Journal of Experimental Psychology: Human Perception and Performance

Real-World Interattribute Distances Lead to Inefficient Face Gender Categorization

Nicolas Dupuis-Roy, Daniel Fiset, Kim Dufresne, Laurent Caplette, and Frédéric Gosselin

Online First Publication, June 16, 2014. http://dx.doi.org/10.1037/a0037066

Journal of Experimental Psychology:
Human Perception and Performance

0096-1523/14/$12.00 http://dx.doi.org/10.1037/a0037066

Real-World Interattribute Distances Lead to Inefficient Face Gender Categorization

Nicolas Dupuis-Roy
Université de Montréal

Daniel Fiset
Université du Québec en Outaouais

Kim Dufresne, Laurent Caplette, and Frédéric Gosselin
Université de Montréal

The processing of interattribute distances is believed to be critical for upright face categorization. A recent study by Taschereau-Dumouchel, Rossion, Schyns, and Gosselin (2010) challenged this idea by showing that participants were nearly at chance when asked to identify faces on the sole basis of real-world interattribute distances, while they were nearly perfect when all other facial cues were shown. However, it remains possible that humans are highly tuned to interattribute distances but that the information conveyed by these cues is scarce. We tested this hypothesis by contrasting the efficiencies—a measure of performance that factors out task difficulty—of 60 observers in 6 face gender categorization tasks. Our main finding is that efficiencies for faces that varied only in terms of their interattribute distances were an order of magnitude lower than efficiencies for faces that varied in all respects, except their interattribute distances, or in all respects. These results provide a definitive blow to the idea that real-world interattribute distances are critical for upright face processing.

Keywords: face gender categorization, configural processing, efficiency, color facial cues, achromatic facial cues

Supplemental materials: http://dx.doi.org/10.1037/a0037066.supp

Francis Galton was one of the first to propose that spatial relations between the main facial attributes (i.e., nose, mouth, eyes, and eyebrows) was a reliable distant-invariant cue that could be used for face recognition (see Galton, 1879). Galton even advocated for the use of relational cues as a simple system to certify the identity of a person—just as he did for fingerprints (Galton, 1888). Nearly a century later, Haig (1984) investigated the impact of single attribute displacement on the recognition of unfamiliar faces. He found that the sensitivity to some displacement (e.g., mouth-up) approached the limit of visual acuity and concluded that humans are highly tuned to spatial relations in faces.

Studies of face inversion have also contributed to the idea that relative distances between attributes are fundamental for face recognition. Faces rotated by 180° in the picture-plane lead to lower recognition performances and higher response latencies (e.g., Hochberg & Galper, 1967). Because this impaired performance is disproportionately larger for faces than for other mono-oriented objects such as houses (Yin, 1969), researchers have employed face inversion as a tool to isolate what is special about upright face processing. It was discovered that processing of interattribute distances[1] (IADs) is affected by inversion more than by processing of the local shape or surface-based properties of attributes (Barton, Keenan, & Bass, 2001; Freire, Lee, & Symons, 2000; Le Grand, Mondloch, Maurer, & Brent, 2001; Rhodes et al., 2007; Sergent, 1984; see Rossion, 2008, 2009, for recent reviews; see McKone & Yovel, 2009, for opposite claim).

A recent study by Taschereau-Dumouchel, Rossion, Schyns, and Gosselin (2010) challenged the hypothesis that IADs are important for *real-world* face recognition. Taschereau-Dumouchel et al. used a sample of 515 face photographs to estimate the face recognition information available in real-world IADs. They found that the IADs from 86 face stimuli used in 14 previous studies on

Nicolas Dupuis-Roy, Département de Psychologie, Université de Montréal; Daniel Fiset, Département de Psychoéducation et de Psychologie, Université du Québec en Outaouais; Kim Dufresne, Laurent Caplette, and Frédéric Gosselin, Département de Psychologie, Université de Montréal.

Supported by the Social Sciences and Humanities Research Council (Grant #249974) and the Natural Sciences and Engineering Research Council of Canada.

