Article

PERCEPTION

Perception
1–17

DOI: 10.1177/0301006618754479
journals.sagepub.com/home/pec
SAGE

The Bandwidth of Diagnostic Horizontal Structure for Face Identification

Matthew V. Pachai
Brain Mind Institute, Ecole Polytechnique Fédérale de Lausanne, Switzerland
Department of Psychology, Neuroscience & Behaviour, McMaster University, Hamilton, ON, Canada

Patrick J. Bennett and Allison B. Sekuler
Department of Psychology, Neuroscience & Behaviour, McMaster University, Hamilton, ON, Canada

Abstract

Horizontally oriented spatial frequency components are a diagnostic source of face identity information, and sensitivity to this information predicts upright identification accuracy and the magnitude of the face-inversion effect. However, the bandwidth at which this information is conveyed, and the extent to which human tuning matches this distribution of information, has yet to be characterized. We designed a 10-alternative forced choice face identification task in which upright or inverted faces were filtered to retain horizontal or vertical structure. We systematically varied the bandwidth of these filters in 10° steps and replaced the orientation components that were removed from the target face with components from the average of all possible faces. This manipulation created patterns that looked like faces but contained diagnostic information in orientation bands unknown to the observer on any given trial. Further, we quantified human performance relative to the actual information content of our face stimuli using an ideal observer with perfect knowledge of the diagnostic band. We found that the most diagnostic information for face identification is conveyed by a narrow band of orientations along the horizontal meridian, whereas human observers use information from a wide range of orientations.

Keywords

face identification, face-inversion effect, ideal observer, horizontal selectivity

Date received: 23 August 2017; accepted: 18 December 2017

Corresponding author:
Matthew V. Pachai, 1032 Sherman Health Research Centre, 4700 Keele St, Toronto, ON, Canada, M3J 1P3.
Email: mattpachai@gmail.com

Introduction

A fundamental challenge in the study of face perception is to characterize the information extracted from face stimuli during a variety of tasks including identity recognition. In general, studies suggest that normal human observers use a small amount of the available stimulus information to identify faces. For example, experiments using a variety of techniques have shown that observers focus primarily on the eyes when identifying faces, even though information is available at other spatial locations (Gold, Sekuler, & Bennett, 2004; Gosselin & Schyns, 2001; Peterson & Eckstein, 2012; Sekuler, Gaspar, Gold, Bennett, 2004; Yarbus, 1967). Furthermore, although diagnostic information is available at many spatial scales, several studies suggest that face identification is based on structure at a limited range of spatial frequencies (Gaspar, Sekuler, & Bennett, 2008; Gold, Bennett, & Sekuler, 1999; Näsänen, 1999; Willenbockel et al., 2010).

The constraints on face identification may be the result of perceptual learning, whereby perceptual processing becomes increasingly reliant on the most diagnostic information for the task at hand (Gauthier & Tarr, 1997; Gold et al., 2004; Heisz & Shore, 2008; LeGrand, Mondloch, Maurer, & Brent, 2004; Tanaka, 2001). Indeed, specialized face processing mechanisms, including the selective processing of diagnostic information, appear to continue maturing well into adulthood (Goffaux, Poncin, & Schiltz, 2015; Susilo, Germine, & Duchaine, 2013). The face inversion effect is a particularly compelling example of perceptual learning. That face identification is severely impaired by picture-plane inversion has been known for decades (Valentine, 1988; Yin, 1969). This phenomenon is highly robust and appears across a variety of face-related tasks (Diamond & Carey, 1986; Farah, Wilson, Drain, & Tanaka, 1998; Tanaka & Farah, 1993). Importantly, picture-plane inversion has no effect on the stimulus information available to discriminate among identities, and therefore any change in performance following face inversion must reflect differential processing on the part of the observer. One oft-assumed view is that inversion impairs the holistic processing of internal facial features, resulting in a qualitative shift of processing strategy (Rossion, 2008). Alternatively, the face inversion effect could simply result from qualitatively similar, but less efficient, processing of the available information (Gaspar, Bennett, & Sekuler, 2008). Results from Sekuler et al. (2004) support the latter view: They find that the eye region was used in identification of both upright and inverted faces, but less efficiently when the faces were inverted. However, these results do not reveal the cause of this inefficiency, and other studies have shown that factors such as spatial frequency processing are similar across face orientations (Gaspar, Sekuler, et al., 2008; Willenbockel et al., 2010).

However, recent experiments point to another source of information related to face perception, namely, the selective use of horizontal structure. Dakin and Watt (2009) first noted that horizontal structure is particularly diagnostic for face identification, is unaffected by manipulations that do not commonly affect face identification (e.g., pose variance), and is affected by manipulations that do commonly impair face identification (e.g., contrast reversal). This diagnostic structure also has been linked to the middle spatial frequency band preferred by human observers during face identification tasks (Goffaux,Van Zon, & Schiltz, 2011), and shown to be necessary for a number of robust phenomena such as identity aftereffects and integrative processing (Goffaux & Dakin, 2010). In addition, the extent to which observers selectively utilize horizontal structure during an identification task is the first measure that has been shown to be significantly correlated with both individuals' overall face identification ability and the size of the face inversion effect (Pachai, Sekuler, & Bennett, 2013).

Although several studies have demonstrated the importance of observers' selective use of horizontal orientation structure for face identification, the orientation tuning of this horizontal bias has not been well quantified. In their original study of orientation-filtered faces, Dakin and Watt (2009) used Gaussian-shaped filters with a standard deviation of 23° and subsequent studies followed suit (Huynh & Balas, 2014; Obermeyer, Kolling, Schaich, & Knopf, 2012; Yu & Chung, 2011), with the exception of Goffaux and colleagues (Goffaux & Dakin, 2010; Goffaux et al., 2011) who used a bandwidth of 14° to more closely match the orientation tuning properties of neurons in V1 (Ringach, Shapley, & Hawken, 2002). These studies also reported variations in human performance without evaluating how these variations could be anticipated as a result of the informational manipulation. Pachai et al. (2013) partially addressed these concerns using an ideal observer analysis to quantify the available identification information at different orientations and found that the horizontal band contained the most diagnostic information, but their filter bandwidth also was a relatively large 22.5°, providing only a coarse analysis of orientation tuning. Most recently, Goffaux and Greenwood (2016) used more narrow orientation filters and reported tuning of horizontal selectivity for human observers during upright face discrimination of approximately 25°, but did not compare this result with that of an ideal observer. Comparison with an ideal observer is an important step because it reveals whether the horizontal selectivity exhibited by human observers matches the distribution of information in face stimuli, clarifying the extent to which such selectivity is an optimal strategy for a given task.

