Published in final edited form as:

Atten Percept Psychophys. 2014 July ; 76(5): 1381–1392. doi:10.3758/s13414-014-0669-4.

Emotion recognition (sometimes) depends on horizontal orientations

Carol M Huynh and Benjamin Balas, PhD

Psychology Department Center for Visual and Cognitive Neuroscience North Dakota State University

Abstract

Face recognition depends critically on horizontal orientations (Goffaux & Dakin, 2010). Face images that lack horizontal features are harder to recognize than those that have that information preserved. Presently, we asked if 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-degrees 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. In Experiment 2 we showed that mouth openness (i.e. open versus closed-mouths) also influenced the emotion-dependent reliance on horizontal information. Lastly, 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 does largely 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

Introduction

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, while for other object classes it appears that there are not sub-bands that contribute disproportionately to recognition (Biederman & Kalocsai, 1997; Collin, 2006). With regard to how orientation sub-bands contribute to face recognition, Dakin & 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 discuss 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 non-faces. 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 & Watt. In contrast, both contrast negative and 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 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 & Dakin (2010) showed that the face inversion effect (i.e., failure to recognize a familiar face when 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 better processed than faces containing only vertical information. However, when presented upside down, the horizontal advantage was greatly disrupted while faces containing vertical information remained largely unaffected. Identity after-effects are 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 towards 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 (Goffaux & Dakin, 2010; Dakin & Watt, 2009).

The amount of information for identification has also been shown to be greatest within the horizontal orientation band. Pachai, Sekuler, and Bennet (2013) found that masking of face images was strongest for noise with orientations at or near 0° (i.e., horizontal) and least for orientations at or near 90° (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. (2013) found observers were actually more sensitive to horizontal

information when faces were upright than inverted suggesting observers were more efficient at utilizing orientations at 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, there is preliminary evidence that horizontal information is not completely dominant over all other orientations. Goffaux & Okamoto-Barth (2013) demonstrated that vertical orientation assisted in the processing of gaze information. The authors compared direct with averted gaze by presenting an array of faces filtered to include horizontal, vertical, or a combination of both types of orientation information. By having subjects search for a target face consisting of either gaze, the authors found that detection was better for direct rather than averted gazes but only when arrays were comprised 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 completing 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 classic model of face perception (1986) proposes dissociable processes for identity and facial expressions (Winston, Henson, Fine-Goulden, & Dolan, 2004; Young, McWeeny, Hay, & Ellis, 1986). Bruce & Young (1986) 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 where 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 (Palermo, Willis, Rivolta, McKone, Wilson, Calder, 2011; Duchaine, Parker, Nakayama, 2003) and behavioral studies of neurotypical individuals. For example, Young, McWeeny, Hay, and Ellis (1986) provided evidence for separate processing of identity and emotional expressions by measuring reaction time in a matching task. They presented pairs of familiar or unfamiliar faces simultaneously and had subjects 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 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, reaction time was indeed faster for familiar than for unfamiliar faces while 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, Oruc, & Barton, 2008), the latter providing evidence of independent neural processing of facial emotion. Different emotions (Happy vs. Sad) also appear to be dissociated neurally (Morris, DeGelder, Weiskrantz, & Dolan, 2001; Calder, Lawrence & Young, 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 (Johnson, 2005; Calder, Lawrence, & Young, 2001). Additionally they are also largely dissociable from identity. Therefore, we hypothesized that the observed bias for horizontal information in identity recognition may not obtain in a facial emotion recognition task, and that orientation biases may also depend on appearance variability in how those emotions were 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, & 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°, 90°) 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, there was a bias towards 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 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 multi-scale, oriented Gabor filters (Li & Cottrell, 2012).

In the current 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°) 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-degree angle, 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 contribution 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 (Experiment 1) and faces expressing posed emotions (Experiment 2). This allowed us to examine both an ecologically valid set of emotional faces in one task and 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 are 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, is driving performance in our tasks. Lastly, we submitted our face images for analysis of the energy content within the horizontal or vertical orientation band to determine if 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 larger than vertical orientation energy, but that this effect was not significantly affected by emotional expression or mouth openness. We conclude 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—Seventeen undergraduate students (11 females/6 males) from North Dakota State University participated in this experiment. All participants reported normal or corrected-to-normal vision. Students provided written informed consent and received course credit for their participation.

