



The effect of training with inverted faces on the selective use of horizontal structure



Matthew V. Pachai^{a,*}, Patrick J. Bennett^b, Allison B. Sekuler^{b,c,d}

^a Department of Psychology, York University, Toronto, ON, Canada

^b Department of Psychology, Neuroscience & Behaviour, McMaster University, Hamilton, ON, Canada

^c Rotman Research Institute, Baycrest Health Sciences, Toronto, ON, Canada

^d Department of Psychology, University of Toronto, Toronto, ON, Canada

ARTICLE INFO

Number of Reviews = 2

Keywords:

Face discrimination
Face inversion effect
Perceptual learning
Horizontal selectivity

ABSTRACT

A growing body of evidence demonstrates that selective processing of structure conveyed by horizontally oriented spatial frequency components is associated with upright face discrimination accuracy and the magnitude of the face inversion effect. In this study, we examined whether the increase in discrimination accuracy for inverted faces that is known to result from practice would coincide with more selective processing of horizontal structure in inverted faces. To assess this hypothesis, our observers practiced discrimination of inverted faces for three training sessions and we measured accuracy, efficiency relative to an ideal observer, and horizontal selectivity before and after training. As hypothesized, we observed more efficient discrimination and more selective processing of horizontal structure after training. However, the effects of training did not generalize reliably to novel face exemplars.

1. Introduction

It is commonly believed that most adults are experts at perceiving human faces, and that such expertise is derived, at least in part, from a lifetime of experience with these stimuli (de Heering, Rossion, & Maurer, 2012; Germine, Duchaine, & Nakayama, 2011; LeGrand, Mondloch, Maurer, & Brent, 2001; Susilo, Germine, & Duchaine, 2013). Consistent with the claim that perceptual learning contributes to face expertise is the observation that face perception is degraded for less familiar faces, such as those that differ from the observer in terms of age (Fulton & Bartlett, 1991; Rhodes & Anastasi, 2012) or race (Brigham & Barkowitz, 1978; Malpass & Kravitz, 1969; Meissner & Brigham, 2001). Perhaps the most robust demonstration of expertise in face perception is the severe decrement in performance following picture-plane inversion (Valentine, 1988; Yin, 1969). This face inversion effect (FIE) is particularly interesting because rotation does not alter the physical information available in the stimulus, and therefore the change in performance implies that humans process inverted faces less efficiently than upright faces (Gaspar, Bennett, & Sekuler, 2008a; Sekuler, Gaspar, Gold, & Bennett, 2004).

Previous studies have shown that adults discriminate faces based on information near the eyes (Davies, Ellis, & Shepherd, 1977; Gold, Sekuler, & Bennett, 2004; Gosselin & Schyns, 2001; Haig, 1985, 1986;

Peterson & Eckstein, 2012; Sekuler et al., 2004) using a limited band of spatial frequencies (Näsänen, 1999; Gaspar, Sekuler, & Bennett, 2008b). It is plausible that the FIE might be associated with a change in the spatial frequencies used to discriminate inverted faces, but studies so far have failed to find a meaningful effect of inversion on the spatial frequency tuning properties of face processing (Gaspar, Sekuler, & Bennett, 2008b; Willenbockel et al., 2010, but see Boutet, Collin, & Faubert, 2003; Watier, Collin, & Boutet, 2010). However, recent work has shown that the FIE is associated with differential sensitivity to horizontally-oriented spatial frequency components (Goffaux & Dakin, 2010; Goffaux & Greenwood, 2016; Pachai, Sekuler, & Bennett, 2013b), which are diagnostic for face identity and expression (Dakin & Watt, 2009; Huynh & Balas, 2014; Pachai et al., 2013b). These results suggest that face expertise, as indexed by the FIE, is related to changes in the way observers use horizontal structure to discriminate faces (see also Pachai, Sekuler, Bennett, Schyns, & Ramon, 2017). In this paper, we examine whether perceptual learning of intact exemplars concurrently alters sensitivity to horizontal structure in inverted faces.

Many studies have demonstrated that adults can learn to discriminate upright and inverted faces more effectively (e.g., de Heering & Maurer, 2013; Germine et al., 2011; Gold et al., 2004; Hussain, Sekuler, & Bennett, 2009a; Hussain, Sekuler, & Bennett, 2009b; Lagusse, Dormal, Biervoeye, Kuefner, & Rossion, 2012). The effects of

* Corresponding author.

E-mail addresses: mattpachai@gmail.com (M.V. Pachai), bennett@mcmaster.ca (P.J. Bennett), sekuler@mcmaster.ca (A.B. Sekuler).

<https://doi.org/10.1016/j.visres.2018.04.003>

Received 23 August 2017; Received in revised form 12 February 2018; Accepted 3 April 2018

Available online 27 April 2018

0042-6989/© 2018 Published by Elsevier Ltd.

perceptual learning with faces can be long-lasting (Hussain, Sekuler, & Bennett, 2011), and sometimes (e.g., Hussain et al., 2009a; Hussain et al., 2011), though not always (e.g. de Heering & Maurer, 2013; Lagusse et al., 2012) are specific to the trained faces. Training-based improvements in face discrimination are associated with an increase in calculation efficiency (Gold, Bennett, & Sekuler, 1999b; Gold et al., 2004), but it remains unknown how visual processing changes to become more efficient. In the present study, we examined if learning enhances observers' ability to extract horizontal structure from faces using a filtering technique that on each trial isolated the diagnostic information to a specific orientation band while retaining non-informative structure in all other orientation bands. Furthermore, to generalize our results across stimuli, we trained different groups of observers on two different sets of face identities. To normalize our results for differences between these face sets and to characterize the trained improvements with regard to the information available in the stimuli, we compared the results of our human observers to an ideal observer that optimally utilized all available information. Finally, to examine the extent to which trained improvements in discrimination accuracy transferred to novel exemplars, our observers returned for an additional session in which they were assessed with the untrained face set.

2. Methods

2.1. Observers

Twenty observers (13 male, age 18–30, $M = 23$ years) participated in the experiment. All observers had normal or corrected-to-normal Snellen acuity, and were paid \$10/hour or given partial course credit for their participation. Experimental protocols were approved by the McMaster University Research Ethics Board in accordance with the Declaration of Helsinki, and informed consent was collected prior to the experiment.

2.2. Apparatus

The experiment was run on an Apple Macintosh G4 computer using MATLAB and the Psychophysics and Video Toolboxes (Brainard, 1997; Pelli, 1997). Stimuli were displayed on a 21" Apple Studio monitor with viewable size of 40×30 cm, a resolution of 1280×1024 pixels (32 pixels per cm), and a frame rate of 85 Hz. An average luminance of 31 cd/m^2 was held constant during the experiment. A chin/head rest was used to stabilize viewing distance at 60 cm, and the experimental apparatus provided the only source of light in the room.

2.3. Stimuli

Two sets of ten faces were used in this experiment (Fig. 1). Both sets were generated using front-facing, digital photographs of five male and five female models who had no facial hair, eye glasses, identifiable marks, or visible piercings. Photographs in face set 1 were cropped using a 198×140 pixel oval window and photographs in face set 2 were cropped using a 207×138 pixel oval window. All photographs were centred in a 256×256 pixel matrix, which at the viewing distance of 60 cm subtended $7.6^\circ \times 7.6^\circ$. For more details on the generation of face set 1, see Gold, Bennett, and Sekuler (1999a) and for face set 2 see Gaspar, Bennett, and Sekuler (2008a).

Across trials, we manipulated the orientation information available to observers by filtering the stimuli in the Fourier domain. Specifically, we selectively retained frequency components from the target face using 18 ideal orientation filters with bandwidths ranging from 10° to 180° in 10° steps, centred on 0° (horizontal) or 90° (vertical). Note that 90° is the largest bandwidth at which the horizontal and vertical filters passed completely different (i.e., non-overlapping) frequency components, and that 180° filters removed no frequency components and

therefore yielded unfiltered faces. After filtering, the components that were removed from the target face were replaced with the corresponding components from the average of the 10 possible faces in the relevant face set. Finally, the power of the final, hybrid image was re-scaled to be constant across filter conditions. The resulting stimuli contained power in all orientation bands and resembled an unfiltered face, but contained diagnostic identity information in only a limited orientation band (see Fig. 2). During the experiment, the stimuli always were inverted by rotating images 180° in the picture plane, and presented at an RMS contrast of 0.5. Stimuli were embedded in white noise with an RMS contrast of 0.1 for later comparison with an ideal observer, which requires external noise.

2.4. Design

Observers were randomly assigned to two groups of 10, each trained with a different face set. The training paradigm consisted of a pre-training assessment, three learning sessions, a post-training assessment, and a transfer assessment. Each session was separated by approximately 24 h except for the transfer session, which was completed approximately 72 h following the post-training assessment. Pre- and post-training assessments consisted of 10 trials in each condition (2 filter orientations \times 18 bandwidths, 360 total trials), randomly intermixed. Learning sessions each consisted of 300 trials using unfiltered, inverted faces. The transfer assessment consisted of two blocks of trials: The first block always contained the trained face set and the second block always contained the novel faces. Both blocks were identical in design to the assessment sessions. No practice trials were presented in any of the sessions.

In each session, trials began with a fixation cross presented at the centre of the screen for 500 ms, followed by a 250 ms blank screen. The stimulus then was presented for 500 ms, followed by a 250 ms blank screen and a response window containing front-facing exemplars of the 10 possible faces. Faces in the response window were always unfiltered and inverted. Observers selected their response using a mouse click with no time constraint, and feedback was provided using 600 Hz and 200 Hz tones for correct and incorrect responses, respectively.

2.5. Data analysis

The dependent measure in this experiment was proportion correct in the 10-AFC discrimination task (P_C). Psychometric functions relating P_C to filter bandwidth were estimated using generalized linear models with a probit link function that included one additional free parameter, λ , representing the upper asymptote. The lower asymptote was fixed at 0.1 (i.e., chance performance). Given that 90° is the largest bandwidth at which the orientation filters isolated independent frequency components, our measure of horizontal selectivity was the difference in P_C extracted from these psychometric functions at 90° bandwidth for horizontal and vertical filters, respectively.