In this study, we sought to quantify more precisely the effects of filter bandwidth on horizontal selectivity during face identification. To that end, faces were filtered to retain horizontal or vertical orientation structure using ideal (sharp-edged) orientation filters with bandwidths ranging from 10° to 180° in 10° steps. Furthermore, to better measure the selective processing employed by human observers when identifying intact faces, we replaced the orientation components removed from the target face with components from the average of all 10 possible faces. This procedure ensured that the diagnostic band was unknown to the observer on a given trial and produced faces that appeared intact despite the filtering manipulation. Finally, we compared the results of our human observers with several simulated observers.

Methods

Observers

A total of 32 naive observers (16 men, aged 19–33 years, $M = 22$ years) participated in the experiment. All observers had normal or corrected-to-normal Snellen acuity, as measured using short standardized tests, and were paid $10/hour or given course credit for their participation. All experimental protocols were approved by the McMaster University Research Ethics Board, and informed consent was collected prior to the experiment.

Stimuli

Stimuli were generated using an Apple Macintosh G4 computer with MATLAB and the Psychophysics and Video Toolboxes (Brainard, 1997; Pelli, 1997), and were presented on a 21-inch Apple Studio display with viewable size of 40×30 cm, a resolution of 1280×960 pixels (32 pixels/cm), and a frame rate of 85 Hz. Average luminance was held constant at 31 cd/m^2 throughout the experiment. The face images were generated using front-facing, digitized photographs of five male and five female models with no facial hair, eye glasses,

or visible piercings. These photographs were cropped using a 198 × 140 pixel oval window, and were centered in a 256 × 256 matrix. At the viewing distance of 60 cm, these stimulus matrices subtended 7.6° × 7.6° and the faces subtended 5.9° × 4.2°. To control for differences in detectability between faces, their amplitude spectra were equalized. Specifically, the Fourier transform of each face was computed, the absolute value of the power at each spatial frequency and orientation was averaged across faces, and the phase spectrum of each face was combined with this average amplitude spectrum before computing inverse Fourier transforms to return each face to the spatial domain. This manipulation produced stimuli with equal detectability that differed only in their phase structure. For more details on stimulus generation, see Gold et al. (1999).

The orientation information available to observers during the identification task was manipulated by filtering the stimuli in the spatial frequency domain and combining the filtered target face with a standard face that was created by averaging the 10 possible target faces (Figure 1). On each trial, an ideal band-pass orientation filter with a full bandwidth (w) ranging from 10° to 180° was used to remove orientation information from the average face within $\pm w/2°$ of 0° (horizontal) or 90° (vertical). The orientation information removed from the average face was then replaced by information from a randomly selected target face by filtering the target face to contain only this information and executing a pixelwise sum of the average and target face on each trial. The root mean square (RMS) contrast of the target face was scaled so that the total power of the hybrid image was constant across conditions. Note that only orientation information within the passband of the orientation filter came from the target face, and that a 180° filter yielded a stimulus that was identical to the unfiltered target face. This procedure allowed us to manipulate the band of orientations that conveyed diagnostic stimulus information while

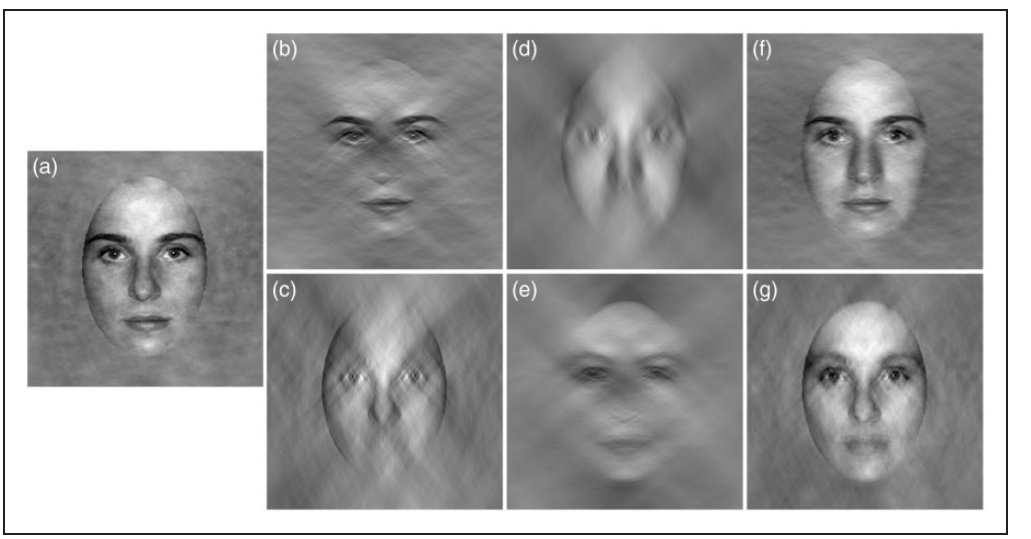

Figure 1. Demonstration of stimulus generation. (a) An unmanipulated target face. The target face filtered to retain horizontal (b) or vertical (c) structure (bandwidth = 90°) The pixelwise average of the 10 possible faces filtered to retain vertical (d) or horizontal (e) structure (bandwidth = 90°). Final stimuli presented in the experiment, drawn from the horizontal (f) or vertical (g) 90° bandwidth conditions. These final stimuli were constructed as the pixelwise sum of the two images to their left, and were equated for RMS contrast after filtering as in the experiment. RMS = root mean square.

ensuring that stimuli in all conditions looked like faces, and contained stimulus-typical power in all orientation bands. Throughout this article, we refer to filters that isolate particular bands of orientation information as isolating those components in the Fourier domain. The final face stimuli, which had an RMS contrast of 0.5, were embedded in a low-contrast white Gaussian noise (RMS contrast = 0.1) prior to presentation. This noise was added to facilitate eventual comparison of data from human observers with that of an ideal observer (for which external noise is required to obtain a threshold).

Design

The experiment consisted of two blocks of trials: face orientation (upright or inverted) varied across blocks, and filter orientation (horizontal or vertical) and bandwidth (10°–180°) varied within blocks. Each block consisted of 360 trials (2 filter orientations × 18 bandwidths × 10 repetitions). Prior to each block, observers completed 10 practice trials with unfiltered faces at the orientation applicable to the current block. The entire experiment took approximately 45 min to complete. The dependent measure was proportion correct on the 10-alternative forced choice (AFC) identification task.

Procedure

Observers viewed the display binocularly, and a chin/head rest stabilized viewing position. Each trial consisted of a fixation cross at the center of the screen for 500 ms, a 250-ms blank screen, a 250-ms stimulus presentation, another 250-ms blank screen, and finally a response window containing the 10 unfiltered faces that remained available until a response was made (Figure 2). Response screen faces were embedded in 192 × 192 pixel matrices subtending 5.7° × 5.7° visual angle (face size 4.4° × 3.1°) and were always presented at the same picture-plane orientation as the target. Figure 2 demonstrates the sequence of a single trial. Observers selected their response using a mouse click, and feedback was provided on every trial using 600-Hz and 200-Hz tones to indicate correct and incorrect responses, respectively.

