

Emotion recognition (sometimes) depends on horizontal orientations

Carol M. Huynh · Benjamin Balas

Published online: 25 March 2014
© Psychonomic Society, Inc. 2014

Abstract Face recognition depends critically on horizontal orientations (Goffaux & Dakin, *Frontiers in Psychology*, 1(143), 1–14, 2010): Face images that lack horizontal features are harder to recognize than those that have this information preserved. We asked whether facial emotional recognition also exhibits this dependency by asking observers to categorize orientation-filtered happy and sad expressions. Furthermore, we aimed to dissociate image-based orientation energy from object-based orientation by rotating images 90 deg in the picture plane. In our first experiment, we showed that the perception of emotional expression does depend on horizontal orientations, and that object-based orientation constrained performance more than image-based orientation did. In Experiment 2, we showed that mouth openness (i.e., open vs. closed mouths) also influenced the emotion-dependent reliance on horizontal information. Finally, we describe a simple computational analysis that demonstrates that the impact of mouth openness was not predicted by variation in the distribution of orientation energy across horizontal and vertical orientation bands. Overall, our results suggest that emotion recognition largely does depend on horizontal information defined relative to the face, but that this bias is modulated by multiple factors that introduce variation in appearance across and within distinct emotions.

Keywords Face recognition · Emotion recognition · Orientation

Face recognition depends on a restricted range of low-level image features. This includes specific spatial-frequency (SF)

bands (Gold, Bennett, & Sekuler, 1999; Yue, Tjan, & Biederman, 2006) and orientation bands. As is the case for all object categories, different spatial frequencies carry different kinds of visual information about face stimuli (Vuilleumier, Armony, Driver, & Dolan, 2003; Goffaux, Hault, Michel, Vuong, & Rossion, 2005; Goffaux & Rossion, 2006), with lower spatial frequencies carrying more information about coarser features (e.g., the face outline), and higher spatial frequencies carrying information about finer details (e.g., texture features or the appearance of the eyes). Mid-range spatial frequencies (~8–16 cycles per face), however, appear to contribute to face recognition disproportionately, whereas for other object classes it appears that subbands do not contribute disproportionately to recognition (Biederman & Kalocsai, 1997; Collin, 2006). With regard to how orientation subbands contribute to face recognition, Dakin and Watt (2009) demonstrated that horizontal orientations appear to contribute disproportionately to famous face identification. The authors asked observers to identify famous faces (celebrities) and found that observers were about 35 % accurate when the orientation was near the vertical axis, which differed significantly from the 56 % accuracy in performance for orientations near the horizontal axis. Additionally, in a separate computational analysis, the authors discussed the possibility that horizontal structures in face images may be a robust cue for face detection. That is, the typical pattern of horizontally oriented features in the face may be a cue that is not disrupted by changes in view or illumination, and may also reliably distinguish faces from nonfaces. The robustness of these structures following typical environmental manipulations may support invariant recognition in many settings. For example, observers are moderately robust to variation in face illumination and viewpoint (Sinha, Balas, Ostrovsky, & Russell, 2006), possibly because neither of these manipulations typically disrupts the structure of the sequence posited by Dakin and Watt. In contrast, both contrast negation and

C. M. Huynh (✉) · B. Balas
Psychology Department, Center for Visual and Cognitive Neuroscience, North Dakota State University,
1210 Albrecht Blvd., Fargo, ND 58108-6050, USA
e-mail: carol.huynh@my.ndsu.edu

face inversion produce stripes that are highly dissimilar to the original image. The disruption of the face sequence in both these circumstances may be the reason that contrast negation and face inversion both disrupt face recognition so profoundly (Galper, 1970; Yin, 1969).

A number of face-processing phenomena depend critically on horizontal orientations within face images. For example, Goffaux and Dakin (2010) showed that the face inversion effect (i.e., failure to recognize a familiar face when it is presented upside down) was preserved for faces containing only horizontal information, but did not obtain for vertically filtered faces. They presented upright and inverted pairs of faces, cars, and scenes that contained horizontal or vertical information, or both. When presented upright, faces containing horizontal information were processed better than faces containing only vertical information. However, when presented upside down, the horizontal advantage was greatly disrupted, whereas faces containing vertical information remained largely unaffected. Identity aftereffects were also driven by horizontal information. Adapting to one of two faces containing horizontal information (i.e., staring at a face for an extended period of time) affected responses to morphed versions of those same two faces (a shift of the psychometric curve toward the adapting face). Finally, the authors showed that masking horizontal information with visual noise disrupted the ability to match faces across different viewpoints. All three manipulations led the authors to conclude that the horizontal structure provides the most useful information about face identity (Dakin & Watt, 2009; Goffaux & Dakin, 2010).

The amount of information for identification has also been shown to be greatest within the horizontal orientation band. Pachai, Sekuler, and Bennett (2013a, b) found that masking of face images was strongest for noise with orientations at or near 0 deg (i.e., horizontal), and least for orientations at or near 90 deg (i.e., vertical). Furthermore, the authors also calculated absolute efficiency scores and proposed that if observers were using masked face information across all orientation bands equally, then no differences should be observed for faces embedded in noise fields with different orientation energy distributions. However, this was not the case; rather, Pachai et al. (2013a, b) found that observers were actually more sensitive to horizontal information when faces were upright than inverted, suggesting that observers were more efficient at utilizing orientation information in the horizontal band and less efficient at using information in the vertical band. Finally, the authors showed that sensitivity to horizontal information in upright faces correlated significantly with the size of the face inversion effect, suggesting that horizontal face information was used efficiently, but only for faces presented upright rather than inverted.