2.6. Ideal observer analysis

An ideal observer uses an optimal decision rule to perform a given task and is limited by the diagnostic information available in the stimulus (Bennett & Banks, 1987; Geisler, 1989, 2011; Gold et al., 2004; Kersten, 1987; Tjan, Braje, Legge, & Kersten, 1995). It has been shown that the ideal decision rule in an discrimination task is to correlate the noisy stimulus on a given trial with templates representing each possible response, then to select the identity that maximizes the correlation (Gold et al., 2004; Tjan et al., 1995). When the stimuli are band-pass filtered, the ideal observer bases its decision on only the spatial frequency components carrying diagnostic information. The performance of such an observer can be compared to human observers, quantifying the extent to which they optimally used the available information. This relationship is termed efficiency, and is defined as the squared ratio of

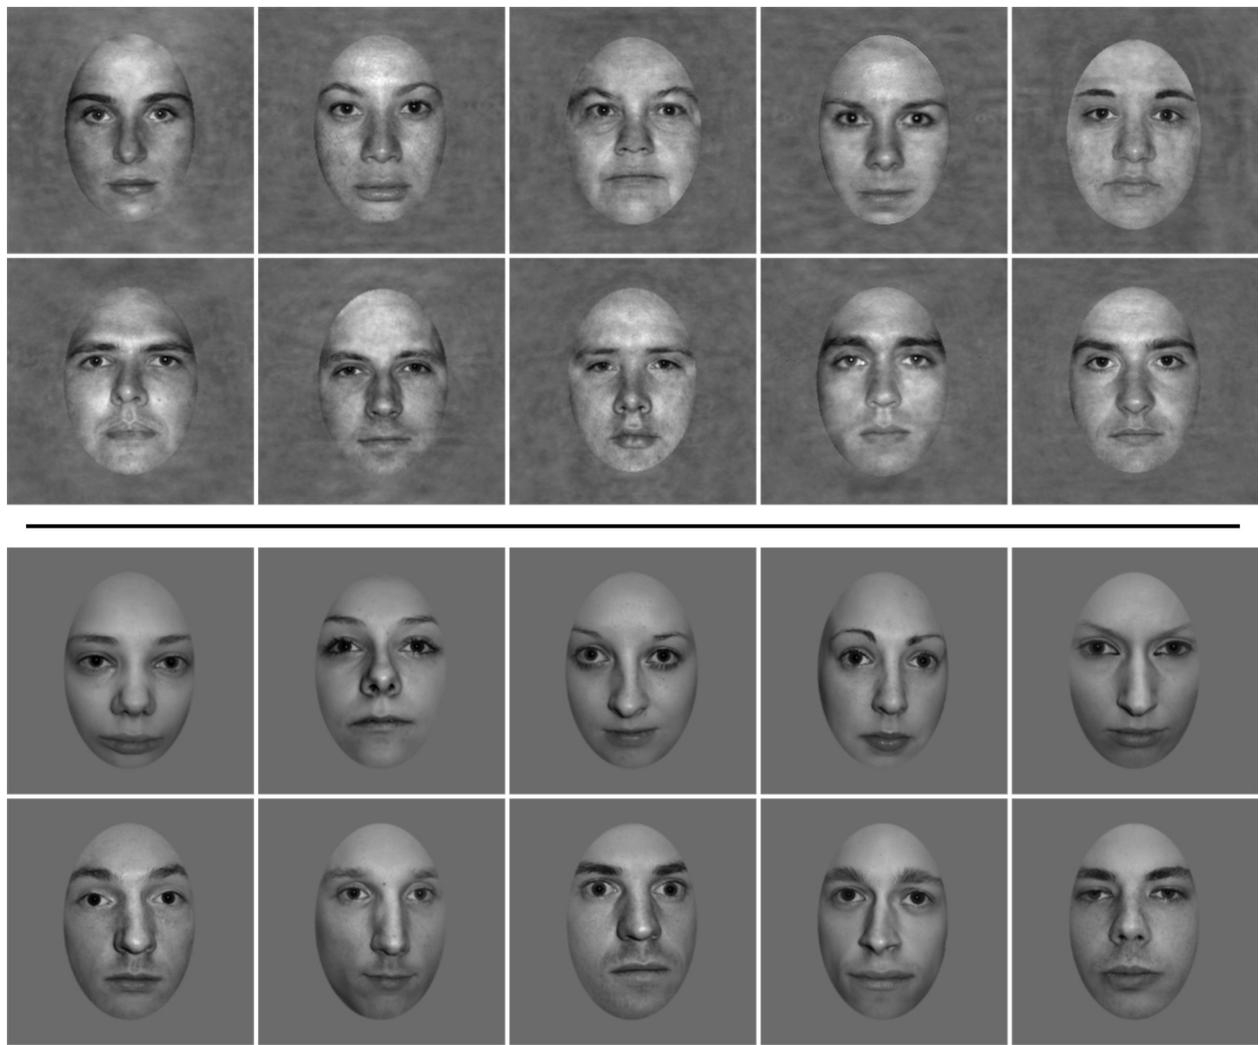


Fig. 1. Unfiltered examples of the identities included in face set 1 (top) and face set 2 (bottom).

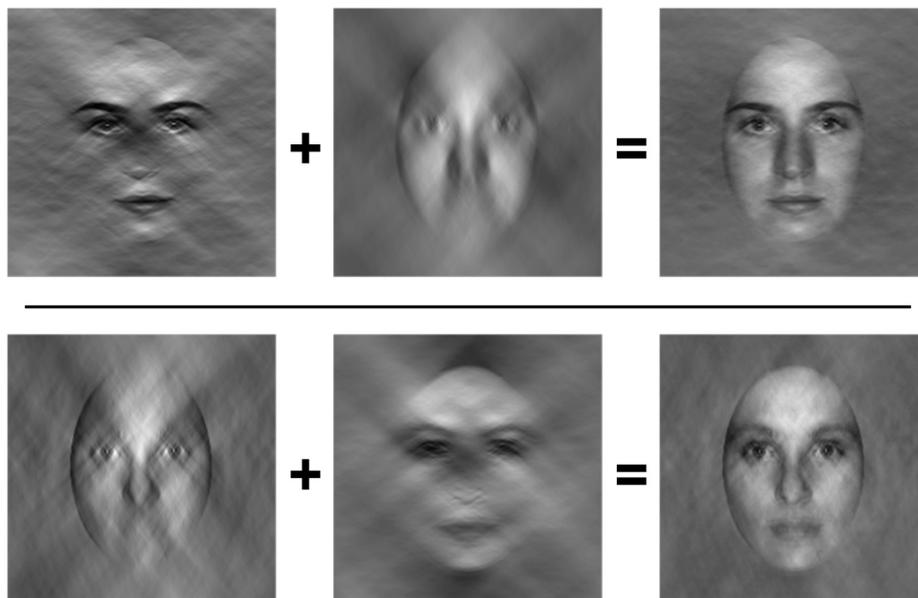


Fig. 2. Demonstration of the filtering technique using an identity from face set 1. Stimuli on left were passed through a 90° filter centred on horizontal (top) or vertical (bottom). The removed frequency components were replaced with the corresponding components from the average of the 10 faces (centre) to produce the final stimuli presented during the experiment (right).

human to ideal d' (Tanner & Birdsall, 1958). Therefore, to compute efficiency, human P_C was transformed to d' using the procedures outlined by Macmillan and Creelman (2004). However, at the RMS contrast shown to human observers, the ideal observer would achieve ceiling P_C in all filter conditions, precluding such a transformation. Therefore, for each filter orientation and bandwidth, we simulated ideal performance at seven RMS contrasts ranging from 0.001 to 0.0025, using 5000 trials per contrast level. We then computed best fitting least squares regression lines relating d' to log-transformed RMS contrast, and extrapolated to the contrast presented to human observers. The resulting d' values were used to compute efficiency based on the relationship described above.

3. Results

All statistical analyses were conducted using R (R Core Team, 2017). Section 3.1 describes the performance of an ideal observer on the 10-AFC discrimination task, establishing the inherent differences between our two face sets. Section 3.2 describes the effect of our training regimen on overall performance and horizontal selectivity. Finally, Section 3.3 describes the results of our transfer assessment, in which we examined whether trained improvements in performance would generalize to untrained identities. All data for human observers are reported as P_C and efficiency.

3.1. Ideal observer results

Fig. 3 plots d' as a function of filter orientation and bandwidth for an ideal observer performing the 10-AFC discrimination task. The performance of such an observer quantifies the diagnostic information available in the stimulus, and this analysis revealed less total information available in face set 1 than face set 2 [$d'_{180^\circ(\text{set1})} = 672; d'_{180^\circ(\text{set2})} = 830$]. We next examined the differential information available in the horizontal band for each face set by computing best-fitting least squares regression lines relating d' to bandwidth from 10° to 90° for horizontal and vertical filters, respectively. These data were fit well by linear regression (all $R^2 > 0.98$), and the resulting parameters are presented in Table 1. This analysis revealed similar slopes between horizontal and vertical filters for both face sets [$\Delta S_{H-V(\text{set1})} = 0.631; \Delta S_{H-V(\text{set2})} = -0.097$] as well as a large difference between intercepts for horizontal and vertical filters

Table 1

Parameters of the best fitting least-squares regression lines fit to d' as a function of filter bandwidth from 10° to 90° for an ideal observer tested with each face set.