Results

Human Observers

All statistical analyses were conducted with R (R Core Team, 2017). Figure 3 plots proportion correct in the 10-AFC identification task as a function of filter bandwidth,

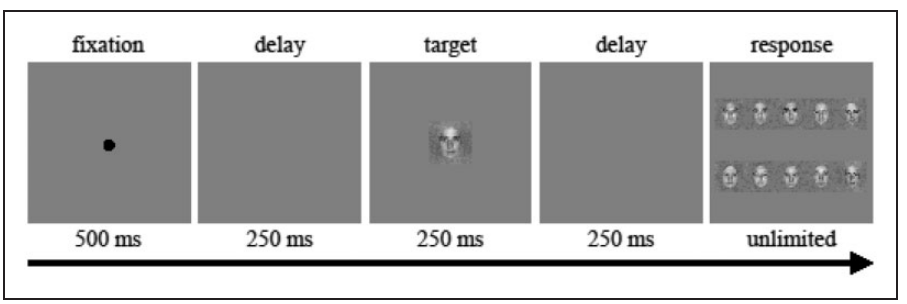

Figure 2. Temporal structure of a single trial. Target faces were presented for 250 ms, and the response screen was presented until response. Exemplars on the response screen were always unfiltered.

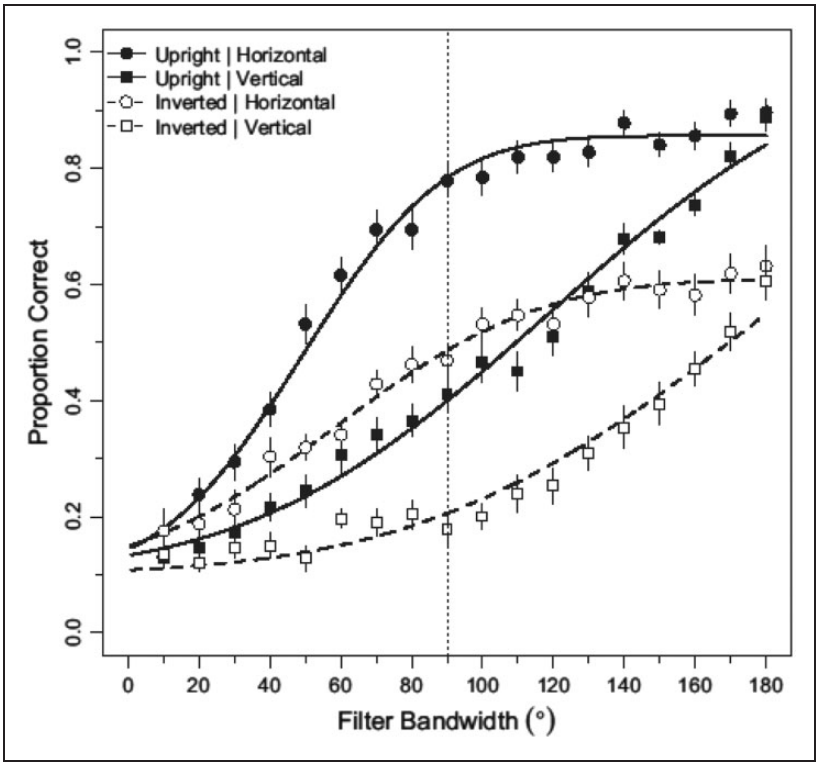

Figure 3. Proportion correct in the 10-AFC identification task plotted as a function of filter bandwidth for each face and filter orientation. A full bandwidth of 90° is indicated by the vertical dotted line. Lines indicate best-fitting psychometric functions fit to the data using generalized linear models with a probit link function with an additional free parameter λ corresponding to the upper asymptote. When the filter bandwidth was 90°, the horizontal and vertical filters isolated the largest, nonoverlapping sets of orientations. As bandwidth increased beyond 90°, the sets of orientations passed by the horizontal and vertical filters became progressively similar. When the filter bandwidth was 180°, the stimuli in the horizontal and vertical filter conditions were identical and corresponded to original (i.e., unfiltered) faces. Error bars represent ±1 SEM. AFC = alternative forced choice; SEM = standard error of the mean.

separately for each filter and face orientation. Lines represent psychometric functions fit to these data using generalized linear models with a probit (i.e., inverse cumulative normal) link function and an additional free parameter, λ, corresponding to the upper asymptote. The lower asymptote was fixed at 0.1. Note that the largest bandwidth at which the horizontal and vertical filters isolate nonoverlapping sets of frequency components is 90°: As bandwidth increases beyond 90°, the frequency components passed by the horizontal and vertical filters become increasingly similar and so, not surprisingly, proportion correct in the two conditions converges to a common point. Visual inspection of Figure 3 indicates that performance improved more rapidly as a function of bandwidth for horizontal than vertical filters, and that this difference was larger for upright relative to inverted faces. The maximum difference between accuracy in the horizontal and vertical filter conditions was observed with filter bandwidths of 90° with differences of 0.37 and 0.29 for upright and inverted faces, respectively.

For the subsequent analyses, we fit psychometric functions separately for each observer and extracted several measures from these fits. First, we computed the mean proportion

