implicitly by using *a priori* assumptions based on environmental statistics (Kersten, Mamassian, & Yuille, 2004).

1. Do visual processes of high-level vision take advantage of environmental statistics to facilitate and support recognition of faces? Empirical studies of the face diet have uncovered many statistical regularities in face exposure. For example, daily exposure of both infants and adults is predominantly to faces that are familiar, and seen up-close, i.e., visually large (Jayaraman et al., 2015; Oruc et al., 2019). This aligns well with the finding that face recognition is more efficient at visual sizes typical of social interaction (Yang, Shafai, & Oruc, 2014). Beyond statistical regularities caused by the observers' interactions with the environment (such as the predominance of visually large faces due to social interaction), the geometry of optics is also a basic determinant of image statistics. For example, faces seen from a distance are visually small. Also, due to the resolution limits and sensitivity characteristics of the visual system, *neural images of distant faces are blurry*. Indeed, critical spatial frequencies for face recognition change with size such that coarser features are utilized for smaller sizes and finer details are used for larger sizes (Oruc & Barton, 2010a; Willenbockel et al., 2010). Based on this, are blurry faces better recognized at small sizes? What is the optimal size at which to view a blurry image? What image manipulations might improve recognition of blurry images?
2. How can we use this knowledge to help those with disordered vision? This population may include people with high-level impairments, such as those with prosopagnosia, and those with low vision, such as resulting from age-related macular degeneration. What are the best approaches to recognizing severely blurry faces? Coarser features are sufficient for supporting recognition of faces under some conditions, e.g., when they mimic viewing at a distance (Mousavi & Oruc, 2019; Oruc & Barton, 2010a; Shahangian & Oruc, 2014). Can we extrapolate from those conditions to help devise visual aids for people with blurry vision? What are the differences between peripherally viewed faces and centrally viewed blurry faces? Can peripheral face-recognition ability be trained?

4.4. Social evaluation of face images

What else is there to think about with regard to the social evaluation of faces? We suggest that the following are likely to be important questions for future work.

1. Are fundamental social evaluations robust to appearance variability? The work of Sutherland et al. (2013) suggests that using ambient images rather than more uniform face stimuli leads to a similar low-dimensional model of social face space, but in other cases, variation in pose affects specific aspects of appearance, which in turn affect social face evaluation. More generally, are social inferences stable across variation in pose, expression, lighting, etc., or is it relatively easy to modulate these judgments by modulating appearance?
2. Are fundamental social evaluations really fundamental? Although

valence/dominance axes frequently emerge as low-dimensional basis functions that explain the variance across a set of social face evaluations, to our knowledge, there hasn't been a great deal of work examining how effectively they can be used to extrapolate to new social inferences. For example, if you know where faces live in a valence/dominance space, can you predict how competent a face looks? What about the presumed sexual orientation of a face, or the presumed mental health of a face? Again, this is not a question about establishing validity, but rather linking a useful low-dimensional model to a broad range of social judgments.

3. What low-visual features support social evaluation? Face recognition has been characterized as depending on mid-range spatial frequencies, horizontal orientations, and holistic face processing. Does social face evaluation depend on the same visual information? Alternatively, are there independent features for social judgments that differ from those that support face recognition? Recent results by Goffaux and Dakin (2010) demonstrate that vertical orientation energy is recruited for gaze judgments, for example, while horizontal orientation energy is recruited for identification. This is one intriguing piece of evidence that social inferences may rely on different features than other face recognition tasks.

5. In closing

The last few decades have seen a flurry of research on face perception. Beyond all that we have learned, we have also come to recognize face perception as a key visual function central to social interaction and relevant to survival. Demands on improved face-perception skills may have shaped the human brain – the way it is connected, organized, and the way other stimuli (e.g., orthography) are processed and recognized. Face research has been, and will continue to be, a key arena that allows us to learn about processing of faces in the brain, as well as the organizational principles of the brain in general. This review serves as a short overview of a limited selection of a vast literature.