Orientation	Face Set 1			Face Set 2		
	R^2	Intercept	Slope	R^2	Intercept	Slope
Horizontal	0.98	109.67	3.12	0.98	118.48	3.44
Vertical	0.98	22.87	2.49	0.99	42.16	3.53

[$\Delta I_{H-V(\text{set1})} = 86.8; \Delta I_{H-V(\text{set2})} = 76.3$]. Together, these parameters demonstrate that the horizontal advantage is conveyed by information with a bandwidth of 10° or less, because a difference at the intercept that does not increase with bandwidth (i.e., approximately equal slopes) indicates an advantage for horizontal structure present at the narrowest bandwidth we tested or smaller. We also observed a smaller difference between intercepts for face set 2, which suggests that the overall informational advantage for this set may be due, in part, to face set 2 having relatively more information in the vertical band of orientations. To visualize more clearly the different horizontal advantages conveyed by each face set, we computed horizontal selectivity as $d'_{\text{horizontal}} - d'_{\text{vertical}}$ and plotted the resulting values in Fig. 4. Finally, to reveal the face features revealed by the horizontal and vertical bands from 10° to 90° , we computed maps of the information conveyed at each bandwidth, which are shown in Fig. 5. These maps quantify the diagnostic information available to observers by representing, for each pixel, the contrast variance across face identities at that location. Visual inspection of these maps demonstrates the extent to which horizontal spatial frequency components convey information around the eyes and eyebrows, as well as the greater availability of diagnostic information in the vertical band of face set 2.

3.2. Training results

Before considering the effect of our training regimen, we examined P_C across the three learning sessions, with each session divided into five 60-trial bins (Fig. 6). Visual inspection of these data indicates that much of the learning occurred by the end of the first training session, consistent with previous results on face learning (Gold et al., 1999b, 2004). We verified this observation using a 2 (face set) \times 3 (session) \times 5 (block)

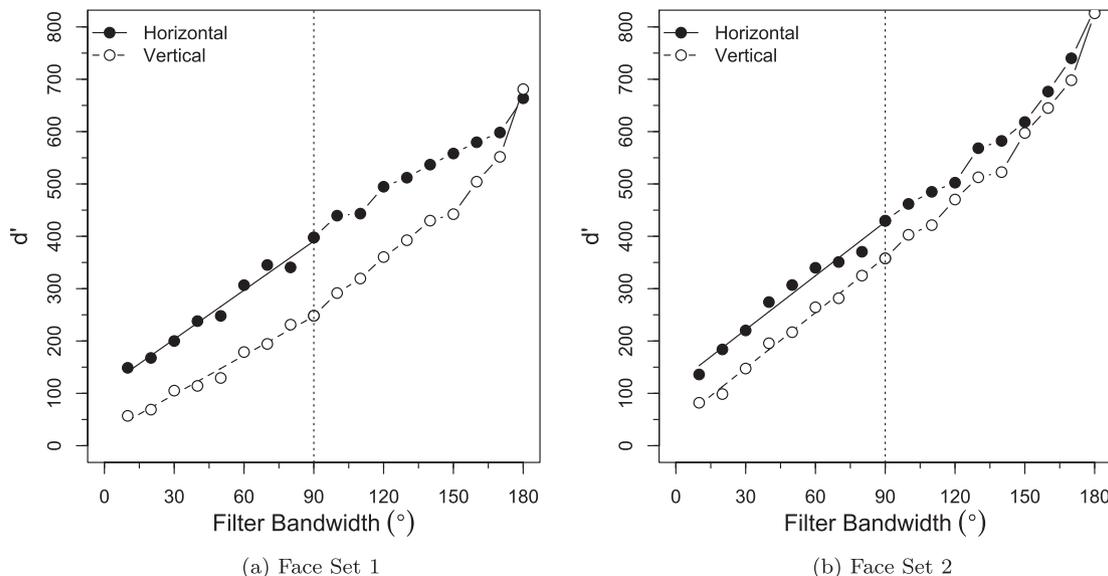


Fig. 3. d' on the 10AFC discrimination task for an ideal observer using a template matching decision rule that weighs all orientations optimally (see Section 2.6 for details). Lines represent the best fitting least squares regression fit to the data from 10° to 90° bandwidth, where 90° is indicated by the vertical dotted line.

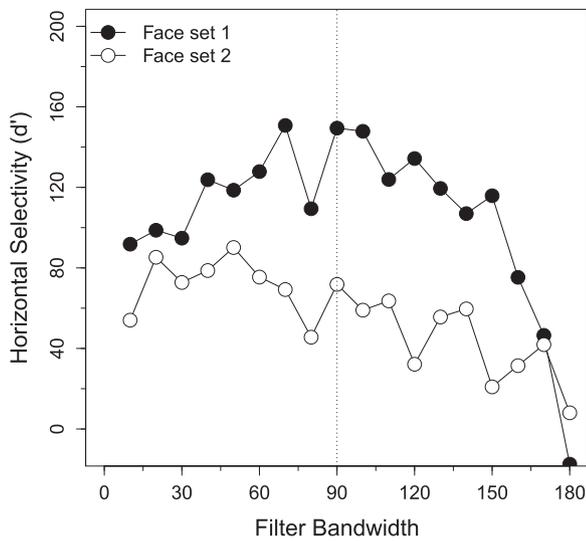


Fig. 4. Horizontal selectivity for the ideal observer, defined as $d'_{horizontal} - d'_{vertical}$ plotted separately for face sets 1 and 2.

mixed ANOVA with face set as a between-subjects factor. This analysis revealed significant main effects of session [$F(2,36) = 22.70, p < 0.0001$] and block [$F(4,72) = 15.20, p < 0.0001$]. These main effects were qualified by a session \times block interaction [$F(8,144) = 5.12, p < 0.0001$]. No other main effects or interactions approached significance.

3.2.1. Proportion correct

The effect of our training regimen on P_c for face set 1 is demonstrated in Fig. 7a. To examine directly the effect of training on horizontal selectivity, we extracted proportion correct at 90° for each observer from psychometric functions fit to the entire range of bandwidths and plotted the resulting values in Fig. 7b. We submitted these data to a 2 (training) \times 2 (filter orientation) repeated measures ANOVA, which revealed main effects of training [$F(1,9) = 58.69, p < 0.0001$], filter orientation [$F(1,9) = 80.03, p < 0.0001$], and a significant training \times filter orientation interaction [$F(1,9) = 13.34, p = 0.0053$]. To further analyze this interaction, we conducted paired t-tests comparing horizontal to vertical performance at 90° before and after training. Before training,

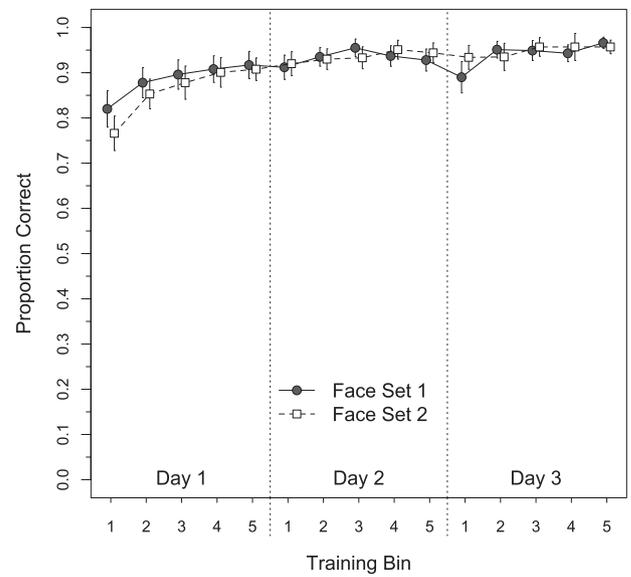


Fig. 6. Proportion correct during the three learning sessions, divided into 60-trial bins. Error bars represent ± 1 SEM.

horizontal performance was greater than vertical performance [$M_H = 0.57; M_V = 0.29; t(9) = 5.7, p = 0.0002$]. After training, horizontal performance was again greater than vertical performance [$M_H = 0.9; M_V = 0.41; t(9) = 9.1, p < 0.0001$], where the aforementioned interaction reveals that this difference was greater after training than before. Together, the results of these analyses are consistent with the observations that (i) accuracy was higher after training; (ii) accuracy was higher in the horizontal filter condition than the vertical filter condition; and (iii) that the effect of training was greater in the horizontal condition than the vertical condition. Hence, training improved performance and increased horizontal selectivity.

The effect of our training regimen on discrimination of face set 2 is plotted in Fig. 8. As with face set 1, we extracted proportion correct at 90° from the psychometric functions and submitted the resulting data to a 2 (training) \times 2 (filter orientation) repeated-measures ANOVA, which revealed significant main effects of training [$F(1,9) = 48.36, p < 0.0001$],

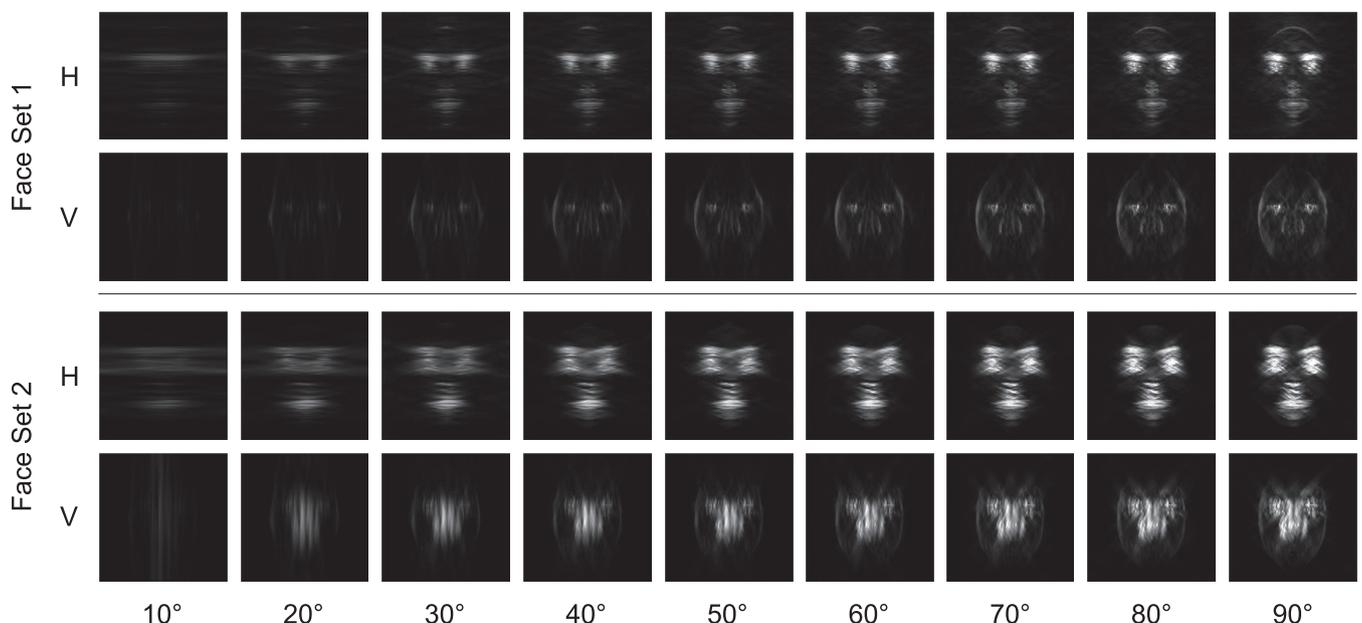


Fig. 5. Maps of the diagnostic information conveyed by horizontal (H) and vertical (V) spatial frequency components at each bandwidth from 10° to 90° for each face set. Each pixel in each image represents the variance in contrast at that location across face identities filtered with a given orientation and bandwidth.