However, preliminary evidence has suggested that horizontal information is not completely dominant over all other

orientations. Goffaux and Okamoto-Barth (2013) demonstrated that vertical orientation assisted in the processing of gaze information. The authors compared direct with averted gazes by presenting an array of faces filtered to include horizontal, vertical, or a combination of both types of orientation information. By having participants search for a target face consisting of either gaze, the authors found that detection was better for direct than for averted gazes, but only when arrays were composed of vertically filtered faces. This suggests that specific facial regions can be useful for communicating relevant social information, and that these orientation bands carry social cues that are distinct from those that carry useful information for individuation. Indeed, the eyes, in particular (relative to the nose and mouth), carry important horizontally oriented information for individuation (Pachai, Sekuler, & Bennett, 2013a), but clearly also carry important vertically oriented information for gaze perception. These reports also suggest that multiple subsets of orientations may be more critical than others, depending on the region of focus within the face and the specific cues that observers require to complete different perceptual tasks. However, for recognition of whole faces, horizontal information appears to be most important.

Since identification critically depends on horizontal information, but some social cues may depend on a broader or different range of orientations, we chose to investigate how facial emotion recognition depends on orientation information. Bruce and Young's (1986) classic model of face perception proposes dissociable processes for identity and facial expressions (Winston, Henson, Fine-Goulden, & Dolan, 2004; Young, McWeeny, Hay, & Ellis, 1986). Bruce and Young identified several distinct types of information that can be derived from viewing faces. This perceptual process is broken up into stages, with the first being the encoding of structural information, from which abstract descriptions of features are obtained. Following this initial stage, they proposed that expression and identity are analyzed independently from one another by separate systems (the expression analysis and face recognition units). This model has received support from clinical studies of prosopagnosic patients (Duchaine, Parker, & Nakayama, 2003; Palermo et al., 2011) and from behavioral studies of neurotypical individuals. For example, Young, McWeeny, Hay, and Ellis provided evidence for separate processing of identity and emotional expressions by measuring response times in a matching task. They presented pairs of familiar or unfamiliar faces simultaneously and had participants decide whether faces were of the same person (identity-matching) or emotion (expression-matching). According to the Bruce and Young model, recognizing expression does not depend on face recognition units, and so performance should be similar across familiar and unfamiliar faces. In contrast, identity matching should result in faster responses to familiar than to unfamiliar faces, due to the rapid and

automatic operation of face recognition units. The results of Young et al.'s task revealed that for identity matching, response times were indeed faster for familiar than for unfamiliar faces, whereas no differences were observed for matching emotional expressions.

Evidence from neuroimaging studies and visual-adaptation paradigms also supports the possibility that identity and emotion are processed independently. For example, Winston et al. (2004) were able to distinguish between neural representations for emotion and identity processing using an fMRI adaptation paradigm. Behavioral face adaptation paradigms similarly reveal that identity adaptation depends on both an expression-dependent mechanism and an expression-independent mechanism (Fox, Oruç, & Barton, 2008), the latter providing evidence of independent neural processing of facial emotion. Different emotions (happy vs. sad) also appear to be dissociated neurally (Calder, Lawrence, & Young, 2001; Morris, DeGelder, Weiskrantz, & Dolan, 2001), suggesting that not only is emotion processing neurally distinct from identity processing, but that distinct emotions may be processed by distinct mechanisms.

Altogether, different emotional expressions appear to be processed by distinct neuroanatomical structures (Calder et al., 2001; Johnson, 2005). Additionally, they are also largely dissociable from identity. Therefore, we hypothesized that the observed bias for horizontal information in identity recognition might not obtain in a facial emotion recognition task, and that orientation biases may also depend on appearance variability in how those emotions are expressed. Indeed, prior reports suggest that not all emotion categories are equally dependent on the same spatial frequencies or orientations. Happy and sad emotion recognition appear to be supported by low (<8 cycles per face) and high (>32 cycles per face) spatial frequencies, respectively (Kumar & Srinivasan, 2011). Yu, Chai, and Chung (2011) measured performance on the categorization of four facial expressions (anger, fear, happiness, and sadness) using multiple orientation filters (i.e., -60, -30, 0, 30, 60, and 90 deg) and concluded that horizontal information is critical for the recognition of most emotions, with the exception of fear expressions. When the degree of orientation reached near vertical, a bias emerged toward labeling faces as "fearful," suggesting that diagnostic cues for recognizing fear may be embedded within the vertical rather than the horizontal component, or at least may be more equally distributed between the two. This result was also borne out by computational modeling simulations that help to explain the differing orientation biases observed by Yu et al. using a model of visual processing based on multiscale oriented Gabor filters (Li & Cottrell, 2012).

In the present study, we asked participants to categorize happy and sad faces that were filtered to include information that was predominantly vertical, predominantly horizontal, or both. Furthermore, we used picture-plane rotation (0 or 90 deg)

to dissociate image-based from object-based orientation. For instance, when the horizontally filtered face image (stimuli containing predominantly horizontal information) is rotated at a 90-deg angle, the raw visual orientation becomes vertical although information along the horizontal structure of the face remains present. This manipulation thus allowed us to determine the relative contributions of a putative bottom-up bias for horizontal orientations and higher-level biases for particular facial features. We conducted two experiments, using faces expressing genuine emotions (Exp. 1) and faces expressing posed emotions (Exp. 2). This allowed us both to examine an ecologically valid set of emotional faces in one task and to complement this analysis with a controlled set of images that made it possible to control for confounds between emotional expression and specific features (e.g., open mouths) that are present in naturally evoked expressions. We hypothesized that the reliance on horizontal orientations for emotion recognition may depend on the appearance of distinct emotional faces, since particular diagnostic features vary substantially by emotion category. Overall, our results were consistent with this hypothesis, insofar as we found that emotion recognition does largely depend on horizontal orientation information, but that this bias is modulated by factors influencing the appearance of specific emotions, including mouth openness (i.e., open vs. closed). Furthermore, we found that structural orientation relative to the face image, as opposed to raw orientation, was driving performance in our tasks. Finally, we submitted our face images for analysis of the energy content within the horizontal and vertical orientation bands, to determine whether our behavioral effects were driven by the relative amounts of orientation energy in the target bands. Our results revealed that, overall, horizontal orientation energy was consistently greater than vertical orientation energy, but that this effect was not significantly affected by emotional expression or mouth openness. We concluded that the extent to which emotional expressions are recognized with a horizontal orientation bias depends on a number of stimulus factors, suggesting that observers are capable of adopting a flexible strategy for recognition that is not constrained by a front-end horizontal bias.