6. For the novice reader: Glossary of important concepts in face recognition research

For the novice reader we provide here a glossary of the important concepts mentioned in this review.

6.1. Ambient faces

This term refers to faces that incorporate the natural variability associated with viewing people in complex, real-world environments. This means that in some ways it is easier to define ambient faces by what they are *not*: They are not constrained to be in a particular pose, not aligned with respect to eye position or other features, not manipulated to equalize luminance or contrast, and otherwise not changed in terms of their appearance. The use of such faces has grown in recent years, leading to a range of results demonstrating how challenging face recognition tasks are under ordinary viewing conditions.

6.2. Detection

Detection refers to the ability to indicate the presence (vs. absence) of a face. In a typical psychophysical experiment this can be tested using a two-interval two-alternative forced choice (2I-2AFC) task, in which one of two possible stimulus intervals (or locations) features a face and the other is blank (or an alternative stimulus category, see below for *categorization*). In a given trial the observer's task is to indicate which one of the two intervals (or locations) contained a face. A detection threshold for signal strength (e.g., contrast or display duration) is estimated over successive trials. Alternatively, reaction time can be used to assess detection ability. For example, Dalrymple and

Duchaine (2016) showed children with developmental prosopagnosia display arrays of 25 images and asked them to indicate whether a face was present or not. Slower reaction times were observed in several of the prosopagnosic participants indicating impairments in face-detection ability.

6.3. Categorization

One step up from detection, categorization refers to the ability to indicate the presence of a face vs. a stimulus of another category, such as a house. In a typical psychophysical experiment, a single interval 2AFC paradigm may be used, where the observer views a single stimulus and indicates whether it was a face or a stimulus from the alternative category. In this definition, categorization is closely related to detection (see above). Alternatively, a categorization task may require the observer to classify a face as one of a fixed number of sub-categories, such as female vs. male. For example, Zhao and Bentin (2008) presented observers with Israeli and Chinese faces and asked them to indicate whether each face was female or male, old or young, and Israeli or Chinese, in separate blocks. They found an *other-race advantage* (ORA) in categorization of race, where participants were faster to correctly respond to faces of the other-race compared to own-race. An ORA was not observed for categorization of age and gender, which allowed the authors to conclude that the attributes that determine age and gender in faces are universal.

6.4. Discrimination

One step further from categorization, discrimination refers to the ability to tell apart one face from another. A same/different task may be used to assess this ability, where a trial features two of the same face or two different faces and accuracy across all trials is measured. Alternatively, in a typical psychophysical design, a discrimination threshold is measured in a 2I-2AFC paradigm, which marks the minimal difference between two faces that enable observers to tell them apart. For example, Oruc et al. (2018) measured discrimination thresholds between various pairs of facial expressions in a group of adults with ASD. A morphing technique was used to generate faces with gradual fine-grained variation of expression strength, e.g., from 100% happy to 0% happy (neutral). On each trial, observers were shown two faces, each displaying the same level of one of two expressions (e.g., 5% happy and 5% angry), and asked which one of the two faces was “happier” (or “angrier”). They found that discrimination thresholds were elevated in the ASD group compared to the control group, indicating expression perception difficulties in this population (also see, Oruc & Barton, 2011, for a similar protocol to measure identity discrimination thresholds.).

6.5. Individuation

One step further from discrimination, individuation refers to the ability to identify a face as a distinct and specific exemplar among other faces. Key here is visual identification – recalling the name or biographical information associated with the face is not required. In a typical psychophysical experiment, an identification task can take the form of an *m*-AFC protocol in which one face randomly selected out of *m* alternatives is shown in any given trial. The observer indicates which one of *m* faces was shown and an identification threshold is measured (Gold, Bennett, & Sekuler, 1999b; Nasanen, 1999; Yang et al., 2014). Alternatively, accuracy (e.g., percent correct) can be used to assess individuation performance in a similar procedure. For example, Peterson and Eckstein (2012) asked observers to identify faces in a 10-AFC task and measured identification accuracy. They found that observers chose to fixate on locations on the face (e.g., just below the eye) that maximize identification accuracy.