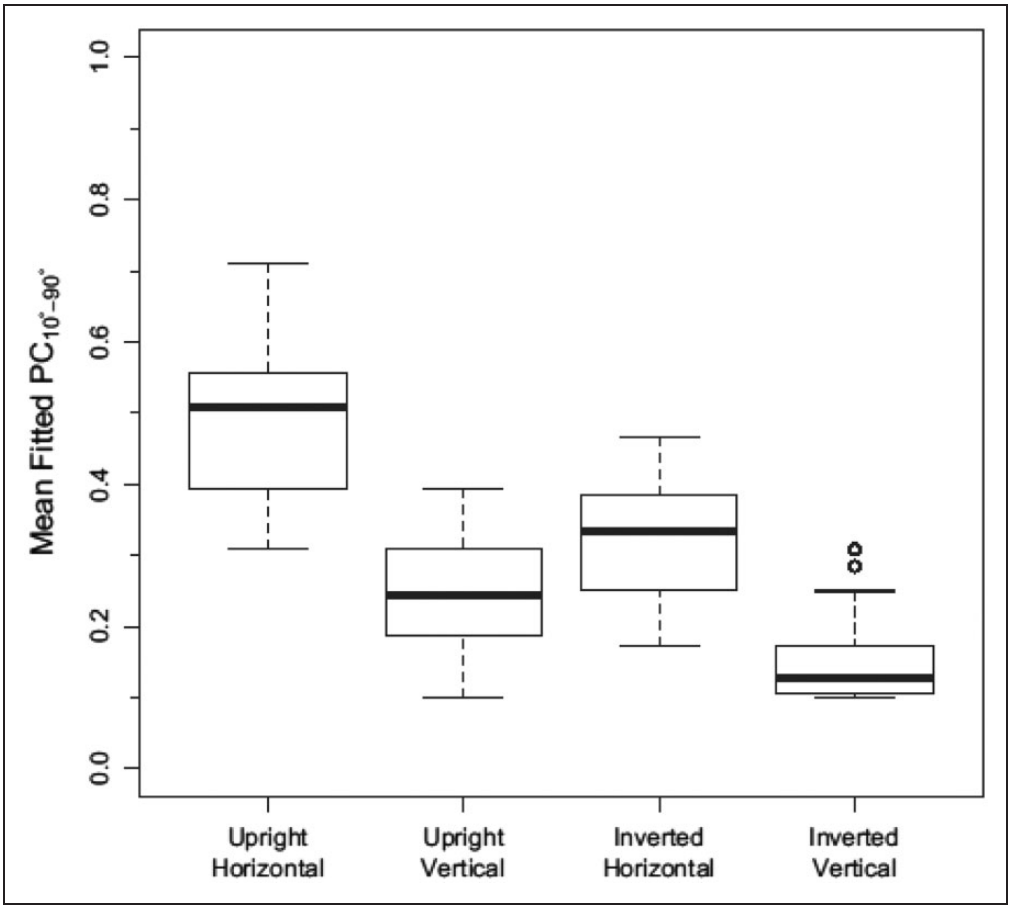

Figure 4. Boxplots representing mean proportion correct from 10° to 90° averaged from best-fitting psychometric functions fit separately to each observer's data for each face and filter orientation. The horizontal line in each box represents the median, whereas the upper and lower edges represent the 75th and 25th percentiles, respectively.

correct for each filter and face orientation from 10° to 90° bandwidth. These values, plotted in Figure 4, quantify differences in performance across the full range of orthogonal bandwidths. A 2 (face orientation) × 2 (filter orientation) analysis of variance (ANOVA) on these values revealed significant main effects of face orientation, $F(1,31) = 112.97$, $p < .0001$, filter orientation $F(1,31) = 160.42$, $p < .0001$, and a significant interaction $F(1,31) = 11.25$, $p = .002$. Together, these results reflect higher performance for upright faces than inverted (main effect of face orientation), higher performance for horizontal structure than vertical (main effect of filter orientation), and more horizontal selectivity for upright faces than inverted (significant interaction). These results are consistent with previous studies (Goffaux & Dakin, 2010; Pachai et al., 2013).

Next, to explore the bandwidth at which horizontal selectivity emerges, we computed separately for each face orientation and filter bandwidth a measure of horizontal selectivity ($PC_{horizontal} - PC_{vertical}$). The resulting values are plotted in Figure 5. Visual inspection of Figure 5 indicates that horizontal selectivity increases as a function of bandwidth up to 90°, after which it begins to decrease for both face orientations as the two orientation filters grow

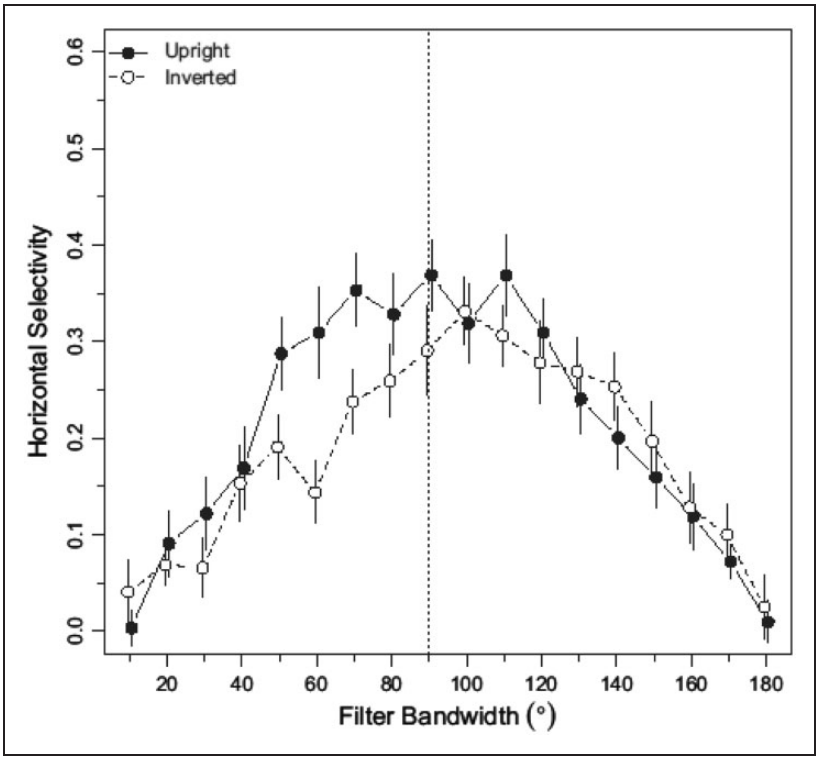

Figure 5. Horizontal selectivity, defined as ($PC_{horizontal} - PC_{vertical}$), plotted as a function of filter bandwidth separately for each face orientation. A full bandwidth of 90° is indicated by the vertical dotted line. Error bars represent ±1 SEM. Points have been offset slightly for visibility. SEM = standard error of the mean.

increasingly similar with bandwidths beyond this point. Note that horizontal selectivity differs as a function of face orientation primarily in the 50°-to-90° bandwidth range, reflecting the fact that horizontal selectivity begins to reach its peak and level off at narrower bandwidths for upright faces than inverted faces. To quantify these patterns, we evaluated the linear and quadratic trends of horizontal selectivity as a function of filter bandwidth from 10° to 90°. The linear trend was significant for both upright, $t(31) = 9.91$, $p < .0001$, and inverted, $t(31) = 6.74$, $p < .0001$ faces, reflecting the linear increase in horizontal selectivity with bandwidth in both conditions. However, the quadratic trend was significant for upright faces, $t(31) = -2.58$, $p = .0148$, but not inverted faces, $t(31) = 0.17$, $p = .8624$, a difference that was itself significant, $t(31) = -2.36$, $p = .0246$. This result reflects the earlier plateau in horizontal selectivity for upright compared with inverted faces.