Experiment 1

In Experiment 1, we investigated whether genuine emotion recognition depended on the horizontal structure of the human face. We also wanted to determine whether image-based orientation or object-based orientation was more relevant to differential performance as a function of orientation energy.

Method

Participants A group of 17 undergraduate students (11 females/six males) from North Dakota State University

participated in this experiment. All participants reported normal or corrected-to-normal vision, provided written informed consent, and received course credit for their participation.

Stimuli Face images of 29 individuals (12 male/17 female) expressing genuine emotions of both happy and sad were taken from the Tarrlab Face Place database (www.face-place.org) and were 250×250 pixels in size. Faces containing certain artifacts (e.g., extensive facial hair) were not chosen—hence, the unequal sample of male and female stimuli. We normalized the images by subtracting the mean luminance value from each image. We filtered these faces in MATLAB 2010A by applying a Fourier transform to each image and multiplying the Fourier energy with either a horizontal or vertical Gaussian filter with a standard deviation of 20 deg. Our stimuli were then created by taking the inverse of the Fourier-transformed image back into the spatial domain (Fig. 1). Following this inverse transformation, all images were readjusted so that the mean luminance and contrast were matched (Dakin & Watt, 2009). In addition to the vertically and horizontally filtered images, we also included a third condition (broadband) composed of faces that contained broadband orientation information (Fig. 1). We had intended this set of images to include only the combination of horizontal and vertical orientation energy from our first two filtered-image conditions, but due to an error in our image-filtering code, these images instead had information from all orientations at low spatial frequencies, and were only restricted to horizontal and vertical orientations at high spatial frequencies. As a result, these control images are effectively composed of orientation energy at all orientations, and therefore primarily allowed us to compare performance with largely unfiltered images to performance in the horizontal and vertical conditions.

Design We used a $2 \times 2 \times 3$ within-subjects design with the factors Emotion (happy, sad), Image Orientation (upright or sideways), and Filter Orientation (vertical, horizontal, broadband). Image orientation was varied in separate blocks, whereas emotion and filter orientation were pseudorandomized within each block. Participants completed a total of 348 trials (174 in the upright condition and 174 in the rotated condition). Block order was counterbalanced, so that half of the participants began with the upright-images and the remaining half with the rotated-images condition.

Procedure Participants viewed the stimuli on a 13-in. MacBook with a 2.4-GHz Intel Core 2 Duo Processor. We recorded participants' responses using an eight-bit USB controller. Stimuli were presented using Psychophysics Toolbox 3.0.10 on a MacOS 10.7.4 system. The participants' task was to label each face according to the expressed emotion (happy/sad). Participants responded by pressing the "B" button on our controller for "sad," and the "A" button for "happy." We asked participants to respond as quickly and as accurately as possible. Each trial began with a fixation cross at the center of a gray screen for 500 ms, followed by a face stimulus that replaced the fixation cross. The face stimulus remained on the screen until either participants had made a response or 2,000 ms had elapsed. Short breaks were offered in between blocks, and the experiment resumed only when participants indicated they were ready.

Results

Sensitivity We computed estimates of sensitivity (d') using hits (the correct classifications of happy faces) and false alarms (the incorrect classifications of sad faces) in each




Fig. 1 Examples of Experiment 1 faces, depicting the same individual filtered vertically, horizontally, and with broadband orientation information

condition. We submitted these sensitivity measures to a 2 (image rotation) \times 3 (filter orientation) repeated measures analysis of variance (ANOVA) and observed a main effect of image rotation, $F(1, 16) = 24.62, p < .0001, \eta^2 = .11$, such that discrimination was better for upright faces ($M = 3.26$) than for faces rotated sideways at a 90-deg angle ($M = 2.87$). We also observed a main effect of filter orientation, $F(2, 32) = 73.40, p < .0001, \eta^2 = .68$, such that discrimination was poorest for vertically filtered faces ($M = 2.43$), followed by horizontally filtered faces ($M = 3.17$) and faces containing broadband information ($M = 3.59$). Bonferroni post-hoc pairwise comparisons revealed that all three of these values differed from one another ($p < .001$). These two factors also interacted, $F(2, 32) = 7.57, p = .002, \eta^2 = .04$, such that image rotation significantly impacted discrimination for vertically filtered faces, but had no effect on horizontally filtered faces and faces containing broadband orientation information (Fig. 2).

Response time A 2 (emotion) \times 2 (image rotation) \times 3 (filter orientation) repeated measures ANOVA of median correct response latencies revealed significant main effects of emotion, $F(1, 16) = 27.9, p < .001$, and filter orientation, $F(2, 32) = 42.74, p < .001$. These effects were driven by longer response latencies for sad (774 ms) than for happy (665 ms) faces, and by slower response latencies for vertically filtered faces (793 ms) than for horizontally filtered faces (700 ms) and faces with broadband orientation information ($M = 665$ ms). We also observed a significant interaction between these two factors, $F(2, 32) = 11.60, p < .001$, such that vertical filtering greatly impaired the recognition of sad faces (Fig. 3), but happy-face latencies did not differ in the horizontal and vertical filtering conditions. Unlike in our analysis of sensitivity, we observed no interaction between image orientation and orientation filter, $F(2, 32) = 1.29, p = .29$.

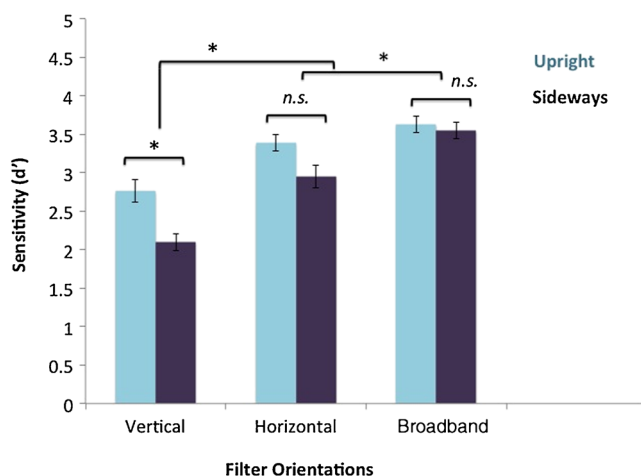


Fig. 2 Sensitivity measures for responses in Experiment 1. We found that image rotation significantly affected discrimination for categorization, especially for vertically filtered faces. Error bars represent ± 1 SEM

Criterion We also ran a 2 (image rotation) \times 3 (filter orientation) repeated measures ANOVA of response bias, C , and found no significant biases in the way participants were responding in any of our conditions, $F < 1$.