Correspondence concerning this article should be addressed to Nicolas Dupuis-Roy, Département de Psychologie, Université de Montréal, P.O. Box 6128, Succursale Centre-Ville, Montréal, QC H3C 3J7, Canada. E-mail: nicolas@dupuis.ca

[1] By "interattribute distances," we mean relative distances between facial attributes that can be manipulated independently from the shapes of these attributes (e.g., the center of gravity to center of gravity interocular distance; e.g., Haig, 1984; Sergent, 1984; Hosie et al., 1988; Rhodes, Brake, & Atkinson, 1993; Tanaka & Sengco, 1997; Leder & Bruce, 1998, 2000; Freire, Lee, & Symons, 2000; Barton, Keenan, & Bass, 2001; Leder, Candrian, & Huber, 2001; Le Grand, Mondloch, Maurer, & Brent, 2001; Bhatt, Bertin, Hayden, & Reed, 2005; Goffaux, Hault, Michel, Vuong, & Rossion, 2005; Hayden, Bhatt, Reed, Corbly, & Joseph, 2007). This excludes, for example, the nasal-corner-to-nasal-corner interocular distance and the temporal-corner-to-temporal-corner interocular distance that cannot be manipulated jointly and independently from attribute size.

IAD had exaggerated this information by 376% compared with real-world IADs. Furthermore, they showed that when their participants resolved a matching-to-sample (ABX) face identification task solely on the basis of real-world IADs, they performed poorly (< 65%) across a broad range of viewing distances. In contrast, recognition was almost perfect when observers recognized faces on the basis of real-world information other than interattribute distances such as attribute shapes and skin properties.

The Taschereau-Dumouchel et al. (2010) study results raise serious doubts about the importance of IADs for face recognition, but they fail to provide a definitive blow to the idea. A low performance with real-world IADs could be due to two possibly interacting causes: (a) scarce information could be available in real-world IADs to resolve face categorization or (b) observers might be inept at using this real-world IAD information. Therefore, it remains possible that real-world IADs are important to face processing in the sense that, although little real-world IAD information could be available, observers may exploit a high proportion of this information. In the present study, we tested this hypothesis for the first time in the context of a face gender categorization task, one of the most biologically relevant facial judgments.

To test our hypothesis, we used the efficiency measure, that is, a performance index that factors out task difficulty (Tanner & Birdsall, 1958; Tjan, Braje, Legge, & Kersten, 1995). Efficiency can be expressed as the ratio of the quantity of noise required by a human observer to reach a given performance to the quantity of noise required by the ideal observer to reach the same performance. The ideal observer is a mathematical model that uses optimally all of the information available for the task at hand. Thus, the ideal observer provides a benchmark for the highest possible performance, and the efficiency measure offers a grasp on the ability of humans to exploit the available information. Within this framework, the hypothesized importance of IADs for upright face gender processing should translate into high efficiency. In other words, human observers should exploit the most from the real-world IAD information. The corollary of this hypothesis is that the efficiency for IADs should be relatively high compared with other face gender cues. This was examined by contrasting the efficiency in three facial-cue conditions: the IAD condition, in which faces varied only in terms of their IADs; the noIAD condition, in which faces varied in all respects except their IADs; and the ALL condition, which featured the original and unmodified faces. Finally, we explored possible interactions between color and efficiencies by contrasting two *color* conditions: the LUM condition, which only exhibited the luminance of faces; and the COL condition, which showed the luminance and chrominance of faces.

Method

Participants

Sixty healthy participants (30 women; $M_{\text{age}} = 22.49$ years, $SD = 3.65$), with normal color vision and normal or corrected-to-normal visual acuity were recruited. Informed consent was obtained prior to the experiment, and a monetary compensation was provided on its completion.

Stimuli

We used the Taschereau-Dumouchel et al. (2010) face database, which comprises 515 White, frontal-view, color portraits (256 females) presenting a neutral expression. These portraits come from many sources: the entire 300-face set of Dupuis-Roy, Fortin, Fiset, and Gosselin (2009), the 146 neutral faces from the Karolinska Directed Emotional Faces (Lundqvist, Flykt, & Öhman, 1998), the 16 neutral faces from Schyns and Oliva (1999), the 10 neutral faces from the CAFE set, six neutral faces from the Ekman and Friesen (1975) set, and 40 additional neutral faces.

The six main attributes of these faces (i.e., eyes, eyebrows, nose, and mouth) were manually segmented; these attributes were aligned across faces using a linear conformal transformation, which preserves IADs.