Simulated Observers

Ideal observer. Our human observers demonstrate a clear reliance on information conveyed by horizontal structure to identify upright, and to a lesser extent, inverted faces. However, we know that there is more diagnostic information available for identification in the horizontal band (Dakin & Watt, 2009; Pachai et al., 2013). To quantify the extent to which our results reflect the horizontal structure inherent to our stimuli, we simulated the

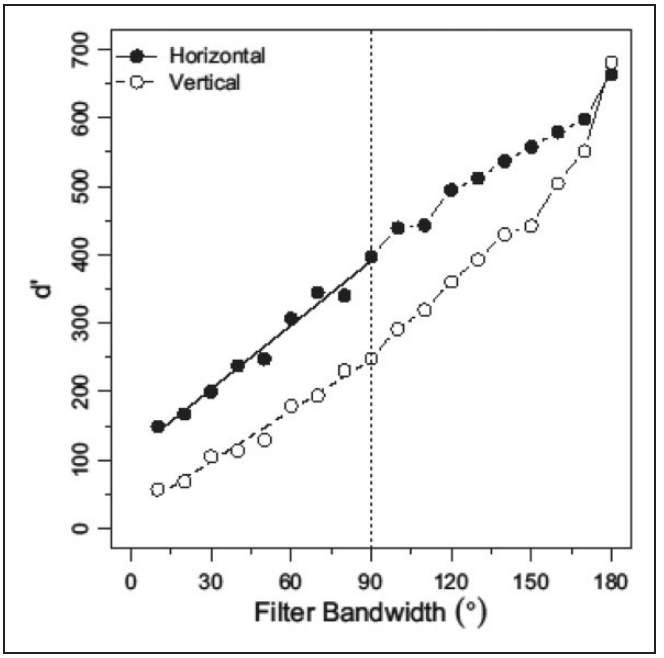

Figure 6. d' as a function of filter bandwidth for a simulated observer using an ideal template matching decision rule. Each point is based on 5,000 simulated trials. Lines represent the best fitting least-squares regression from 10° to 90° bandwidth. A full bandwidth of 90° is indicated by the vertical dotted line.

performance of an ideal observer that optimally utilized the available information. For an object identification task such as ours, an ideal observer correlates the noisy, filtered stimulus with optimally tuned templates of each identity and selects the template yielding the highest correlation (Tjan, Braje, Legge, & Kersten, 1995). This observer isolates precisely the diagnostic orientation band on a given trial and optimally uses the information therein.

The dependent measure of our simulation was proportion correct on the 10-AFC identification task. However, using the same signal-to-noise ratio presented to human observers would result in ceiling performance in all conditions. Therefore, we simulated performance with each filter orientation and bandwidth at seven stimulus RMS contrasts ranging from 0.001 to 0.0025. These contrast values were chosen in pilot simulations to produce neither ceiling nor floor performance in any filter condition. As in the human experiment, Gaussian white noise (RMS contrast = 0.1) was added to the stimulus on every trial. We transformed the resulting proportion correct values into d' using the procedures outlined in Macmillan and Creelman (2004), where each d' value was based on 5,000 simulated trials. Then, separately for each filter orientation and bandwidth, we calculated the best fitting least-squares regression line relating d' to log-transformed RMS contrast (all $R^2 \geq .99$) and extrapolated to the RMS contrast of 0.5 presented to human observers.

Extrapolated d' values for our simulated observer are plotted as a function of bandwidth in Figure 6. Note that the data from human observers (Figures 3 and 5) are qualitatively similar when expressed as proportion correct or d', permitting comparison between human and ideal observers. From the d' measured for the ideal observer, we calculated, separately for each filter orientation, the best fitting least-squares regression lines relating d' to filter bandwidth from 10° to 90° bandwidth (both $R^2 \geq .98$). This analysis revealed

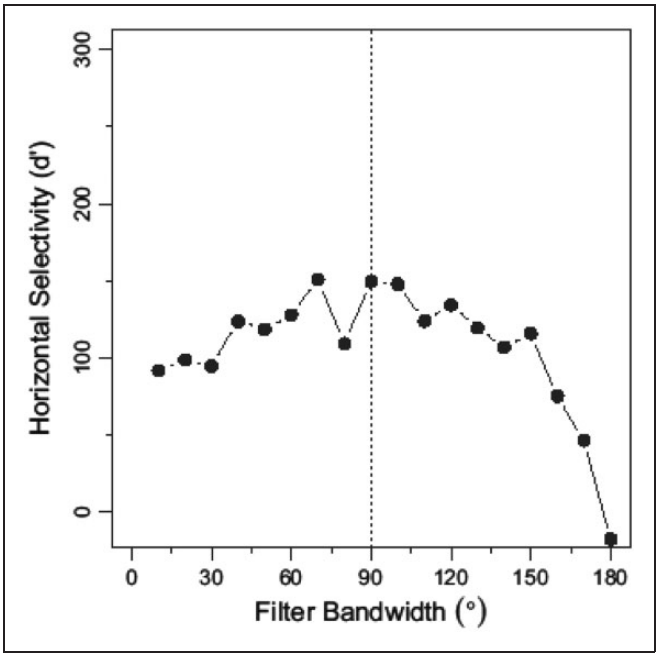

Figure 7. Horizontal selectivity for the ideal observer, calculated as $d'_{horizontal} - d'_{vertical}$, plotted as a function of filter bandwidth. A full bandwidth of 90° is indicated by the vertical dotted line.

a comparatively small difference in slopes, $H = 3.12$, $V = 2.49$, and a large difference between intercepts, $H = 109.67$, $V = 22.87$. The intercepts indicate that the differential value of information in the horizontal and vertical bands lies in the narrowest bandwidth range ($< 10°$), while the slopes indicate the informational gain of increasing bandwidth is roughly equal for the two filters, although slightly higher in the horizontal band. To demonstrate this effect more directly, we computed horizontal selectivity for the ideal observer as the difference in d' between horizontal and vertical filters, similar to the human data plotted in Figure 5. The resulting values, plotted in Figure 7, reveal a relatively high and constant horizontal bias from 10° to 90° reflecting the value of horizontal information conveyed at bandwidths narrower than 10°.

Absolute efficiency. We next examined the relationship between the performance of our human observers and the ideal observer. We transformed individual observers' data to d' values, converting accuracy of 0 and 1 to 0.05 and 0.95, respectively, based on the procedures described in Macmillan and Creelman (2004). We then calculated, separately for each observer, efficiency (η) which is equal to $(d'_{human}/d'_{ideal})^2$ (Tanner & Birdsall, 1958). On this measure, a value of 1 indicates that observers are performing optimally, while lower values indicate less optimal performance.

Efficiency is plotted as a function of bandwidth separately for each face and filter orientation in Figure 8. We first note the low overall values of efficiency, common in face-identification tasks (Gaspar, Bennett, et al., 2008; Gold et al., 1999; Pachai et al., 2013). To quantify these results further, we conducted a 2 (face orientation) × 2 (filter orientation) × 9 (filter bandwidth, 10°–90°) repeated-measures ANOVA. Note that the ANOVA was conducted on bandwidths from 10° to 90° because beyond 90°, the filters begin to