Discussion

Our results demonstrate that the facial cues important for emotion recognition were preferentially carried by horizontal orientation. Discrimination ability (as indexed by d' values) was poorer for vertically filtered images than for horizontally filtered images and images with broadband orientation. We also found that discrimination of emotional faces was worse when images were rotated 90 deg in the picture plane, but that sideways rotation did not lead to a “flipped” orientation bias. That is, horizontal orientations relative to the object (not the image) yielded better performance than did vertical orientations. This suggests that the horizontal bias that we have observed is not driven solely by a front-end bias for horizontally tuned neurons in early vision. Were this the case, we would have expected horizontal information relative to the image to be a better predictor of superior performance. Instead, our effect of planar rotation is largely consistent with previous reports on the effect of inversion on face processing (Freire, Lee, & Symons, 2000; Maurer, Grand, & Mondloch, 2002). A change in orientation disrupts the efficiency of face processing—in our case, impacting vertically filtered faces significantly more than faces with horizontal or with both horizontal and vertical energy. This disruption in performance for vertically rather than horizontally filtered faces rotated sideways is inconsistent with previous results showing that rotation impacts both types of information (Jacques, d’Arripe, & Rossion, 2007; Yin, 1969). However, one major difference between the results that we obtained from previous studies is that our picture-plane rotation did not involve a complete 180-deg rotation. Thus, our face images were not completely inverted, but were presented sideways instead. Furthermore, according to Goffaux and Rossion (2007), face inversion does not equally disrupt vertical and horizontal facial information. Although Goffaux and Rossion (2007) did not actually investigate orientation bands, they did find differences between the extractions of different orientation information within facial features. Specifically, they observed the poorest performance for recognizing vertical rather than horizontal facial relations rotated sideways at a 90-deg angle. Again, orientation information was not the target of their investigation; however, the report that performance differences could arise for various facial information rotated within the picture plane might possibly explain the discrepancies between our results and previous findings.

In terms of response latency, our main effect of emotion category is consistent with previous reports (Elfenbein &




Fig. 3 Average response latencies for correct responses in Experiment 1. For both upright and sideways faces, we found that vertically filtered faces were recognized more slowly than our other filtered images, but only when sad faces were recognized. Error bars represent ± 1 SEM

Ambady, 2003; Kirita & Endo, 1995) that happy-face categorization was carried out faster than sad-face categorization. In terms of our initial hypotheses regarding the potential for different emotions to exhibit a differential horizontal bias, we also found that the preference for horizontal information depended on emotion category, since happy-face response latencies revealed a reduced horizontal bias, relative to sad faces.

Together, our sensitivity and response time data suggest that (a) orientation biases for emotion recognition may manifest in an emotion-dependent manner, and (b) structural orientation relative to the face image (not raw orientation on the retina) drives differential performance in our task.

One important limitation of our first experiment, however, was that our use of genuine emotional expressions may have introduced confounding factors that could underlie the interaction we observed between emotion category and filter orientation. Specifically, mouth openness varies substantially in genuine happy and sad faces, and the prevalence of open mouths in happy faces may be the basis of the interaction that we observed here. To examine the emotion dependence of the horizontal orientation bias in more depth, we continued in Experiment 2 by using stimuli from a database of posed emotional expressions that permitted systematic control of mouth openness and emotional expression.

Experiment 2

In Experiment 2, we wished to replicate and extend the results of Experiment 1 using a controlled set of face stimuli within which mouth openness could be manipulated. Specifically, we chose to use a set of posed emotions from the NimStim Face Set (Tottenham et al., 2009), in which the position of the mouth (open vs. closed) was systematically varied in happy and sad emotional expressions.

Method

Participant A group of 21 undergraduate students (11 females/10 males) from North Dakota State University participated in this experiment. All reported normal or corrected-to-normal vision, provided written informed consent, and received course credit for their participation.

Stimuli Twenty-six individual face images (13 male/13 female) expressing posed emotions (26 happy/26 sad) were taken from the NimStim Face Set (Tottenham et al., 2009) and were presented 650×506 pixels in size. Each emotion was presented in two versions (Fig. 4): One expressed each emotion with a closed mouth, and the other with an opened mouth (156 closed mouth/156 open mouth). The process of filtering our stimuli to obtain horizontally filtered and vertically filtered images was identical to that described in Experiment 1. We note, however, that our third condition in this task (horizontal + vertical) was composed of images containing the combination of the horizontal and vertical orientation subbands from the horizontal and vertical conditions. As compared to Experiment 1 (in which energy at all orientations was included at low spatial frequencies), this third set of images contained restricted orientation information at all spatial frequencies, allowing us to more closely examine the additive effects of horizontal and vertical orientation information. To distinguish between the two types of control images used in Experiment 1, we refer to the third orientation condition in this task as “horizontal + vertical,” to contrast these stimuli with the “broadband” stimuli employed in Experiment 1.

Design We used a within-subjects design with the factors Emotion (happy, sad), Mouth Openness (open, closed), Image Orientation (upright or sideways), and Filter Orientation (vertical, horizontal, horizontal + vertical). Participants completed a total of 624 trials, broken up into four blocks—156 upright




Fig. 4 Examples of happy and sad faces with closed and open mouths, used in Experiment 2. Filtering operations were applied to these images in the same manner described for Experiment 1

with closed mouth, 156 upright with open mouth, 156 rotated with closed mouth, and 156 rotated with open mouth—with the factors Emotion and Filter Orientation randomized within each block. The block order was counterbalanced across participants.