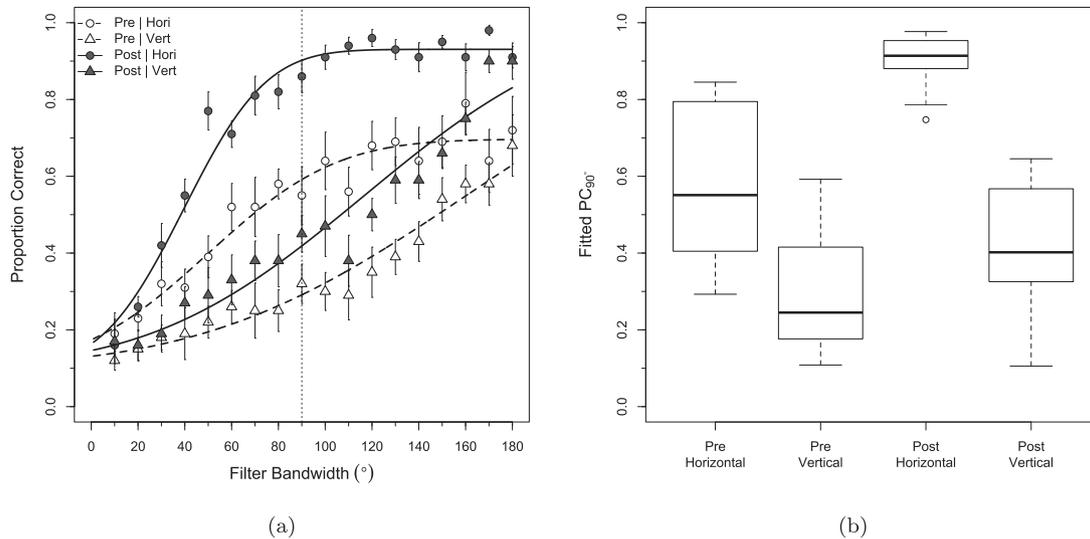


Fig. 7. (a) Proportion correct on the 10-AFC discrimination task before and after three days of training with face set 1. Fitted lines are psychometric functions calculated for the mean data. 90° bandwidth is indicated by the vertical dotted line. Error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to data from individual observers that spanned the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

filter orientation [$F(1,9) = 50.96, p < 0.0001$] and a significant interaction [$F(1,9) = 12.35, p = 0.0066$]. To further analyze this interaction, we again conducted paired t-tests comparing horizontal to vertical performance at 90° before and after training. Before training, horizontal performance was greater than vertical performance [$M_H = 0.39; M_V = 0.21; t(9) = 6.5, p = 0.0001$]. After training, horizontal performance was again greater than vertical performance [$M_H = 0.73; M_V = 0.38; t(9) = 6.2, p = 0.0001$], where the aforementioned interaction reveals that this difference was greater after training than before. Together, the results of these analyses are qualitatively similar to face set 1: training improved overall discrimination accuracy (main effect of training) as well as horizontal selectivity (training \times filter orientation interaction).

3.2.2. Efficiency

To rescale the effect of our training regimen with regard to the differential information available in the two face sets, we transformed

P_C to efficiency using the relationship described in Section 2.6. Efficiency is plotted for both face sets in Fig. 9. We submitted these data, separately for each face set, to 2 (training) \times 2 (filter orientation) \times 9 (bandwidth) repeated-measures ANOVAs. Note that for these analyses, we include only the non-overlapping range of bandwidths (10–90°).

For face set 1, visual inspection of Fig. 9 suggests an overall improvement in efficiency following training, and an increase in horizontal selectivity that emerges at intermediate bandwidths. Quantitatively, our omnibus ANOVA revealed a significant main effect of training [$F(1,9) = 24.59, p = 0.0008$] but no significant main effect of filter [$F(1,9) = 0.85, p = 0.3805$] or bandwidth [$F(8,72) = 1.94, p = 0.0668$]. These main effects were qualified by a significant filter orientation \times filter bandwidth interaction [$F(8,72) = 4.59, p = 0.0002$] and a significant training \times filter orientation \times filter bandwidth interaction [$F(8,72) = 3.06, p = 0.0051$]. The interactions of training \times filter orientation [$F(1,9) = 0.13, p = 0.7226$] and training \times filter bandwidth [$F(8,72) = 1.41, p = 0.2076$] were not significant. Together, these results

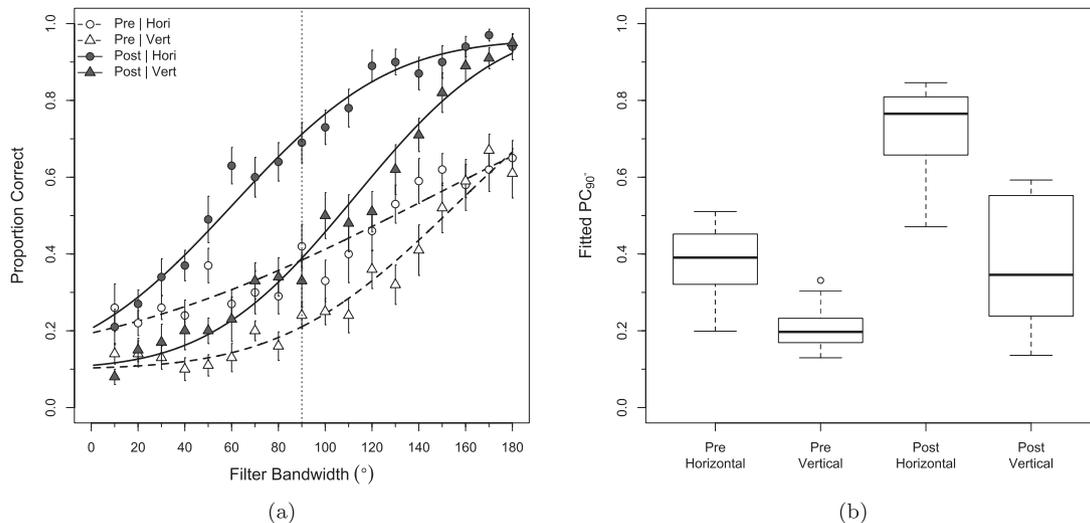


Fig. 8. (a) Proportion correct on the 10-AFC discrimination task before and after three days of training with face set 2. Fitted lines are psychometric functions calculated using the mean data. 90° bandwidth is indicated by the vertical dotted line. Error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to data from individual observers that spanned the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

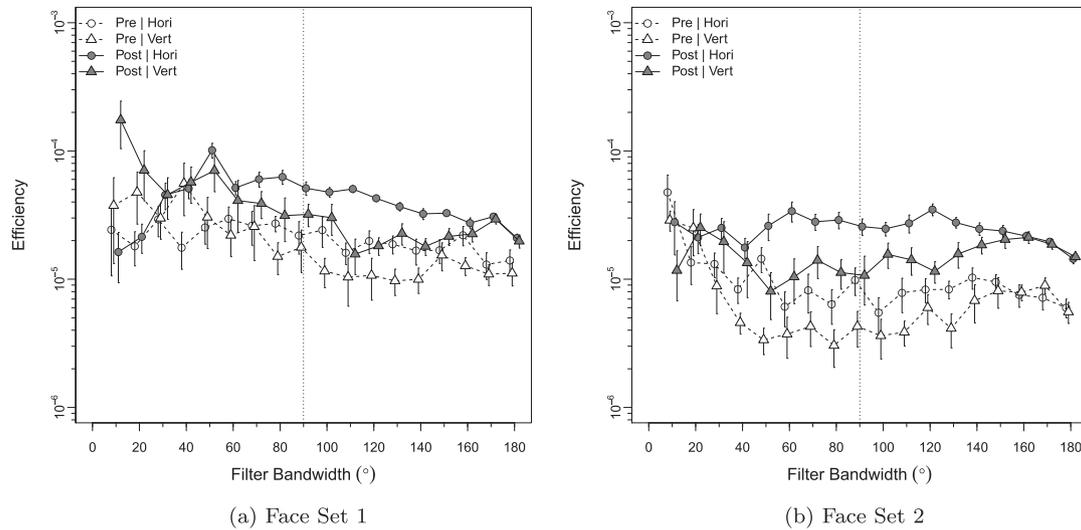


Fig. 9. Efficiency relative to an ideal observer that weighs all frequency components optimally before training (open symbols) and after training (closed symbols) for (a) face set 1 and (b) face set 2. See Section 2.6 for details on efficiency calculation. Points have been offset slightly for visibility. Error bars represent ± 1 SEM.