6.6. Memory

Face memory refers to the ability to recall whether a given face has been seen before. This has often been assessed in an old/new paradigm in which observers first view (or learn) a set of faces and later are asked whether a face was among those viewed before (e.g., Chance, Goldstein, & McBride, 1975) or which one of *m* (two or more) faces is the one seen before (e.g., Hancock & Rhodes, 2008). The Cambridge Face Memory Test (Duchaine & Nakayama, 2006a), a standardized protocol that adopts this latter approach, is one of the most commonly used contemporary tests of face memory.

6.7. Recognition

Recognition is a tricky term in the human face-perception literature as it is used to describe two very distinct concepts. Following the usage in computer vision, some face perception studies use the term *recognition* to refer to individuation (see definition above) (e.g., Guo, Oruc, & Barton, 2009; Nasanen, 1999; Royer, Blais, Gosselin, Duncan, & Fiset, 2015; Shafai & Oruc, 2018). In other studies, face recognition refers to face memory (see definition above) (e.g., Chiroro & Valentine, 1995; Hancock & Rhodes, 2008).

6.8. Naming

Naming requires individuation as well as memory of the face but also includes the additional task of recalling biographical information (e.g., the name) that is attached to the face. For example, to examine the diagnostic features that enable recognition of familiar individuals using the Bubbles technique (Gosselin & Schyns, 2001), Butler, Blais, Gosselin, Bub, and Fiset (2010) presented observers with images of celebrity faces and asked them to verbally state the name of the celebrity. Alternative semantic information indicating recognition of the identity was also accepted as correct. Their study revealed similarities between recognition of familiar faces and unfamiliar ones (e.g., the use of the eye region in both) as well as some differences.

6.9. Telling faces apart

All faces look alike – they share the same configuration with two eyes above the nose above the mouth. Subtle variations in the appearance of facial features and the spacing among them set different faces apart. A good face recognition system must be sensitive enough to detect these subtle variations to tell faces apart.

6.10. Telling faces together

Different facial images of the same individual may vary drastically. In fact, images of two different faces may be more similar than those of the same one (Andrews et al., 2015; Jenkins et al., 2011). A good face-recognition system must be robust enough to ignore variations across images of the same individual and assign the same identity to these images.

References

- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nature Neuroscience*, 7(10), 1057–1058.
- Andrews, S., Jenkins, R., Cursiter, H., & Burton, A. M. (2015). Telling faces together: Learning new faces through exposure to multiple instances. *The Quarterly Journal of Experimental Psychology (Hove)*, 68(10), 2041–2050.
- Baker, K. A., Laurence, S., & Mondloch, C. J. (2017). How does a newly encountered face become familiar? The effect of within-person variability on adults' and children's perception of identity. *Cognition*, 161, 19–30.
- Balas, B., & Pacella, J. (2017). Trustworthiness perception is disrupted in artificial faces. *Computers in Human Behavior*, 77, 240–248.
- Balas, B., & Pearson, H. (2017). Intra- and extra-personal variability in person recognition. *Visual Cognition*, 25, 456–469.