Procedure All stimulus display parameters and response collection routines were identical to those described in Experiment 1.

Results

Sensitivity We computed estimates of sensitivity (d') by using hits (correct classifications of happy faces) and false alarms (incorrect classifications of sad faces) in each condition. We submitted these sensitivity measures to a 2 (planar rotation) \times 2 (mouth openness) \times 3 (filter orientation) repeated measures ANOVA and observed significant main effects of image rotation, $F(1, 20) = 22.99$, $p < .001$, $\eta^2 = .09$; mouth

openness, $F(1, 20) = 12.62, p = .002, \eta^2 = .016$; and filter orientation, $F(2, 40) = 86.75, p < .001, \eta^2 = .40$. Discrimination was better for upright faces ($M = 2.98$) than for sideways faces ($M = 2.46$), and for open-mouthed ($M = 2.81$) than for close-mouthed ($M = 2.62$) faces. Discrimination was worse for vertically filtered faces ($M = 1.94$) than for horizontally filtered faces ($M = 3.04$) and faces containing both types of orientation information ($M = 3.17$). Bonferroni-corrected pairwise comparisons revealed that horizontally filtered faces did not differ from faces containing both types of orientation information ($p = .58$)—an effect we had also obtained in Experiment 1, possibly due to the difference in appearance between the “horizontal + vertical” stimuli used in these two tasks. We also observed significant two-way interactions between image rotation and filter orientation, $F(2, 40) = 3.88, p = .029, \eta^2 = .01$, and between mouth openness and filter orientation, $F(2, 40) = 9.12, p = .001, \eta^2 = .04$. As in the results from Experiment 1, vertically filtered faces were greatly affected by image rotation, such that discrimination was worst for vertically filtered faces rotated sideways at a 90-deg angle. Furthermore, discrimination for vertically filtered faces was significantly worse when mouths were closed than when they were open (Fig. 5).

Response time We analyzed the median response latencies for correct responses in all experimental conditions using a 2 (emotion) \times 2 (image rotation) \times 3 (filter orientation) \times 2 (mouth openness) repeated measures ANOVA. This analysis revealed significant main effects of all factors: emotion, $F(1, 20) = 29.03, p < .001$; image orientation, $F(1, 20) = 46.08, p < .001$; mouth openness, $F(1, 20) = 6.67, p = .018$; and orientation filter, $F(2, 40) = 140.36, p < .001$. Participants were faster to respond to happy than to sad faces (688 vs. 735 ms), to upright than to sideways faces (679 vs. 743 ms), and to faces displaying open rather than closed mouths (695 vs. 728 ms). Participants were also slowest to respond to faces

containing vertical information ($M = 794$ ms), which differed significantly from their response latencies to both horizontally filtered faces ($M = 675$ ms) and faces containing both orientations ($M = 665$ ms). We also observed a significant two-way interaction between emotion and mouth openness, $F(1, 20) = 18.71, p < .001$, and also an interaction between emotion and orientation filter, $F(2, 40) = 13.21, p < .001$. These two-way interactions were qualified by a three-way interaction between emotion, mouth openness, and orientation filter, $F(2, 40) = 10.10, p < .001$. In this case, we observed a difference between happy and sad response latencies for open-mouthed, vertically filtered faces that did not hold for the other combinations of mouth openness and filter orientation (Fig. 6).

Criterion We ran a 2 (image rotation) \times 2 (mouth openness) \times 3 (filter orientation) repeated measures ANOVA of response bias and found significant main effects of mouth openness, $F(1, 20) = 16.83, p = .001$, and filter orientation, $F(2, 40) = 55.56, p < .001$. These main effects were qualified by a two-way interaction between mouth openness and filter orientation, $F(2, 40) = 8.58, p < .001$, indicating that for vertically filtered faces, responses were significantly more biased toward responding “happy” when mouths were open than when they were closed (Table 1). We suggest that this criterion shift may have resulted from participants’ tendency to infer that observers were expressing happiness when teeth (which contain vertical edges) are visible.

Discussion

Overall, our results from Experiment 2 largely replicated the main results of Experiment 1 with regard to the roles of image orientation, filter orientation, and emotion category on emotion recognition. We found that a 90-deg rotation negatively impacted performance, but did not induce the “flip” between horizontal and vertical orientations that one would expect if

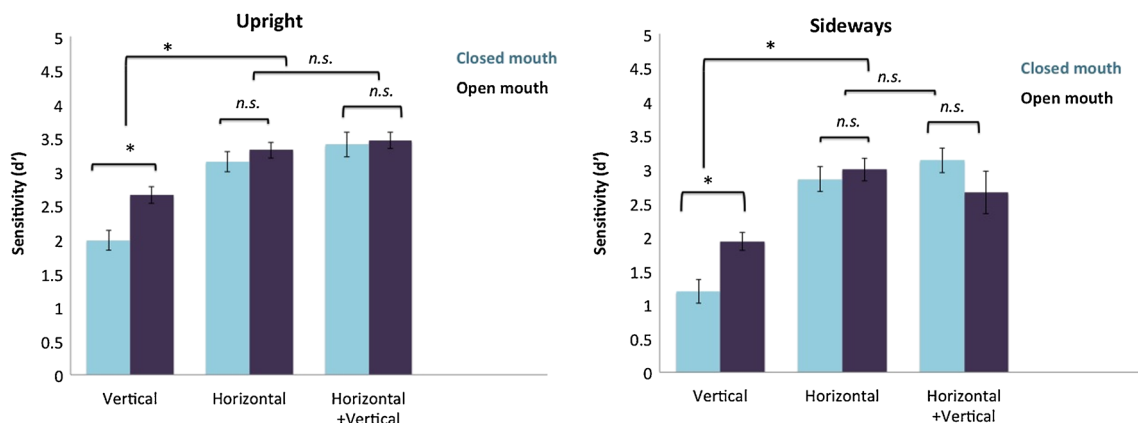


Fig. 5 Sensitivity measures for responses in Experiment 2. We found that vertically filtered faces were greatly affected by sideways rotation, especially for closed-mouth faces. Furthermore, we observed worse

discrimination for vertically filtered faces with closed than with open mouths. Error bars represent ± 1 SEM