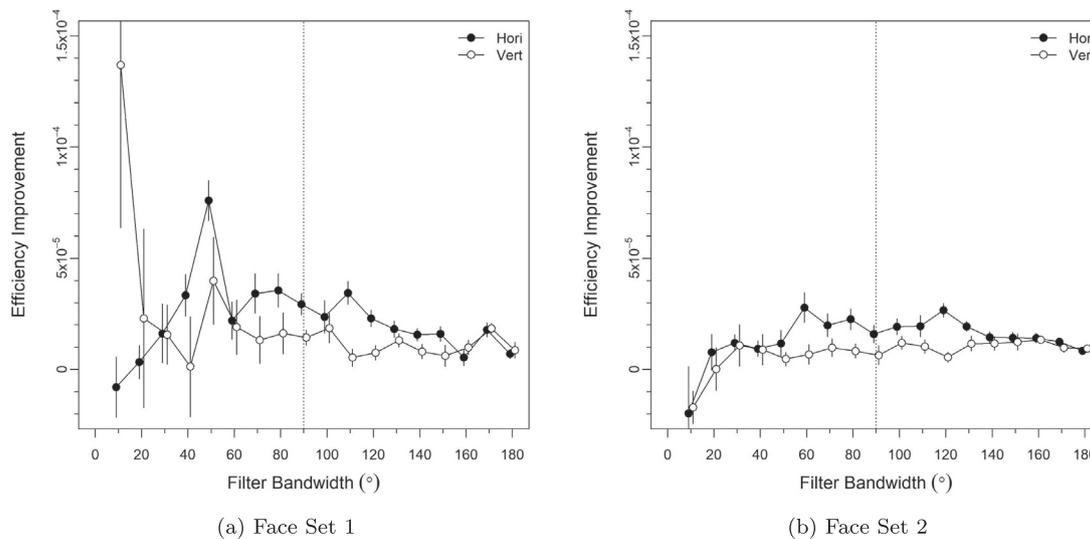


Fig. 10. The improvement in efficiency resulting from our training regimen plotted separately for (a) face set 1 and (b) face set 2. See Section 2.6 for details on efficiency calculation. Points have been offset slightly for visibility. Error bars represent ± 1 SEM.

suggest that training improved overall performance and that the effect of training on horizontal selectivity varied as a function of bandwidth. The effect of bandwidth can be visualized more directly in Fig. 10a, which plots the effect of training as the difference between post-training and pre-training efficiency. Fig. 10a reveals a greater improvement for horizontal than vertical orientations that stabilizes after approximately 60°. That horizontal orientations improved more than vertical can also be characterized as increased horizontal selectivity in this bandwidth range following training.

For face set 2, visual inspection of Fig. 9 reveals an effect of training on overall efficiency, particularly for the horizontal band. Quantification of these patterns with an omnibus ANOVA revealed significant main effects of training [$F(1,9) = 6.68, p = 0.0295$], filter orientation [$F(1,9) = 9.85, p = 0.012$], and bandwidth [$F(8,72) = 4.04, p = 0.0005$]. Further, all of the two-way interactions were significant, specifically training \times filter orientation [$F(1,9) = 12.13, p = 0.0069$], training \times bandwidth [$F(8,72) = 4.50, p = 0.0002$], and filter orientation \times bandwidth [$F(8,72) = 2.51, p = 0.0181$]. The training \times filter orientation \times bandwidth interaction was not significant [$F(8,72) = 0.51, p = 0.8451$]. These effects of training on efficiency can be visualized more directly in

Fig. 10b, which again reveals a greater improvement for horizontal than vertical orientations that stabilizes after approximately 60°. Again, that horizontal orientations improved more than vertical can also be characterized as increased horizontal selectivity in this bandwidth range following training.

3.3. Transfer results

During the transfer session, observers completed one block with the trained face set and one block with the untrained face set. Fig. 11 replots the post-training data from Section 3.2.1 along with proportion correct with trained faces from the transfer session. This comparison demonstrates the extent to which trained improvements degraded in the 72 h prior to the transfer session. These results show clearly that trained improvements in performance remained intact, so we proceeded to examine transfer to untrained faces.

3.3.1. Proportion correct

For clarity in the following analyses, we define Groups 1 and 2 as observers who were trained with face set 1 and 2, respectively. First, we

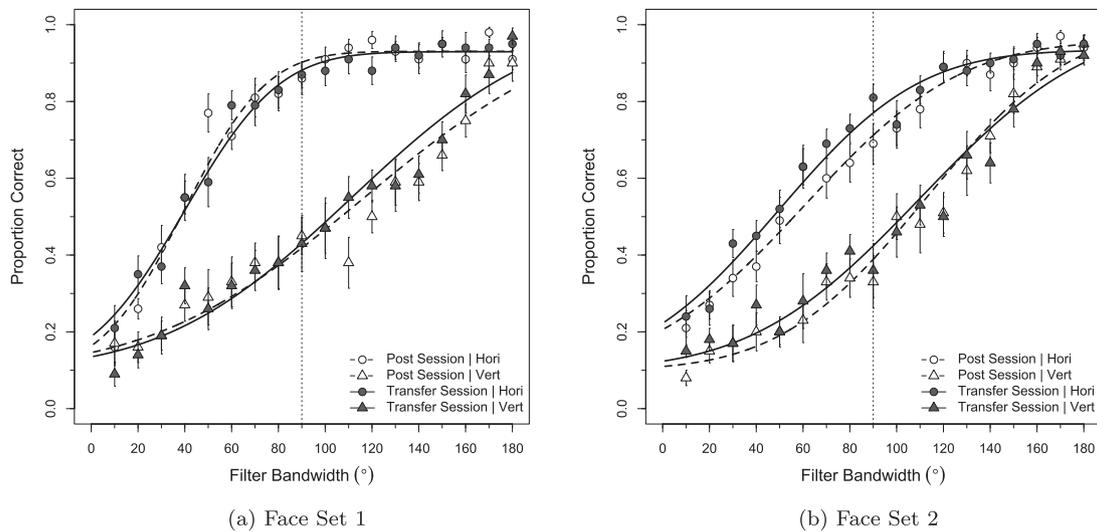


Fig. 11. Proportion correct on the 10-AFC discrimination task after training, assessed during the post-training session (open symbols) and the transfer session 72 h later (closed symbols). The dotted line represents 90° bandwidth, and the error bars represent ± 1 SEM.

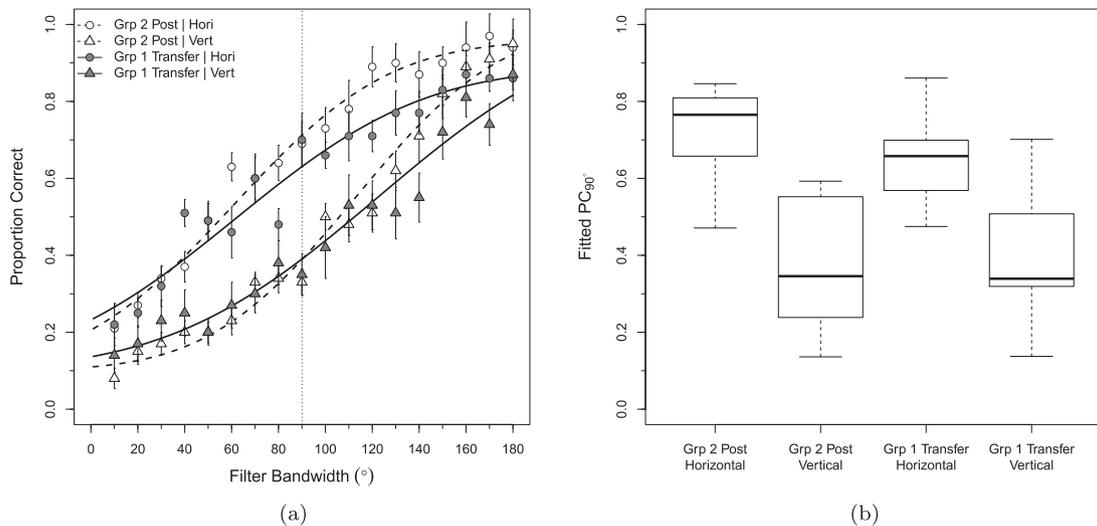


Fig. 12. Proportion correct on the 10-AFC discrimination task with face set 2. Group 1 indicates those observers trained with face set 1 and group 2 indicates those observers trained with face set 2. (a) Post-training assessment for group 2 (replotted from Fig. 8a) plotted with the transfer assessment for group 1, following their training with face set 1. The dotted line represents 90° bandwidth, and all error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

compared the transfer assessment for Group 1 with the post-training assessment for Group 2. This comparison, shown in Fig. 12a, quantifies the effect of training directly with face set 2 to the effect of training with face set 1. Note that both groups had roughly equal experience with the task at the time of this comparison. Visual inspection of Fig. 12a suggests that training with face set 1 produced comparable effects on discrimination of face set 2 to training directly with those stimuli, as overall performance and horizontal selectivity appear similar. To quantify these effects, we extracted proportion correct at 90° from psychometric functions fit to the entire range of bandwidths (see Fig. 12b), and submitted the output to a 2 (group) \times 2 (filter orientation) mixed ANOVA with group as a between-subjects factor. This analysis revealed a significant main effect of filter orientation [$F(1,18) = 57.02, p < 0.0001$] but no main effect of group [$F(1,18) = 0.44, p = 0.5154$] or group \times filter orientation interaction [$F(1,18) = 1.32, p = 0.2648$].

Next, we examined whether the effect of training with face set 2 transferred to face set 1. This comparison, plotted in Fig. 13a, reveals

asymmetrical transfer. To quantify this result, we again submitted proportion correct at 90° (see Fig. 13b) to a 2 (group) \times 2 (filter orientation) mixed ANOVA, which revealed significant main effects of group [$F(1,18) = 59.33, p < 0.0001$] and filter orientation [$F(1,18) = 132.74, p < 0.0001$] but no significant group \times filter orientation interaction [$F(1,18) = 2.28, p = 0.1482$]. Here, the main effect of group demonstrates that training on face set 2 had less of an effect on overall performance than training with face set 1 directly. To determine whether training with face set 2 had any effect on performance with face set 1, we next compared this group to the pre-training assessment for group 1.