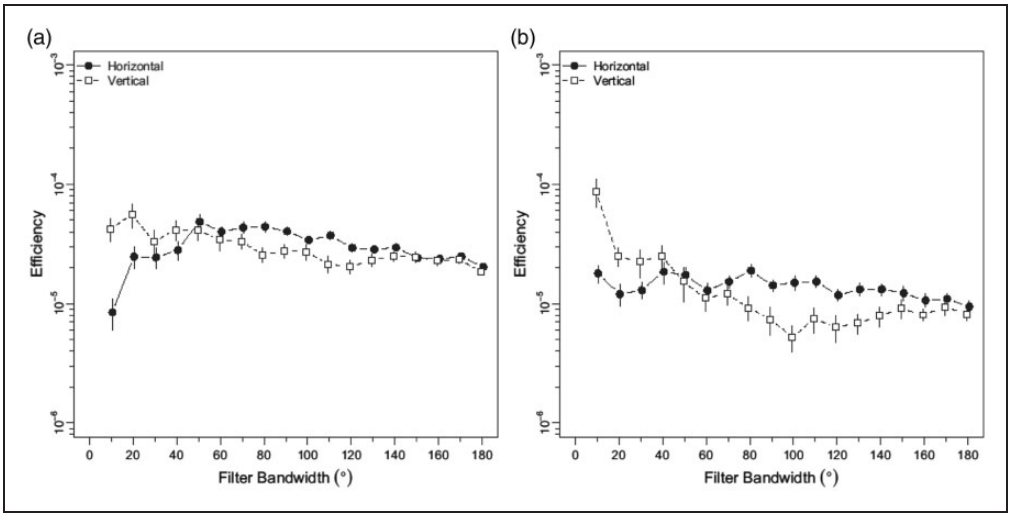

Figure 8. Human efficiency relative to an ideal observer is plotted as a function of filter bandwidth for upright faces (a) and inverted faces (b). See text for details about the calculation of these values. Error bars represent ±1 SEM, and points have been offset slightly for visibility. SEM = standard error of the mean.

Table 1. Summary of a 2 (face orientation) × 2 (filter orientation) × 9 (filter bandwidth) repeated-measures ANOVA conducted on human efficiency with bandwidths from 10° to 90°.

Effect	*df*	*F*	*p*
Face orientation	1,31	41.69	<0001
Filter orientation	1,31	3.94	.0559
Filter bandwidth	8,428	3.11	.0023
Face × filter	1,31	0.79	.3811
Face × bandwidth	8,248	7.28	< .0001
Filter × bandwidth	8,248	11.91	< .0001
Face × filter × bandwidth	8,248	1.36	.2156

ANOVA = analysis of variance.

overlap; however, an ANOVA conducted on the entire data set yielded similar results. The results of the ANOVA, which are summarized in Table 1, are consistent with the trends evident in Figure 8. The significant main effects suggest that efficiency was higher with upright than inverted faces, and with horizontal than vertical filters. The significant face orientation and filter orientation interaction suggests that the effect of filter orientation was larger for inverted faces than upright faces. The significant interaction between filter orientation and bandwidth indicates that the effect of filter orientation depended on bandwidth, with efficiency being higher for vertical filters at narrow bandwidths (i.e., $< \approx 40°$) and higher for horizontal filters at wider bandwidths (e.g., 60°–90°). As no previous study of horizontal selectivity has parametrically manipulated bandwidth as we have, this finding is a novel contribution to the literature. Together, our behavioral and ideal observer results show that human observers perform better with horizontal structure, but they make inefficient use of the extensive information available in that band. In particular,

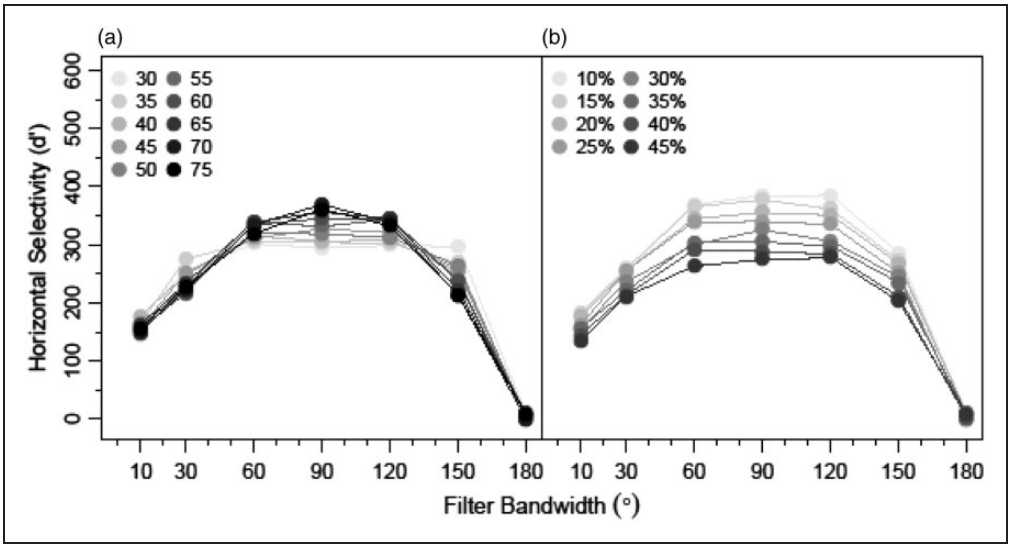

Figure 9. Summary of how the bandwidth and weight parameters affect the horizontal selectivity (calculated as $d'_{horizontal} - d'_{vertical}$) of a fixed-bandwidth simulated observer. (a) The effect of bandwidth, averaged across all stimulated weights. (b) The effect of weight, averaged across all simulated bandwidths. All points are based on 10,000 simulated trials. See text for more details.

observers fail to make efficient use of the information contained within a narrow bandwidth ($\approx 10^\circ$) along the horizontal meridian.

Fixed-bandwidth observers. Finally, we simulated the performance of suboptimal observers who were unable to isolate the diagnostic information on a given trial. These observers optimally utilized the horizontal structure conveyed in one of the 10 bandwidths (30°–75° in 5° steps) while utilizing all other orientation components suboptimally with one of the eight weights (10%–45% in 5% steps). These observers used the same decision rule described for the ideal observer analysis. As with the ideal observer analysis, we simulated the performance of these observers at seven stimulus RMS contrasts ranging from 0.001 to 0.0025 before extrapolating to the RMS contrast shown to human observers. Figure 9 summarizes the effects of bandwidth and weight on simulated selectivity by representing the effect of bandwidth averaged across weights (Figure 9(a)) and the effect of weight averaged across bandwidths (Figure 9(b)). Visual inspection of Figure 9 reveals that increasing bandwidth produces a more peaked horizontal selectivity curve, whereas increasing the weight applied to nonhorizontal bandwidths both lowers and flattens the horizontal selectivity curve because horizontal and nonhorizontal structures are treated more equally.