The experiment comprised six stimulus sets: 3 × 2 ([IADs, noIAD, ALL] × [LUM, COL]). Figure 1A depicts sample stimuli from each of these tasks.

IAD stimuli. Each of the 515 stimuli of this task was created as follows: First, the six facial attributes of all faces in the database were displaced using cubic interpolation to the real-world attribute locations of one face of the database. Second, two gender prototypes were created by dot-averaging all faces belonging to the same gender. Third, the two gender prototypes were blended according to each subject point of subjective equality (see Procedure subsection). Note that we measured one point of subjective equality for COL stimuli and another for LUM stimuli.

noIAD stimuli. The 515-face stimuli used in this task were created by displacing the six facial attributes of every face in the database using cubic interpolation to their average locations.

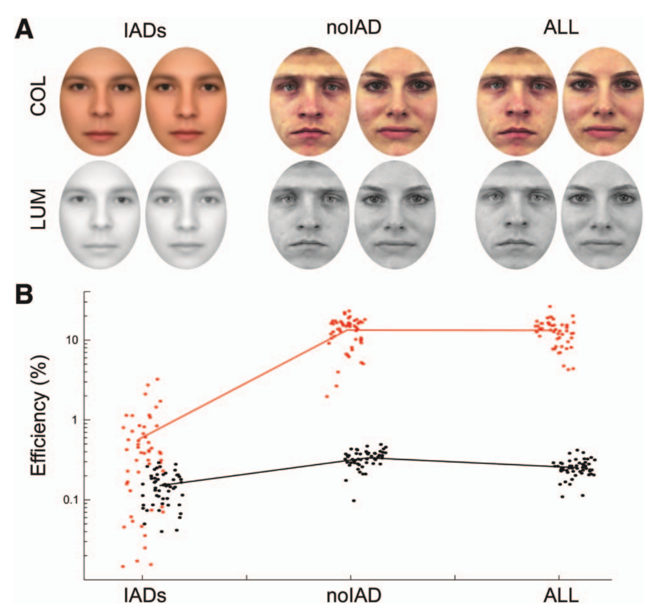

Figure 1. (A) Example of masculine (left), feminine (right), COL (up), and LUM (bottom) face images used in the interattribute distance (IAD), noIAD, and ALL tasks. All images of a given gender correspond to a single face identity taken from the Karolinska Directed Emotional Faces database (Lundqvist, Flykt, & Öhman, 1998). (B) Average efficiency recorded in the six tasks: the color condition is represented by the color of the line (red for COL and black for LUM). Red and black dots represent individual efficiencies.

ALL stimuli. The 515 original faces in the database were used in this task.

Apparatus

The experimental programs were run on four computers in the Matlab (MathWorks Inc., Natick, MA) environment, using functions from the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). The four high-resolution CRT monitors used for stimuli presentation were set to display 1024 $\times$ 768 pixels at a refresh rate of 85 Hz. The relationship between RGB values and luminance levels (measured with a Samsung SyncMaster 753 *df* photometer) was computed for each color channel independently; the three best-fitted gamma functions were then used to create luminance noise and to analyze the luminance content of the stimuli. Participants were seated in a dimly lit room. Their viewing distance was maintained by a chin-rest so that the stimuli interocular width was 1.05° of visual angle.

Signal and Noise Adjustments

The computation of the efficiency measure requires the addition of visual noise to the stimuli to ensure that the ideal observer makes errors. Although using different external noise levels across conditions is not usually recommended, we were forced to do so in the present case due to the large discrepancy between the sensitivities of the IAD tasks and the other tasks. The noise and signal contrast levels were adjusted once and for all in a series of preliminary experiments (see supplemental materials online). We added luminance Gaussian noise to the LUM stimuli while we added luminance Gaussian noise independently to each of the RGB channels of the COL stimuli. The root mean square contrast of the face images and the power of the noise (expressed in contrast units) were, respectively, 27.58 and 7.32×10^{-3} in the IADs−COL condition; 69.73 and 1.08×10^{-3} in the IADs−LUM condition; 26.63 and 107.04×10^{-3} in the noIAD−COL condition; and 61.38 and 2.99×10^{-3} in the noIAD−LUM condition. For the ALL−COL and ALL−LUM conditions, we used the same signal contrast and noise power as in the noIAD−COL and noIAD−LUM conditions, respectively.