Fig. 14a re-plots the pre-training assessment for Group 1 along with the transfer assessment for Group 2, comparing two groups of observers with no training on face set 1. We submitted these data to a 2 (group) \times 2 (filter orientation) mixed ANOVA on proportion correct at 90° (Fig. 14b), which revealed a significant main effect of filter orientation [$F(1,18) = 83.96, p < 0.0001$] but no main effect of group [$F(1,18) = 2.56, p = 0.1269$] or interaction [$F(1,18) = 1.66, p = 0.2146$].

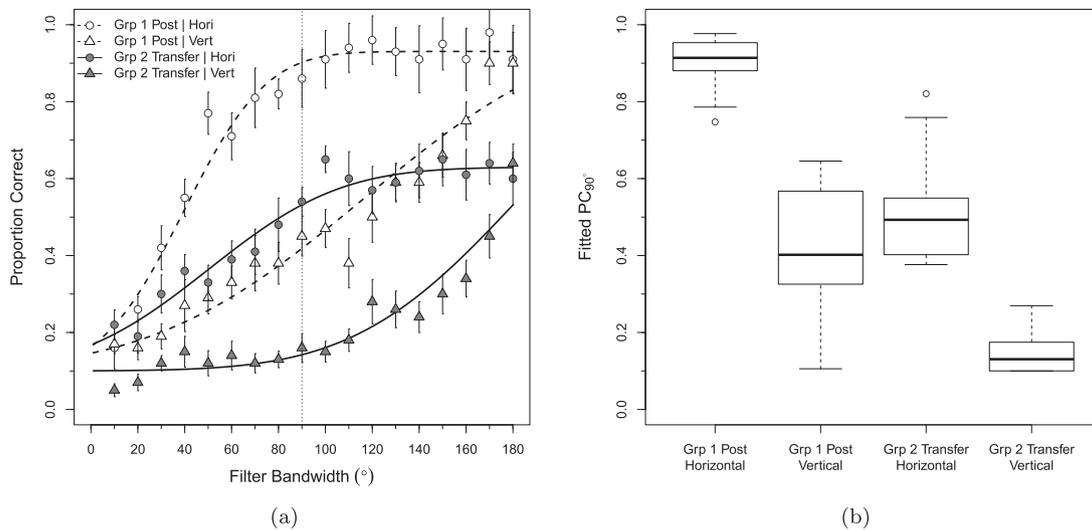


Fig. 13. Proportion correct on the 10-AFC discrimination task with face set 1. Group 1 indicates those observers trained with face set 1 and group 2 indicates those observers trained with face set 2. (a) Post-training assessment for group 1 (replotted from Fig. 7a) plotted with the transfer assessment for group 2, following their training with face set 2. The dotted line represents 90° bandwidth, and all error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

Together, these results demonstrate a failure to transfer overall accuracy (no main effect of group) or horizontal selectivity (no interaction) from face set 2 to face set 1.

Finally, Fig. 15a re-plots the pre-training assessment for Group 2 with the transfer assessment for Group 1, comparing two groups of observers with no training on face set 2. We submitted these data to a 2 (group) × 2 (filter orientation) mixed ANOVA on proportion correct at 90° (Fig. 15b), which revealed significant main effects of group [$F(1,18) = 25.96, p < 0.0001$] and filter orientation [$F(1,18) = 46.90, p < 0.0001$], but no significant interaction [$F(1,18) = 1.43, p = 0.2468$]. Together, these results reveal transfer of overall accuracy (main effect of group), but no transfer of horizontal selectivity (no interaction) from face set 1 to face set 2.

Together, these results demonstrate asymmetrical transfer between face sets. Specifically, improvements resulting from training with face

set 1 transferred to face set 2, but improvements resulting from training with face set 2 did not transfer to set 1.

4. Discussion

In the present study, we found that training with inverted faces (300 trials/session for 3 sessions) improved both discrimination accuracy and selective processing of horizontal structure in the trained stimuli. Further, an ideal observer analysis confirmed that the horizontal band contained more information diagnostic for identity than vertical (see also Pachai, Bennett, & Sekuler, 2018), and that learning increased observers' ability to capitalize on that information difference. Our observed improvements were reflected in both accuracy (Figs. 7 & 8) and efficiency relative to an ideal observer (Fig. 9), demonstrating that training resulted in more efficient sampling of the diagnostic

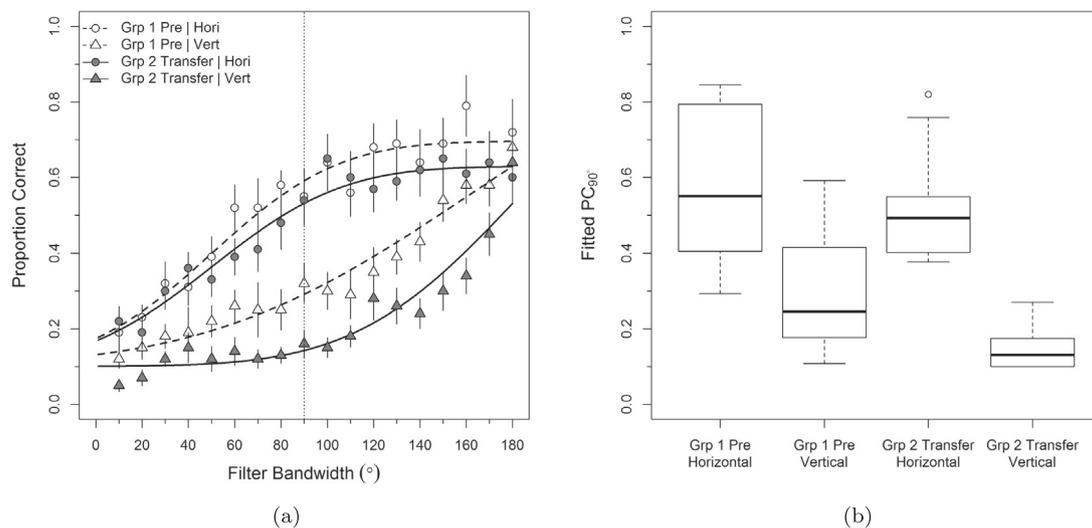


Fig. 14. Proportion correct on the 10-AFC discrimination task with face set 1. Group 1 indicates those observers that would go on to be trained with face set 1 and group 2 indicates those observers trained with face set 2. (a) Pre-training assessment for group 1 (replotted from Fig. 7a) plotted with the transfer assessment for group 2, following their training with face set 2. The dotted line represents 90° bandwidth, and all error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

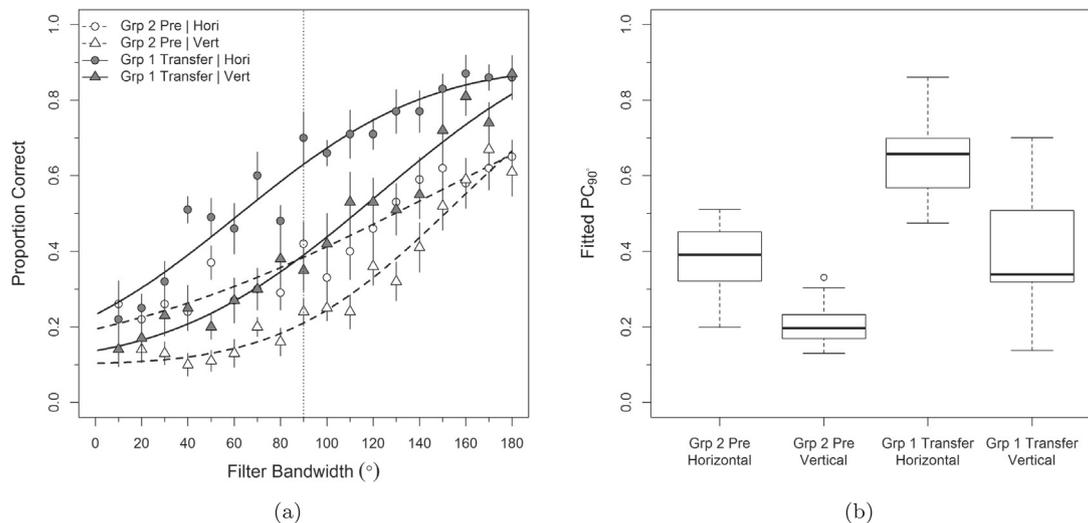


Fig. 15. Proportion correct on the 10-AFC discrimination task with face set 2. Group 1 indicates those observers trained with face set 1 and group 2 indicates those observers trained with face set 2. (a) Pre-training assessment for group 2 (replotted from Fig. 8a) plotted with the transfer assessment for group 1, following their training with face set 1. The dotted line represents 90° bandwidth, and all error bars represent ± 1 SEM. (b) Boxplots of orientation tuning, defined as proportion correct at 90° bandwidth extracted from psychometric functions fit to the entire bandwidth range. The central line represents the median, and the box frame represents the 25th and 75th percentiles, respectively.

information inherent to face stimuli. Importantly, however, the trained improvements in horizontal selectivity were observable particularly with broad bandwidth stimuli (Fig. 10), suggesting that the improvements resulting from our particular training paradigm were not narrowly tuned to the most diagnostic horizontal band. We also observed an improvement in efficiency for vertical structure, suggesting that our training regimen did not exclusively affect the processing of horizontal spatial frequency components.