The 10 bandwidths and eight weights we simulated produced 80 selectivity curves, each of which was then compared with human selectivity to determine the best fitting parameters. To this end, we computed the Euclidean distance between each simulated observer and the average of our human observers, after normalizing both to have a mean of 0 and standard deviation of 1. This normalization ensured comparison of relative curve shape rather than absolute similarity of selectivity values. Our human data were best fit by a simulated observer with a horizontal bandwidth of 70° and a 10% weight applied to all other components. The results for this simulated observer are plotted in Figure 10. Finally, to quantify further the bandwidth utilized by our human observers, we fit mean performance with a bilinear

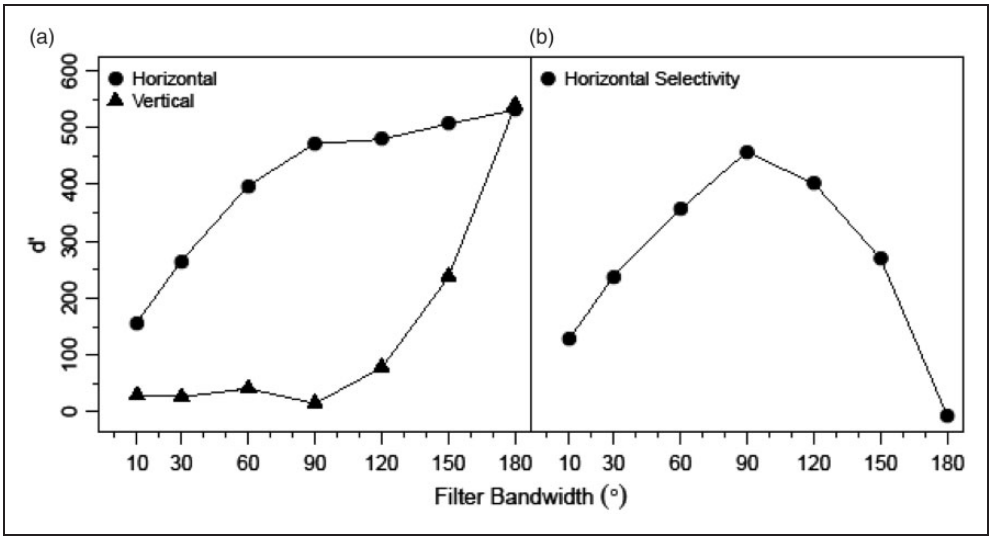

Figure 10. The best-fitting fixed-bandwidth simulation. This observer optimally utilized horizontal structure with a bandwidth of 70°, applying a 10% weight to all other components. (a) d' as a function of bandwidth for horizontal and vertical filters. (b) Horizontal selectivity, calculated as $d'_{horizontal} - d'_{vertical}$. Each point is based on 10,000 trials. See text for more details.

regression using the *segmented* package in *R* (Muggeo, 2008). Although individual human data were noisy, the average data were well fit by this model (all $R^2 \geq .98$). For upright faces, the estimated break point for horizontal filters was 73.6 (95% confidence interval, CI: 67.3–79.8), matching the 70% bandwidth revealed by our Euclidean distance measure. For vertical filters, the estimated breakpoint was 115.8 (95% CI: 99.2–132.5).

Discussion

The results of this experiment are consistent with recent reports that preferential use of horizontal structure supports accurate identification of upright faces, and that this preference is reduced by picture-plane inversion (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Greenwood, 2016; Goffaux et al., 2011; Pachai et al., 2013). Furthermore, we used an ideal observer to quantify the bandwidth at which the horizontal information bias inherent to face stimuli is conveyed, and revealed the relationship between the horizontal selectivity demonstrated by human observers and the availability of this information in the stimulus. We demonstrated that an ideal observer performs better with horizontal structure than vertical structure even at 10° bandwidth, and that the magnitude of this difference increases only slightly as bandwidth increases to 90° (Figure 6). This result suggests that the orientation structure diagnostic for face identification is conveyed primarily by spatial frequency components that are within ±5° of horizontal. In contrast, when target faces were filtered with very narrow bandwidth filters, we found that human observers perform near chance in all conditions (see Figure 3)—presumably because real observers integrate information over a much larger orientation bandwidth—and consequently absolute efficiency was *lower* for horizontal structure than vertical in this range. However, horizontal selectivity increases rapidly with bandwidth (see Figure 5) until the information contained in horizontal and vertical bandwidths converges to a common point. In other words, our

results suggest that human observers integrate information across a wider orientation range than the narrow band in which most of the diagnostic horizontal structure is conveyed. Indeed, when we simulated constrained observers, the pattern of results matched human performance most closely with an estimated horizontal bandwidth of 70°. Recently, Goffaux and Greenwood (2016) quantified the bandwidth of horizontal selectivity for face identification to be narrower, at approximately 25°. It is important to note, however, that in addition to using a different task (same/different vs. 10- AFC) Goffaux and Greenwood removed orientation components from their face stimuli, whereas we replaced these components with nondiagnostic structure. It is possible that presenting face-like information in every orientation band, as we have, encourages use of a wider bandwidth. However, the precise manner in which early orientation signals are combined for face identification, and the manner in which different filtering manipulations may affect this process, remains a question for future research.

As expected, we observed a significant face-inversion effect, with lower overall identification accuracy for inverted faces than for upright faces. As seen in Figure 3, the magnitude of the face-inversion effect increased with orientation bandwidth until about 90°, and then remained relatively constant. Efficiency for both horizontal and vertical structures also was reduced for inverted faces, except at the very smallest orientation bandwidth (Figure 8). Given that perceptual learning has been shown to alter the magnitude of the face-inversion effect (Hussain, Sekuler, & Bennett, 2009; Laguesse, Dormal, Biervoye, Kuefner, & Rossion, 2012), future research should examine whether training increases orientation channel flexibility, specifically targeting the selective processing of horizontal structure, or whether overall efficiency for all orientations simply increases with training.

Selective extraction of horizontal structure has been shown to affect other robust phenomena associated with face perception. For example, Goffaux and Dakin (2010) demonstrated that when horizontal structure is rendered unavailable, identity aftereffects and holistic processing are significantly reduced. Face perception also is associated with neural signals such as the N170, a negative-going peak in the Event-related potential (ERP) that begins to emerge 130 ms after presentation of a face (Rousselet, Pernet, Bennett, & Sekuler, 2008). Jacques, Schiltz, and Goffaux (2014) reported delayed N170 peaks and an attenuated N170 inversion effect when horizontal structure was phase-randomized, which they attributed to impaired recruitment of face-specific neural representations when horizontal information is rendered unavailable. Further, Huynh and Balas (2014) demonstrated that horizontal structure is required for accurate discrimination of happy and sad facial expressions. Interestingly, in the same study, when the faces were rotated 90° such that the horizontal structure became vertical in image-centered coordinates, observers remained tuned to the horizontal information in object-centered coordinates, suggesting some flexibility in the orientation channels used during face perception. Future research should determine more precisely the limits of that flexibility, and the behavioral timeframe over which selectivity for horizontal structure emerges.

Overall, the current results further confirm the important role of horizontal structure for face identification. We demonstrate that a narrow band of orientation components along the horizontal meridian is maximally diagnostic for face identification, but that human observers utilize a much wider range of information when performing such tasks.

Author Note

Portions of this work appeared previously at the Annual Meeting of the Vision Sciences Society and were completed in partial fulfillment of the requirements for a PhD at McMaster University.