Procedure

Subjects completed the experiment on two separate days. On Day 1, they first did two preliminary tasks and, then, two of the six experimental conditions (15–20 min each). On Day 2, they did the remaining experimental conditions. The order of experimental conditions was randomized across participants.

In the preliminary tasks, each participant judged five blends of the gender prototypes (from 80% male−20% female to 20% male−80% female) 20 times on a male-to-female continuous scale. There were two preliminary tasks: one displaying LUM stimuli, and the other COL stimuli. The points of subjective equality were interpolated from fitted psychophysical curves, and were used to create subject-specific stimuli in the IAD tasks (see IAD stimuli subsection).

Each experimental condition consisted of the unique presentation of the 512 stimuli created for that condition in a random order. Each trial was constructed as follows: A black fixation cross was shown for 750 ms on a gray background; a uniform gray screen was displayed for 250 ms; and, finally, a noisy face stimulus revealed through an ellipse was presented at the center of a uniform gray screen and remained there until the participant had indicated its gender by pressing the appropriate keyboard key. No feedback was provided.

Results

Statistical analyses were conducted on 53 subjects (25 men): two subjects were excluded during the experiment because they had complained of visual fatigue, and five others were rejected because their performance was at least 3 *SD* below the group average in at least one experimental condition.

The average performance in the IADs−COL and IADs−LUM conditions—the most difficult ones—was above chance level: respectively 56.63%, $t(52) = 16.51$, $p < .001$ and 56.08%, $t(52) = 15.55$, $p < .001$. Binomial tests performed on individual performance revealed that 34 participants reached significance on the IADs−COL condition and 31 on the IADs−LUM condition (see supplemental materials, Table 1).

Efficiencies

The ideal observer maximizes the a posteriori probability $P(G_i \mid S)$ of selecting the proper gender G_i when a face exemplar S is embedded in Gaussian noise. Tjan et al. (1995) demonstrated that this ideal decision rule is equivalent to selecting the gender category i that maximizes the following term:

$$\sum_{jk} \exp\left(-\frac{1}{2\sigma^2}\|S - T_{ijk}\|^2\right), \tag{1}$$

the weighted sum of the square of the Euclidean distance between the noisy face stimulus (S) and the face template (T_{ijk}—the jth face exemplar of gender i with k independent chromatic channels). The ideal observer went through the same six tasks human observers did, except that each stimulus was presented 10 times at 10 different noise levels. The noise power needed by the ideal observer to reach a given subject's sensitivity index (d') was interpolated from power curves fitted to the average d' recorded at each noise level. The R^2 values of the fitted curves were high, ranging from 96−100% ($M = 98.5\%$).

Finally, an efficiency measure was computed following Tanner and Birdsall's (1958) formulation:

$$\eta = \frac{N_h}{N_i} \tag{2}$$

where N_h is the noise power required by a human observer to reach a given performance, and N_i is the noise power required by the ideal observer to reach the same performance. Figure 1B shows the individual and group efficiencies for the different conditions.

A multivariate analysis of variance indicated two significant main effects—facial-cue conditions, Wilks's $\lambda = .09$, $F(2, 50) = 261.75$, $p < .001$, and color conditions, Wilks's $\lambda = .09$, $F(2, 50) = 494.32$, $p < .001$—and a significant interaction, Wilks's $\lambda = .09$, $F(2, 50) = 257.10$, $p < .001$. The interaction was decomposed into simple effects: the facial-cue conditions contained significant differences for both COL, Wilks's λ = .09, $F(2, 50) = 259.51$, $p < .001$, and LUM, Wilks's λ = .15, $F(2, 50) = 146.44$, $p < .001$. Pairwise comparisons indicated that the efficiencies for noIAD−COL (M = 13.4%) and ALL−COL (M = 13.2%) were significantly higher than the efficiency for IADs−COL (M = 0.276%; $p < .001$, Bonferroni-corrected). However, no significant difference was found between the noIAD−COL and ALL−COL conditions. Similarly, the efficiencies for noIAD−LUM (M = 0.338%) and ALL−LUM (M = 0.255%) were significantly higher than the efficiency for IADs−LUM (M = 0.152%). The efficiency for noIAD−LUM was also significantly higher than for ALL−LUM ($p < .001$).