The absolute levels of efficiency in our task were notably lower than those observed for other visual tasks (Banks, Geisler, & Bennett, 1987; Banks & Bennett, 1988; Banks, Sekuler, & Anderson, 1991; Bennett, Sekuler, & Ozin, 1999; Braje, Tjan, & Legge, 1995; Gold et al., 2004; Pelli & Farell, 1999; Taylor, Bennett, & Sekuler, 2009; Tjan et al., 1995), but were qualitatively similar to the efficiency for 10-AFC face discrimination measured by Pachai et al. (2013b). This low efficiency suggests that sampling of the available stimulus information was highly suboptimal. Such suboptimal sampling is unsurprising, as response classification techniques have clearly established that face discrimination is based on a spatially limited sample of information from the eyes and eyebrows where stimulus differences are most diagnostic, rather than on the broad range of information spread across the entire face including less diagnostic regions (Gold et al., 2004; Gosselin & Schyns, 2001; Haig, 1985; Sekuler et al., 2004; Vinette, Gosselin, & Schyns, 2004). It is important to note that such conclusions hold true, in particular, for stimuli such as ours that are repeated several times and are tightly controlled (i.e., only one viewpoint, constant lighting, cropped into a uniform oval). However, using such stimuli, Sekuler et al. (2004) demonstrated that small differences in spatial sampling are predictive of the face inversion effect. Further, Gold et al. (2004) demonstrated that spatial sampling becomes more efficient with perceptual learning, but that information is selectively sampled from the eye region both before and after training. Finally, recent findings from our lab suggest that observers selectively sample horizontal structure from the eye region during face discrimination (Pachai, Sekuler, & Bennett, 2013a). The current results further elucidate the nature of how stimulus processing improves with training, showing that – even if increased efficiency is constrained to a relatively localized spatial region, as suggested by previous studies – observers improve in their ability to extract diagnostic structure from this region. Although we note again that these effects are based on tightly controlled stimuli with reduced naturalistic

variation, we believe that the evidence to date, including the present result, suggests that face inversion and perceptual training modulate the efficiency with which diagnostic horizontal structure is sampled from the region around the eyes and eyebrows.

Although we observed robust effects of perceptual training, these improvements did not reliably transfer to untrained identities. Specifically, we observed stimulus specificity after training with face set 2, in that trained improvements did not transfer to face set 1, but stimulus generalization after training with face set 1. In general, stimulus specificity has been the norm in perceptual learning research (e.g., Ball & Sekuler, 1987; Fiorentini & Berardi, 1981; Husk, Bennett, & Sekuler, 2007; Hussain et al., 2009b; Hussain, McGraw, Sekuler, & Bennett, 2012; Schoups, Vogels, & Orban, 1995; Yi, Olson, & Chun, 2006). However, recent results have suggested that specificity may depend on the training regimen employed. For example, in tasks where learning is typically retinotopic, significant transfer to minimally-trained locations can be induced by intermittent practice at those locations (Xiao et al., 2008; Zhang et al., 2010). Further, Hussain, Bennett, and Sekuler (2012) showed that increasing stimulus variability across trials leads to greater generalization of learning in a texture identification task. Also, de Heering and Maurer (2013) demonstrated transfer to novel identities after training on face stimuli with high viewpoint variability and Lagusse et al. (2012) demonstrated robust transfer after training on a variety of face-related tasks with a large number of identities. Together these results suggest that training with a large set of identities or viewpoints may be required to induce generalized learning in face-related tasks. In our experiment, only one image of each identity was presented, and these images contained more head and gaze variability in face set 2 (see Fig. 1). Therefore, observers trained with face set 2 could have learned incidental discrimination cues unrelated to identity features that were unavailable in face set 1, leading learning with face set 2 to be more specific. To what extent such factors may have affected our transfer results, and whether more generalized improvements in horizontal selectivity can indeed be induced, remain open questions for future research.

It is important to note that while designing this experiment we ensured observers never were cued to the diagnostic band of orientation information. During the training sessions observers viewed only unfiltered (full-face) stimuli and during the assessment sessions we embedded the diagnostic orientation band in a full-face context. However,

it remains unclear whether other training regimens would lead to even greater enhancements of sensitivity to horizontal structure in face stimuli. For example, paradigms in which observers are overtly trained to utilize horizontal structure might produce more learning. This is a worthwhile avenue for future research, as studies aimed at maximizing training-based improvements in face identification have important implications for amelioration of the face identification deficits experienced by a number of populations, including older observers (Grady, 2002; Konar, Bennett, & Sekuler, 2013; Maylor & Valentine, 1992; Obermeyer, Kolling, Schaich, & Knopf, 2012; Rousselet et al., 2009) and individuals with prosopagnosia (Barton, Radcliffe, Cherkasova, & Edelman, 2007; Barton, 2008; Behrmann, Avidan, Marotta, & Kimchi, 2005; Busigny & Rossion, 2010; Davies-Thompson et al., 2017; Duchaine & Nakayama, 2006, but see DeGutis, Chiu, Grosso, & Cohan, 2014 for evidence that training may have limited effects on acquired prosopagnosia), schizophrenia (Archer, Hay, & Young, 1992; Christensen, Spencer, King, Sekuler, & Bennett, 2013; Williams, Loughland, Gordon, & Davidson, 1999), or autism (Barton et al., 2007; Jiang et al., 2013; Langdell, 1978; Nagai et al., 2013; Rutherford, Clements, & Sekuler, 2007).

5. Conclusion

We trained observers to discriminate intact, inverted faces, measuring discrimination accuracy and selective processing of horizontal structure before and after training. This study was motivated by demonstrations that horizontal structure is a highly diagnostic cue to face identity (Dakin & Watt, 2009), and that differences in the selective processing of this structure are associated with the magnitude of the FIE (Pachai et al., 2013b). Our results further support the role of horizontal structure in accurate face discrimination, as trained improvements in face discrimination accuracy were particularly associated with improvements for this orientation band. Although these improvements did not reliably transfer to untrained identities, further studies using modified training regimens may extend the present results to produce long-lasting, generalized improvements in face discrimination performance.

Acknowledgements

We thank Donna Waxman for her assistance with data collection and two anonymous reviewers for their helpful suggestions on a previous version of this manuscript. Portions of this work appeared previously at the Annual Meeting of the Vision Sciences Society and were completed in partial fulfillment of the requirements for a PhD at McMaster University. Financial support for this work was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) and Canada Research Chairs program.

References

Archer, J., Hay, D. C., & Young, A. W. (1992). Face processing in psychiatric conditions. *British Journal of Clinical Psychology, 31*(1), 45–61.

Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research, 27*(6), 953–965.

Banks, M. S., & Bennett, P. J. (1988). Optical and photoreceptor immaturities limit the spatial and chromatic vision of human neonates. *Journal of the Optical Society of America, A, Optics and Image Science, 5*(12), 2059–2079.

Banks, M. S., Geisler, W. S., & Bennett, P. J. (1987). The physical limits of grating visibility. *Vision Research, 27*(11), 1915–1924.

Banks, M. S., Sekuler, A. B., & Anderson, S. J. (1991). Peripheral spatial vision: Limits imposed by optics, photoreceptors, and receptor pooling. *Journal of the Optical Society of America, A, Optics and Image Science, 8*, 1775–1787.

Barton, J. J. S. (2008). Structure and function in acquired prosopagnosia: Lessons from a series of 10 patients with brain damage. *Journal of Neuropsychology, 21*(1), 197–225.

Barton, J. J. S., Radcliffe, N., Cherkasova, M. V., & Edelman, J. A. (2007). Scan patterns during the processing of facial identity in prosopagnosia. *Experimental Brain Research, 181*(2), 199–211.

Behrmann, M., Avidan, G., Marotta, J. J., & Kimchi, R. (2005). Detailed exploration of face-related processing in congenital prosopagnosia: 1. Behavioral findings. *Journal of*

Cognitive Neuroscience, 17(7), 1130–1149.

Bennett, P. J., & Banks, M. S. (1987). Sensitivity loss in odd-symmetric mechanisms and phase anomalies in peripheral vision. *Nature, 326*, 873–876.

Bennett, P. J., Sekuler, A. B., & Ozin, L. (1999). Effects of aging on calculation efficiency and equivalent noise. *Journal of the Optical Society of America, A, Optics, Image Science, and Vision, 16*(3), 654–668.

Boutet, I., Collin, C. A., & Faubert, J. (2003). Configural face encoding and spatial frequency information. *Perception & Psychophysics, 65*(7), 1078–1093.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*(4), 433–436.

Braje, W. L., Tjan, B. S., & Legge, G. E. (1995). Human efficiency for recognizing and detecting low-pass filtered objects. *Vision Research, 35*(21), 2955–2966.

Brigham, J., & Barkowitz, P. (1978). Do “they all look alike?” The Effect of Race, Sex, Experience, and Attitudes on the Ability to Recognize Faces. *Journal of Applied Social Psychology, 8*(4), 306–318.

Busigny, T., & Rossion, B. (2010). Acquired prosopagnosia abolishes the face inversion effect. *Cortex, 46*(8), 965–981.

Christensen, B. K., Spencer, J. M. Y., King, J. P., Sekuler, A. B., & Bennett, P. J. (2013). Noise as a mechanism of anomalous face processing among persons with Schizophrenia. *Frontiers in Psychology, 4*, 1–10.

Dakin, S. C., & Watt, R. J. (2009). Biological bar codes in human faces. *Journal of Vision, 9*(4), 1–10.

Davies, G., Ellis, H., & Shepherd, J. (1977). Cue saliency in faces as assessed by the ‘photofit’ technique. *Perception, 6*, 263–269.

Davies-Thompson, J., Fletcher, K., Hills, C., Pancaroglu, R., Corrow, S. L., & Barton, J. J. S. (2017). Perceptual learning of faces: A rehabilitative study of acquired prosopagnosia. *Journal of Cognitive Neuroscience, 29*(3), 573–591.

de Heering, A., & Maurer, D. (2013). The effect of spatial frequency on perceptual learning of inverted faces. *Vision Research, 86*, 107–114.

de Heering, A., Rossion, B., & Maurer, D. (2012). Developmental changes in face recognition during childhood: Evidence from upright and inverted faces. *Cognitive Development, 27*(1), 17–27.

DeGutis, J. M., Chiu, C., Grosso, M. E., & Cohan, S. (2014). Face processing improvements in prosopagnosia: Successes and failures over the last 50 years. *Frontiers in Human Neuroscience, 8*(August), 1–14.