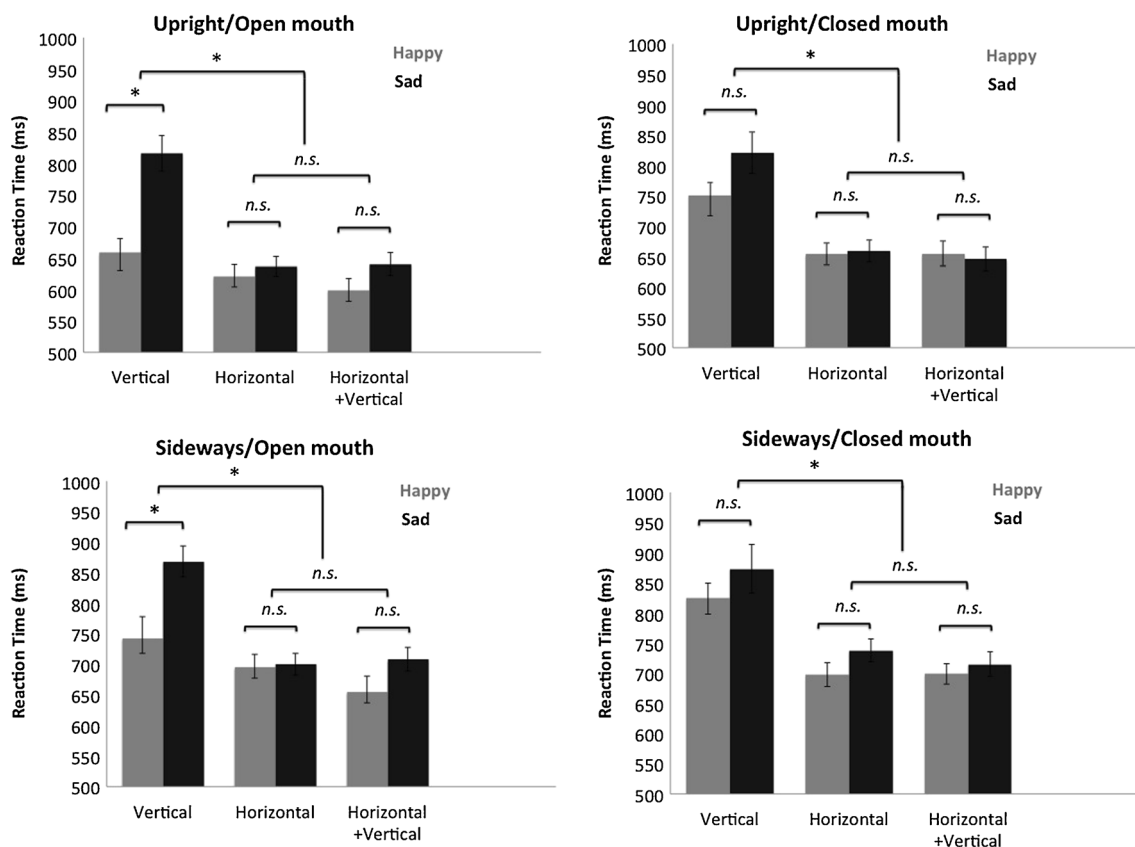


Fig. 6 Average response latencies for correct responses in Experiment 2, for upright and sideways faces with open or closed mouths. Error bars represent ± 1 SEM

image-based information were the relevant factor constraining performance. We also found that horizontally filtered images were recognized more accurately than vertically filtered images, but additionally demonstrated that this depends on the openness of the mouth. Specifically, face expressions with open mouths were more robustly recognized after vertical filtering than faces with closed mouths. This suggests that the effect of emotion that we observed in the response latency data from Experiment 1 may have largely been driven by the confound between mouth openness and emotion category in genuine emotional expressions. With this confound removed in Experiment 2, we could see that the openness of the mouth modulates the magnitude of the horizontal bias. The larger point, however, is that the horizontal bias is impacted by

stimulus appearance: Observers do not exhibit an unwavering bias for horizontal relative to vertical orientations, but exhibit varying levels of bias as a function of stimulus appearance. We will conclude below by discussing the results of both experiments in more detail, discussing the role of low-level image statistics in determining performance in both of our tasks, and suggesting important avenues for further research. Finally, people were significantly biased toward responding “happy” when face expressions displayed open rather than closed mouths. This could be due to the visibility of the teeth, however slight it might be. A possible reason for this bias could be due to the association between teeth and happy expressions (e.g., smiles). Thus, it is possible that sad expressions that exhibited the slightest amount of teeth information may have led observers to assume that the current expression was happy, and to respond accordingly.

Table 1 Response bias, *C*, for Experiment 2

	Upright		Sideways	
	Closed Mouth	Open Mouth	Closed Mouth	Open Mouth
Vertical filter	-.29	-.48	-.25	-.57
Horizontal filter	.03	-.01	-.04	.04
Horizontal + vertical filter	.06	.00	.04	-.07

Experiment 3

Given that we found that the horizontal bias for emotion recognition is affected by stimulus manipulations, a natural question to ask was whether our results were driven largely by the amounts of horizontal versus vertical information in our different categories of images. That is, were observers simply

using more horizontal orientation energy because some images include more of it than others? In principle, more orientation energy in a particular band does not necessarily mean more *diagnostic* information in that band, but it is possible that our behavioral effects might simply reflect the low-level statistics of our emotional faces. To examine this, we computed the power spectrum of our face stimuli in all conditions in order to make comparisons of the summed energies within horizontal and vertical orientation bands.

Method

Stimuli The face stimuli were the same as those used in Experiment 1 and 2.

Design For Experiment 1, we used a 2 (happy, sad) \times 2 (horizontal, vertical) design. For Experiment 2, we created a 2 (happy, sad) \times 2 (open-mouth, closed-mouth) \times 2 (horizontal, vertical) design.