Stimuli—Face images of 29 individuals (12 male/17 females) 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 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-degrees. Our stimuli were then created by taking the inverse of the Fourier transformed image back into the spatial domain (**Figure 1**). Following this inverse transformation, all images were readjusted so that mean luminance and contrast were matched (Dakin & Watt, 2009). In addition to the vertically (V) and horizontally (H) filtered images, we also included a third condition (Broadband) comprised of faces that contain

broadband orientation information (Figure 1). We had intended this set of images to only include 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 have information from all orientations at low spatial frequencies and are only restricted to horizontal and vertical orientations at high spatial frequencies. As a result, these control images are effectively comprised of orientation energy at all orientations, and therefore primarily allow 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 of Emotion (happy, sad), Image Orientation (upright or sideways), and Filter Orientation (vertical, horizontal, broadband). Image orientation was varied in separate blocks, while emotion and filter orientation were pseudo-randomized within each block. Participants completed a total of 348 trials (174 trials in the upright condition and 174 trials in the rotated condition). Block order was counterbalanced so that half of the participants began with the upright images and the remaining half began with the rotated images condition.

Procedure—Participants viewed the stimuli on a 13-inch MacBook with a 2.4 GHz Intel Core 2 Duo Processor. We recorded participants' responses using an 8-bit USB controller. Stimuli were presented using PsychToolbox 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 500ms, followed by a face stimulus that replaced the fixation cross. The face stimulus remained on the screen until participants either made a response or 2000ms 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 classification of happy) and false alarms (the incorrect classification of sad) in each condition. We submitted these sensitivity measures to a 2 (image rotation) $\times$ 3 (filter orientation) Repeated Measures ANOVA and observed a main effect of image rotation ($F(1,16) = 24.62, p < .0001, \eta^2 = 0.11$) such that discrimination was better for upright faces ($M = 3.26$) than for faces rotated sideways at a 90-degree angle ($M = 2.87$). We also observed a main effect of filter-orientation ($F(2,32) = 73.40, p < .0001, \eta^2 = 0.68$) such that discrimination was poorest for vertically-filtered faces ($M = 2.43$) followed by horizontally-filtered faces ($M = 3.17$) and faces containing both horizontal and vertical 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 ($F(2,32) = 7.57, p = .002, \eta^2 = 0.04$) such that image rotation significantly impacted discrimination for vertically-filtered but had no effect on horizontally-filtered faces and faces containing both broadband orientation information (**Figure 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 faces (774ms) than for happy faces (665ms), and by slower response latencies for vertically-filtered faces (793ms) relative to horizontal (700ms) and faces with broadband orientation ($M=665$ ms). There was also 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 (**Figure 3**) but happy face latencies did not differ in the horizontal and vertical filtering conditions. Unlike our analysis of sensitivity, we observed no interaction between image orientation and orientation filter ($F(2,32) = 1.29, p = 0.29$).

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 relative to horizontally-filtered images and images with broadband orientation. We also found that discrimination of emotional faces was worse when images were rotated 90-degrees 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 vertical orientations. This suggests that the horizontal bias we have observed is not solely driven 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, & Mondloc, 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 both horizontal and vertical energy. This disruption in performance for vertically-filtered faces rather than horizontally- filtered faces rotated sideways is inconsistent with previous results showing that rotation impacts both types of information (Yin, 1969; Goffaux & Rossion, 2011; Jacques, d'Arripe, & Rossion). However, one major difference between the results we obtained from previous studies is that our picture-plane rotation does not involve a complete 180-degree rotation. Thus, our face images are not completely inverted but are presented sideways instead. Furthermore, according to Goffaux and Rossion (2007), face inversion does not disrupt the same amount of facial information, whether vertical or horizontal. Although Goffaux & 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 poorest performance for recognizing vertical rather than horizontal facial relations rotated sideways at a 90-degree angle. Again, orientation information was not the target of their investigation but the report that performance

differences could arise for various facial information rotated along the picture-plane may possibly explain the discrepancies between our results from those of previous findings.

In terms of response latency, our main effect of emotion category is consistent with previous reports (Elfenbein & 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 RT data suggest that (a) orientation biases for emotion recognition may manifest in an emotion-dependent manner and that (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, is that our use of genuine emotional expressions may have introduced confounding factors that 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 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 where mouth openness could be manipulated. Specifically, we chose to use a set of posed emotions from the NimStim Face Set (Tottenham, Tanaka, Leon, McCarry, Nurse, Hare, Marcus, Westerlund, Casey, & Nelson, 2009) in which the position of the mouth (open vs. closed) was systematically varied in happy and sad emotional expressions.