- Balas, B., & Saville, A. (2015). N170 face specificity and face memory depend on hometown size. *Neuropsychologia*, 69, 211–217.
- Balas, B., & Saville, A. (2017). Hometown size affects the processing of naturalistic face variability. *Vision Research*, 141, 228–236.
- Bar-Haim, Y., Ziv, T., Lamy, D., & Hodes, R. M. (2006). Nature and nurture in own-race face processing. *Psychological Science*, 17(2), 159–163.
- Barton, J. J. (2008). Structure and function in acquired prosopagnosia: Lessons from a series of 10 patients with brain damage. *Journal of Neuropsychology*, 2(Pt 1), 197–225.
- Batres, C., Re, D. E., & Perrett, D. I. (2015). Influence of perceived height, masculinity, and age on each other and on perceptions of dominance in male faces. *Perception*, 44(11), 1293–1309.
- Bonnar, L., Gosselin, F., & Schyns, P. G. (2002). Understanding Dali's Slave market with the disappearing bust of Voltaire: A case study in the scale information driving perception. *Perception*, 31(6), 683–691.
- Borkenau, P., Brecke, S., Mottig, C., & Paelecke, M. (2009). Extraversion is accurately perceived after a 50-ms exposure to a face. *Journal of Research in Personality*, 43, 703–706.
- 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, 339–360.
- Bruce, V., Henderson, Z., Newman, C., & Burton, A. M. (2001). Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied*, 7(3), 207–218.
- Burton, A. M., Kramer, R. S., Ritchie, K. L., & Jenkins, R. (2016). Identity from variation: Representations of faces derived from multiple instances. *Cognitive Science*, 40(1), 202–223.
- Butler, S., Blais, C., Gosselin, F., Bub, D., & Fiset, D. (2010). Recognizing famous people. *Attention, Perception, & Psychophysics*, 72(6), 1444–1449.
- Calder, A. J., Jenkins, R., Cassel, A., & Clifford, C. W. (2008). Visual representation of eye gaze is coded by a nonopponent multichannel system. *Journal of Experimental Psychology: General*, 137(2), 244–261.
- Carey, S., Diamond, R., & Woods, B. (1980). Development of face recognition – A maturational component? *Developmental Psychology*, 16(4), 257–269.
- Carre, J. M., & McCormick, C. M. (2008). In your face: Facial metrics predict aggressive behaviour in the laboratory and in varsity and professional hockey players. *Proceedings of the Royal Society, Series B: Biological Sciences*, 275(1651), 2651–2656.
- Cassia, V. M., Kuefner, D., Picozzi, M., & Vesco, E. (2009). Early experience predicts later plasticity for face processing: Evidence for the reactivation of dormant effects. *Psychological Science*, 20(7), 853–859.
- Cassia, V. M., Picozzi, M., Kuefner, D., & Casati, M. (2009). Why mix-ups don't happen in the nursery: Evidence for an experience-based interpretation of the other-age effect. *The Quarterly Journal of Experimental Psychology (Hove)*, 62(6), 1099–1107.
- Chance, J., Goldstein, A. G., & McBride, L. (1975). Differential experience and recognition memory for faces. *Journal of Social Psychology*, 97(2), 243–253.
- Chiroro, P., & Valentine, T. (1995). An investigation of the contact hypothesis of the own-race bias in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 48(4), 4879–4894.
- Clifford, C. W., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., & Schwartz, O. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research*, 47(25), 3125–3131.
- Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9(4), 2.
- Dalrymple, K. A., & Duchaine, B. (2016). Impaired face detection may explain some but not all cases of developmental prosopagnosia. *Developmental Science*, 19(3), 440–451.