To estimate the available information in the different conditions, we computed the ratio of contrast energy and threshold noise power required for the ideal observer to reach a d' of 1.60, that is, the average of all observed d' values in the experiment: 1.26×10^5 (IADs−COL), 7.87×10^3 (IADs−LUM), 97.07 (noIAD−COL), 164.29 (noIAD−LUM), 91.51 (ALL−COL), and 151.10 (ALL−LUM).

Discussion

The idea that the processing of spatial relations between the main internal features of faces (e.g., nose, mouth, eyes, eyebrows) is a distinctive and critical aspect of upright face recognition is widely accepted in the face perception literature (e.g., Maurer, Le Grand, & Mondloch, 2002). A recent study by Taschereau-Dumouchel et al. (2010) challenged this idea by showing that participants were nearly at chance level when asked to identify faces on the sole basis of real-world interattribute distances (IADs), and that they were nearly perfect when all other facial cues were shown while IADs were kept constant across faces. However, a low performance with real-world IADs could be due to two possibly interacting causes: (a) the information to resolve the task might be scarce or (b) observers might be inept at using the available information. Therefore, real-world IADs could be important for face processing inasmuch as observers exploit a high proportion of their meager information. We compared the efficiency—a performance index that factors out task difficulty—for face gender categorization in six conditions: 3 facial-cue conditions (IADs [faces varied only in terms of their IADs], noIAD [faces varied in all respects except their IADs], ALL [original faces]) × 2 *color* conditions (LUM [only luminance of faces], COL [luminance and chrominance of faces]).

The average accuracy recorded for the IADs−COL and IADs−LUM conditions was just above chance level, which confirms the main result of Taschereau-Dumouchel et al. (2010). More importantly, however, we found low efficiencies (≈0.2%) in the IAD conditions, indicating that participants did poorly using real-world IADs. In other words, human observers appear unable to exploit a high proportion of the real-world IAD information to categorize facial gender even when it is the only available information. Furthermore, the efficiencies found in the ALL conditions were much higher (i.e., ≈6.8%) and comparable to those reported for other face recognition tasks (e.g., Gold, Bennett, & Sekuler, 1999; Gold et al., 2013; Hammal, Gosselin, & Fortin, 2009). These findings provide a definitive blow to the idea that real-world IADs are critical for face recognition mechanisms in the real world.

We also explored the impact of color on face gender discrimination efficiencies and found significantly higher efficiencies for COL than for LUM faces, regardless of the facial-cue condition. This is consistent with studies on face gender discrimination showing that human observers are sensitive to chromatic differences between men's and women's facial pigmentation, and that they include chroma in their representation of the gender of a face (see Dupuis-Roy et al., 2009; Nestor & Tarr, 2008a, 2008b; Tarr, Kersten, Cheng, & Rossion, 2001, 2002; but see Bruce et al., 1999). In addition to the color face gender information, two mid-level vision mechanisms might benefit from color and, in turn, improve gender discrimination. First, it has been proposed that color can enhance the parsing of a face into its constituent parts using surface segmentation and edge localization (see Yip & Sinha, 2002). Second, it is known that humans' contrast sensitivity to low spatial frequencies is higher for chromatic than achromatic variations (Mullen, 1985). Therefore, it is possible that color enhances the signal in the spatial frequency channels that are the most relevant to face gender categorization. This last hypothesis is also consistent with the findings of Yip and Sinha (2002), who observed that the facilitating role of color in a face individuation task was further amplified as faces were increasingly low-passed (see also Tarr et al., 2001). In any case, our findings provide the first direct demonstration that the gains associated with chromatic variations in the face also lead to higher efficiencies.