Duchaine, B. C., & Nakayama, K. (2006). Developmental prosopagnosia: A window to content-specific face processing. *Current Opinion in Neurobiology, 16*, 166–173.

Fiorienti, A., & Berardi, N. (1981). Learning in grating waveform discrimination: Specificity for orientation and spatial frequency. *Vision Research, 21*, 1149–1158.

Fulton, A., & Bartlett, J. C. (1991). Young and old faces in young and old heads: The factor of age in face recognition. *Psychology and Aging, 6*(4), 623–630.

Gaspar, C. M., Bennett, P. J., & Sekuler, A. B. (2008a). The effects of face inversion and contrast-reversal on efficiency and internal noise. *Vision Research, 48*(8), 1084–1095.

Gaspar, C. M., Sekuler, A. B., & Bennett, P. J. (2008b). Spatial frequency tuning of upright and inverted face identification. *Vision Research, 48*(28), 2817–2826.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review, 96*(2), 267–314.

Geisler, W. S. (2011). Contributions of ideal observer theory to vision research. *Vision Research, 51*(7), 771–781.

Germine, L. T., Duchaine, B., & Nakayama, K. (2011). Where cognitive development and aging meet: Face learning ability peaks after age 30. *Cognition, 118*, 201–210.

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

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

Gold, J. M., Bennett, P. J., & Sekuler, A. B. (1999a). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research, 39*(21), 3537–3560.

Gold, J. M., Bennett, P. J., & Sekuler, A. B. (1999b). Signal but not noise changes with perceptual learning. *Nature, 402*, 176–178.

Gold, J. M., Sekuler, A. B., & Bennett, P. J. (2004). Characterizing perceptual learning with external noise. *Cognitive Science, 28*(2), 167–207.

Gosselin, F., & Schyns, P. G. (2001). Bubbles: A technique to reveal the use of information in recognition tasks. *Vision Research, 41*(17), 2261–2271.

Grady, C. L. (2002). Age-related differences in face processing: a meta-analysis of three functional neuroimaging experiments. *Canadian Journal of Experimental Psychology, 56*(3), 208–220.

Haig, N. D. (1985). How faces differ—a new comparative technique. *Perception, 14*(5), 601–615.

Haig, N. D. (1986). Exploring recognition with interchanged facial features. *Perception, 15*(3), 235–247.

Husk, J. S., Bennett, P. J., & Sekuler, A. B. (2007). Inverting houses and textures: Investigating the characteristics of learned inversion effects. *Vision Research, 47*(27), 3350–3359.

Hussain, Z., Bennett, P. J., & Sekuler, A. B. (2012). Versatile perceptual learning of textures after variable exposures. *Vision Research, 61*, 89–94.

Hussain, Z., McGraw, P. V., Sekuler, A. B., & Bennett, P. J. (2012). The rapid emergence of stimulus specific perceptual learning. *Frontiers in psychology, 3*(July), 226.

Hussain, Z., Sekuler, A. B., & Bennett, P. J. (2009a). How much practice is needed to produce perceptual learning? *Vision Research, 49*(21), 2624–2634.

Hussain, Z., Sekuler, A. B., & Bennett, P. J. (2009b). Perceptual learning modifies inversion effects for faces and textures. *Vision Research, 49*(18), 2273–2284.

Hussain, Z., Sekuler, A. B., & Bennett, P. J. (2011). Superior identification of familiar visual patterns a year after learning. *Psychological Science, 22*(6), 724–730.

Huynh, C. M., & Balas, B. (2014). Emotion recognition (sometimes) depends on horizontal orientations. *Attention, Perception & Psychophysics*.

Jiang, X., Bollich, A., Cox, P., Hyder, E., James, J., Gowani, S. A., ... Hadjikhani, N. (2013). A quantitative link between face discrimination deficits and neuronal

- selectivity for faces in autism. *NeuroImage: Clinical*, 2(1), 320–331.
- Kersten, D. (1987). Statistical efficiency for the detection of visual noise. *Vision Research*, 27(6), 1029–1040.
- Konar, Y., Bennett, P. J., & Sekuler, A. B. (2013). Effects of aging on face identification and holistic face processing. *Vision Research*, 88, 38–46.
- Laguerre, R., Dormal, G., Biervoeye, A., Kuefner, D., & Rossion, B. (2012). Extensive visual training in adulthood significantly reduces the face inversion effect. *Journal of Vision*, 12, 1–13.
- Langdell, T. (1978). Recognition of faces: An approach to the study of autism. *Journal of Child Psychology and Psychiatry*, 19(3), 255–268.
- LeGrand, R., Mondloch, C. J., Maurer, D., & Brent, H. P. (2001). Early visual experience and face processing. *Science*, 410, 890.
- Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: a user's guide* (2nd ed.). Psychology Press.
- Malpass, R., & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology*, 13(4), 330–334.
- Maylor, E., & Valentine, T. (1992). Linear and nonlinear effects of aging on categorizing and naming faces. *Psychology and Aging*, 7(2), 317–323.
- Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law*, 7(1), 3–35.
- Nagai, M., Bennett, P. J., Rutherford, M. D., Gaspar, C. M., Kumada, T., & Sekuler, A. B. (2013). Comparing face processing strategies between typically-developed observers and observers with autism using sub-sampled-pixels presentation in response classification technique. *Vision Research*, 79, 27–35.
- Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, 39(23), 3824–3833.
- Obermeyer, S., Kolling, T., Schaich, A., & Knopf, M. (2012). Differences between old and young adults' ability to recognize human faces underlie processing of horizontal information. *Frontiers in Aging Neuroscience*, 4(April), 3.
- Pachai, M. V., Bennett, P. J., & Sekuler, A. B. (2017). The Bandwidth of Diagnostic Horizontal Structure for Face Identification. *Perception* 030100661875447.
- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013a). Masking of individual facial features reveals the use of horizontal structure in the eyes. *Journal of Vision*, 13(9), 411.
- 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(February), 1–9.
- Pachai, M. V., Sekuler, A. B., Bennett, P. J., Schyns, P. G., & Ramon, M. (2017). Personal familiarity enhances sensitivity to horizontal structure during processing of face identity. *Journal of Vision*, 17(6), 5.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pelli, D. G., & Farell, B. (1999). Why use noise? *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, 16(3), 647–653.
- Peterson, M. F., & Eckstein, M. P. (2012). Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences*, 109(48).
- R Core Team. (2017). R: A Language and Environment for Statistical Computing.
- Rhodes, M. G., & Anastasi, J. S. (2012). The own-age bias in face recognition: A meta-analytic and theoretical review. *Psychological Bulletin*, 138(1), 146–174.
- Rousselet, G. A., Husk, J. S., Pernet, C. R., Gaspar, C. M., Bennett, P. J., & Sekuler, A. B. (2009). Age-related delay in information accrual for faces: evidence from a parametric, single-trial EEG approach. *BMC Neuroscience*, 10, 114.
- Rutherford, M. D., Clements, K. A., & Sekuler, A. B. (2007). Differences in discrimination of eye and mouth displacement in autism spectrum disorders. *Vision Research*, 47(15), 2099–2110.
- Schoups, A. A., Vogels, R., & Orban, G. A. (1995). Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity and monocularity. *The Journal of Physiology*, 483, 797–810.
- Sekuler, A. B., Gaspar, C. M., Gold, J. M., & Bennett, P. J. (2004). Inversion leads to quantitative, not qualitative, changes in face processing. *Current Biology*, 14(5), 391–396.
- Susilo, T., Germine, L., & Duchaine, B. (2013). Face recognition ability matures late: Evidence from individual differences in young adults. *Journal of Experimental Psychology: Human Perception and Performance*, 39(5), 1212–1217.
- Tanner, W. P., & Birdsall, T. G. (1958). Definitions of d' and eta as psychophysical measures. *The Journal of the Acoustical Society of America*, 30(10), 922–928.
- Taylor, C. P., Bennett, P. J., & Sekuler, A. B. (2009). Spatial frequency summation in visual noise. *Journal of the Optical Society of America. A, Optics and Image Science and Vision*, 26(11), B84–B93.
- Tjan, B. S., Braje, W. L., Legge, G. E., & Kersten, D. (1995). Human efficiency for recognizing 3-D objects in luminance noise. *Vision Research*, 35(21), 3053–3069.
- Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, 79(4), 471–491.
- Vinette, C., Gosselin, F., & Schyns, P. G. (2004). Spatio-temporal dynamics of face recognition in a flash: It's in the eyes. *Cognitive Science*, 28(2), 289–301.
- Watier, N. N., Collin, C. A., & Boutet, I. (2010). Spatial-frequency thresholds for configural and featural discriminations in upright and inverted faces. *Perception*, 39(4), 502–513.
- Willenbockel, V., Fiset, D., Chauvin, A., Blais, C., Arguin, M., Tanaka, J. W., ... Gosselin, F. (2010). Does face inversion change spatial frequency tuning? *Journal of Experimental Psychology: Human Perception and Performance*, 36(1), 122–135.
- Williams, L. M., Loughland, C. M., Gordon, E., & Davidson, D. (1999). Visual scanpaths in schizophrenia: Is there a deficit in face recognition? *Schizophrenia Research*, 40(3), 189–199.
- Xiao, L.-Q., Zhang, J.-Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, 18(24), 1922–1926.
- Yi, D.-J., Olson, I. R., & Chun, M. M. (2006). Shape-specific perceptual learning in a figure-ground segregation task. *Vision Research*, 46, 914–924.
- Yin, R. K. (1969). Looking at Upside-Down Faces. *Journal of Experimental Psychology*, 81(1), 141–145.
- Zhang, J.-Y., Zhang, G.-L., Xiao, L.-Q., Klein, S. A., Levi, D. M., & Yu, C. (2010). Rule-based learning explains visual perceptual learning and its specificity and transfer. *The Journal of Neuroscience*, 30(37), 12323–12328.