- Damasio, A. R., Tranel, D., & Damasio, H. (1990). Face agnosia and the neural substrates of memory. *Annual Review of Neuroscience*, 13, 89–109.
- de Heering, A., Rossion, B., & Maurer, D. (2012). Developmental changes in face recognition during childhood: Evidence from upright and inverted faces. *Cognitive Development*, 27, 17–27.
- Duchaine, B. C., & Nakayama, K. (2006a). 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.
- Duchaine, B. C., & Nakayama, K. (2006b). Developmental prosopagnosia: A window to content-specific face processing. *Current Opinion in Neurobiology*, 16(2), 166–173.
- Fowler, K. A., Lilienfeld, S. O., & Patrick, C. J. (2009). Detecting psychopathy from thin slices of behavior. *Psychological Assessment*, 21(1), 68–78.
- Gauthier, I., McGugin, R. W., Richler, J. J., Herzmann, G., Speegle, M., & Van Gulick, A. E. (2014). Experience moderates overlap between object and face recognition, suggesting a common ability. *Journal of Vision*, 14(8), 7.
- 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.
- Germin, L. T., Duchaine, B., & Nakayama, K. (2011). Where cognitive development and aging meet: Face learning ability peaks after age 30. *Cognition*, 118(2), 201–210.
- Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7), 926–932.
- Goffaux, V., & Dakin, S. C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*, 1, 143.
- Gold, J., Bennett, P. J., & Sekuler, A. B. (1999a). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research*, 39(21), 3537–3560.
- Gold, J., Bennett, P. J., & Sekuler, A. B. (1999b). Signal but not noise changes with perceptual learning. *Nature*, 402(6758), 176–178.
- Gosselin, F., & Schyns, P. G. (2001). Bubbles: A technique to reveal the use of information in recognition tasks. *Vision Research*, 41(17), 2261–2271.
- Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, 15(8), 536–548.
- Guo, X. M., Oruc, I., & Barton, J. J. (2009). Cross-orientation transfer of adaptation for facial identity is asymmetric: A study using contrast-based recognition thresholds. *Vision Research*, 49(18), 2254–2260.
- Hancock, K. J., & Rhodes, G. (2008). Contact, configural coding and the other-race effect in face recognition. *British Journal of Psychology*, 99(Pt 1), 45–56.
- Harrison, S. A., Gauthier, I., Hayward, W. G., & Richler, J. J. (2014). Other-race effects manifest in overall performance, not qualitative processing style. *Visual Cognition*, 22(6), 843–864.
- 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.
- Hill, R. A., & Dunbar, R. I. (2003). Social network size in humans. *Human Nature*, 14(1), 53–72.
- Jayaraman, S., Fausey, C. M., & Smith, L. B. (2015). The faces in infant-perspective scenes change over the first year of life. *PLoS One*, 10(5), e0123780.
- Jenkins, R., Dowsett, A. J., & Burton, A. M. (2018). How many faces do people know? *Proceedings of the Royal Society, Series B: Biological Sciences*, 285(1888).
- Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of the same face. *Cognition*, 121(3), 313–323.
- Jiang, F., Blanz, V., & O'Toole, A. J. (2006). Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, 17(6), 493–500.
- Johnston, R. A., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A review. *Memory*, 17(5), 577–596.
- Jones, B. C., DeBruine, L., Flake, J. K., Aczel, B., Adamkovic, M., Alaei, R., ... Francesco, F. (2019). To which world regions does the valence-dominance model of social perception apply? *PsyArxiv*.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17(11), 4302–4311.