Procedure For the stimuli in each experiment, we submitted the summed energies to repeated measures ANOVAs with Filter Orientation and Emotion as factors, for our Experiment 1 images, and with Filter Orientation, Emotion, and Mouth Openness as factors, for our Experiment 2 images.

Results

For both Experiments 1 and 2, we found that the horizontal structure contained significantly more energy than did the vertical structure in both Experiment 1, $F(1, 28) = 19.87$, $p < .001$, and Experiment 2, $F(1, 26) = 64$, $p < .001$. We did not observe any other main effects or interactions.

Discussion

By itself, our main effect of filter orientation is not surprising, given that faces inherently possess more features that are driven by horizontal content, such as the eyes, eyebrows, and mouth (Dakin & Watt, 2009). Critically, however, we found that this main effect of filter orientation did not interact with any of our other factors, nor did these factors significantly affect the summed energies. Our results thus suggest that the variation in the horizontal bias that we observed behaviorally was not driven solely by the low-level statistics of our stimulus categories (emotion category and mouth openness). We suggest, therefore, that our observers were not simply making use of horizontal information on the basis of the extent to which it dominated our images, but were potentially using a flexible perceptual strategy adapted to the diagnosticity of information in different orientation bands.

General discussion

Overall, our results demonstrate that emotion processing generally depends on horizontal orientations. We observed better discrimination and faster responses to horizontally filtered faces (as compared to vertically filtered faces) in both of our experiments. However, other factors, such as the emotion category displayed by our stimuli (Exps. 1 and 2) and the position of the mouth (Exp. 2), also play critical roles in determining the magnitude of this bias. Overall, we suggest that our results demonstrate that the orientation biases that affect emotion recognition are relatively flexible—the extent to which horizontal orientation is favored is not “locked down” by preferential connections to early vision, and it is malleable in response to variations in stimulus appearance.

To speak to the impact of early orientation processing on emotion recognition, we demonstrated that rotating the image 90 deg in the image plane did not substantially impact the horizontal bias. Instead, the pattern that we observed for both vertical and horizontal orientations, defined relative to the object, were not reversed, even though their image-based orientations were reversed in the sideways condition (i.e., vertical orientation became horizontal, and horizontal became vertical). Orientation biases for emotion recognition are thus defined relative to the face, suggesting that critical features in the face image (rather than raw orientation) dictate performance, but that the specific features that are diagnostic for emotion recognition vary subject to multiple sources of appearance variation. The encoding of face viewpoint, by contrast, has recently been shown to have a substantial image-based orientation component (Balas & Valente, 2012). Balas and Valente observed adaption to the image axis, rather than the object, when participants adapted to upright or sideways (90 deg in the picture plane) faces rotated in depth. Our results are inconsistent with this possibility for mechanisms of emotion recognition. Instead, we suggest that the orientation biases observed for face recognition likely do not result from a purely feed-forward bias for horizontal orientations that propagates downstream to face-sensitive neural loci.

To the best of our knowledge, ours is also the first study to take into account the possible variation of expressions within each emotion category and to examine the impact of this variation on information biases for recognition. Here, we made the distinction between open and closed mouths (Exp. 2) to see whether mouth openness affected recognition across orientation bands, on the basis of our observation in Experiment 1 that these factors were typically confounded in genuine emotional expressions. Our results indicated that mouth openness did play a key role, especially in happy faces. This suggests that the presence of specific diagnostic features in particular variants of some emotional expressions modulates the overall dependence on basic features such as spatial frequency and/or orientation. In our case, the visibility of the teeth may have

provided additional insight into the type of emotion being presented, particularly in the absence of other facial cues. In this study, observers may have taken the presence of vertical lines located within the mouth region as an indication of a happy expression. Therefore, different orientation information becomes more or less important, depending on the kind of features that are available for viewing. In general, this suggests that emotion recognition (and potentially other tasks) may in fact be supported by flexible processes that are not completely constrained by global biases affecting low-level feature extraction, but may instead be constrained by specific diagnostic micro-patterns.

We note, however, that discrimination between emotional faces was consistently best for faces containing both horizontal and vertical orientations (Exp. 2) or broadband orientation information (Exp. 1). This sensitivity advantage for faces with broader orientation information is also coupled with the fastest response during categorization. Thus, having access to a wide range of orientation information is certainly useful for face recognition, indicating that vertical and oblique orientations do contribute (although to a lesser degree) to recognition (Pachai et al., 2013a, b). However, for faces that are limited to a narrow band of orientations, there is clearly an advantage for the presence of horizontal over vertical information for the categorization of facial expressions.

Our results also raise a number of additional intriguing issues for further study. For example, although a 90-deg planar rotation did not affect the orientation biases that we observed here, face inversion (a 180-deg rotation) is known to disrupt the processing of facial information in general (Yin, 1969). In addition, it has been shown that observers encode information in the horizontal band differently for upright than for inverted faces, indicating differences in observer sensitivity to faces rotated in the picture plane (Pachai et al., 2013b) as a function of the orientations that comprise the image. Extending these results to include emotion recognition (and possibly other face recognition tasks) would further clarify how information biases are affected by task demands and the impact of image transformations that are known to disrupt face processing. Determining how the neural response to emotional faces is modulated by orientation filtering would also complement recent results demonstrating that the behavioral effects observed by Dakin and Watt (2009) are also manifested in the responses measured from putative face-sensitive cortical loci (Jacques, Schiltz, Collet, Oever, & Goffaux, 2011). Our demonstration that the position of the mouth affects the relative bias for horizontal orientations over vertical orientations is potentially interesting to examine in this context, especially given the face-sensitive N170 component's robust responses to isolated eyes (Itier, Alain, Sedore, & McIntosh, 2007). Finally, to our knowledge no results have revealed how biases for orientation bands in face patterns emerge during typical development. Understanding how emotion recognition

proceeds as a function of specific emotion categories (and confusions between these) and the biases for spatial frequency and orientation bands could yield important insights into how the statistical regularities that define diagnostic image features are learned and applied as the visual system gains experience with complex and socially relevant patterns.