We believe that the low efficiency for processing IADs could originate from multiple cue integration processes in the brain. Studies on the integration of multiple cues within a single sensory modality (e.g., Hillis, Watt, Landy, & Banks, 2004; Young, Landy, & Maloney, 1993) and between sensory modalities (e.g., Alais & Burr, 2004; Battaglia, Jacobs, & Aslin, 2003; Ernst & Banks, 2002) have shown that humans tend to weigh multiple cues quasi-optimally as a function of their reliability. If humans also weigh the multiple facial cues quasi-optimally as a function of their reliability (e.g., Gold, Mundy, & Tjan, 2012), then the scarce information conveyed by IADs should be largely underweighted compared with all the information contained in other local cues, such as the 2-dimensional shape of the features, the shape-from-shading cues, the skin texture, and pigmentation cues. Interestingly, multiple cue integration hypothesis might help explain the results of previous studies, which have stated the importance of IADs for upright face recognition. Taschereau-Dumouchel et al. (2010) surveyed 14 IAD studies of face identification (Barton et al., 2001; Bhatt, Bertin, Hayden, & Reed, 2005; Freire et al., 2000; Goffaux, Hault, Michel, Vuong, & Rossion, 2005; Haig, 1984; Hayden, Bhatt, Reed, Corbly, & Joseph, 2007; Hosie et al., 1988; Leder & Bruce, 1998, 2000; Leder, Candrian, & Huber, 2001; Le Grand et al., 2001; Rhodes, Brake, & Atkinson, 1993; Sergent, 1984; Tanaka & Sengco, 1997) and found that most of them had greatly exaggerated the signal in their IADs, up to a factor of 376% compared with real-world IADs. By amplifying the signal in IADs, these studies might have artificially promoted their use. In conclusion, our results provide a definitive blow to the idea that real-world IADs are critical for face recognition mechanisms in the real world.

References

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology, 14,* 257–262. doi:10.1016/j.cub.2004.01.029

Barton, J. J. S., Keenan, J. P., & Bass, T. (2001). Discrimination of spatial relations and features in faces: Effects of inversion and viewing duration. *British Journal of Psychology, 92,* 527–549. doi:10.1348/000712601162329

Battaglia, P. W., Jacobs, R. A., & Aslin, R. N. (2003). Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America, A: Optics, Image Science & Vision, 20,* 1391–1397. doi:10.1364/JOSAA.20.001391

Bhatt, R. S., Bertin, E., Hayden, A., & Reed, A. (2005). Face processing in infancy: Developmental changes in the use of different kinds of relational information. *Child Development, 76,* 169–181. doi:10.1111/j.1467-8624.2005.00837.x

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

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

Dupuis-Roy, N., Fortin, I., Fiset, D., & Gosselin, F. (2009). Uncovering gender discrimination cues in a realistic setting. *Journal of Vision, 9*(2), 10. doi:10.1167/9.2.10

Ekman, P., & Friesen, W. V. (1975). *Unmasking the face: A guide to recognizing emotions from facial clues*. Englewood Cliffs, NJ: Prentice Hall.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433. doi:10.1038/415429a

Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception, 29,* 159–170. doi:10.1068/p3012

Galton, F. (1879). Composite portraits, made by combining those of many different persons into a single, resultant figure. *Journal of the Anthropological Institute, 8,* 132–144.

Galton, F. (1888). Personal identification and description. *Nature, 38,* 173–177. doi:10.1038/038173b0

Goffaux, V., Hault, B., Michel, C., Vuong, Q. C., & Rossion, B. (2005). The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception, 34,* 77–86. doi:10.1068/p5370

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature, 402,* 176–178. doi:10.1038/46027

Gold, J. M., Barker, J. D., Barr, S., Bittner, J. L., Bromfield, W. D., Chu, N., . . . Srinath, A. (2013). The efficiency of dynamic and static facial expression recognition. *Journal of Vision, 13*(5), 23. doi:10.1167/13.5.23

Gold, J. M., Mundy, P. J., & Tjan, B. S. (2012). The perception of a face is no more than the sum of its parts. *Psychological Science, 23*(4), 427–434. doi:10.1177/0956797611427407

Haig, N. D. (1984). The effect of feature displacement on face recognition. *Perception, 13*, 505–512. doi:10.1068/p130505

Hammal, Z., Gosselin, F., & Fortin, I. (2009). How efficient are the recognition of dynamic and static facial expressions? *Journal of Vision, 9*(8), 499. doi:10.1167/9.8.499

Hayden, A., Bhatt, R. S., Reed, A., Corbly, C. R., & Joseph, J. E. (2007). The development of expert face processing: Are infants sensitive to normal differences in second-order relational information? *Journal of Experimental Child Psychology, 97,* 85–98. doi:10.1016/j.jecp.2007.01.004

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision, 4,* 967–992. doi:10.1167/4.12.1

Hochberg, J., & Galper, R. E. (1967). Recognition of faces: I. An exploratory study. *Psychonomic Science, 9,* 619–620. doi:10.3758/BF03327918