- Kelly, D. J., Quinn, P. C., Slater, A. M., Lee, K., Ge, L., & Pascalis, O. (2007). The other-race effect develops during infancy: Evidence of perceptual narrowing. *Psychological Science*, 18(12), 1084–1089.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271–304.
- Kriegeskorte, N., Formisano, E., Sorger, B., & Goebel, R. (2007). Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proceedings of the National Academy of Sciences USA*, 104(51), 20600–20605.
- Laurence, S., Zhou, X., & Mondloch, C. J. (2016). The flip side of the other-race coin: They all look different to me. *British Journal of Psychology*, 107(2), 374–388.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, 4(1), 89–94.
- Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, 129(4), 559–574.
- Malpass, R. S., & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology*, 13(4), 330–334.
- Matthews, C. M., Davis, E. E., & Mondloch, C. J. (2018). Getting to know you: The development of mechanisms underlying face learning. *Journal of Experimental Child Psychology*, 167, 295–313.
- Maurer, D., Grand, R. L., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255–260.
- 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*, 7(1), 3–35.
- Michel, C., Caldara, R., & Rossion, B. (2006). Same-race faces are perceived more holistically than other-race faces. *Visual Cognition*, 14, 55–73.
- Michel, C., Rossion, B., Han, J., Chung, C. S., & Caldara, R. (2006). Holistic processing is finely tuned for faces of one's own race. *Psychological Science*, 17(7), 608–615.
- Mignault, A., & Chaudhuri, A. (2003). The many faces of a neutral face: Head tilt and the perception of dominance and emotion. *Journal of Nonverbal Behavior*, 27, 111–132.
- Mondloch, C. J., Elms, N., Maurer, D., Rhodes, G., Hayward, W. G., Tanaka, J. W., & Zhou, G. (2010). Processes underlying the cross-race effect: An investigation of holistic, featural, and relational processing of own-race versus other-race faces. *Perception*, 39(8), 1065–1085.
- Moshontz, H., Campbell, L., Ebersole, C. R., Ijzerman, H., Urry, H. L., Forscher, P. S., & Chartier, C. R. (2018). The psychological science accelerator: Advancing psychology through a distributed collaborative network. *Advances in Methods and Practices in Psychological Science*, 1(4), 501–515.
- Mousavi, S. M., & Oruc, I. (2019). Size effects in the recognition of blurry faces. *Perception*, Under Review.
- Na, J., Kim, S., Oh, H., Choi, I., & O'Toole, A. (2015). Competence judgments based on facial appearance are better predictors of American elections than of Korean elections. *Psychological Science*, 26(7), 1107–1113.
- Nasanen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, 39(23), 3824–3833.
- Natu, V. S., Jiang, F., Narvekar, A., Keshvari, S., Blanz, V., & O'Toole, A. J. (2010). Dissociable neural patterns of facial identity across changes in viewpoint. *Journal of Cognitive Neuroscience*, 22(7), 1570–1582.
- Neil, L., Cappagli, G., Karaminis, T., Jenkins, R., & Pellicano, E. (2016). Recognizing the same face in different contexts: Testing within-person face recognition in typical development and in autism. *Journal of Experimental Child Psychology*, 143, 139–153.
- Nelson, C. A. (2001). The development and neural bases of face recognition. *Infant and Child Development*, 10, 3–18.
- Noyes, E., Phillips, P. J., & O'Toole, A. J. (2017). What is a super-recogniser? In M.