Methods

Participant—Twenty-one undergraduate students (11 females/10 males) from North Dakota State University participated in this experiment, all reported normal or corrected-to-normal vision. Students provided written informed consent and received course credit for their participation.

Stimuli—Twenty-six individual face images (13 male/13 females) expressing posed emotions (26 happy/26 sad) were taken from the NimStim Face Set (Tottenham et. al., 2009) and were 650 × 506 pixels in size. There were two versions of each emotion (**Figure 4**); one expressed each emotion with a closed mouth and the other with an opened mouth (156 closed-mouth/156 opened-mouth). The process of filtering our stimuli to obtain horizontally-filtered (H) and vertically-filtered (V) images was identical to that described in Experiment 1. We note however, that our third condition in this task (HV) was comprised of

images containing the combination of the horizontal and vertical orientation subbands from the H and V conditions. Compared to Experiment 1 (in which energy at all orientations was included at low spatial frequencies) this third set of images contains 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 “HV” to contrast these stimuli with the “Broadband” stimuli employed in Experiment 1.

Design—We used a within-subjects design with the factors of 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 with closed mouth, 156 upright with opened mouth, 156 rotated with closed mouth, and 156 rotated with opened mouth with the factors of Emotion and Filter Orientation randomized within each block. 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 classification of happy) and false alarms (incorrect classification of sad) in each condition. We submitted these sensitivity measures to a 2 (planar rotation) $\times$ 2 (mouth openness) $\times$ 3 (filter orientation) Repeated Measures and observed significant main effects of image rotation, ($F(1,20)= 22.99, p < .001, \eta^2= 0.09$), mouth openness ($F(1,20)= 12.62, p = 0.002, \eta^2 = 0.016$) and filter orientation, ($F(2,40) = 86.75, p < .001, \eta^2= 0.40$). Discrimination was better for upright faces ($M=2.98$) than for sideways faces ($M=2.46$) and for open ($M=2.81$) than 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 ($p=0.58$) – an effect we did obtain 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, ($F(2,40)= 3.88, p = .029, \eta^2= 0.01$) and between mouth openness and filter orientation, ($F(2,40)= 9.12, p = .001, \eta^2= 0.04$). Similar to results from Experiment 1, vertically- filtered faces were greatly impacted by image rotation such that discrimination was worst for vertically-filtered faces rotated sideways at a 90-degree angle. Furthermore, discrimination for vertically- filtered faces was significantly worse when mouths were closed than when they were open (**Figure 5**).

Response Time—We analyzed the median response latency 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 (688ms vs. 735ms), upright than sideways (679ms vs. 743ms), and to faces displaying open rather than closed-mouths (695ms vs. 728ms). 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 obtain for other combinations of mouth openness and filter orientation (**Figure 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 towards responding “happy” when mouths were open than when they were closed (**Table 1**). We suggest that this criterion shift may result from participants’ tendency to infer that observers are 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 role of image orientation, filter orientation, and emotion category on emotion recognition. We found that a 90-degree rotation negatively impacted performance, but did not induce the “flip” between horizontal and vertical orientations that one would expect if 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 are more robustly recognized after vertical filtering than faces with closed mouths. This suggests that the effect of emotion 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 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 orientations relative to vertical orientations, but exhibit varying levels of bias as a function of stimulus appearance. We 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 towards responding “happy” when face

expressions displayed open than closed mouths. This could be due to the visibility of the teeth however slight it may 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 to observers to assume the current expression is happy and are therefore responding accordingly.

Experiment 3

Given that we have found that the horizontal bias for emotion recognition is affected by stimulus manipulations, a natural question to ask is whether our results are largely driven by the amount of horizontal vs. vertical information in our different categories of images. That is, are observers simply using more horizontal orientation energy because there is more of it in some images 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 may simply reflect the low-level statistics of our emotional faces. To examine this, we computed the power spectrum of face stimuli in all conditions in order to make comparisons of the summed energy within horizontal and vertical orientation bands.

Methods

Stimuli—Face stimuli were the same as those used in Experiment 1 and 2.