Author note This research was supported by NIGMS Grant No. GM103505. Special thanks to Dan Gu for technical support, and to Steve Dakin for providing code to filter the images used in this study. Thanks also to Alyson Saville and Christopher Tonsager for their assistance with participant recruitment and data collection.

References

- Balas, B., & Valente, N. (2012). View-adaptation reveals coding of face pose along image, not object, axes. *Vision Research*, *67*, 22–27.
- Biederman, I., & Kalocsai, P. (1997). Neurocomputational bases of object and face recognition. *Philosophical Transactions of the Royal Society B*, *352*, 1203–1219. doi:10.1098/rstb.1997.0103
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*, 305–327.
- Calder, A. J., Lawrence, A. D., & Young, A. W. (2001). Neuropsychology of fear and loathing. *Nature Reviews Neuroscience*, *2*, 252–363.
- Collin, C. A. (2006). Spatial-frequency thresholds for object categorization at basic and subordinate levels. *Perception*, *35*, 41–52.
- Dakin, S., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, *9*(4), 2.1–10. doi:10.1167/9.4.2
- Duchaine, B., Parker, H., & Nakayama, K. (2003). Normal recognition of emotion in a prosopagnosic. *Perception*, *32*, 827–838.
- Elfenbein, H., & Ambady, N. (2003). When familiarity breeds accuracy: Cultural exposure and facial emotion recognition. *Journal of Personality and Social Psychology*, *85*, 276–290.
- Fox, C. J., Oruç, I., & Barton, J. J. (2008). It doesn't matter how you feel: The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision*, *8*(3), 11. doi:10.1167/8.3.11
- 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.
- Galper, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science*, *19*, 207–208.
- Goffaux, V., & Dakin, S. C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*, *1*(143), 1–14. doi:10.3389/fpsyg.2010.00143
- 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.
- Goffaux, V., & Okamoto-Barth, S. (2013). Contribution of cardinal orientations to the “stare-in-the-crowd” effect. Poster presented at the Vision Sciences Society 2013 Annual Meeting, Naples, Florida.
- Goffaux, V., & Rossion, B. (2006). Faces are “spatial”—Holistic face perception is supported by low spatial frequencies. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 1023–1039. doi:10.1037/0096-1523.32.4.1023
- Goffaux, V., & Rossion, B. (2007). Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 995–1002. doi:10.1037/0096-1523.33.4.995

- Gold, J., 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.
- Itier, R. J., Alain, C., Sedore, K., & McIntosh, A. R. (2007). Early face processing specificity: It's in the eyes! *Journal of Cognitive Neuroscience*, *19*, 1815–1826.
- Jacques, C., d'Arripe, O., & Rossion, B. (2007). The time course of the inversion effect during individual face discrimination. *Journal of Vision*, *7*(8), 3. doi:10.1167/7.8.3
- Jacques, C., Schiltz, C., Collet, K., Oever, S., & Goffaux, V. (2011). *Scrambling horizontal face structure: behavioral and electrophysiological evidence for a tuning of visual face processing to horizontal information*. Poster presented at the Vision Sciences Society 2011 Annual Meeting, Naples, Florida.
- Johnson, M. H. (2005). Subcortical face processing. *Nature Reviews Neuroscience*, *6*, 766–774.
- Kirita, T., & Endo, M. (1995). Happy face advantage in recognizing facial expressions. *Acta Psychologica*, *89*, 149–163.
- Kumar, D., & Srinivasan, N. (2011). Emotion perception is mediated by spatial frequency content. *Emotion*, *11*, 114–1151.
- Li, R., & Cottrell, G. (2012). A new angle on the EMPATH model: Spatial frequency orientation in recognition of facial expressions. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Building bridges across cognitive sciences around the world: Proceeding of the 34th Annual Meeting of the Cognitive Science Society* (pp. 1894–1899). Austin, TX: Cognitive Science Society.
- Maurer, D., Grand, R. L., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, *6*, 255–260.
- Morris, J. S., DeGelder, B., Weiskrantz, L., & Dolan, R. L. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain*, *124*, 1241–1252.
- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013a). *Masking of individual facial features reveals the use of horizontal structure in the eyes*. Poster presented at the Vision Sciences Society 2013 Annual Meeting, Naples, Florida.
- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013b). Sensitivity to information conveyed by horizontal contours is correlated with face identification accuracy. *Frontiers in Psychology*, *4*(74), 1–8. doi:10.3389/fpsyg.2013.00074
- Palermo, R., Willis, M. L., Rivolta, D., McKone, E., Wilson, C. E., & Calder, A. J. (2011). Impaired holistic coding of facial expression and facial identity in congenital prosopagnosia. *Neuropsychologia*, *49*, 1226–1235.
- Sinha, P., Balas, B., Ostrovsky, Y., & Russell, R. (2006). Face recognition by humans: 20 results all computer vision researchers should know about. *Proceedings of the IEEE*, *94*, 1948–1962.
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., & Nelson, C. (2009). The Nimstim set of facial expressions: Judgments from typical research participants. *Psychiatry Research*, *168*, 242–249.
- Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience*, *6*, 624–631. doi:10.1038/nm1057
- Winston, J. S., Henson, R. N. A., Fine-Goulden, M. R., & Dolan, R. J. (2004). fMRI-adaptation reveals dissociable neural representation of identity and expression in face perception. *Journal of Physiology*, *92*, 1830–1839.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141–145.
- Young, A. W., McWeeny, K. H., Hay, D. C., & Ellis, A. W. (1986). Matching familiar and unfamiliar faces on identity and expression. *Psychological Research*, *48*, 63–68.
- Yu, D., Chai, A., & Chung, S. T. L. (2011). *Orientation information in encoding facial expressions*. Poster presented at the Vision Sciences Society 2011 Annual Meeting, Naples, Florida.
- Yue, X., Tjan, B. S., & Biederman, I. (2006). What makes faces special? *Vision Research*, *46*, 3802–3811.