- Bindemann, A. M. Megreya (Eds.). *Face Processing: Systems, Disorders, and Cultural Differences* (pp. 173–201). New York: Nova.
- O'Neil, S. F., & Webster, M. A. (2011). Adaptation and the perception of facial age. *Visual Cognition*, 19(4), 534–550.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences USA*, 105(32), 11087–11092.
- Oruc, I., & Barton, J. J. (2010a). Critical frequencies in the perception of letters, faces, and novel shapes: Evidence for limited scale invariance for faces. *Journal of Vision*, 10(12), 20.
- Oruc, I., & Barton, J. J. (2010b). A novel face aftereffect based on recognition contrast thresholds. *Vision Research*, 50(18), 1845–1854.
- Oruc, I., & Barton, J. J. (2011). Adaptation improves discrimination of face identity. *Proceedings of the Royal Society, Series B: Biological Sciences*, 278(1718), 2591–2597.
- Oruc, I., Shafai, F., & Iarocci, G. (2018). Link between facial identity and expression abilities suggestive of origins of face impairments in autism: Support for the social-motivation hypothesis. *Psychological Science*, 29(11), 1859–1867.
- Oruc, I., Shafai, F., Murthy, S., Lages, P., & Ton, T. (2019). The adult face-diet: A naturalistic observation study. *Vision Research*, 157, 222–229.
- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013). Sensitivity to information conveyed by horizontal contours is correlated with face identification accuracy. *Frontiers in Psychology*, 4, 74.
- Palmeri, T. J., & Gauthier, I. (2004). Visual object understanding. *Nature Reviews Neuroscience*, 5(4), 291–303.
- Pascalis, O., de Haan, M., & Nelson, C. A. (2002). Is face processing species-specific during the first year of life? *Science*, 296(5571), 1321–1323.
- Perrett, D. I., Burt, D. M., Penton-Voak, I. S., Lee, K. J., Rowland, D. A., & Edwards, R. (1999). Symmetry and human facial attractiveness. *Evolution and Human Behavior*, 20, 295–307.
- Peterson, M. F., & Eckstein, M. P. (2012). Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences USA*, 109(48), E3314–E3323.
- Rhodes, G., Ewing, L., Hayward, W. G., Maurer, D., Mondloch, C. J., & Tanaka, J. W. (2009). Contact and other-race effects in configural and component processing of faces. *British Journal of Psychology*, 100(Pt 4), 717–728.
- Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research*, 46(18), 2977–2987.
- Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W., & Nakayama, K. (2003). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science*, 14(6), 558–566.
- Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychol (Amst)*, 128(2), 274–289.
- Rossion, B., Dricot, L., Devolder, A., Bodart, J. M., Crommelinck, M., De Gelder, B., & Zoonjies, R. (2000). Hemispheric asymmetries for whole-based and part-based face processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 12(5), 793–802.
- Rossion, B., & Michel, C. (2011). An experience-based holistic account of the other-race effect. In G. Rhodes, A. J. Calder, M. Johnson, & J. V. Haxby (Eds.). *Oxford Handbook of Face Perception*. Oxford: Oxford University Press.
- Royer, J., Blais, C., Gosselin, F., Duncan, J., & Fiset, D. (2015). When less is more: Impact of face processing ability on recognition of visually degraded faces. *Journal of Experimental Psychology: Human Perception and Performance*, 41(5), 1179–1183.
- Ruiz-Soler, M., & Beltran, F. S. (2006). Face perception: An integrative review of the role of spatial frequencies. *Psychological Research*, 70(4), 273–292.
- Rule, N. O., Krendl, A. C., Ivcevic, Z., & Ambady, N. (2013). Accuracy and consensus in judgments of trustworthiness from faces: Behavioral and neural correlates. *Journal of Personality and Social Psychology*, 104(3), 409–426.
- Russell, R. (2003). Sex, beauty, and the relative luminance of facial features. *Perception*, 32(9), 1093–1107.
- Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16(2), 252–257.
- Russell, R., Sinha, P., Biederman, I., & Nederhouser, M. (2006). Is pigmentation important for face recognition? Evidence from contrast negation. *Perception*, 35(6), 749–759.
- Said, C. P., Sebe, N., & Todorov, A. (2009). Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion*, 9(2), 260–264.
- Sangrigoli, S., Pallier, C., Argenti, A. M., Ventureyra, V. A., & de Schonen, S. (2005). Reversibility of the other-race effect in face recognition during childhood. *Psychological Science*, 16(6), 440–444.
- Schweinberger, S. R., Zasker, R., Walther, C., Golle, J., Kovacs, G., & Wiese, H. (2010). Young without plastic surgery: Perceptual adaptation to the age of female and male faces. *Vision Research*, 50(23), 2570–2576.
- Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science*, 13(5), 402–409.
- Shafai, F., & Oruc, I. (2018). Qualitatively similar processing for own- and other-race faces: Evidence from efficiency and equivalent input noise. *Vision Research*, 143, 58–65.
- Shahangian, K., & Oruc, I. (2014). Looking at a blurry old family photo? Zoom out! *Perception*, 43(1), 90–98.
- Sugden, N. A., Mohamed-Ali, M. I., & Moulson, M. C. (2014). I spy with my little eye: Typical, daily exposure to faces documented from a first-person infant perspective. *Developmental Psychobiology*, 56(2), 249–261.
- Sugita, Y. (2008). Face perception in monkeys reared with no exposure to faces. *Proceedings of the National Academy of Sciences USA*, 105(1), 394–398.
- Sun, G., Song, L., Bentin, S., Yang, Y., & Zhao, L. (2013). Visual search for faces by race: A cross-race study. *Vision Research*, 89, 39–46.
- Sun, G., Zhang, G., Yang, Y., Bentin, S., & Zhao, L. (2014). Mapping the time course of other-race face classification advantage: A cross-race ERP study. *Brain Topography*, 27(5), 663–671.
- Sussman, A. B., Petkova, K., & Todorov, A. (2013). Competence ratings in US predict presidential election outcomes in Bulgaria. *Journal of Experimental Social Psychology*, 49, 771–775.
- Sutherland, C. A., Oldmeadow, J. A., Santos, I. M., Towler, J., Burt, D. M., & Young, A. W. (2013). Social inferences from faces: Ambient images generate a three-dimensional model. *Cognition*, 127(1), 105–118.
- Sutherland, C. A., Rowley, L. E., Amoaku, U. T., Daguzan, E., Kidd-Rossiter, K. A., Maceviciute, U., & Young, A. W. (2015). Personality judgments from everyday images of faces. *Frontiers in Psychology*, 6, 1616.
- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology: A, Human Experimental Psychology*, 46(2), 225–245.
- Tanaka, J. W., & Gauthier, I. (1997). Expertise in object and face categorization. *The Psychology of Learning and Motivation*, 36, 83–125.
- Tanaka, J. W., Kiefer, M., & Bukach, C. M. (2004). A holistic account of the own-race effect in face recognition: Evidence from a cross-cultural study. *Cognition*, 93(1), B1–9.
- Tardif, J., Morin Duchesne, X., Cohan, S., Royer, J., Blais, C., Fiset, D., ... Gosselin, F. (2019). Use of face information varies systematically from developmental prosopagnosics to super-recognizers. *Psychological Science*, 30(2), 300–308.
- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623–1626.
- Todorov, A., Pakrashi, M., & Oosterhof, N. N. (2009). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*, 27, 813–833.
- Wang, H., Han, C., Hahn, A. C., Fasolt, V., Morrison, D. K., Holzleitner, I. J., ... Jones, B. C. (2019). A data-driven study of Chinese participants' social judgments of Chinese faces. *PLoS One*, 14(1), e0210315.
- Webster, M. A. (2015). Visual adaptation. *Annual Review of Vision Science*, 1, 547–567.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557–561.
- Webster, M. A., & MacLeod, D. I. (2011). Visual adaptation and face perception. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 366(1571), 1702–1725.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6), 598–604.
- Willenbockel, V., Fiset, D., Chauvin, A., Blais, C., Arguin, M., Tanaka, J. W., ... Gosselin, F. (2010). Does face inversion change spatial frequency tuning? *Journal of Experimental Psychology: Human Perception and Performance*, 36(1), 122–135.
- Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science*, 17(7), 592–598.
- Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., ... Duchaine, B. (2010). Human face recognition ability is specific and highly heritable. *Proceedings of the National Academy of Sciences USA*, 107(11), 5238–5241.
- Yan, X., Andrews, T. J., Jenkins, R., & Young, A. W. (2016). Cross-cultural differences and similarities underlying other-race effects for facial identity and expression. *The Quarterly Journal of Experimental Psychology*, 69(7), 1247–1254.
- Yang, H., Shen, J., Chen, J., & Fang, F. (2011). Face adaptation improves gender discrimination. *Vision Research*, 51(1), 105–110.
- Yang, N., Shafai, F., & Oruc, I. (2014). Size determines whether specialized expert processes are engaged for recognition of faces. *Journal of Vision*, 14(8), 17.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81, 141–145.
- Young, A. W., & Burton, A. M. (2017). Recognizing faces. *Current Directions in Psychological Science*, 26(3), 212–217.
- Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, 16(6), 747–759.
- Zhao, L., & Bentin, S. (2008). Own- and other-race categorization of faces by race, gender, and age. *Psychonomic Bulletin & Review*, 15(6), 1093–1099.