Hosie, J. A., Ellis, H. D., & Haig, N. D. (1988). The effect of feature displacement on the perception of well-known faces, *Perception, 17,* 461–474. doi:10.1068/p170461

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

Leder, H., & Bruce, V. (1998). Local and relational aspects of face distinctiveness. *The Quarterly Journal of Experimental Psychology. A: Human Experimental Psychology, 51,* 449–473. doi:10.1080/713755777

Leder, H., & Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *The Quarterly Journal of Experimental Psychology. A: Human Experimental Psychology, 53,* 513–536. doi:10.1080/713755889

Leder, H., Candrian, G., & Huber, O. (2001). Configural features in the context of upright and inverted faces. *Perception, 30,* 73–83. doi:10.1068/p2911

Lundqvist, D., Flykt, A., & Öhman, A. (1998). *The Karolinska Directed Emotional Faces (KDEF)* [CD-ROM]. Stockholm, Sweden: Karolinska Institutet, Department of Clinical Neuroscience, Psychology section.

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

McKone, E., & Yovel, G. (2009). Why does picture-plane inversion sometimes dissociate perception of features and spacing in faces, and sometimes not? Toward a new theory of holistic processing. *Psychonomic Bulletin & Review, 16,* 778–797. doi:10.3758/PBR.16.5.778

Mullen, K. T. (1985). The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *The Journal of Physiology, 359*(1), 381–400.

Nestor, A., & Tarr, M. J. (2008a). Gender recognition of human faces using color. *Psychological Science, 19,* 1242–1246. doi:10.1111/j.1467-9280.2008.02232.x

Nestor, A., & Tarr, M. J. (2008b). The segmental structure of faces and its use in gender recognition. *Journal of Vision, 8*(7), 7.1–7.12. doi:10.1167/8.7.7

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. doi:10.1163/156856897X00366

Rhodes, G., Brake, S., & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition, 47,* 25–57. doi:10.1016/0010-0277(93)90061-Y

Rhodes, G., Hayward, W. G., & Winkler, C. (2007). Expert face coding: Configural and component coding of own-race and other-race faces. *Psychonomic Bulletin and Review, 13,* 499–505.

Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica, 128,* 274–289. doi:10.1016/j.actpsy.2008.02.003

Rossion, B. (2009). Distinguishing the cause and consequence of face inversion: The perceptual field hypothesis. *Acta Psychologica, 132,* 300–312. doi:10.1016/j.actpsy.2009.08.002

Schyns, P. G., & Oliva, A. (1999). Dr Angry and Mr Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition, 69,* 243–265. doi:10.1016/S0010-0277(98)00069-9

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology, 75,* 221–242. doi:10.1111/j.2044-8295.1984.tb01895.x

Tanaka, J. W., & Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory & Cognition, 25,* 583–592. doi:10.3758/BF03211301

Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d′ and η as psychophysical measures. *Journal of the Acoustical Society of America, 30,* 922–928. doi:10.1121/1.1909408

Tarr, M. J., Kersten, D., Cheng, Y., & Rossion, B. (2001). It's Pat! Sexing faces using only red and green. *Journal of Vision, 1*(3), 337. doi:10.1167/1.3.337

Taschereau-Dumouchel, V., Rossion, B., Schyns, P. G., & Gosselin, F. (2010). Interattribute distances do not represent the identity of real world faces. *Frontiers in Psychology, 1,* 159. doi:10.3389/fpsyg.2010.00159

Tjan, B. S., Braje, W. L., Legge, G. E., & Kersten, D. (1995). Human efficiency for recognizing 3-D objects in luminance noise. *Vision Research, 35,* 3053–3069. doi:10.1016/0042-6989(95)00070-G

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics, 33,* 113–120. doi:10.3758/BF03202828

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

Yip, A. W., & Sinha, P. (2002). Contribution of color to face recognition. *Perception, 31,* 995–1003. doi:10.1068/p3376

Young, M. J., Landy, M. S., & Maloney, L. T. (1993). A perturbation analysis of depth perception from combinations of texture and motion cues. *Vision Research, 33,* 2685–2696. doi:10.1016/0042-6989(93)90228-O

Received September 11, 2013
Revision received April 9, 2014
Accepted April 10, 2014 ■