Design—In Experiment 1, we performed a 2 (happy, sad) $\times$ 2 (horizontal, vertical) design. In Experiment 2, we performed 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 Experiment 1 and 2, we found that the horizontal structure contained significantly more energy than 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 we observed behaviorally was not driven solely by the low-level statistics of our stimulus categories (emotion category, mouth openness). We suggest therefore, that our observers were not simply making use of horizontal information based on 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 (compared to vertically-filtered faces) in both of our experiments. However, other factors, such as the emotion category displayed by our stimuli (Experiments 1 and 2) and the position of the mouth (Experiment 2) also play a critical role in determining the magnitude of this bias. Overall, we suggest that our results demonstrate that orientation biases that affect emotion recognition are relatively flexible – the extent to which horizontal orientation is favored is not “locked-down” either by preferential connections to early vision and it is malleable in response to variation in stimulus appearance.

To speak to the impact of early orientation processing on emotion recognition, we demonstrated that rotating the image 90-degrees in the image plane did not substantially impact the horizontal bias. Instead, the pattern 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. Encoding of face viewpoint, by contrast, has recently been shown to have a substantial image-based orientation component (Balas & Valente, 2012). Balas and Valente (2012) observed adaption to the image axis, rather than the object, when subjects adapted to upright or sideways (90° 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 examine the impact of this variation on information biases for recognition. Here, we have made the distinction between open and closed mouths (Experiment 2) to see if mouth openness affected recognition across orientation bands, based on 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 like 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., 2013). 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, though a 90-degree planar rotation did not affect the orientation biases we observed here, face inversion (a 180-degree 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 inverted faces, indicating differences in observer sensitivity to faces rotated in the picture-plane (Pachai, Sekuler, & Bennett, 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 & Watt (2009) are also manifested in the responses measured from putative face-sensitive cortical loci (Jacques, Schlitz, & 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 there are no results describing 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 statistical regularities that define diagnostic image features are learned and applied as the visual system gains experience with complex and socially relevant patterns.

Acknowledgments

This research was supported by NIGMS grant 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 subject recruitment and data collection.

References

- Balas B, Valente N. View-adaptation reveals coding of face pose along image, not object, axes. *Vision Research*. 2012; 67:22–27. [PubMed: 22796427]

- Biederman I, Kolcsai P. Neurocomputational bases of object and face recognition. *Philosophical Transactions of the Royal Society of London, Series B*. 1997; 352:1203–1219. [PubMed: 9304687]
- Bruce V, Young A. Understanding face recognition. *British Journal of Psychology*. 1986; 77:305–327. [PubMed: 3756376]
- Calder AJ, Lawrence AD, Young AW. Neuropsychology of fear and loathing. *Nature Reviews Neuroscience*. 2001; 2:252–363.
- Collin CA. Spatial-frequency thresholds for object categorization at basic and subordinate levels. *Perception*. 2006; 35:41–52. [PubMed: 16491707]
- Dakin S, Watt RJ. Biological “bar codes” in human faces. *Journal of Vision*. 2009; 9:1–10.
- Duchaine B, Parker H, Nakayama K. Normal recognition of emotion in a prosopagnosic. *Perception*. 2003; 32:827–838. [PubMed: 12974568]
- Elfenbein H, Ambady N. When familiarity breeds accuracy: cultural exposure and facial emotion recognition. *Journal of Personality and Social Psychology*. 2003; 85:276–290. [PubMed: 12916570]
- Fox CJ, Oruç I, Barton JJ. It doesn't matter how you feel. The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision*. 2008; 8(3):11. [PubMed: 18484817]
- Freire A, Lee K, Symons LA. The face-inversion effect as a deficit in the encoding of configural information: direct evidence. *Perception*. 2000; 29:159–170. [PubMed: 10820599]
- Galper RE. Recognition of faces in photographic negative. *Psychonomic Science*. 1970; 19(4):207–208.
- Goffaux V, Dakin S. Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*. 2010; 1:1–14. [PubMed: 21833184]
- Goffaux V, Hault B, Michel C, Vuong QC, Rossion B. The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception-London*. 2005; 34(1):77–86. [PubMed: 15773608]
- Goffaux, V.; Okamoto-Barth, S. Contribution of cardinal orientations to the “Stare-in-the-crowd” effect.. Poster presented at the Vision Sciences Society 2013 Annual Meeting; Naples, Florida. 2013.
- Goffaux V, Rossion B. Faces are “spatial”—holistic face perception is supported by low spatial frequencies. *Journal of Experimental Psychology: Human Perception and Performance*. 2006; 32(4):1023. [PubMed: 16846295]
- Goffaux V, Rossion B. Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology*. 2007; 33:995–1001. [PubMed: 17683242]
- Gold J, Bennett PJ, Sekuler AB. Identification of band-pass filtered letters and faces by human and ideal observers. *Vision research*. 1999; 39(21):3537–3560. [PubMed: 10746125]
- Itier RJ, Alain C, Sedore K, McIntosh AR. Early face processing specificity: it's in the eyes. *Journal of Cognitive Neuroscience*. 2007; 19(11):1815–1826. [PubMed: 17958484]
- Jacques, C.; Schlitz, C.; Collet, K.; Oever, S.; Goffaux, V. 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. 2011.
- Johnson MH. Subcortical face processing. *Nature Reviews Neuroscience*. 2005; 6:766–774.
- Kirita T, Endo M. Happy face advantage in recognizing facial expressions. *Acta Psychologica*. 1995; 89:149–163.
- Kumar D, Srinivasan N. Emotion perception is mediated by spatial frequency content. *Emotion*. 2011; 11(5):114–1151.
- Li, R.; Cottrell, G. A New Angle on the EMPATH Model: Spatial Frequency Orientation in Recognition of Facial Expressions.. In: Miyake, N.; Peebles, D.; Cooper, RP., editors. *Proceeding of the 34th Annual Conference of the Cognitive Science Society*. Cognitive Science Society; Austin, TX: 2012.
- Maurer D, Grand RL, Mondloch CJ. The many faces of configural processing. *Trends in cognitive sciences*. 2002; 6(6):255–260. [PubMed: 12039607]