Author Matthew V. Pachai is now affiliated to Department of Psychology, York University, Toronto, ON, Canada.
Author Allison B. Sekuler is also affiliated to Rotman Research Institute, Baycrest Health Sciences, Toronto, ON, Canada and Department of Psychology, University of Toronto, Toronto, ON, Canada.

Acknowledgment

We thank Donna Waxman for her assistance with data collection.

Declaration of Conflicting Interests

The author(s) declared no potential conflict of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Financial support for this work was provided by the Natural Sciences and Engineering Research Council of Canada.

References

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

Dakin, S. C., & Watt, R. J. (2009). Biological "bar codes" in human faces. *Journal of Vision*, *9*, 1–10. doi: 10.1167/9.4.2

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107–117. doi: 10.1037/0096-3445.115.2.107

Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. W. (1998). What is "special" about face perception? *Psychological Review*, *105*, 482–498. doi: 10.1037/0033-295X.105.3.482

Gaspar, C. M., Bennett, P. J., & Sekuler, A. B. (2008). The effects of face inversion and contrast-reversal on efficiency and internal noise. *Vision Research*, *48*, 1084–1095. doi: 10.1016/j.visres.2007.12.014

Gaspar, C. M., Sekuler, A. B., & Bennett, P. J. (2008). Spatial frequency tuning of upright and inverted face identification. *Vision Research*, *48*, 2817–2826. doi: 10.1016/j.visres.2008.09.015

Gauthier, I., & Tarr, M. J. (1997). Becoming a "greeble" expert: Exploring mechanisms for face recognition. *Vision Research*, *37*, 1673–1682. doi: 10.1016/S0042-6989(96)00286-6

Goffaux, V., & Dakin, S. C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Perception*, *1*, 1–14. doi: 10.3389/fpsyg.2010.00143

Goffaux, V., & Greenwood, J. A. (2016). The orientation selectivity of face identification. *Scientific Reports*, *6*, 34204. doi: 10.1038/srep34204

Goffaux, V., Poncin, A., & Schiltz, C. (2015). Selectivity of face perception to horizontal information over lifespan (from 6 to 74 year old). *PLoS ONE*, *10*, 1–17. doi: 10.1371/journal.pone.0138812

Goffaux, V., van Zon, J., & Schiltz, C. (2011). The horizontal tuning of face perception relies on the processing of intermediate and high spatial frequencies. *Journal of Vision*, *11*, 1–9. doi: 10.1167/11.10.1

Gold, J. M., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research*, *39*, 3537–3560. doi: 10.1016/S0042-6989(99)00080-2

Gold, J. M., Sekuler, A. B., & Bennett, P. J. (2004). Characterizing perceptual learning with external noise. *Cognitive Science*, *28*, 167–207. doi: 10.1016/j.cogsci.2003.10.005

Gosselin, F., & Schyns, P. G. (2001). Bubbles: A technique to reveal the use of information in recognition tasks. *Vision Research*, *41*, 2261–2271. doi: 10.1016/S0042-6989(01)00097-9

Heisz, J. J., & Shore, D. I. (2008). More efficient scanning for familiar faces. *Journal of Vision*, *8*, 1–10. doi: 10.1167/8.1.9

Hussain, Z., Sekuler, A. B., & Bennett, P. J. (2009). Perceptual learning modifies inversion effects for faces and textures. *Vision Research*, *49*, 2273–2284. doi: 10.1016/j.visres.2009.06.014

Huynh, C. M., & Balas, B. (2014). Emotion recognition (sometimes) depends on horizontal orientations. *Attention, Perception & Psychophysics*, *76*, 1381–1392. doi: 10.3758/s13414-014-0669-4

Jacques, C., Schiltz, C., & Goffaux, V. (2014). Face perception is tuned to horizontal orientation in the N170 time window. *Journal of Vision*, *14*, 1–18. doi: 10.1167/14.2.5

Laguesse, R., Dormal, G., Biervoye, A., Kuefner, D., & Rossion, B. (2012). Extensive visual training in adulthood significantly reduces the face inversion effect. *Journal of Vision*, *12*, 1–13. doi: 10.1167/12.10.14

LeGrand, R., Mondloch, C. J., Maurer, D., & Brent, H. P. (2004). Impairment in holistic face impairment early processing following visual deprivation. *Psychological Science*, *15*, 762–768. doi: 10.1111/j.0956-7976.2004.00753.x

Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: A user's guide* 2nd ed. Mahwah, New Jersey: Psychology Press.

Muggeo, V. M. (2008). Segmented: An R package to fit regression models with broken-line relationships. *R News*, *8*, 20–25.

Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, *39*, 3824–3833. doi: 10.1016/S0042-6989(99)00096-6

Obermeyer, S., Kolling, T., Schaich, A., & Knopf, M. (2012). Differences between old and young adults' ability to recognize human faces underlie processing of horizontal information. *Frontiers in Aging Neuroscience*, *4*, 3. doi: 10.3389/fnagi.2012.00003

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*, 1–9. doi: 10.3389/fpsyg.2013.00074

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

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*, *109*, E3314–E3323. doi: 10.1073/pnas.1214269109

R Core Team. (2017). *R: A language and environment for statistical computing*. Vienna, Austria: Author.

Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation selectivity in macaque V1: Diversity and laminar dependence. *The Journal of Neuroscience*, *22*, 5639–5651. doi: 20026567

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

Rousselet, G. A., Pernet, C. R., Bennett, P. J., & Sekuler, A. B. (2008). Parametric study of EEG sensitivity to phase noise during face processing. *BMC Neuroscience*, *9*, 98. doi: 10.1186/1471-2202-9-98

Sekuler, A. B., Gaspar, C. M., Gold, J. M., & Bennett, P. J. (2004). Inversion leads to quantitative, not qualitative, changes in face processing. *Current Biology*, *14*, 391–396. doi: 10.1016/j.cub.2004.02.028

Susilo, T., Germine, L., & Duchaine, B. (2013). Face recognition ability matures late: Evidence from individual differences in young adults. *Journal of Experimental Psychology: Human Perception & Performance*, *39*, 1212–1217. doi: 10.1037/a0033469

Tanaka, J. W. (2001). The entry point of face recognition: Evidence for face expertise. *Journal of Experimental Psychology: General*, *130*, 534–543. doi: 10.1037/0096-3445.130.3.534

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Journal of Experimental Psychology*, *46*, 225–245. doi: 10.1080/14640749308401045

Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d' and eta as psychophysical measures. *The Journal of the Acoustical Society of America*, *30*, 922–928. doi: 10.1121/1.1909408

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

Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, *79*, 471–491. doi: 10.1111/j.2044-8295.1988.tb02747.x

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*, 122–135. doi: 10.1037/a0016465

Yarbus, A. L. (1967). *Eye movements during perception of complex objects*. New York, NY: Plenum Press.

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

Yu, D., & Chung, S. T. L. (2011). Critical orientation for face identification in central vision loss. *Optometry and Vision Science*, *88*, 724–732. doi: 10.1097/OPX.0b013e318213933c.