- Morris JS, DeGelder B, Weiskrantz L, Dolan RL. Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain*. 2001; 124:1241–1252. [PubMed: 11353739]
- Pachai, M.; Sekuler, A.; Bennett, P. 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. 2013a.
- Pachia M, Sekuler A, Bennett P. Sensitivity to information conveyed by horizontal contours is correlated with face identification accuracy. *Frontiers in Psychology*. 2013b; 4:1–8. [PubMed: 23382719]
- Palermo R, Willis ML, Rivolta D, McKone E, Wilson CE, Calder AJ. Impaired holistic coding of facial expression and facial identity in congenital prosopagnosia. *Neuropsychologia*. 2011; 49:1226–1235. [PubMed: 21333662]
- Sinha P, Balas B, Ostrovsky Y, Russell R. Face recognition by humans: 20 results all computer vision researchers should know about. *Proceedings of the IEEE*. 2006; 94:1948–1962.
- Tottenham N, Tanaka JW, Leon AC, McCarry T, Nurse M, Hare TA, Marcus DJ, Westerlund A, Casey BJ, Nelson C. The Nimstim Set of facial expressions: Judgments from typical research participants. *Psychiatry Research*. 2009; 168(3):242–249. [PubMed: 19564050]
- Vuilleumier P, Armony JL, Driver J, Dolan RJ. Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature neuroscience*. 2003; 6(6):624–631.
- Winston JS, Henson RNA, Fine-Goulden MR, Dolan RJ. fMRI- Adaptation reveals dissociable neural representation of identity and expression in face perception. *Journal of Physiology*. 2004; 92:1830–1839.
- Yin RK. Looking at upside-down faces. *Journal of Experimental Psychology*. 1969; 81(1):141–145.
- Young AW, McWeeny KH, Hay DC, Ellis AW. Matching familiar and unfamiliar faces on identity and expression. *Psychological Research*. 1986; 48:63–68. [PubMed: 3774919]
- Yu, D.; Chai, A.; Chung, STL. Orientation information in encoding facial expressions.; Poster presented at the Vision Sciences Society 2011 Annual Meeting. Naples, Florida. 2011.
- Yue X, Tjan BS, Biederman I. What makes faces special?. *Vision research*. 2006; 46(22):3802–3811. [PubMed: 16938328]

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

Fig. 2.

Sensitivity measures for responses in Experiment 1. We find that image rotation significantly impacted discrimination for categorization, especially for vertically- filtered faces. Error bars represent +/- 1 s.e.m.

Fig. 3.

Average response latency for correct responses in Experiment 1. For both upright and sideways faces, we find that vertically- filtered faces are more slowly recognized than our other filtered images, but only when sad faces are recognized. Error bars represent ± 1 s.e.m.

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

Fig. 5.

Sensitivity measures for responses in Experiment 2. We find that vertically- filtered faces were greatly impacted by sideways rotation, especially for closed mouth openness.

Furthermore, we find worse discrimination for vertically- filtered faces with closed than open mouths. Error bars represent ± 1 s.e.m.

Fig. 6.

Average response latency for correct responses in Experiment 2 for upright and sideways faces with open or closed mouths. Error bars represent ± 1 s.e.m.

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